

PDF Reference

third edition

Adobe Portable Document Format
Version 1.4

Adobe Systems Incorporated



ADDISON-WESLEY

Boston • San Francisco • New York • Toronto • Montreal
London • Munich • Paris • Madrid
Capetown • Sydney • Tokyo • Singapore • Mexico City

Library of Congress Cataloging-in-Publication Data

PDF reference : Adobe portable document format version 1.4 / Adobe Systems Incorporated. — 3rd ed.

p. cm.

Includes bibliographical references and index.

ISBN 0-201-75839-3 (alk. paper)

1. Adobe Acrobat. 2. Portable document software. I. Adobe Systems.

QA76.76.T49 P38 2001

005.7'2—dc21

2001053899

© 1985–2001 Adobe Systems Incorporated. All rights reserved.

NOTICE: All information contained herein is the property of Adobe Systems Incorporated.

No part of this publication (whether in hardcopy or electronic form) may be reproduced or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior written consent of the publisher.

Adobe, the Adobe logo, Acrobat, the Acrobat logo, Acrobat Capture, Acrobat Reader, Adobe Garamond, Aldus, Distiller, ePaper, Extreme, FrameMaker, Illustrator, InDesign, Minion, Myriad, PageMaker, Photoshop, Poetica, PostScript, and XMP are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States and/or other countries. Microsoft and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Apple, Mac, Mac OS, Macintosh, QuickDraw, and TrueType are trademarks of Apple Computer, Inc., registered in the United States and other countries. KanjiTalk is a trademark of Apple Computer, Inc. UNIX is a registered trademark of The Open Group. Unicode is a registered trademark of Unicode, Inc. Java is a trademark of Sun Microsystems, Inc. JavaScript is a registered trademark of Netscape Communications Corporation. QuarkXPress is a trademark of Quark, Inc. and/or certain of the Quark Affiliated Companies, Reg. U.S. Pat. & Tm. Off. and in many other countries. PANTONE is a registered trademark and Hexachrome is a trademark of Pantone, Inc. PANOSE is a trademark of Hewlett-Packard Company. OEB is a trademark of the Open eBook Forum. Helvetica and Times are registered trademarks of Linotype-Hell AG and/or its subsidiaries. Arial and Times New Roman are trademarks of The Monotype Corporation registered in the U.S. Patent and Trademark Office and may be registered in certain other jurisdictions. ITC Zapf Dingbats is a registered trademark of International Typeface Corporation. Ryumin Light is a trademark of Morisawa & Co., Ltd. All other trademarks are the property of their respective owners.

All instances of the name PostScript in the text are references to the PostScript language as defined by Adobe Systems Incorporated unless otherwise stated. The name PostScript also is used as a product trademark for Adobe Systems' implementation of the PostScript language interpreter. Except as otherwise stated, any mention of a "PostScript output device," "PostScript printer," "PostScript software," or similar item refers to a product that contains PostScript technology created or licensed by Adobe Systems Incorporated, not to one that purports to be merely compatible.

This publication and the information herein are furnished AS IS, are subject to change without notice, and should not be construed as a commitment by Adobe Systems Incorporated. Adobe Systems Incorporated assumes no responsibility or liability for any errors or inaccuracies, makes no warranty of any kind (express, implied, or statutory) with respect to this publication, and expressly disclaims any and all warranties of merchantability, fitness for particular purposes, and noninfringement of third-party rights.

1 2 3 4 5 6 7 8 9 CRS 04030201

First printing, December 2001

Contents

Preface xix

Chapter 1: Introduction 1

- 1.1 About This Book 1
- 1.2 Introduction to PDF 1.4 Features 4
- 1.3 Related Publications 5
- 1.4 Intellectual Property 6

Chapter 2: Overview 9

- 2.1 Imaging Model 10
- 2.2 Other General Properties 14
- 2.3 Using PDF 19
- 2.4 PDF and the PostScript Language 21

Chapter 3: Syntax 23

- 3.1 Lexical Conventions 24
- 3.2 Objects 27
- 3.3 Filters 41
- 3.4 File Structure 61
- 3.5 Encryption 71
- 3.6 Document Structure 81
- 3.7 Content Streams and Resources 92
- 3.8 Common Data Structures 98
- 3.9 Functions 106
- 3.10 File Specifications 118

Chapter 4: Graphics 131

- 4.1 Graphics Objects 132
- 4.2 Coordinate Systems 136
- 4.3 Graphics State 147
- 4.4 Path Construction and Painting 161
- 4.5 Color Spaces 172
- 4.6 Patterns 219
- 4.7 External Objects 261
- 4.8 Images 262
- 4.9 Form XObjects 281
- 4.10 PostScript XObjects 289

Chapter 5: Text 291

- 5.1 Organization and Use of Fonts 292
- 5.2 Text State Parameters and Operators 300
- 5.3 Text Objects 308
- 5.4 Introduction to Font Data Structures 314
- 5.5 Simple Fonts 316
- 5.6 Composite Fonts 334
- 5.7 Font Descriptors 355
- 5.8 Embedded Font Programs 364
- 5.9 ToUnicode CMaps 368

Chapter 6: Rendering 373

- 6.1 CIE-Based Color to Device Color 374
- 6.2 Conversions among Device Color Spaces 376
- 6.3 Transfer Functions 380
- 6.4 Halftones 382
- 6.5 Scan Conversion Details 403

Chapter 7: Transparency 409

- 7.1 Overview of Transparency 410
- 7.2 Basic Compositing Computations 412
- 7.3 Transparency Groups 425
- 7.4 Soft Masks 439
- 7.5 Specifying Transparency in PDF 441
- 7.6 Color Space and Rendering Issues 454

Chapter 8: Interactive Features 471

- 8.1 Viewer Preferences 471
- 8.2 Document-Level Navigation 474
- 8.3 Page-Level Navigation 481
- 8.4 Annotations 488
- 8.5 Actions 513
- 8.6 Interactive Forms 528
- 8.7 Sounds 568
- 8.8 Movies 570

Chapter 9: Document Interchange 573

- 9.1 Procedure Sets 574
- 9.2 Metadata 575
- 9.3 File Identifiers 580
- 9.4 Page-Piece Dictionaries 581
- 9.5 Marked Content 583
- 9.6 Logical Structure 588
- 9.7 Tagged PDF 612
- 9.8 Accessibility Support 651
- 9.9 Web Capture 659
- 9.10 Prepress Support 676

Appendix A: Operator Summary 699**Appendix B: Operators in Type 4 Functions** 703

- B.1 Arithmetic Operators 703
- B.2 Relational, Boolean, and Bitwise Operators 704
- B.3 Conditional Operators 704
- B.4 Stack Operators 704

Appendix C: Implementation Limits 705

- C.1 General Implementation Limits 706
- C.2 Implementation Limits Affecting Web Capture 708

Appendix D: Character Sets and Encodings 709

- D.1 Latin Character Set and Encodings 711
- D.2 Expert Set and MacExpertEncoding 715
- D.3 Symbol Set and Encoding 718
- D.4 ZapfDingbats Set and Encoding 721

Appendix E: PDF Name Registry 723**Appendix F: Linearized PDF** 725

- F.1 Background and Assumptions 726
- F.2 Linearized PDF Document Structure 728
- F.3 Hint Tables 741
- F.4 Access Strategies 751

Appendix G: Example PDF Files 757

- G.1 Minimal PDF File 757
- G.2 Simple Text String Example 760
- G.3 Simple Graphics Example 762
- G.4 Page Tree Example 765
- G.5 Outline Tree Example 770
- G.6 Updating Example 774

Appendix H: Compatibility and Implementation Notes 783

- H.1 PDF Version Numbers 783
- H.2 Feature Compatibility 786
- H.3 Implementation Notes 787

Bibliography 811**Index** 817

Figures

- 2.1 Creating PDF files using PDF Writer 20
- 2.2 Creating PDF files using Acrobat Distiller 21

- 3.1 PDF components 24
- 3.2 Initial structure of a PDF file 62
- 3.3 Structure of an updated PDF file 70
- 3.4 Structure of a PDF document 82
- 3.5 Inheritance of attributes 92
- 3.6 Mapping with the **Decode** array 112

- 4.1 Graphics objects 135
- 4.2 Device space 138
- 4.3 User space 140
- 4.4 Relationships among coordinate systems 142
- 4.5 Effects of coordinate transformations 143
- 4.6 Effect of transformation order 144
- 4.7 Miter length 154
- 4.8 Cubic Bézier curve generated by the **c** operator 165
- 4.9 Cubic Bézier curves generated by the **v** and **y** operators 166
- 4.10 Nonzero winding number rule 170
- 4.11 Even-odd rule 171
- 4.12 Color specification 174
- 4.13 Color rendering 175
- 4.14 Component transformations in a CIE-based *ABC* color space 182
- 4.15 Component transformations in a CIE-based *A* color space 183
- 4.16 Starting a new triangle in a free-form Gouraud-shaded triangle mesh 245
- 4.17 Connecting triangles in a free-form Gouraud-shaded triangle mesh 246
- 4.18 Varying the value of the edge flag to create different shapes 247
- 4.19 Lattice-form triangle meshes 248
- 4.20 Coordinate mapping from a unit square to a four-sided Coons patch 251
- 4.21 Painted area and boundary of a Coons patch 252
- 4.22 Color values and edge flags in Coons patch meshes 254
- 4.23 Edge connections in a Coons patch mesh 255
- 4.24 Control points in a tensor-product patch 257
- 4.25 Typical sampled image 262
- 4.26 Source image coordinate system 265
- 4.27 Mapping the source image 266

5.1	Glyphs painted in 50% gray	295
5.2	Glyph outlines treated as a stroked path	296
5.3	Graphics clipped by a glyph path	297
5.4	Glyph metrics	298
5.5	Metrics for horizontal and vertical writing modes	300
5.6	Character spacing in horizontal writing	303
5.7	Word spacing in horizontal writing	303
5.8	Horizontal scaling	304
5.9	Leading	304
5.10	Text rise	307
5.11	Operation of the TJ operator in horizontal writing	312
5.12	Output from Example 5.9	327
5.13	Characteristics represented in the Flags entry of a font descriptor	359
6.1	Various halftoning effects	389
6.2	Halftone cell with a nonzero angle	396
6.3	Angled halftone cell divided into two squares	396
6.4	Halftone cell and two squares tiled across device space	397
6.5	Tiling of device space in a type 16 halftone	399
6.6	Flatness tolerance	404
6.7	Rasterization without stroke adjustment	407
8.1	Presentation timing	487
8.2	Open annotation	489
8.3	Coordinate adjustment with the NoRotate flag	494
8.4	Square and circle annotations	505
8.5	QuadPoints specification	506
8.6	FDF file structure	559
9.1	Simple Web Capture file structure	662
9.2	Complex Web Capture file structure	663
9.3	Page boundaries	678
9.4	Trapping example	688
G.1	Output of Example G.3	763
G.2	Page tree for Example G.4	765
G.3	Document outline as displayed in Example G.5	770
G.4	Document outline as displayed in Example G.6	772

- Plate 1** Additive and subtractive color
- Plate 2** Uncalibrated color
- Plate 3** Lab color space
- Plate 4** Color gamuts
- Plate 5** Rendering intents
- Plate 6** Duotone image
- Plate 7** Quadtone image
- Plate 8** Colored tiling pattern
- Plate 9** Uncolored tiling pattern
- Plate 10** Axial shading
- Plate 11** Radial shadings depicting a cone
- Plate 12** Radial shadings depicting a sphere
- Plate 13** Radial shadings with extension
- Plate 14** Radial shading effect
- Plate 15** Coons patch mesh
- Plate 16** Transparency groups
- Plate 17** Isolated and knockout groups
- Plate 18** RGB blend modes
- Plate 19** CMYK blend modes
- Plate 20** Blending and overprinting

Tables

3.1	White-space characters	26
3.2	Escape sequences in literal strings	30
3.3	Examples of literal names using the # character	33
3.4	Entries common to all stream dictionaries	38
3.5	Standard filters	42
3.6	Typical LZW encoding sequence	48
3.7	Optional parameters for LZWDecode and FlateDecode filters	49
3.8	Predictor values	51
3.9	Optional parameters for the CCITTFaxDecode filter	53
3.10	Optional parameter for the JBIG2Decode filter	57
3.11	Optional parameter for the DCTDecode filter	60
3.12	Entries in the file trailer dictionary	68
3.13	Entries common to all encryption dictionaries	72
3.14	Additional encryption dictionary entries for the standard security handler	76
3.15	User access permissions	77
3.16	Entries in the catalog dictionary	83
3.17	Required entries in a page tree node	86
3.18	Entries in a page object	88
3.19	Entries in the name dictionary	93
3.20	Compatibility operators	95
3.21	Entries in a resource dictionary	97
3.22	PDF data types	99
3.23	Entries in a name tree node dictionary	102
3.24	Example of a name tree	103
3.25	Entries in a number tree node dictionary	106
3.26	Entries common to all function dictionaries	108
3.27	Additional entries specific to a type 0 function dictionary	110
3.28	Additional entries specific to a type 2 function dictionary	113
3.29	Additional entries specific to a type 3 function dictionary	114
3.30	Operators in type 4 functions	116
3.31	Examples of file specifications	121
3.32	Entries in a file specification dictionary	122
3.33	Additional entries in an embedded file stream dictionary	124
3.34	Entries in an embedded file parameter dictionary	125
3.35	Entries in a Mac OS file information dictionary	125

4.1	Operator categories	134
4.2	Device-independent graphics state parameters	148
4.3	Device-dependent graphics state parameters	150
4.4	Line cap styles	153
4.5	Line join styles	154
4.6	Examples of line dash patterns	155
4.7	Graphics state operators	156
4.8	Entries in a graphics state parameter dictionary	157
4.9	Path construction operators	163
4.10	Path-painting operators	167
4.11	Clipping path operators	172
4.12	Color space families	176
4.13	Entries in a CalGray color space dictionary	183
4.14	Entries in a CalRGB color space dictionary	185
4.15	Entries in a Lab color space dictionary	188
4.16	Additional entries specific to an ICC profile stream dictionary	190
4.17	ICC profile types	191
4.18	Ranges for typical ICC color spaces	192
4.19	Rendering intents	198
4.20	Entry in a DeviceN color space attributes dictionary	207
4.21	Color operators	216
4.22	Additional entries specific to a type 1 pattern dictionary	221
4.23	Entries in a type 2 pattern dictionary	231
4.24	Shading operator	232
4.25	Entries common to all shading dictionaries	234
4.26	Additional entries specific to a type 1 shading dictionary	237
4.27	Additional entries specific to a type 2 shading dictionary	238
4.28	Additional entries specific to a type 3 shading dictionary	240
4.29	Additional entries specific to a type 4 shading dictionary	244
4.30	Additional entries specific to a type 5 shading dictionary	249
4.31	Additional entries specific to a type 6 shading dictionary	253
4.32	Data values in a Coons patch mesh	256
4.33	Data values in a tensor-product patch mesh	260
4.34	XObject operator	261
4.35	Additional entries specific to an image dictionary	267
4.36	Default Decode arrays	272
4.37	Entries in an alternate image dictionary	274
4.38	Inline image operators	278
4.39	Entries in an inline image object	279
4.40	Additional abbreviations in an inline image object	280
4.41	Additional entries specific to a type 1 form dictionary	284
4.42	Entries common to all group attributes dictionaries	287
4.43	Entries in a reference dictionary	288
4.44	Additional entries specific to a PostScript XObject dictionary	290

5.1	Text state parameters	301
5.2	Text state operators	302
5.3	Text rendering modes	306
5.4	Text object operators	308
5.5	Text-positioning operators	310
5.6	Text-showing operators	311
5.7	Font types	315
5.8	Entries in a Type 1 font dictionary	317
5.9	Entries in a Type 3 font dictionary	323
5.10	Type 3 font operators	326
5.11	Entries in an encoding dictionary	330
5.12	Entries in a CIDSystemInfo dictionary	337
5.13	Entries in a CIDFont dictionary	338
5.14	Predefined CJK CMap names	343
5.15	Character collections for predefined CMaps, by PDF version	346
5.16	Additional entries in a CMap dictionary	349
5.17	Entries in a Type 0 font dictionary	353
5.18	Entries common to all font descriptors	356
5.19	Font flags	358
5.20	Additional font descriptor entries for CIDFonts	361
5.21	Character classes in CJK fonts	362
5.22	Embedded font organization for various font types	365
5.23	Additional entries in an embedded font stream dictionary	366
6.1	Predefined spot functions	385
6.2	PDF halftone types	391
6.3	Entries in a type 1 halftone dictionary	393
6.4	Additional entries specific to a type 6 halftone dictionary	395
6.5	Additional entries specific to a type 10 halftone dictionary	398
6.6	Additional entries specific to a type 16 halftone dictionary	400
6.7	Entries in a type 5 halftone dictionary	401
7.1	Variables used in the basic compositing formula	414
7.2	Standard separable blend modes	417
7.3	Standard nonseparable blend modes	419
7.4	Variables used in the source shape and opacity formulas	422
7.5	Variables used in the result shape and opacity formulas	423
7.6	Revised variables for the basic compositing formulas	427
7.7	Arguments and results of the group compositing function	429
7.8	Variables used in the group compositing formulas	431
7.9	Variables used in the page group compositing formulas	437
7.10	Entries in a soft-mask dictionary	446
7.11	Restrictions on the entries in a soft-mask image dictionary	448

- 7.12** Additional entry in a soft-mask image dictionary 449
- 7.13** Additional entries specific to a transparency group attributes dictionary 450
- 7.14** Overprinting behavior in the opaque imaging model 464
- 7.15** Overprinting behavior in the transparent imaging model 464

- 8.1** Entries in a viewer preferences dictionary 472
- 8.2** Destination syntax 475
- 8.3** Entries in the outline dictionary 478
- 8.4** Entries in an outline item dictionary 478
- 8.5** Outline item flags 479
- 8.6** Entries in a page label dictionary 483
- 8.7** Entries in a thread dictionary 484
- 8.8** Entries in a bead dictionary 484
- 8.9** Entries in a transition dictionary 486
- 8.10** Entries common to all annotation dictionaries 490
- 8.11** Annotation flags 493
- 8.12** Entries in a border style dictionary 495
- 8.13** Entries in an appearance dictionary 497
- 8.14** Annotation types 499
- 8.15** Additional entries specific to a text annotation 500
- 8.16** Additional entries specific to a link annotation 501
- 8.17** Additional entries specific to a free text annotation 502
- 8.18** Additional entries specific to a line annotation 503
- 8.19** Line ending styles 504
- 8.20** Additional entries specific to a square or circle annotation 505
- 8.21** Additional entries specific to markup annotations 506
- 8.22** Additional entries specific to a rubber stamp annotation 507
- 8.23** Additional entries specific to an ink annotation 508
- 8.24** Additional entries specific to a pop-up annotation 509
- 8.25** Additional entries specific to a file attachment annotation 509
- 8.26** Additional entries specific to a sound annotation 510
- 8.27** Additional entries specific to a movie annotation 511
- 8.28** Additional entries specific to a widget annotation 512
- 8.29** Entries common to all action dictionaries 514
- 8.30** Entries in an annotation's additional-actions dictionary 515
- 8.31** Entries in a page object's additional-actions dictionary 515
- 8.32** Entries in a form field's additional-actions dictionary 516
- 8.33** Entries in the document catalog's additional-actions dictionary 516
- 8.34** Action types 518
- 8.35** Additional entries specific to a go-to action 519
- 8.36** Additional entries specific to a remote go-to action 520
- 8.37** Additional entries specific to a launch action 521

8.38	Entries in a Windows launch parameter dictionary	521
8.39	Additional entries specific to a thread action	522
8.40	Additional entries specific to a URI action	523
8.41	Entry in a URI dictionary	524
8.42	Additional entries specific to a sound action	525
8.43	Additional entries specific to a movie action	526
8.44	Additional entries specific to a hide action	527
8.45	Named actions	527
8.46	Additional entries specific to named actions	528
8.47	Entries in the interactive form dictionary	529
8.48	Signature flags	530
8.49	Entries common to all field dictionaries	531
8.50	Field flags common to all field types	532
8.51	Additional entries common to all fields containing variable text	534
8.52	Entries in an appearance characteristics dictionary	536
8.53	Field flags specific to button fields	538
8.54	Additional entry specific to a checkbox field	540
8.55	Additional entry specific to a radio button field	543
8.56	Field flags specific to text fields	543
8.57	Additional entry specific to a text field	544
8.58	Field flags specific to choice fields	546
8.59	Additional entries specific to a choice field	546
8.60	Entries in a signature dictionary	549
8.61	Additional entries specific to a submit-form action	551
8.62	Flags for submit-form actions	551
8.63	Additional entries specific to a reset-form action	555
8.64	Flag for reset-form actions	555
8.65	Additional entries specific to an import-data action	556
8.66	Additional entries specific to a JavaScript action	556
8.67	Entry in the FDF trailer dictionary	560
8.68	Entries in the FDF catalog dictionary	561
8.69	Entries in the FDF dictionary	561
8.70	Additional entry in an embedded file stream dictionary for an encrypted FDF file	563
8.71	Entries in the JavaScript dictionary	563
8.72	Entries in an FDF field dictionary	564
8.73	Entries in an icon fit dictionary	566
8.74	Entries in an FDF page dictionary	567
8.75	Entries in an FDF template dictionary	567
8.76	Entries in an FDF named page reference dictionary	568
8.77	Additional entry for annotation dictionaries in an FDF file	568
8.78	Additional entries specific to a sound object	569
8.79	Entries in a movie dictionary	571
8.80	Entries in a movie activation dictionary	571

9.1	Predefined procedure sets	574
9.2	Entries in the document information dictionary	576
9.3	Additional entries in a metadata stream dictionary	578
9.4	Additional entry for components having metadata	578
9.5	PDF constructs that do not take metadata	579
9.6	Entries in a page-piece dictionary	582
9.7	Entries in an application data dictionary	582
9.8	Marked-content operators	584
9.9	Entries in the structure tree root	590
9.10	Entries in a structure element dictionary	591
9.11	Entries in a marked-content reference dictionary	594
9.12	Entries in an object reference dictionary	599
9.13	Additional dictionary entries for structure element access	602
9.14	Entry common to all attribute objects	605
9.15	Entry in the mark information dictionary	614
9.16	Property list entries for artifacts	616
9.17	Derivation of font characteristics	622
9.18	Standard structure types for grouping elements	627
9.19	Block-level structure elements	629
9.20	Standard structure types for paragraphlike elements	630
9.21	Standard structure types for list elements	630
9.22	Standard structure types for table elements	631
9.23	Standard structure types for inline-level structure elements	633
9.24	Standard structure types for illustration elements	637
9.25	Standard attribute owners	639
9.26	Standard layout attributes	640
9.27	Standard layout attributes common to all standard structure types	641
9.28	Additional standard layout attributes specific to block-level structure elements	643
9.29	Standard layout attributes specific to inline-level structure elements	646
9.30	Standard list attribute	650
9.31	Standard table attributes	651
9.32	Entries in the Web Capture information dictionary	660
9.33	Entries common to all Web Capture content sets	668
9.34	Additional entries specific to a Web Capture page set	668
9.35	Additional entries specific to a Web Capture image set	669
9.36	Entries in a source information dictionary	670
9.37	Entries in a URL alias dictionary	671
9.38	Entries in a Web Capture command dictionary	672
9.39	Web Capture command flags	673
9.40	Entries in a Web Capture command settings dictionary	674
9.41	Entries in a box color information dictionary	681
9.42	Entries in a box style dictionary	681

9.43	Additional entries specific to a printer's mark annotation	682
9.44	Additional entries specific to a printer's mark form dictionary	683
9.45	Entries in a separation dictionary	684
9.46	Entries in a PDF/X output intent dictionary	685
9.47	Additional entries specific to a trap network annotation	691
9.48	Additional entries specific to a trap network appearance stream	692
9.49	Entry in an OPI version dictionary	694
9.50	Entries in a version 1.3 OPI dictionary	694
9.51	Entries in a version 2.0 OPI dictionary	697
A.1	PDF content stream operators	699
C.1	Architectural limits	706
D.1	Latin-text encodings	710
F.1	Entries in the linearization parameter dictionary	733
F.2	Standard hint tables	736
F.3	Page offset hint table, header section	743
F.4	Page offset hint table, per-page entry	744
F.5	Shared object hint table, header section	746
F.6	Shared object hint table, shared object group entry	747
F.7	Thumbnail hint table, header section	748
F.8	Thumbnail hint table, per-page entry	749
F.9	Generic hint table	750
F.10	Interactive form or logical structure hint table	750
G.1	Objects in minimal example	758
G.2	Objects in simple text string example	760
G.3	Objects in simple graphics example	762
G.4	Object usage after adding four text annotations	775
G.5	Object usage after deleting two text annotations	778
G.6	Object usage after adding three text annotations	780
H.1	Abbreviations for standard filter names	789
H.2	Acrobat behavior with unknown filters	789
H.3	Names of standard fonts	795

Preface

THE ORIGINS OF THE Portable Document Format and the Adobe® Acrobat® product family date to early 1990. At that time, the PostScript® page description language was rapidly becoming the worldwide standard for the production of the printed page. PDF builds on the PostScript page description language by layering a document structure and interactive navigation features on PostScript's underlying imaging model, providing a convenient, efficient mechanism enabling documents to be reliably viewed and printed anywhere.

The PDF specification was first published at the same time the first Acrobat products were introduced in 1993. Since then, updated versions of the specification have been and continue to be available from Adobe via the World Wide Web. This book is the third professionally published edition of the specification. Like its predecessor, it is completely self-contained, including the precise documentation of the underlying imaging model from PostScript along with the PDF-specific features that are combined in version 1.4 of the PDF standard.

Over the past eight years, aided by the explosive growth of the Internet, PDF has become the *de facto* standard for the electronic exchange of documents. Well over 200 million copies of the free Acrobat Reader® application have been distributed around the world, facilitating efficient sharing of digital content. In addition, PDF is now the industry standard for the intermediate representation of printed material in electronic prepress systems for conventional printing applications. As major corporations, government agencies, and educational institutions streamline their operations by replacing paper-based workflow with electronic exchange of information, the impact and opportunity for the application of PDF will continue to grow at a rapid pace.

PDF is the file format that underlies Adobe ePaper® Solutions, a family of products supporting Adobe's vision for Network Publishing—the process of creating, managing, and accessing digital content on diverse platforms and devices. ePaper

fulfills a set of requirements related to business process needs for the global desktop user, including:

- Preservation of document fidelity across the enterprise, independently of the device, platform, and software
- Merging of content from diverse sources—Web sites, word processing and spreadsheet programs, scanned documents, photos, and graphics—into one self-contained document while maintaining the integrity of all original source documents
- Real-time collaborative editing of documents from multiple locations or platforms
- Digital signatures to certify authenticity
- Security and permissions to allow the creator to retain control of the document and associated rights
- Accessibility of content to those with disabilities
- Extraction and reuse of content using other file formats and applications

A significant number of third-party developers and systems integrators offer customized enhancements and extensions to Adobe's core family of products. Adobe publishes the PDF specification in order to encourage the development of such third-party applications.

The emergence of PDF as a standard for electronic information exchange is the result of concerted effort by many individuals in both the private and public sectors. Without the dedication of Adobe employees, our industry partners, and our customers, the widespread acceptance of PDF could not have been achieved. We thank all of you for your continuing support and creative contributions to the success of PDF.

*Chuck Geschke and John Warnock
May 2001*

CHAPTER 1

Introduction

THE ADOBE PORTABLE DOCUMENT FORMAT (PDF) is the native file format of the Adobe® Acrobat® family of products. The goal of these products is to enable users to exchange and view electronic documents easily and reliably, independently of the environment in which they were created. PDF relies on the same imaging model as the PostScript® page description language to describe text and graphics in a device-independent and resolution-independent manner. To improve performance for interactive viewing, PDF defines a more structured format than that used by most PostScript language programs. PDF also includes objects, such as annotations and hypertext links, that are not part of the page itself but are useful for interactive viewing and document interchange.

1.1 About This Book

This book provides a description of the PDF file format and is intended primarily for application developers wishing to develop *PDF producer* applications that create PDF files directly. It also contains enough information to allow developers to write *PDF consumer* applications that read existing PDF files and interpret or modify their contents.

Although the PDF specification is independent of any particular software implementation, some PDF features are best explained by describing the way they are processed by a typical application program. In such cases, this book uses the Adobe Acrobat family of PDF viewer applications as its model. (The prototypical viewer is the fully capable Acrobat product, not the limited Acrobat Reader® product.) Similarly, Appendix C discusses some implementation limits in the Acrobat viewer applications, even though these limits are not part of the file format itself. To provide guidance to implementors of PDF producer and consumer applications, compatibility and implementation notes in Appendix H describe

the behavior of Acrobat viewer applications when they encounter newer features they do not understand, as well as areas in which the Acrobat products diverge from the specification presented in this book.

This third edition of the *PDF Reference* describes version 1.4 of PDF. (See implementation note 1 in Appendix H.) Throughout the book, information specific to particular versions of PDF is marked as such—for example, with indicators like (*PDF 1.3*) or (*PDF 1.4*). Features so marked may be new in the indicated version or may have been substantially redefined in that version. Features designated (*PDF 1.0*) have generally been superseded in later versions; unless otherwise stated, features identified as specific to other versions are understood to be available in later versions as well. (PDF viewer applications designed for a specific PDF version generally ignore newer features they do not recognize; implementation notes in Appendix H point out exceptions.)

The rest of the book is organized as follows:

- Chapter 2, “Overview,” briefly introduces the overall architecture of PDF and the design considerations behind it, compares it with the PostScript language, and describes the underlying imaging model that they share.
- Chapter 3, “Syntax,” presents the syntax of PDF at the object, file, and document level. It sets the stage for subsequent chapters, which describe how that information is interpreted as page descriptions, interactive navigational aids, and application-level logical structure.
- Chapter 4, “Graphics,” describes the graphics operators used to describe the appearance of pages in a PDF document.
- Chapter 5, “Text,” discusses PDF’s special facilities for presenting text in the form of character shapes, or *glyphs*, defined by fonts.
- Chapter 6, “Rendering,” considers how device-independent content descriptions are matched to the characteristics of a particular output device.
- Chapter 7, “Transparency,” discusses the operation of the *transparent imaging model*, introduced in PDF 1.4, in which objects can be painted with varying degrees of opacity, allowing the previous contents of the page to show through.
- Chapter 8, “Interactive Features,” describes those features of PDF that allow a user to interact with a document on the screen, using the mouse and keyboard.

- Chapter 9, “Document Interchange,” shows how PDF documents can incorporate higher-level information that is useful for the interchange of documents among applications.

The appendices contain useful tables and other auxiliary information.

- Appendix A, “Operator Summary,” lists all the operators used in describing the visual content of a PDF document.
- Appendix B, “Operators in Type 4 Functions,” summarizes the PostScript operators that can be used in PostScript calculator functions, which contain code written in a small subset of the PostScript language.
- Appendix C, “Implementation Limits,” describes typical size and quantity limits imposed by the Acrobat viewer applications.
- Appendix D, “Character Sets and Encodings,” lists the character sets and encodings that are assumed to be predefined in any PDF viewer application.
- Appendix E, “PDF Name Registry,” discusses a registry, maintained for developers by Adobe Systems, that contains private names and formats used by PDF producers or Acrobat plug-in extensions.
- Appendix F, “Linearized PDF,” describes a special form of PDF file organization designed to work efficiently in network environments.
- Appendix G, “Example PDF Files,” presents several examples showing the structure of actual PDF files, ranging from one containing a minimal one-page document to one showing how the structure of a PDF file evolves over the course of several revisions.
- Appendix H, “Compatibility and Implementation Notes,” provides details on the behavior of Acrobat viewer applications and describes how viewer applications should handle PDF files containing features that they do not recognize.

A color plate section provides illustrations of some of PDF’s color-related features. References in the text of the form “see Plate 1” refer to the contents of this section.

The book concludes with a Bibliography and an Index.

The enclosed CD-ROM contains the entire text of this book in PDF form.

1.2 Introduction to PDF 1.4 Features

The most significant addition in PDF 1.4 is the new *transparent imaging model*, which extends the opaque imaging model of earlier versions to include the ability to paint objects with varying degrees of opacity, allowing previously painted objects to show through. Transparency is covered primarily in Chapter 7, with implications reflected in other chapters as well. Other new features introduced in PDF 1.4 include the following:

- A filter for decoding JBIG2-encoded data (Section 3.3.6, “JBIG2Decode Filter”)
- Enhancements to encryption (Section 3.5, “Encryption”)
- Specification of the PDF version in the document catalog (Section 3.6.1, “Document Catalog”)
- Embedding of data in a file stream (Section 3.10.3, “Embedded File Streams”)
- The ability to import content from one PDF document into another (Section 4.9.3, “Reference XObjects”)
- New predefined CMaps (“Predefined CMaps” on page 343)
- Additional viewer preferences for controlling the area of a page to be displayed or printed (Section 8.1, “Viewer Preferences”)
- Specification of a color and font style for text in an outline item (Section 8.2.2, “Document Outline”)
- Annotation names (Section 8.4, “Annotations”) and new entries in specific annotation dictionaries (Section 8.4.1, “Annotation Dictionaries”)
- Additional trigger events for actions affecting the document as a whole (Section 8.5.2, “Trigger Events”)
- Assorted enhancements to interactive forms and Forms Data Format (FDF), including multiple-selection list boxes, file-select controls, XML form submission, embedded FDF files, Unicode[®] specification of field export values, and support for remote collaboration and digital signatures in FDF files (Section 8.6, “Interactive Forms”)
- Metadata streams, a new architecture for attaching descriptive information to PDF documents and their constituent parts (Section 9.2.2, “Metadata Streams”)

- Standardized structure types and attributes for describing the logical structure of a document (Section 9.7, “Tagged PDF”)
- Support for accessibility to disabled users, including specification of the language used for text (Section 9.8, “Accessibility Support”)
- Support for the display and preview of production-related page boundaries such as the crop box, bleed box, and trim box (Section 9.10.1, “Page Boundaries”)
- Facilities for including printer’s marks such as registration targets, gray ramps, color bars, and cut marks to assist in the production process (Section 9.10.2, “Printer’s Marks”)
- Output intents for matching the color characteristics of a document with those of a target output device or production environment in which it will be printed (Section 9.10.4, “Output Intents”)

1.3 Related Publications

PDF and the PostScript page description language share the same underlying Adobe imaging model. A document can be converted straightforwardly between PDF and the PostScript language; the two representations produce the same output when printed. However, PostScript includes a general-purpose programming language framework not present in PDF. The *PostScript Language Reference* is the comprehensive reference for the PostScript language and its imaging model.

PDF and PostScript support several standard formats for font programs, including Adobe Type 1, CFF (Compact Font Format), TrueType[®], and CID-keyed fonts. The PDF manifestations of these fonts are documented in this book. However, the specifications for the font files themselves are published separately, because they are highly specialized and are of interest to a different user community. A variety of Adobe publications are available on the subject of font formats, most notably the following:

- *Adobe Type 1 Font Format* and Adobe Technical Note #5015, *Type 1 Font Format Supplement*
- Adobe Technical Note #5176, *The Compact Font Format Specification*
- Adobe Technical Note #5177, *The Type 2 Charstring Format*
- Adobe Technical Note #5014, *Adobe CMap and CID Font Files Specification*

See the Bibliography for additional publications related to PDF and the contents of this book.

1.4 Intellectual Property

The general idea of using an interchange format for electronic documents is in the public domain. Anyone is free to devise a set of unique data structures and operators that define an interchange format for electronic documents. However, Adobe Systems Incorporated owns the copyright for the particular data structures and operators and the written specification constituting the interchange format called the Portable Document Format. Thus, these elements of the Portable Document Format may not be copied without Adobe's permission.

Adobe will enforce its copyright. Adobe's intention is to maintain the integrity of the Portable Document Format standard. This enables the public to distinguish between the Portable Document Format and other interchange formats for electronic documents. However, Adobe desires to promote the use of the Portable Document Format for information interchange among diverse products and applications. Accordingly, Adobe gives anyone copyright permission, subject to the conditions stated below, to:

- Prepare files whose content conforms to the Portable Document Format
- Write drivers and applications that produce output represented in the Portable Document Format
- Write software that accepts input in the form of the Portable Document Format and displays, prints, or otherwise interprets the contents
- Copy Adobe's copyrighted list of data structures and operators, as well as the example code and PostScript language function definitions in the written specification, to the extent necessary to use the Portable Document Format for the purposes above

The conditions of such copyright permission are:

- Software that accepts input in the form of the Portable Document Format must respect the access permissions specified in that document. Accessing the docu-

ment in ways not permitted by the document's access permissions is a violation of the document author's copyright.

- Anyone who uses the copyrighted list of data structures and operators, as stated above, must include an appropriate copyright notice.

This limited right to use the copyrighted list of data structures and operators does not include the right to copy this book, other copyrighted material from Adobe, or the software in any of Adobe's products that use the Portable Document Format, in whole or in part, nor does it include the right to use any Adobe patents, except as may be permitted by an official Adobe Patent Clarification Notice (see the Bibliography).

Acrobat, Acrobat Capture, Acrobat Reader, ePaper, the "Get Acrobat Reader" Web logo, the "Adobe PDF" Web logo, and all other trademarks, service marks, and logos used by Adobe (the "Marks") are the registered trademarks or trademarks of Adobe Systems Incorporated in the United States and other countries. Nothing in this book is intended to grant you any right or license to use the Marks for any purpose.

CHAPTER 2

Overview

THE ADOBE *PORTABLE DOCUMENT FORMAT (PDF)* is a file format for representing documents in a manner independent of the application software, hardware, and operating system used to create them and of the output device on which they are to be displayed or printed. A *PDF document* consists of a collection of *objects* that together describe the appearance of one or more *pages*, possibly accompanied by additional interactive elements and higher-level application data. A *PDF file* contains the objects making up a PDF document along with associated structural information, all represented as a single self-contained sequence of bytes.

A document's pages (and other visual elements) may contain any combination of text, graphics, and images. A page's appearance is described by a PDF *content stream*, which contains a sequence of *graphics objects* to be painted on the page. This appearance is fully specified; all layout and formatting decisions have already been made by the application generating the content stream.

In addition to describing the static appearance of pages, a PDF document may contain interactive elements that are possible only in an electronic representation. PDF supports *annotations* of many kinds for such things as text notes, hypertext links, markup, file attachments, sounds, and movies. A document can define its own user interface; keyboard and mouse input can trigger *actions* that are specified by PDF objects. The document can contain *interactive form* fields to be filled in by the user, and can export the values of these fields to or import them from other applications.

Finally, a PDF document can contain higher-level information that is useful for interchange of content among applications. In addition to specifying appearance, a document's content can include identification and logical structure informa-

tion that allows it to be searched, edited, or extracted for reuse elsewhere. PDF is particularly well suited for representing a document as it moves through successive stages of a prepress production workflow.

2.1 Imaging Model

At the heart of PDF is its ability to describe the appearance of sophisticated graphics and typography. This is achieved through the use of the *Adobe imaging model*, the same high-level, device-independent representation used in the PostScript page description language.

Although application programs could theoretically describe any page as a full-resolution pixel array, the resulting file would be bulky, device-dependent, and impractical for high-resolution devices. A high-level *imaging model* enables applications to describe the appearance of pages containing text, graphical shapes, and sampled images in terms of abstract graphical elements rather than directly in terms of device pixels. Such a description is economical and device-independent, and can be used to produce high-quality output on a broad range of printers, displays, and other output devices.

2.1.1 Page Description Languages

Among its other roles, PDF serves as a *page description language*: a language for describing the graphical appearance of pages with respect to an imaging model. An application program produces output through a two-stage process:

1. The application generates a device-independent description of the desired output in the page description language.
2. A program controlling a specific output device interprets the description and *renders* it on that device.

The two stages may be executed in different places and at different times; the page description language serves as an interchange standard for the compact, device-independent transmission and storage of printable or displayable documents.

2.1.2 Adobe Imaging Model

The Adobe imaging model is a simple and unified view of two-dimensional graphics borrowed from the graphic arts. In this model, “paint” is placed on a page in selected areas.

- The painted figures may be in the form of character shapes (*glyphs*), geometric shapes, lines, or sampled images such as digital representations of photographs.
- The paint may be in color or in black, white, or any shade of gray; it may also take the form of a repeating *pattern* (PDF 1.2) or a smooth transition between colors (PDF 1.3).
- Any of these elements may be *clipped* to appear within other shapes as they are placed onto the page.

A page’s content stream contains *operands* and *operators* describing a sequence of graphics objects. A PDF viewer application maintains an implicit *current page* that accumulates the marks made by the painting operators. Initially, the current page is completely blank. For each graphics object encountered in the content stream, the viewer places marks on the current page, which replace or combine with any previous marks they may overlay. Once the page has been completely composed, the accumulated marks are rendered on the output medium and the current page is cleared to blank again.

Versions of PDF up to and including PDF 1.3 use an *opaque imaging model* in which each new graphics object painted onto a page completely obscures the previous contents of the page at those locations (subject to the effects of certain optional parameters that may modify this behavior; see Section 4.5.6, “Overprint Control”). No matter what color an object has—white, black, gray, or color—it is placed on the page as if it were applied with opaque paint. PDF 1.4 introduces a new *transparent imaging model* in which objects painted on the page are not required to be fully opaque. Instead, newly painted objects are *composited* with the previously existing contents of the page, producing results that combine the colors of the object and its backdrop according to their respective opacity characteristics. The transparent imaging model is described in Chapter 7.

The principal graphics objects (among others) are as follows:

- A *path object* consists of a sequence of connected and disconnected points, lines, and curves that together describe shapes and their positions. It is built up through the sequential application of *path construction operators*, each of which appends one or more new elements. The path object is ended by a *path-painting operator*, which paints the path on the page in some way. The principal path-painting operators are **S** (stroke), which paints a line along the path, and **f** (fill), which paints the interior of the path.
- A *text object* consists of one or more glyph shapes representing characters of text. The glyph shapes for the characters are described in a separate data structure called a *font*. Like path objects, text objects can be stroked or filled.
- An *image object* is a rectangular array of *sample values*, each representing a color at a particular position within the rectangle. Such objects are typically used to represent photographs.

The painting operators require various parameters, some explicit and others implicit. Implicit parameters include the current color, current line width, current font (typeface and size), and many others. Together, these implicit parameters make up the *graphics state*. There are operators for setting the value of each implicit parameter in the graphics state; painting operators use the values currently in effect at the time they are invoked.

One additional implicit parameter in the graphics state modifies the results of painting graphics objects. The *current clipping path* outlines the area of the current page within which paint can be placed. Although painting operators may attempt to place marks anywhere on the current page, only those marks falling within the current clipping path will affect the page; those falling outside it will not. Initially, the current clipping path encompasses the entire imageable area of the page. It can temporarily be reduced to the shape defined by a path or text object, or to the intersection of multiple such shapes. Marks placed by subsequent painting operators will then be confined within that boundary.

2.1.3 Raster Output Devices

Much of the power of the Adobe imaging model derives from its ability to deal with the general class of *raster output devices*. These encompass such technologies as laser, dot-matrix, and ink-jet printers, digital imagesetters, and raster-scan displays. The defining property of a raster output device is that a printed or dis-

played image consists of a rectangular array, or *raster*, of dots called *pixels* (picture elements) that can be addressed individually. On a typical *bilevel* output device, each pixel can be made either black or white. On some devices, pixels can be set to intermediate shades of gray or to some color. The ability to set the colors of individual pixels makes it possible to generate printed or displayed output that can include text, arbitrary graphical shapes, and reproductions of sampled images.

The *resolution* of a raster output device measures the number of pixels per unit of distance along the two linear dimensions. Resolution is typically—but not necessarily—the same horizontally and vertically. Manufacturers' decisions on device technology and price/performance tradeoffs create characteristic ranges of resolution:

- Computer displays have relatively low resolution, typically 75 to 110 pixels per inch.
- Dot-matrix printers generally range from 100 to 250 pixels per inch.
- Ink-jet and laser-scanned xerographic printing technologies achieve medium-level resolutions of 300 to 1400 pixels per inch.
- Photographic technology permits high resolutions of 2400 pixels per inch or more.

Higher resolution yields better quality and fidelity of the resulting output, but is achieved at greater cost. As the technology improves and computing costs decrease, products evolve to higher resolutions.

2.1.4 Scan Conversion

An abstract graphical element (such as a line, a circle, a character glyph, or a sampled image) is rendered on a raster output device by a process known as *scan conversion*. Given a mathematical description of the graphical element, this process determines which pixels to adjust and what values to assign to those pixels to achieve the most faithful rendition possible at the available device resolution.

The pixels on a page can be represented by a two-dimensional array of pixel values in computer memory. For an output device whose pixels can only be black or white, a single bit suffices to represent each pixel. For a device that can reproduce gray levels or colors, multiple bits per pixel are required.

Note: Although the ultimate representation of a printed or displayed page is logically a complete array of pixels, its actual representation in computer memory need not consist of one memory cell per pixel. Some implementations use other representations, such as display lists. The Adobe imaging model has been carefully designed not to depend on any particular representation of raster memory.

For each graphical element that is to appear on the page, the scan converter sets the values of the corresponding pixels. When the interpretation of the page description is complete, the pixel values in memory represent the appearance of the page. At this point, a raster output process can *render* this representation (make it visible) on a printed page or display screen.

Scan-converting a graphical shape, such as a rectangle or circle, entails determining which device pixels lie inside the shape and setting their values appropriately (for example, to black). Because the edges of a shape do not always fall precisely on the boundaries between pixels, some policy is required for deciding how to set the pixels along the edges. Scan-converting a glyph representing a text character is conceptually the same as scan-converting an arbitrary graphical shape; however, character glyphs are much more sensitive to legibility requirements and must meet more rigid objective and subjective measures of quality.

Rendering grayscale elements on a bilevel device is accomplished by a technique known as *halftoning*. The array of pixels is divided into small clusters according to some pattern (called the *halftone screen*). Within each cluster, some pixels are set to black and some to white in proportion to the level of gray desired at that location on the page. When viewed from a sufficient distance, the individual dots become imperceptible and the perceived result is a shade of gray. This enables a bilevel raster output device to reproduce shades of gray and to approximate natural images such as photographs. Some color devices use a similar technique.

2.2 Other General Properties

This section describes other notable general properties of PDF, aside from its imaging model.

2.2.1 Portability

PDF files are represented as sequences of 8-bit binary bytes. A PDF file is designed to be portable across all platforms and operating systems. The binary rep-

resentation is intended to be generated, transported, and consumed directly, without translation between native character sets, end-of-line representations, or other conventions used on various platforms.

Any PDF file can also be represented in a form that uses only 7-bit ASCII (American Standard Code for Information Interchange) character codes. This is useful for the purpose of exposition, as in this book. However, this representation is not recommended for actual use, since it is less efficient than the normal binary representation. Regardless of which representation is used, PDF files must be transported and stored as *binary* files, not as *text* files; inadvertent changes, such as conversion between text end-of-line conventions, will damage the file and may render it unusable.

2.2.2 Compression

To reduce file size, PDF supports a number of industry-standard compression filters:

- JPEG compression of color and grayscale images
- CCITT (Group 3 or Group 4), run-length, and (in PDF 1.4) JBIG2 compression of monochrome images
- LZW (Lempel-Ziv-Welch) and (beginning with PDF 1.2) Flate compression of text, graphics, and images

Using JPEG compression, color and grayscale images can be compressed by a factor of 10 or more. Effective compression of monochrome images depends on the compression filter used and the properties of the image, but reductions of 2:1 to 8:1 are common (or 20:1 to 50:1 for JBIG2 compression of an image of a page full of text). LZW or Flate compression of the content streams describing all other text and graphics in the document results in compression ratios of approximately 2:1. All of these compression filters produce binary data, which can then be further converted to ASCII base-85 encoding if a 7-bit ASCII representation is desired.

2.2.3 Font Management

Managing fonts is a fundamental challenge in document interchange. Generally, the receiver of a document must have the same fonts that were originally used to

create it. If a different font is substituted, its character set, glyph shapes, and metrics may differ from those in the original font. This can produce unexpected and undesirable results, such as lines of text extending into margins or overlapping with graphics.

PDF provides various means for dealing with font management:

- The original font programs can be embedded in the PDF file. PDF supports various font formats, including Type 1, TrueType[®], and CID-keyed fonts. This ensures the most predictable and dependable results.
- To conserve space, a font subset can be embedded, containing just the glyph descriptions for those characters that are actually used in the document. Also, Type 1 fonts can be represented in a special compact format.
- PDF prescribes a set of 14 standard fonts that can be used without prior definition. These include four faces each of three Latin text typefaces (Courier, Helvetica*, and Times*), as well as two symbolic fonts (Symbol and ITC Zapf Dingbats[®]). These fonts, or suitable substitute fonts with the same metrics, are guaranteed to be available in all PDF viewer applications.
- A PDF file can refer by name to fonts that are not embedded in the PDF file. In this case, a viewer application will use those fonts if they are available in the viewer's environment. This approach suffers from the uncertainties noted above.
- A PDF file contains a *font descriptor* for each font that it uses (other than the standard 14). The font descriptor includes font metrics and style information, enabling a viewer application to select or synthesize a suitable substitute font if necessary. Although the glyphs' shapes will differ from those intended, their placement will be accurate.

Font management is primarily concerned with producing the correct appearance of text—that is, the shape and placement of glyphs. However, it is sometimes necessary for a PDF application to extract the meaning of the text, represented in some standard information encoding such as Unicode. In some cases, this information can be deduced from the encoding used to represent the text in the PDF file. Otherwise, the PDF producer application should specify the mapping explicitly by including a special object, the **ToUnicode CMap**.

2.2.4 Single-Pass File Generation

Because of system limitations and efficiency considerations, it may be necessary or desirable for an application program to generate a PDF file in a single pass. For example, the program may have limited memory available or be unable to open temporary files. For this reason, PDF supports single-pass generation of files. Although some PDF objects must specify their length in bytes, a mechanism is provided allowing the length to follow the object itself in the PDF file. In addition, information such as the number of pages in the document can be written into the file after all pages have been generated.

A PDF file that is generated in a single pass is generally not ordered for most efficient viewing, particularly when accessing the contents of the file over a network. When generating a PDF file that is intended to be viewed many times, it is worthwhile to perform a second pass to optimize the order in which objects occur in the file. PDF specifies a particular file organization, *Linearized PDF*, which is documented in Appendix F. Other optimizations are also possible, such as detecting duplicated sequences of graphics objects and collapsing them to a single shared sequence that is specified only once.

2.2.5 Random Access

A PDF file should be thought of as a flattened representation of a data structure consisting of a collection of objects that can refer to each other in any arbitrary way. The order of the objects' occurrence in the PDF file has no semantic significance. In general, a viewer application should process a PDF file by following references from object to object, rather than by processing objects sequentially. This is particularly important for interactive document viewing or for any application in which pages or other objects in the PDF file are accessed out of sequence.

To support such random access to individual objects, every PDF file contains a *cross-reference table* that can be used to locate and directly access pages and other important objects within the file. The cross-reference table is stored at the end of the file, allowing applications that generate PDF files in a single pass to store it easily and those that read PDF files to locate it easily. Using the cross-reference table makes the time needed to locate a page or other object nearly independent of the length of the document. This allows PDF documents containing hundreds or thousands of pages to be accessed efficiently.

2.2.6 Security

PDF has two security features that can be used, separately or together, in any document:

- The document can be *encrypted* so that only authorized users can access it. There is separate authorization for the owner of the document and for all other users; the users' access can be selectively restricted to allow only certain operations, such as viewing, printing, or editing.
- The document can be digitally *signed* to certify its authenticity. The signature may take many forms, including a document digest that has been encrypted with a public/private key, a biometric signature such as a fingerprint, and others. Any subsequent changes to a signed PDF file will invalidate the signature.

2.2.7 Incremental Update

Applications may allow users to modify PDF documents. Users should not have to wait for the entire file—which can contain hundreds of pages or more—to be rewritten each time modifications to the document are saved. PDF allows modifications to be appended to a file, leaving the original data intact. The addendum appended when a file is incrementally updated contains only those objects that were actually added or modified, and includes an update to the cross-reference table. Incremental update allows an application to save modifications to a PDF document in an amount of time proportional to the size of the modification rather than the size of the file.

In addition, because the original contents of the document are still present in the file, it is possible to undo saved changes by deleting one or more addenda. The ability to recover the exact contents of an original document is critical when digital signatures have been applied and subsequently need to be verified.

2.2.8 Extensibility

PDF is designed to be extensible. Not only can new features be added, but applications based on earlier versions of PDF can behave reasonably when they encounter newer features that they do not understand. Appendix H describes how a PDF viewer application should behave in such cases.

Additionally, PDF provides means for applications to store their own private information in a PDF file. This information can be recovered when the file is imported by the same application, but is ignored by other applications. This allows PDF to serve as an application's native file format while allowing its documents to be viewed and printed by other applications. Application-specific data can be stored either as *marked content* annotating the graphics objects in a PDF content stream or as entirely separate objects unconnected with the PDF content.

2.3 Using PDF

PDF files may be produced either directly by application programs or indirectly by conversion from other file formats or imaging models. As PDF documents and applications that process them become more prevalent, new ways of creating and using PDF will be invented. One of the goals of this book is to make the file format accessible so that application developers can expand on the ideas behind PDF and the applications that initially support it.

Many applications can generate PDF files directly, and some can import them as well. This is the most desirable approach, since it gives the application access to the full capabilities of PDF, including the imaging model and the interactive and document interchange features. Alternatively, existing applications that do not generate PDF directly can still be used to produce PDF output by indirect methods. There are two principal ways of doing this:

- The application describes its printable output by making calls to an application programming interface (API) such as GDI in Microsoft® Windows® or Quick-Draw® in the Apple® Mac® OS. A software component called a *printer driver* intercepts these calls and interprets them to generate output in PDF form.
- The application produces printable output directly in some other file format, such as PostScript, PCL, HPGL, or DVI, which is then converted into PDF by a separate translation program.

Note, however, that while these indirect strategies are often the easiest way to obtain PDF output from an existing application, the resulting PDF files may not make the best use of the high-level Adobe imaging model. This is because the information embodied in the application's API calls or in the intermediate output file often describes the desired results at too low a level; any higher-level information maintained by the original application has been lost and is not available to the printer driver or translator.

Figures 2.1 and 2.2 show how Adobe Acrobat products support these indirect approaches. PDF Writer (Figure 2.1), available on the Windows and Mac OS platforms, acts as a printer driver, intercepting graphics and text operations generated by a running application program through the operating system's API. Instead of converting these operations into printer commands and transmitting them directly to a printer, PDF Writer converts them to equivalent PDF operators and embeds them in a PDF file. The result is a platform-independent file that can be viewed and printed by a PDF viewer application, such as Adobe Acrobat, running on any supported platform—even a different platform from the one on which the file was originally generated.

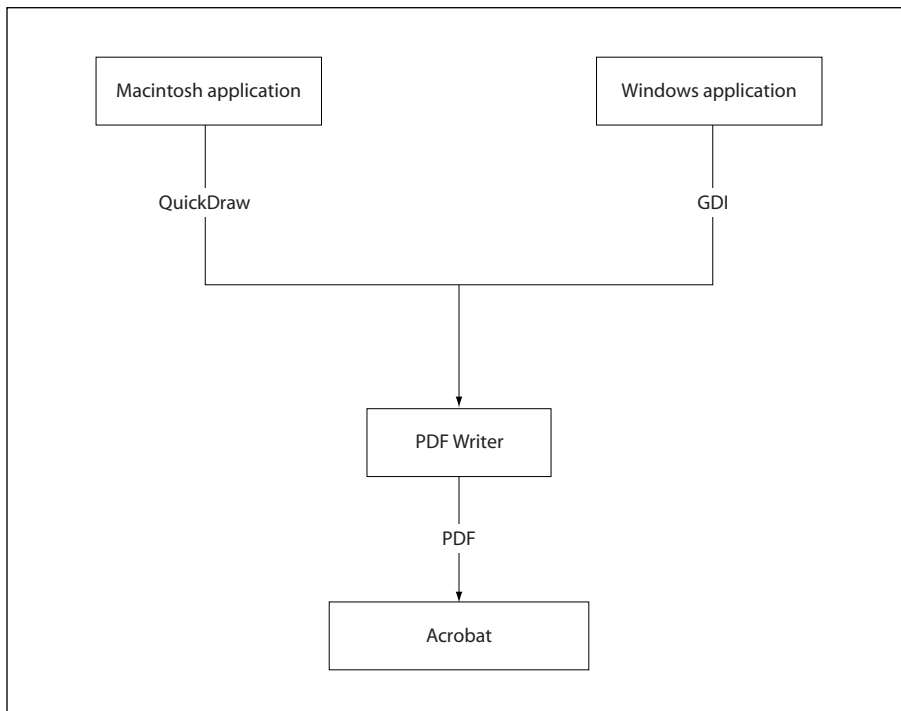


FIGURE 2.1 *Creating PDF files using PDF Writer*

Instead of describing their printable output via API calls, some applications produce PostScript page descriptions directly—either because of limitations in the QuickDraw or GDI imaging models or because the applications run on platforms such as DOS or UNIX[®], where no system-level printer driver exists. PostScript

files generated by such applications can be converted into PDF files using the Acrobat Distiller[®] application (see Figure 2.2). Because PostScript and PDF share the same Adobe imaging model, Acrobat Distiller can preserve the exact graphical content of the PostScript file in the translation to PDF. Additionally, Distiller supports a PostScript language extension, called **pdfmark**, that allows the producing application to embed instructions in the PostScript file for creating hypertext links, logical structure, and other interactive and document interchange features of PDF. Again, the resulting PDF file can be viewed with a viewer application, such as Adobe Acrobat, on any supported platform.




FIGURE 2.2 *Creating PDF files using Acrobat Distiller*

2.4 PDF and the PostScript Language

The PDF operators for setting the graphics state and painting graphics objects are similar to the corresponding operators in the PostScript language. Unlike PostScript, however, PDF is not a full-scale programming language; it trades reduced flexibility for improved efficiency and predictability. PDF therefore differs from PostScript in the following significant ways:

- PDF enforces a strictly defined file structure that allows an application to access parts of a document in arbitrary order.
- To simplify the processing of content streams, PDF does not include common programming language features such as procedures, variables, and control constructs.

- PDF files contain information such as font metrics to ensure viewing fidelity.
- A PDF file may contain additional information that is not directly connected with the imaging model, such as hypertext links for interactive viewing and logical structure information for document interchange.

Because of these differences, a PDF file generally cannot be transmitted directly to a PostScript output device for printing (although a few such devices do also support PDF directly). An application printing a PDF document to a PostScript device must carry out the following steps:

1. Insert *procedure sets* containing PostScript procedure definitions to implement the PDF operators.
2. Extract the content for each page. Each content stream is essentially the script portion of a traditional PostScript program using very specific procedures, such as **m** for **moveto** and **l** for **lineto**.
3. Decode compressed text, graphics, and image data as necessary. The compression filters used in PDF are compatible with those used in PostScript; they may or may not be supported, depending on the LanguageLevel of the target output device.
4. Insert any needed resources, such as fonts, into the PostScript file. These can be either the original fonts themselves or suitable substitute fonts based on the font metrics in the PDF file. Fonts may need to be converted to a format that the PostScript interpreter recognizes, such as Type 1 or Type 42.
5. Put the information in the correct order. The result is a traditional PostScript program that fully represents the visual aspects of the document but no longer contains PDF elements such as hypertext links, annotations, and bookmarks.
6. Transmit the PostScript program to the output device.

CHAPTER 3

Syntax

THIS CHAPTER COVERS everything about the syntax of PDF at the object, file, and document level. It sets the stage for subsequent chapters, which describe how the contents of a PDF file are interpreted as page descriptions, interactive navigational aids, and application-level logical structure.

PDF syntax is best understood by thinking of it in four parts, as shown in Figure 3.1:

- *Objects.* A PDF document is a data structure composed from a small set of basic types of data object. Section 3.1, “Lexical Conventions,” describes the character set used to write objects and other syntactic elements. Section 3.2, “Objects,” describes the syntax and essential properties of the objects themselves. Section 3.2.7, “Stream Objects,” provides complete details of the most complex data type, the stream object.
- *File structure.* The PDF file structure determines how objects are stored in a PDF file, how they are accessed, and how they are updated. This structure is independent of the semantics of the objects. Section 3.4, “File Structure,” describes the file structure. Section 3.5, “Encryption,” describes a file-level mechanism for protecting a document’s contents from unauthorized access.
- *Document structure.* The PDF document structure specifies how the basic object types are used to represent components of a PDF document: pages, fonts, annotations, and so forth. Section 3.6, “Document Structure,” describes the overall document structure; later chapters address the detailed semantics of the components.
- *Content streams.* A PDF *content stream* contains a sequence of instructions describing the appearance of a page or other graphical entity. These instructions, while also represented as objects, are conceptually distinct from the objects that

represent the document structure and are described separately. Section 3.7, “Content Streams and Resources,” discusses PDF content streams and their associated resources.




FIGURE 3.1 *PDF components*

In addition, this chapter describes some data structures, built from basic objects, that are so widely used that they can almost be considered basic object types in their own right. These objects are covered in Sections 3.8, “Common Data Structures”; 3.9, “Functions”; and 3.10, “File Specifications.”

PDF’s object and file syntax is also used as the basis for other file formats. These include the Forms Data Format (FDF), described in Section 8.6.6, “Forms Data Format,” and the Portable Job Ticket Format (PJTF), described in Adobe Technical Note #5620, *Portable Job Ticket Format*.

3.1 Lexical Conventions

At the most fundamental level, a PDF file is a sequence of 8-bit bytes. These bytes can be grouped into *tokens* according to the syntax rules described below. One or more tokens are then assembled to form higher-level syntactic entities, principally *objects*, which are the basic data values from which a PDF document is constructed.

PDF can be entirely represented using byte values corresponding to the visible printable subset of the ASCII character set, plus characters that appear as “white space,” such as space, tab, carriage return, and line feed characters. ASCII is the American Standard Code for Information Interchange, a widely used convention

for encoding a specific set of 128 characters as binary numbers. However, a PDF file is not restricted to the ASCII character set; it can contain arbitrary 8-bit bytes, subject to the following considerations:

- The tokens that delimit objects and that describe the structure of a PDF file are all written in the ASCII character set, as are all the reserved words and the names used as keys in standard dictionaries.
- The data values of certain types of object—strings and streams—can be but need not be written entirely in ASCII. For the purpose of exposition (as in this book), ASCII representation is preferred. However, in actual practice, data that is naturally binary, such as sampled images, is represented directly in binary for the sake of compactness and efficiency.
- A PDF file containing binary data must be transported and stored by means that preserve all bytes of the file faithfully; that is, as a binary file rather than a text file. Such a file is not portable to environments that impose reserved character codes, maximum line lengths, end-of-line conventions, or other restrictions.

Note: In this chapter, the term *character* is synonymous with *byte* and merely refers to a particular 8-bit value. This is entirely independent of any logical meaning that the value may have when it is treated as data in specific contexts, such as representing human-readable text or selecting a glyph from a font.

3.1.1 Character Set

The PDF character set is divided into three classes, called *regular*, *delimiter*, and *white-space* characters. This classification determines the grouping of characters into tokens, except within strings, streams, and comments; different rules apply in those contexts.

White-space characters (see Table 3.1) separate syntactic constructs such as names and numbers from each other. All white-space characters are equivalent, except in comments, strings, and streams. In all other contexts, PDF treats any sequence of consecutive white-space characters as if there were just one.

TABLE 3.1 White-space characters

DECIMAL	HEXADECIMAL	OCTAL	NAME
0	00	000	Null (NUL)
9	09	011	Tab (HT)
10	0A	012	Line feed (LF)
12	0C	014	Form feed (FF)
13	0D	015	Carriage return (CR)
32	20	040	Space (SP)

The carriage return (CR) and line feed (LF) characters, also called *newline characters*, are treated as *end-of-line* (EOL) markers. The combination of a carriage return followed immediately by a line feed is treated as one EOL marker. For the most part, EOL markers are treated the same as any other white-space characters. However, there are certain instances in which an EOL marker is required or recommended—that is, the following token must appear at the beginning of a line.

Note: *The examples in this book illustrate a recommended convention for arranging tokens into lines. However, the examples' use of white space for indentation is purely for clarity of exposition and is not recommended for practical use.*

The *delimiter characters* (,), <, >, [,], {, }, /, and % are special. They delimit syntactic entities such as strings, arrays, names, and comments. Any of these characters terminates the entity preceding it and is not included in the entity.

All characters besides the white-space characters and delimiters are referred to as *regular characters*. These include 8-bit binary characters that are outside the ASCII character set. A sequence of consecutive regular characters comprises a single token.

Note: *PDF is case-sensitive; corresponding uppercase and lowercase letters are considered distinct.*

3.1.2 Comments

Any occurrence of the percent sign character (%) outside a string or stream introduces a *comment*. The comment consists of all characters between the percent sign and the end of the line, including regular, delimiter, space, and tab characters. PDF ignores comments, treating them as if they were single white-space characters. That is, a comment separates the token preceding it from the one following; thus the PDF fragment

```
abc% comment {/%} blah blah blah
123
```

is syntactically equivalent to just the tokens `abc` and `123`.

Comments (other than the `%PDF-1.4` and `%%EOF` comments described in Section 3.4, “File Structure”) have no semantics. They are not necessarily preserved by applications that edit PDF files (see implementation note 2 in Appendix H). In particular, there is no PDF equivalent of the PostScript document structuring conventions (DSC).

3.2 Objects

PDF supports eight basic types of object:

- Boolean values
- Integer and real numbers
- Strings
- Names
- Arrays
- Dictionaries
- Streams
- The null object

Objects may be labeled so that they can be referred to by other objects. A labeled object is called an *indirect object*.

The following sections describe each object type, as well as how to create and refer to indirect objects.

3.2.1 Boolean Objects

PDF provides *boolean objects* identified by the keywords **true** and **false**. Boolean objects can be used as the values of array elements and dictionary entries, and can also occur in PostScript calculator functions as the results of boolean and relational operators and as operands to the conditional operators **if** and **ifelse** (see Section 3.9.4, “Type 4 (PostScript Calculator) Functions”).

3.2.2 Numeric Objects

PDF provides two types of numeric object: integer and real. *Integer objects* represent mathematical integers within a certain interval centered at 0. *Real objects* approximate mathematical real numbers, but with limited range and precision; they are typically represented in fixed-point, rather than floating-point, form. The range and precision of numbers are limited by the internal representations used in the machine on which the PDF viewer application is running; Appendix C gives these limits for typical implementations.

An integer is written one or more decimal digits optionally preceded by a sign:

```
123 43445 +17 -98 0
```

The value is interpreted as a signed decimal integer and is converted to an integer object. If it exceeds the implementation limit for integers, it is converted to a real object.

A real value is written as one or more decimal digits with an optional sign and a leading, trailing, or embedded period (decimal point):

```
34.5 -3.62 +123.6 4. -.002 0.0
```

The value is interpreted as a real number and is converted to a real object. If it exceeds the implementation limit for real numbers, an error occurs.

Note: PDF does not support the PostScript syntax for numbers with nondecimal radices (such as `16#FFFE`) or in exponential format (such as `6.02E23`).

Throughout this book, the term *number* refers to an object whose type may be either integer or real. Wherever a real number is expected, an integer may be used instead and will be automatically converted to an equivalent real value. For example, it is not necessary to write the number 1.0 in real format; the integer 1 will suffice.

3.2.3 String Objects

A *string object* consists of a series of bytes—unsigned integer values in the range 0 to 255. The string elements are not integer objects, but are stored in a more compact format. The length of a string is subject to an implementation limit; see Appendix C.

There are two conventions, described in the following sections, for writing a string object in PDF:

- As a sequence of literal characters enclosed in parentheses ()
- As hexadecimal data enclosed in angle brackets < >

This section describes only the basic syntax for writing a string as a sequence of bytes. Strings can be used for many purposes and can be formatted in a variety of ways. When a string is used for a specific purpose (to represent a date, for example), it is useful to have a standard format for that purpose (see Section 3.8.2, “Dates”). Such formats are merely conventions for interpreting the contents of a string and are not in themselves separate object types. The use of a particular format is described with the definition of the string object that uses that format.

Literal Strings

A *literal string* is written as an arbitrary number of characters enclosed in parentheses. Any characters may appear in a string except unbalanced parentheses and the backslash, which must be treated specially. Balanced pairs of parentheses within a string require no special treatment.

The following are valid literal strings:

```
(This is a string)
(Strings may contain newlines
and such.)
```

(Strings may contain balanced parentheses () and special characters (*!&}^% and so on).)
 (The following is an empty string.)
 ()
 (It has zero (0) length.)

Within a literal string, the backslash (\) is used as an escape character for various purposes, such as to include newline characters, nonprinting ASCII characters, unbalanced parentheses, or the backslash character itself in the string. The character immediately following the backslash determines its precise interpretation (see Table 3.2). If the character following the backslash is not one of those shown in the table, the backslash is ignored.

TABLE 3.2 Escape sequences in literal strings

SEQUENCE	MEANING
<code>\n</code>	Line feed (LF)
<code>\r</code>	Carriage return (CR)
<code>\t</code>	Horizontal tab (HT)
<code>\b</code>	Backspace (BS)
<code>\f</code>	Form feed (FF)
<code>\(</code>	Left parenthesis
<code>\)</code>	Right parenthesis
<code>\\</code>	Backslash
<code>\ddd</code>	Character code <i>ddd</i> (octal)

If a string is too long to be conveniently placed on a single line, it may be split across multiple lines by using the backslash character at the end of a line to indicate that the string continues on the following line. The backslash and the end-of-line marker following it are not considered part of the string. For example:

(These \
 two strings \
 are the same.)
 (These two strings are the same.)

If an end-of-line marker appears within a literal string without a preceding backslash, the result is equivalent to `\n` (regardless of whether the end-of-line marker itself was a carriage return, a line feed, or both). For example:

```
(This string has an end-of-line at the end of it.  
)  
(So does this one.\n)
```

The `\ddd` escape sequence provides a way to represent characters outside the printable ASCII character set. For example:

```
(This string contains \245two octal characters\307.)
```

The number *ddd* may consist of one, two, or three octal digits, with high-order overflow ignored. It is required that three octal digits be used, with leading zeros as needed, if the next character of the string is also a digit. For example, the literal

```
(\0053)
```

denotes a string containing two characters, `\005` (Control-E) followed by the digit 3, whereas both

```
(\053)
```

and

```
(\53)
```

denote strings containing the single character `\053`, a plus sign (+).

This notation provides a way to specify characters outside the 7-bit ASCII character set using ASCII characters only. However, any 8-bit value may appear in a string. In particular, when a document is encrypted (see Section 3.5, “Encryption”), all of its strings are encrypted and often contain arbitrary 8-bit values. Note that the backslash character is still required as an escape to specify unbalanced parentheses or the backslash character itself.

Hexadecimal Strings

Strings may also be written in hexadecimal form; this is useful for including arbitrary binary data in a PDF file. A hexadecimal string is written as a sequence of hexadecimal digits (0–9 and either A–F or a–f) enclosed within angle brackets (< and >):

```
<4E6F762073686D6F7A206B6120706F702E>
```

Each pair of hexadecimal digits defines one byte of the string. White-space characters (such as space, tab, carriage return, line feed, and form feed) are ignored.

If the final digit of a hexadecimal string is missing—that is, if there is an odd number of digits—the final digit is assumed to be 0. For example,

```
<901FA3>
```

is a 3-byte string consisting of the characters whose hexadecimal codes are 90, 1F, and A3, but

```
<901FA>
```

is a 3-byte string containing the characters whose hexadecimal codes are 90, 1F, and A0.

3.2.4 Name Objects

A *name object* is an atomic symbol uniquely defined by a sequence of characters. *Uniquely defined* means that any two name objects made up of the same sequence of characters are identically the same object. *Atomic* means that a name has no internal structure; although it is defined by a sequence of characters, those characters are not “elements” of the name.

A slash character (/) introduces a name. The slash is not part of the name itself, but a prefix indicating that the following sequence of characters constitutes a name. There can be no white-space characters between the slash and the first character in the name. The name may include any regular characters, but not delimiter or white-space characters (see Section 3.1, “Lexical Conventions”).

Uppercase and lowercase letters are considered distinct: /A and /a are different names. The following are examples of valid literal names:

```
/Name1
/ASomewhatLongerName
/A;Name_With~Various***Characters?
/1.2
/$$
/@pattern
/.notdef
```

Note: The token / (a slash followed by no regular characters) is a valid name.

Beginning with PDF 1.2, any character except null (character code 0) may be included in a name by writing its 2-digit hexadecimal code, preceded by the number sign character (#); see implementation notes 3 and 4 in Appendix H. This syntax is required in order to represent any of the delimiter or white-space characters or the number sign character itself; it is recommended but not required for characters whose codes are outside the range 33 (!) to 126 (~). The examples shown in Table 3.3 are valid literal names in PDF 1.2 and higher.

TABLE 3.3 Examples of literal names using the # character

LITERAL NAME	RESULT
/Adobe#20Green	Adobe Green
/PANTONE#205757#20CV	PANTONE 5757 CV
/paired#28#29parentheses	paired()parentheses
/The_Key_of_F#23_Minor	The_Key_of_F#_Minor
/A#42	AB

The length of a name is subject to an implementation limit; see Appendix C. The limit applies to the number of characters in the name's internal representation. For example, the name /A#20B has four characters (/ , A , space , B), not six.

As stated above, name objects are treated as atomic symbols within a PDF file. Ordinarily, the bytes making up the name are never treated as text to be presented to a human user or to an application external to a PDF viewer. However, occa-

sionally the need arises to treat a name object as text, such as one that represents a font name (see the **BaseFont** entry in Table 5.8 on page 317) or a structure type (see Section 9.6.2, “Structure Types”).

In such situations, it is recommended that the sequence of bytes (after expansion of # sequences, if any) be interpreted according to UTF-8, a variable-length byte-encoded representation of Unicode in which the printable ASCII characters have the same representations as in ASCII. This enables a name object to represent text in any natural language, subject to the implementation limit on the length of a name. (See implementation note 5 in Appendix H.)

***Note:** PDF does not prescribe what UTF-8 sequence to choose for representing any given piece of externally specified text as a name object. In some cases, there are multiple UTF-8 sequences that could represent the same logical text. Name objects defined by different sequences of bytes constitute distinct name objects in PDF, even though the UTF-8 sequences might have identical external interpretations.*

In PDF, name objects always begin with the slash character (/), unlike keywords such as **true**, **false**, and **obj**. This book follows a typographic convention of writing names without the leading slash when they appear in running text and tables. For example, **Type** and **FullScreen** denote names that would actually be written in a PDF file (and in code examples in this book) as `/Type` and `/FullScreen`.

3.2.5 Array Objects

An *array object* is a one-dimensional collection of objects arranged sequentially. Unlike arrays in many other computer languages, PDF arrays may be heterogeneous; that is, an array’s elements may be any combination of numbers, strings, dictionaries, or any other objects, including other arrays. The number of elements in an array is subject to an implementation limit; see Appendix C.

An array is written as a sequence of objects enclosed in square brackets ([and]):

```
[549 3.14 false (Ralph) /SomeName]
```

PDF directly supports only one-dimensional arrays. Arrays of higher dimension can be constructed by using arrays as elements of arrays, nested to any depth.

3.2.6 Dictionary Objects

A *dictionary object* is an associative table containing pairs of objects, known as the dictionary's *entries*. The first element of each entry is the *key* and the second element is the *value*. The key must be a name (unlike dictionary keys in PostScript, which may be objects of any type). The value can be any kind of object, including another dictionary. A dictionary entry whose value is **null** (see Section 3.2.8, "The Null Object") is equivalent to an absent entry. (Note that this differs from PostScript, where **null** behaves like any other object as the value of a dictionary entry.) The number of entries in a dictionary is subject to an implementation limit; see Appendix C.

Note: *No two entries in the same dictionary should have the same key. If a key does appear more than once, its value is undefined.*

A dictionary is written as a sequence of key-value pairs enclosed in double angle brackets (<<...>>). For example:

```
<< /Type /Example
  /Subtype /DictionaryExample
  /Version 0.01
  /IntegerItem 12
  /StringItem (a string)
  /Subdictionary << /Item1 0.4
                    /Item2 true
                    /LastItem (not!)
                    /VeryLastItem (OK)
  >>
>>
```

Note: *Do not confuse the double angle brackets with single angle brackets (< and >), which delimit a hexadecimal string (see "Hexadecimal Strings" on page 32).*

Dictionary objects are the main building blocks of a PDF document. They are commonly used to collect and tie together the attributes of a complex object, such as a font or a page of the document, with each entry in the dictionary specifying the name and value of an attribute. By convention, the **Type** entry of such a dictionary identifies the type of object the dictionary describes. In some cases, a **Subtype** entry (sometimes abbreviated **S**) is used to further identify a specialized subcategory of the general type. The value of the **Type** or **Subtype** entry is always

a name. For example, in a font dictionary, the value of the **Type** entry is always **Font**, whereas that of the **Subtype** entry may be **Type1**, **TrueType**, or one of several other values.

The value of the **Type** entry can almost always be inferred from context. The operand of the **Tf** operator, for example, must be a font object, so the **Type** entry in a font dictionary serves primarily as documentation and as information for error checking. The **Type** entry is not required unless so stated in its description; however, if the entry is present, it must have the correct value. In addition, the value of the **Type** entry in any dictionary, even in private data, must be either a name defined in this book or a registered name; see Appendix E for details.

3.2.7 Stream Objects

A *stream object*, like a string object, is a sequence of bytes. However, a PDF application can read a stream incrementally, while a string must be read in its entirety. Furthermore, a stream can be of unlimited length, whereas a string is subject to an implementation limit. For this reason, objects with potentially large amounts of data, such as images and page descriptions, are represented as streams.

Note: As with strings, this section describes only the syntax for writing a stream as a sequence of bytes. What those bytes represent is determined by the context in which the stream is referenced.

A stream consists of a dictionary that describes a sequence of bytes, followed by zero or more lines of bytes bracketed between the keywords **stream** and **endstream**:

```
dictionary
stream
...Zero or more lines of bytes...
endstream
```

All streams must be indirect objects (see Section 3.2.9, “Indirect Objects”) and the stream dictionary must be a direct object. The keyword **stream** that follows the stream dictionary should be followed by either a carriage return and a line feed or by just a line feed, and not by a carriage return alone. The sequence of bytes that make up a stream lie between the **stream** and **endstream** keywords; the

stream dictionary specifies the exact number of bytes. Alternatively, in PDF 1.2 the bytes may be contained in an external file, in which case the stream dictionary specifies the file and any bytes between **stream** and **endstream** are ignored. (See implementation note 6 in Appendix H.)

Note: Without the restriction against following the keyword **stream** by a carriage return alone, it would be impossible to differentiate a stream that uses carriage return as its end-of-line marker and has a line feed as its first byte of data from one that uses a carriage return–line feed sequence to denote end-of-line.

Table 3.4 lists the entries common to all stream dictionaries; certain types of stream may have additional dictionary entries, as indicated where those streams are described. The optional entries regarding *filters* for the stream indicate whether and how the data in the stream must be transformed (“decoded”) before it is used. Filters are described further in Section 3.3, “Filters.”

Stream Extent

Every stream dictionary has a **Length** entry that indicates how many bytes of the PDF file are used for the stream’s data. (If the stream has a filter, **Length** is the number of bytes of *encoded* data.) In addition, most filters are defined so that the data is self-limiting; that is, they use an encoding scheme in which an explicit *end-of-data* (EOD) marker delimits the extent of the data. Finally, streams are used to represent many objects from whose attributes a length can be inferred. *All of these constraints must be consistent.*

For example, an image with 10 rows and 20 columns, using a single color component and 8 bits per component, requires exactly 200 bytes of image data. If the stream uses a filter, there must be enough bytes of encoded data in the PDF file to produce those 200 bytes. An error occurs if **Length** is too small, if an explicit EOD marker occurs too soon, or if the decoded data does not contain 200 bytes.

It is also an error if the stream contains too *much* data, with the exception that there may be an extra end-of-line marker in the PDF file before the keyword **endstream**.

TABLE 3.4 Entries common to all stream dictionaries

KEY	TYPE	VALUE
Length	integer	<i>(Required)</i> The number of bytes from the beginning of the line following the keyword stream to the last byte just before the keyword endstream . (There may be an additional EOL marker, preceding endstream , that is not included in the count and is not logically part of the stream data.) See “Stream Extent,” above, for further discussion.
Filter	name or array	<i>(Optional)</i> The name of a filter to be applied in processing the stream data found between the keywords stream and endstream , or an array of such names. Multiple filters should be specified in the order in which they are to be applied.
DecodeParms	dictionary or array	<i>(Optional)</i> A parameter dictionary, or an array of such dictionaries, used by the filters specified by Filter . If there is only one filter and that filter has parameters, DecodeParms must be set to the filter’s parameter dictionary unless all the filter’s parameters have their default values, in which case the DecodeParms entry may be omitted. If there are multiple filters and any of the filters has parameters set to non-default values, DecodeParms must be an array with one entry for each filter: either the parameter dictionary for that filter, or the null object if that filter has no parameters (or if all of its parameters have their default values). If none of the filters have parameters, or if all their parameters have default values, the DecodeParms entry may be omitted. (See implementation note 7 in Appendix H.)
F	file specification	<i>(Optional; PDF 1.2)</i> The file containing the stream data. If this entry is present, the bytes between stream and endstream are ignored, the filters are specified by FFilter rather than Filter , and the filter parameters are specified by FDecodeParms rather than DecodeParms . However, the Length entry should still specify the number of those bytes. (Usually there are no bytes and Length is 0.)
FFilter	name or array	<i>(Optional; PDF 1.2)</i> The name of a filter to be applied in processing the data found in the stream’s external file, or an array of such names. The same rules apply as for Filter .
FDecodeParms	dictionary or array	<i>(Optional; PDF 1.2)</i> A parameter dictionary, or an array of such dictionaries, used by the filters specified by FFilter . The same rules apply as for DecodeParms .

3.2.8 Null Object

The *null object* has a type and value that are unequal to those of any other object. There is only one object of type null, denoted by the keyword **null**. An indirect object reference (see Section 3.2.9, “Indirect Objects”) to a nonexistent object is treated the same as a null object; specifying the null object as the value of a dictionary entry (Section 3.2.6, “Dictionary Objects”) is equivalent to omitting the entry entirely.

3.2.9 Indirect Objects

Any object in a PDF file may be labeled as an *indirect object*. This gives the object a unique *object identifier* by which other objects can refer to it (for example, as an element of an array or as the value of a dictionary entry). The object identifier consists of two parts:

- A positive integer *object number*. Indirect objects are often numbered sequentially within a PDF file, but this is not required; object numbers may be assigned in any arbitrary order.
- A nonnegative integer *generation number*. In a newly created file, all indirect objects have generation numbers of 0. Nonzero generation numbers may be introduced when the file is later updated; see Sections 3.4.3, “Cross-Reference Table,” and 3.4.5, “Incremental Updates.”

Together, the combination of an object number and a generation number uniquely identifies an indirect object. The object retains the same object number and generation number throughout its existence, even if its value is modified.

The definition of an indirect object in a PDF file consists of its object number and generation number, followed by the value of the object itself bracketed between the keywords **obj** and **endobj**. For example, the definition

```
12 0 obj
  (Brillig)
endobj
```

defines an indirect string object with an object number of 12, a generation number of 0, and the value Brillig. The object can then be referred to from elsewhere in

the file by an *indirect reference* consisting of the object number, the generation number, and the keyword **R**:

```
12 0 R
```

An indirect reference to an undefined object is not an error; it is simply treated as a reference to the null object. For example, if a file contains the indirect reference

```
17 0 R
```

but does not contain the corresponding definition

```
17 0 obj
...
endobj
```

then the indirect reference is considered to refer to the null object.

Note: *In the data structures that make up a PDF document, certain values are required to be specified as indirect object references. Except where this is explicitly called out, any object (other than a stream) may be specified either directly or as an indirect object reference; the semantics are entirely equivalent. Note in particular that content streams, which define the visible contents of the document, may not contain indirect references (see Section 3.7.1, “Content Streams”).*

Example 3.1 shows the use of an indirect object to specify the length of a stream. The value of the stream’s **Length** entry is an integer object that follows the stream itself in the file. This allows applications that generate PDF in a single pass to defer specifying the stream’s length until after its contents have been generated.

Example 3.1

```
7 0 obj
  << /Length 8 0 R >>                % An indirect reference to object 8
stream
BT
  /F1 12 Tf
  72 712 Td
  (A stream with an indirect length) Tj
ET
endstream
endobj
```



```

8 0 obj
  77                               % The length of the preceding stream
endobj

```

3.3 Filters

Stream filters are introduced in Section 3.2.7, “Stream Objects.” A *filter* is an optional part of the specification of a stream, indicating how the data in the stream must be decoded before it is used. For example, if a stream has an **ASCIIHexDecode** filter, an application reading the data in that stream will transform the ASCII hexadecimal-encoded data in the stream into binary data.

An application program that produces a PDF file can encode certain information (for example, data for sampled images) to compress it or to convert it to a portable ASCII representation. Then an application that reads (“consumes”) the PDF file can invoke the corresponding decoding filter to convert the information back to its original form.

The filter or filters for a stream are specified by the **Filter** entry in the stream’s dictionary (or the **FFilter** entry if the stream is external). Filters can be cascaded to form a *pipeline* that passes the stream through two or more decoding transformations in sequence. For example, data encoded using LZW and ASCII base-85 encoding (in that order) can be decoded using the following entry in the stream dictionary:

```
/Filter [/ASCII85Decode /LZWDecode]
```

Some filters may take parameters to control how they operate. These optional parameters are specified by the **DecodeParms** entry in the stream’s dictionary (or the **FDecodeParms** entry if the stream is external).

PDF supports a standard set of filters that fall into two main categories:

- *ASCII filters* enable decoding of arbitrary 8-bit binary data that has been encoded as ASCII text. (See Section 3.1, “Lexical Conventions,” for an explanation of why this type of encoding might be useful.) Note that ASCII filters serve no useful purpose in a PDF file that is encrypted; see Section 3.5, “Encryption.”
- *Decompression filters* enable decoding of data that has been compressed. The compressed data is always in 8-bit binary format, even if the original data happens to be ASCII text. (Compression is particularly valuable for large sampled

images, since it reduces storage requirements and transmission time. Note that some types of compression are *lossy*, meaning that some data is lost during the encoding, resulting in a loss of quality when the data is decompressed; compression in which no loss of data occurs is called *lossless*.)

The standard filters are summarized in Table 3.5, which also indicates whether they accept any optional parameters. The following sections describe these filters and their parameters (if any) in greater detail, including specifications of encoding algorithms for some filters. (See also implementation notes 8 and 9 in Appendix H.)

TABLE 3.5 Standard filters

FILTER NAME	PARAMETERS?	DESCRIPTION
ASCIHexDecode	no	Decodes data encoded in an ASCII hexadecimal representation, reproducing the original binary data.
ASCII85Decode	no	Decodes data encoded in an ASCII base-85 representation, reproducing the original binary data.
LZWDecode	yes	Decompresses data encoded using the LZW (Lempel-Ziv-Welch) adaptive compression method, reproducing the original text or binary data.
FlateDecode	yes	(<i>PDF 1.2</i>) Decompresses data encoded using the zlib/deflate compression method, reproducing the original text or binary data.
RunLengthDecode	no	Decompresses data encoded using a byte-oriented run-length encoding algorithm, reproducing the original text or binary data (typically monochrome image data, or any data that contains frequent long runs of a single byte value).
CCITTFaxDecode	yes	Decompresses data encoded using the CCITT facsimile standard, reproducing the original data (typically monochrome image data at 1 bit per pixel).
JBIG2Decode	yes	(<i>PDF 1.4</i>) Decompresses data encoded using the JBIG2 standard, reproducing the original monochrome (1 bit per pixel) image data (or an approximation of that data).
DCTDecode	yes	Decompresses data encoded using a DCT (discrete cosine transform) technique based on the JPEG standard, reproducing image sample data that approximates the original data.

Example 3.2 shows a stream, containing the marking instructions for a page, that was compressed using the LZW compression method and then encoded in ASCII base-85 representation. Example 3.3 shows the same stream without any encoding. (The stream's contents are explained in Section 3.7.1, "Content Streams," and the operators used there are further described in Chapter 5.)

Example 3.2

```

1 0 obj
  << /Length 534
    /Filter [/ASCII85Decode /LZWDecode]
  >>
stream
J.)6T`?p&<!J9%_[umg"B7/Z7KNXbN'S+,*Q/&"OLT'F
LIDK#!n`$"<Atdi`\Vn%b%)&'cA*VnK\CJY(sF>c!Jnl@
RM]WM;jjH6Gnc75idkL5]+cPZKEBPWdr>FF(kj1_R%W_d
&/jS!;iuad7h?[L-F$+]]0A3Ck*$I0KZ?;<)CJtqi65Xb
Vc3\n5ua:Q/=0$W<#N3U;H,MQKqfg1?!UpR;6oN[C2E4
ZNr8Udn.'p+?#X+1>0Kuk$bCDF/(3fL5]Oq)^kZ!C2H1
'TO]RI?Q:&'<5&iP!$Rq;BXRecDN[IJB`,)o8XJOSJ9sD
S]hQ;Rj@!ND)bD_q&C\g:inYC%)&u#:u,M6Bm%!Y!Kb1+
":aAa'S`ViJgLLb8<W9k6YI\0McJQkDeLWdPN?9A'jX*
al>iG1p&i;eVoK&juJHs9%;Xomop"5KatWRT"JQ#qYUL,
JD?M$0QP)IKn06l1apKDC@\qJ4B!!(5m+j.7F790m(Vj8
8l8Q:_CZ(Gm1%X\N1&u!FKHMB~>
endstream
endobj

```

Example 3.3

```

1 0 obj
  << /Length 568 >>
stream
2 J
BT
/F1 12 Tf
0 Tc
0 Tw
72.5 712 TD
[(Unencoded streams can be read easily) 65 (,)] TJ
0 -14 TD
[(b) 20 (ut generally tak) 10 (e more space than \311)] TJ
T* (encoded streams.) Tj
0 -28 TD

```

```

[(Se) 25 (v) 15 (eral encoding methods are a) 20 (v) 25 (ailable in PDF) 80 (.)] Tj
0 -14 TD
(Some are used for compression and others simply) Tj
T* [(to represent binary data in an) 55 (ASCII format.)] Tj
T* (Some of the compression encoding methods are \
suitable) Tj
T* (for both data and images, while others are \
suitable only) Tj
T* (for continuous-tone images.) Tj
ET
endstream
endobj

```

3.3.1 ASCIIHexDecode Filter

The **ASCIIHexDecode** filter decodes data that has been encoded in ASCII hexadecimal form. ASCII hexadecimal encoding and ASCII base-85 encoding (described in the next section) convert binary data, such as image data, to 7-bit ASCII characters. In general, ASCII base-85 encoding is preferred to ASCII hexadecimal encoding because it is more compact: it expands the data by a factor of 4:5, compared with 1:2 for ASCII hexadecimal encoding.

For each pair of ASCII hexadecimal digits (0–9 and A–F or a–f), the **ASCIIHexDecode** filter produces one byte of binary data. All white-space characters (see Section 3.1, “Lexical Conventions”) are ignored. A right angle bracket character (>) indicates EOD. Any other characters will cause an error. If the filter encounters the EOD marker after reading an odd number of hexadecimal digits, it will behave as if a 0 followed the last digit.

3.3.2 ASCII85Decode Filter

The **ASCII85Decode** filter decodes data that has been encoded in ASCII base-85 encoding and produces binary data. The following paragraphs describe the process for encoding binary data in ASCII base-85; the **ASCII85Decode** filter reverses this process.

The ASCII base-85 encoding uses the characters ! through u and the character z, with the 2-character sequence ~> as its EOD marker. The **ASCII85Decode** filter ignores all white-space characters (see Section 3.1, “Lexical Conventions”). Any

other characters, and any character sequences that represent impossible combinations in the ASCII base-85 encoding, will cause an error.

Specifically, ASCII base-85 encoding produces 5 ASCII characters for every 4 bytes of binary data. Each group of 4 binary input bytes, $(b_1 b_2 b_3 b_4)$, is converted to a group of 5 output bytes, $(c_1 c_2 c_3 c_4 c_5)$, using the relation

$$(b_1 \times 256^3) + (b_2 \times 256^2) + (b_3 \times 256^1) + b_4 = \\ (c_1 \times 85^4) + (c_2 \times 85^3) + (c_3 \times 85^2) + (c_4 \times 85^1) + c_5$$

In other words, 4 bytes of binary data are interpreted as a base-256 number and then converted into a base-85 number. The five “digits” of the base-85 number are then converted to ASCII characters by adding 33 (the ASCII code for the character !) to each. The resulting encoded data contains only printable ASCII characters with codes in the range 33 (!) to 117 (u). As a special case, if all five digits are 0, they are represented by the character with code 122 (z) instead of by five exclamation points (!!!!!).

If the length of the binary data to be encoded is not a multiple of 4 bytes, the last, partial group of 4 is used to produce a last, partial group of 5 output characters. Given n (1, 2, or 3) bytes of binary data, the encoder first appends $4 - n$ zero bytes to make a complete group of 4. It then encodes this group in the usual way, but without applying the special z case. Finally, it writes only the first $n + 1$ characters of the resulting group of 5. These characters are immediately followed by the ~> EOD marker.

The following conditions (which never occur in a correctly encoded byte sequence) will cause errors during decoding:

- The value represented by a group of 5 characters is greater than $2^{32} - 1$.
- A z character occurs in the middle of a group.
- A final partial group contains only one character.

3.3.3 LZWDecode and FlateDecode Filters

The **LZWDecode** and (in PDF 1.2) **FlateDecode** filters have much in common and so are discussed together in this section. They decode data that has been encoded using the LZW or Flate data compression method, respectively.

- LZW (Lempel-Ziv-Welch) is a variable-length, adaptive compression method that has been adopted as one of the standard compression methods in the *Tag Image File Format* (TIFF) standard. Details on LZW encoding follow in the next section.
- The Flate method is based on the public-domain zlib/deflate compression method, which is a variable-length Lempel-Ziv adaptive compression method cascaded with adaptive Huffman coding. It is fully defined in Internet RFCs 1950, *ZLIB Compressed Data Format Specification*, and 1951, *DEFLATE Compressed Data Format Specification* (see the Bibliography).

Both of these methods compress either binary data or ASCII text but (like all compression methods) always produce binary data, even if the original data was text.

The LZW and Flate compression methods can discover and exploit many patterns in the input data, whether the data is text or images. As described later, both filters support optional transformation by a *predictor function*, which improves the compression of sampled image data. Thanks to its cascaded adaptive Huffman coding, Flate-encoded output is usually much more compact than LZW-encoded output for the same input. Flate and LZW decoding speeds are comparable, but Flate encoding is considerably slower than LZW encoding.

Usually, both Flate and LZW encodings compress their input substantially. However, in the worst case (in which no pair of adjacent characters appears twice), Flate encoding *expands* its input by no more than 11 bytes or a factor of 1.003 (whichever is larger), plus the effects of algorithm tags added by PNG predictors. For LZW encoding, the best case (all zeros) provides a compression approaching 1365:1 for long files, but the worst-case expansion is at least a factor of 1.125, which can increase to nearly 1.5 in some implementations (plus the effects of PNG tags as with Flate encoding).

Details of LZW Encoding

Data encoded using the LZW compression method consists of a sequence of codes that are 9 to 12 bits long. Each code represents a single character of input data (0–255), a clear-table marker (256), an EOD marker (257), or a table entry representing a multiple-character sequence that has been encountered previously in the input (258 or greater).

Initially, the code length is 9 bits and the LZW table contains only entries for the 258 fixed codes. As encoding proceeds, entries are appended to the table, associating new codes with longer and longer sequences of input characters. The encoder and the decoder maintain identical copies of this table.

Whenever both the encoder and the decoder independently (but synchronously) realize that the current code length is no longer sufficient to represent the number of entries in the table, they increase the number of bits per code by 1. The first output code that is 10 bits long is the one following the creation of table entry 511, and similarly for 11 (1023) and 12 (2047) bits. Codes are never longer than 12 bits, so entry 4095 is the last entry of the LZW table.

The encoder executes the following sequence of steps to generate each output code:

1. Accumulate a sequence of one or more input characters matching a sequence already present in the table. For maximum compression, the encoder looks for the longest such sequence.
2. Emit the code corresponding to that sequence.
3. Create a new table entry for the first unused code. Its value is the sequence found in step 1 followed by the next input character.

For example, suppose the input consists of the following sequence of ASCII character codes:

45 45 45 45 45 65 45 45 45 66

Starting with an empty table, the encoder proceeds as shown in Table 3.6.

Codes are packed into a continuous bit stream, high-order bit first. This stream is then divided into 8-bit bytes, high-order bit first. Thus, codes can straddle byte boundaries arbitrarily. After the EOD marker (code value 257), any leftover bits in the final byte are set to 0.

TABLE 3.6 Typical LZW encoding sequence

INPUT SEQUENCE	OUTPUT CODE	CODE ADDED TO TABLE	SEQUENCE REPRESENTED BY NEW CODE
–	256 (clear-table)	–	–
45	45	258	45 45
45 45	258	259	45 45 45
45 45	258	260	45 45 65
65	65	261	65 45
45 45 45	259	262	45 45 45 66
66	66	–	–
–	257 (EOD)	–	–

In the example above, all the output codes are 9 bits long; they would pack into bytes as follows (represented in hexadecimal):

```
80 0B 60 50 22 0C 0C 85 01
```

To adapt to changing input sequences, the encoder may at any point issue a clear-table code, which causes both the encoder and the decoder to restart with initial tables and a 9-bit code length. By convention, the encoder begins by issuing a clear-table code. It must issue a clear-table code when the table becomes full; it may do so sooner.

Note: *The LZW compression method is the subject of U.S. patent number 4,558,302 and corresponding foreign patents owned by the Unisys Corporation. Adobe Systems has licensed this patent for use in its Acrobat products; however, independent software vendors (ISVs) may be required to license this patent directly from Unisys to develop software that uses the LZW method to compress data in PDF files. For information on Unisys licensing policies, send e-mail to <lwz_info@unisys.com>, or visit the Unisys Web site at <<http://www.unisys.com>>.*

LZWDecode and FlateDecode Parameters

The **LZWDecode** and **FlateDecode** filters accept optional parameters to control the decoding process. Most of these parameters are related to techniques that re-

duce the size of compressed sampled images (rectangular arrays of color values, described in Section 4.8, “Images”). For example, image data frequently changes very little from sample to sample; subtracting the values of adjacent samples (a process called *differencing*), and encoding the differences rather than the raw sample values, can reduce the size of the output data. Furthermore, when the image data contains several color components (red-green-blue or cyan-magenta-yellow-black) per sample, taking the difference between the values of corresponding components in adjacent samples, rather than between different color components in the same sample, often reduces the output data size.

Table 3.7 shows the parameters that can optionally be specified for **LZWDecode** and **FlateDecode** filters. Except where otherwise noted, all values supplied to the decoding filter for any optional parameters must match those used when the data was encoded.

TABLE 3.7 Optional parameters for LZWDecode and FlateDecode filters

KEY	TYPE	VALUE
Predictor	integer	A code that selects the predictor algorithm, if any. If the value of this entry is 1, the filter assumes that the normal algorithm was used to encode the data, without prediction. If the value is greater than 1, the filter assumes that the data was differenced before being encoded, and Predictor selects the predictor algorithm. For more information regarding Predictor values greater than 1, see “LZW and Flate Predictor Functions,” below. Default value: 1.
Colors	integer	<i>(Used only if Predictor is greater than 1)</i> The number of interleaved color components per sample. Valid values are 1 to 4 in PDF 1.2 or earlier, and 1 or greater in PDF 1.3 or later. Default value: 1.
BitsPerComponent	integer	<i>(Used only if Predictor is greater than 1)</i> The number of bits used to represent each color component in a sample. Valid values are 1, 2, 4, and 8. Default value: 8.
Columns	integer	<i>(Used only if Predictor is greater than 1)</i> The number of samples in each row. Default value: 1.
EarlyChange	integer	<i>(LZWDecode only)</i> An indication of when to increase the code length. If the value of this entry is 0, code length increases are postponed as long as possible. If it is 1, they occur one code early. This parameter is included because LZW sample code distributed by some vendors increases the code length one code earlier than necessary. Default value: 1.

LZW and Flate Predictor Functions

LZW and Flate encoding compress more compactly if their input data is highly predictable. One way of increasing the predictability of many continuous-tone sampled images is to replace each sample with the difference between that sample and a *predictor function* applied to earlier neighboring samples. If the predictor function works well, the postprediction data will cluster toward 0.

Two groups of predictor functions are supported. The first, the *TIFF* group, consists of the single function that is Predictor 2 in the TIFF standard. (In the TIFF standard, Predictor 2 applies only to LZW compression, but here it applies to Flate compression as well.) TIFF Predictor 2 predicts that each color component of a sample will be the same as the corresponding color component of the sample immediately to its left.

The second supported group of predictor functions, the *PNG* group, consists of the “filters” of the World Wide Web Consortium’s Portable Network Graphics recommendation, documented in Internet RFC 2083, *PNG (Portable Network Graphics) Specification* (see the Bibliography). The term *predictors* is used here instead of *filters* to avoid confusion. There are five basic PNG predictor algorithms (and a sixth that chooses the optimum predictor function separately for each row):

None	No prediction
Sub	Predicts the same as the sample to the left
Up	Predicts the same as the sample above
Average	Predicts the average of the sample to the left and the sample above
Paeth	A nonlinear function of the sample above, the sample to the left, and the sample to the upper left

The predictor algorithm to be used, if any, is indicated by the **Predictor** filter parameter (see Table 3.7), which can have any of the values listed in Table 3.8.

For **LZWDecode** and **FlateDecode**, a **Predictor** value greater than or equal to 10 merely indicates that a PNG predictor is in use; the specific predictor function used is explicitly encoded in the incoming data. The value of **Predictor** supplied by the decoding filter need not match the value used when the data was encoded if they are both greater than or equal to 10.

TABLE 3.8 Predictor values

VALUE	MEANING
1	No prediction (the default value)
2	TIFF Predictor 2
10	PNG prediction (on encoding, PNG None on all rows)
11	PNG prediction (on encoding, PNG Sub on all rows)
12	PNG prediction (on encoding, PNG Up on all rows)
13	PNG prediction (on encoding, PNG Average on all rows)
14	PNG prediction (on encoding, PNG Paeth on all rows)
15	PNG prediction (on encoding, PNG optimum)

The two groups of predictor functions have some commonalities. Both assume the following:

- Data is presented in order, from the top row to the bottom row and, within a row, from left to right.
- A row occupies a whole number of bytes, rounded up if necessary.
- Samples and their components are packed into bytes from high-order to low-order bits.
- All color components of samples outside the image (which are necessary for predictions near the boundaries) are 0.

The predictor function groups also differ in significant ways:

- The postprediction data for each PNG-predicted row begins with an explicit algorithm tag, so different rows can be predicted with different algorithms to improve compression. TIFF Predictor 2 has no such identifier; the same algorithm applies to all rows.
- The TIFF function group predicts each color component from the prior instance of that component, taking into account the number of bits per component and components per sample. In contrast, the PNG function group predicts each byte of data as a function of the corresponding byte of one or

more previous image samples, regardless of whether there are multiple color components in a byte or whether a single color component spans multiple bytes. This can yield significantly better speed at the cost of somewhat worse compression.

3.3.4 RunLengthDecode Filter

The **RunLengthDecode** filter decodes data that has been encoded in a simple byte-oriented format based on run length. The encoded data is a sequence of *runs*, where each run consists of a *length* byte followed by 1 to 128 bytes of data. If the *length* byte is in the range 0 to 127, the following *length* + 1 (1 to 128) bytes are copied literally during decompression. If *length* is in the range 129 to 255, the following single byte is to be copied $257 - \textit{length}$ (2 to 128) times during decompression. A *length* value of 128 denotes EOD.

The compression achieved by run-length encoding depends on the input data. In the best case (all zeros), a compression of approximately 64:1 is achieved for long files. The worst case (the hexadecimal sequence 00 alternating with FF) results in an expansion of 127:128.

3.3.5 CCITTFaxDecode Filter

The **CCITTFaxDecode** filter decodes image data that has been encoded using either Group 3 or Group 4 CCITT facsimile (fax) encoding. CCITT encoding is designed to achieve efficient compression of monochrome (1 bit per pixel) image data at relatively low resolutions, and so is useful only for bitmap image data, not for color images, grayscale images, or general data.

The CCITT encoding standard is defined by the International Telecommunications Union (ITU), formerly known as the Comité Consultatif International Téléphonique et Télégraphique (International Coordinating Committee for Telephony and Telegraphy). The encoding algorithm is not described in detail here, but can be found in ITU Recommendations T.4 and T.6 (see the Bibliography). For historical reasons, we refer to these documents as the CCITT standard.

CCITT encoding is bit-oriented, not byte-oriented. This means that, in principle, encoded or decoded data might not end at a byte boundary. This problem is dealt with in the following ways:

- Unencoded data is treated as complete scan lines, with unused bits inserted at the end of each scan line to fill out the last byte. This is compatible with the PDF convention for sampled image data.
- Encoded data is ordinarily treated as a continuous, unbroken bit stream. The **EncodedByteAlign** parameter (described in Table 3.9) can be used to cause each encoded scan line to be filled to a byte boundary; although this is not prescribed by the CCITT standard and fax machines never do this, some software packages find it convenient to encode data this way.
- When a filter reaches EOD, it always skips to the next byte boundary following the encoded data.

If the **CCITTFaxDecode** filter encounters improperly encoded source data, an error will occur. The filter will not perform any error correction or resynchronization, except as noted for the **DamagedRowsBeforeError** parameter in Table 3.9.

Table 3.9 lists the optional parameters that can be used to control the decoding. Except where noted otherwise, all values supplied to the decoding filter by any of these parameters must match those used when the data was encoded.

TABLE 3.9 Optional parameters for the CCITTFaxDecode filter

KEY	TYPE	VALUE
K	integer	<p>A code identifying the encoding scheme used:</p> <ul style="list-style-type: none"> <0 Pure two-dimensional encoding (Group 4) 0 Pure one-dimensional encoding (Group 3, 1-D) >0 Mixed one- and two-dimensional encoding (Group 3, 2-D), in which a line encoded one-dimensionally can be followed by at most K – 1 lines encoded two-dimensionally <p>The filter distinguishes among negative, zero, and positive values of K to determine how to interpret the encoded data; however, it does not distinguish between different positive K values. Default value: 0.</p>

EndOfLine	boolean	A flag indicating whether end-of-line bit patterns are required to be present in the encoding. The CCITTFaxDecode filter always accepts end-of-line bit patterns, but requires them only if EndOfLine is true . Default value: false .
EncodedByteAlign	boolean	A flag indicating whether the filter expects extra 0 bits before each encoded line so that the line begins on a byte boundary. If true , the filter skips over encoded bits to begin decoding each line at a byte boundary. If false , the filter does not expect extra bits in the encoded representation. Default value: false .
Columns	integer	The width of the image in pixels. If the value is not a multiple of 8, the filter adjusts the width of the unencoded image to the next multiple of 8, so that each line starts on a byte boundary. Default value: 1728.
Rows	integer	The height of the image in scan lines. If the value is 0 or absent, the image's height is not predetermined, and the encoded data must be terminated by an end-of-block bit pattern or by the end of the filter's data. Default value: 0.
EndOfBlock	boolean	A flag indicating whether the filter expects the encoded data to be terminated by an end-of-block pattern, overriding the Rows parameter. If false , the filter stops when it has decoded the number of lines indicated by Rows or when its data has been exhausted, whichever occurs first. The end-of-block pattern is the CCITT end-of-facsimile-block (EOFB) or return-to-control (RTC) appropriate for the K parameter. Default value: true .
BlackIs1	boolean	A flag indicating whether 1 bits are to be interpreted as black pixels and 0 bits as white pixels, the reverse of the normal PDF convention for image data. Default value: false .
DamagedRowsBeforeError	integer	The number of damaged rows of data to be tolerated before an error occurs. This entry applies only if EndOfLine is true and K is nonnegative. Tolerating a damaged row means locating its end in the encoded data by searching for an EndOfLine pattern and then substituting decoded data from the previous row if the previous row was not damaged, or a white scan line if the previous row was also damaged. Default value: 0.

The compression achieved using CCITT encoding depends on the data, as well as on the value of various optional parameters. For Group 3 one-dimensional encoding, in the best case (all zeros), each scan line compresses to 4 bytes, and the

compression factor depends on the length of a scan line. If the scan line is 300 bytes long, a compression ratio of approximately 75:1 is achieved. The worst case, an image of alternating ones and zeros, produces an expansion of 2:9.

3.3.6 JBIG2Decode Filter

The **JBIG2Decode** filter (*PDF 1.4*) decodes monochrome (1 bit per pixel) image data that has been encoded using JBIG2 encoding. JBIG stands for the Joint Bi-Level Image Experts Group, a group within the International Organization for Standardization (ISO) that developed the format; JBIG2 is the second version of a standard originally released as JBIG1 and was approaching standards approval at the time of publication of this book.

JBIG2 encoding, which provides for both lossy and lossless compression, is useful only for monochrome images, not for color images, grayscale images, or general data. The algorithms used by the encoder, and the details of the format, are not described here; a working draft of the JBIG2 specification can be found through the Web site for the JBIG and JPEG (Joint Photographic Experts Group) committees at <<http://www.jpeg.org>>.

In general, JBIG2 provides considerably better compression than the existing CCITT standard (discussed in Section 3.3.5). The compression it achieves depends strongly on the nature of the image. Images of pages containing text in any language will compress particularly well, with typical compression ratios of 20:1 to 50:1 for a page full of text. The JBIG2 encoder builds a table of unique symbol bitmaps found in the image, and other symbols found later in the image are matched against the table. Matching symbols are replaced by an index into the table, and symbols that fail to match are added to the table. The table itself is compressed using other means. This results in high compression ratios for documents in which the same symbol is repeated often, as is typical for images created by scanning text pages. This method also results in high compression of “white space” in the image, which does not need to be encoded because it contains no symbols.

While best compression is achieved for images of text, the JBIG2 standard also includes algorithms for compressing regions of an image that contain dithered half-tone images (for example, photographs).

The JBIG2 compression method can also be used for encoding multiple images into a single JBIG2 bit stream. Typically, these images will be scanned pages of a multiple-page document. Since a single table of symbol bitmaps is used to match symbols across multiple pages, this type of encoding can result in higher compression ratios than if each of the pages had been individually encoded using JBIG2.

In general, an image may be specified in PDF as either an *image XObject* or an *inline image* (as described in Section 4.8, “Images”); however, the **JBIG2Decode** filter can be applied only to image XObjects.

This filter addresses both single-page and multiple-page JBIG2 bit streams, by representing each JBIG2 “page” as a PDF image, as follows:

- The filter uses the embedded file organization of JBIG2. (The details of this and the other types of file organization are provided in an annex of the ISO specification.) The optional 2-byte combination (marker) mentioned in the specification is not used in PDF. JBIG2 bit streams in random-access organization should be converted to the embedded file organization. Bit streams in sequential organization need no reorganization, except for the mappings described below.
- The JBIG2 file header, end-of-page segments, and end-of-file segment are not used in PDF. These should be removed before the PDF objects described below are created.
- The image XObject to which the **JBIG2Decode** filter is applied contains all segments that are associated with the JBIG2 page represented by that image—that is, all segments whose segment page association field contains the page number of the JBIG2 page represented by the image. In the image XObject, however, the segment’s page number should always be 1—that is, when each such segment is written to the XObject, the value of its segment page association field should be set to 1.
- If the bit stream contains global segments (segments whose segment page association field contains 0), these must be placed in a separate PDF stream, and the filter parameter listed in Table 3.10 should refer to that stream. The stream can be shared by multiple image XObjects whose JBIG2 encodings use the same global segments.

TABLE 3.10 Optional parameter for the JBIG2Decode filter

KEY	TYPE	VALUE
JBIG2Globals	stream	A stream containing the JBIG2 global (page 0) segments. Global segments must be placed in this stream even if only a single JBIG2 image XObject refers to it.

Example 3.4 shows an image that was compressed using the JBIG2 compression method and then encoded in ASCII hexadecimal representation. Since the JBIG2 bit stream contains global segments, these are placed in a separate PDF stream, as indicated by the **JBIG2Globals** filter parameter.

Example 3.4

```

5 0 obj
  << /Type /XObject
    /Subtype /Image
    /Width 52
    /Height 66
    /ColorSpace /DeviceGray
    /BitsPerComponent 1
    /Length 224
    /Filter [/ASCIIHexDecode /JBIG2Decode]
    /DecodeParms [null << /JBIG2Globals 6 0 R >>]
  >>
  stream
  000000013000010000001300000034000000420000000000
  000000400000000000002062000010000001e000000340000
  0042000000000000000000200100000000231db51ce51ffac>
  endstream
endobj

6 0 obj
  << /Length 126
    /Filter /ASCIIHexDecode
  >>
  stream
  0000000000010000000032000003fffdff02fefefe000000
  01000000012ae225aea9a5a538b4d9999c5c8e56ef0f872
  7f2b53d4e37ef795cc5506dffac>
  endstream
endobj

```

The JBIG2 bit stream for this example is as follows:

```
97 4A 42 32 0D 0A 1A 0A 01 00 00 00 01 00 00 00 00 00 01 00 00 00 00 32
00 00 03 FF FD FF 02 FE FE FE 00 00 00 01 00 00 00 01 2A E2 25 AE A9 A5
A5 38 B4 D9 99 9C 5C 8E 56 EF 0F 87 27 F2 B5 3D 4E 37 EF 79 5C C5 50 6D
FF AC 00 00 00 01 30 00 01 00 00 00 13 00 00 00 34 00 00 00 42 00 00 00
00 00 00 00 00 40 00 00 00 00 00 02 06 20 00 01 00 00 00 1E 00 00 00 34
00 00 00 42 00 00 00 00 00 00 02 00 10 00 00 00 02 31 DB 51 CE 51
FF AC 00 00 00 03 31 00 01 00 00 00 00 00 00 00 04 33 01 00 00 00 00
```

This bit stream is made up of the parts listed below (in the order listed).

1. The JBIG2 file header

```
97 4A 42 32 0D 0A 1A 0A 01 00 00 00 01
```

Since the JBIG2 file header is not used in PDF, this header is not placed in the JBIG2 stream object, and is discarded.

2. The first JBIG2 segment (segment 0)—in this case, the symbol dictionary segment

```
00 00 00 00 00 01 00 00 00 00 32 00 00 03 FF FD FF 02 FE FE FE 00 00 00
01 00 00 00 01 2A E2 25 AE A9 A5 A5 38 B4 D9 99 9C 5C 8E 56 EF 0F 87
27 F2 B5 3D 4E 37 EF 79 5C C5 50 6D FF AC
```

This is a global segment (segment page association = 0) and so is placed in the **JBIG2Globals** stream.

3. The page information segment

```
00 00 00 01 30 00 01 00 00 00 13 00 00 00 34 00 00 00 42 00 00 00 00
00 00 00 00 40 00 00
```

and the immediate text region segment

```
00 00 00 02 06 20 00 01 00 00 00 1E 00 00 00 34 00 00 00 42 00 00 00
00 00 00 00 00 02 00 10 00 00 00 02 31 DB 51 CE 51 FF AC
```

These two segments constitute the contents of the JBIG2 page, and are placed in the PDF XObject representing this image.

4. The end-of-page segment

```
00 00 00 03 31 00 01 00 00 00 00
```

and the end-of-file segment

```
00 00 00 04 33 01 00 00 00 00
```

Since these are not used in PDF, they are discarded.

The resulting PDF image object, then, contains the page information segment and the immediate text region segment, and refers to a **JBIG2Globals** stream that contains the symbol dictionary segment.

3.3.7 DCTDecode Filter

The **DCTDecode** filter decodes grayscale or color image data that has been encoded in the JPEG baseline format. (JPEG stands for the Joint Photographic Experts Group, a group within the International Organization for Standardization that developed the format; DCT stands for discrete cosine transform, the primary technique used in the encoding.)

JPEG encoding is a lossy compression method, designed specifically for compression of sampled continuous-tone images and not for general data compression. Data to be encoded using JPEG consists of a stream of image samples, each consisting of one, two, three, or four color components. The color component values for a particular sample must appear consecutively. Each component value occupies an 8-bit byte.

During encoding, several parameters control the algorithm and the information loss. The values of these parameters, which include the dimensions of the image and the number of components per sample, are entirely under the control of the encoder and are stored in the encoded data. **DCTDecode** generally obtains the parameter values it requires directly from the encoded data. However, in one instance, the parameter might not be present in the encoded data but must be specified in the filter parameter dictionary; see Table 3.11.

The details of the encoding algorithm are not presented here but can be found in the ISO specification and in *JPEG: Still Image Data Compression Standard*, by Pennebaker and Mitchell (see the Bibliography). Briefly, the JPEG algorithm breaks an image up into blocks 8 samples wide by 8 high. Each color component in an image is treated separately. A two-dimensional DCT is performed on each block. This operation produces 64 coefficients, which are then quantized. Each coefficient may be quantized with a different step size. It is this quantization that results in the loss of information in the JPEG algorithm. The quantized coefficients are then compressed.

TABLE 3.11 Optional parameter for the DCTDecode filter

KEY	TYPE	VALUE
ColorTransform	integer	<p>A code specifying the transformation to be performed on the sample values:</p> <ul style="list-style-type: none"> 0 No transformation. 1 If the image has three color components, transform <i>RGB</i> values to <i>YUV</i> before encoding and from <i>YUV</i> to <i>RGB</i> after decoding. If the image has four components, transform <i>CMYK</i> values to <i>YUVK</i> before encoding and from <i>YUVK</i> to <i>CMYK</i> after decoding. This option is ignored if the image has one or two color components. <p><i>Note:</i> The <i>RGB</i> and <i>YUV</i> used here have nothing to do with the color spaces defined as part of the Adobe imaging model. The purpose of converting from <i>RGB</i> to <i>YUV</i> is to separate luminance and chrominance information (see below).</p> <p>The default value of ColorTransform is 1 if the image has three components and 0 otherwise. In other words, conversion between <i>RGB</i> and <i>YUV</i> is performed for all three-component images unless explicitly disabled by setting ColorTransform to 0. Additionally, the encoding algorithm inserts an Adobe-defined marker code in the encoded data indicating the ColorTransform value used. If present, this marker code overrides the ColorTransform value given to DCTDecode. Thus it is necessary to specify ColorTransform only when decoding data that does not contain the Adobe-defined marker code.</p>

The encoding algorithm can reduce the information loss by making the step size in the quantization smaller at the expense of reducing the amount of compression achieved by the algorithm. The compression achieved by the JPEG algorithm depends on the image being compressed and the amount of loss that is acceptable. In general, a compression of 15:1 can be achieved without perceptible loss of information, and 30:1 compression causes little impairment of the image.

Better compression is often possible for color spaces that treat luminance and chrominance separately than for those that do not. The *RGB*-to-*YUV* conversion provided by the filters is one attempt to separate luminance and chrominance; it conforms to CCIR recommendation 601-1. Other color spaces, such as the CIE 1976 $L^*a^*b^*$ space, may also achieve this objective. The chrominance components can then be compressed more than the luminance by using coarser sampling or quantization, with no degradation in quality.

The JPEG filter implementation in Adobe Acrobat products does not support features of the JPEG standard that are irrelevant to images. In addition, certain choices have been made regarding reserved marker codes and other optional features of the standard. For details, see Adobe Technical Note #5116, *Supporting the DCT Filters in PostScript Level 2*.

In addition to the baseline JPEG format, beginning with PDF 1.3 the **DCTDecode** filter supports the progressive JPEG extension. This extension does not add any entries to the **DCTDecode** parameter dictionary; the distinction between baseline and progressive JPEG is represented in the encoded data.

***Note:** There is no benefit to using progressive JPEG for stream data that is embedded in a PDF file. Decoding progressive JPEG is slower and consumes more memory than baseline JPEG. The purpose of this feature is to enable a stream to refer to an external file whose data happens to be already encoded in progressive JPEG. (See also implementation note 10 in Appendix H.)*

3.4 File Structure

The preceding sections describe the syntax of individual objects. This section describes how objects are organized in a PDF file for efficient random access and incremental update. A canonical PDF file initially consists of four elements (see Figure 3.2):

- A one-line *header* identifying the version of the PDF specification to which the file conforms
- A *body* containing the objects that make up the document contained in the file
- A *cross-reference table* containing information about the indirect objects in the file
- A *trailer* giving the location of the cross-reference table and of certain special objects within the body of the file

This initial structure may be modified by later updates, which append additional elements to the end of the file; see Section 3.4.5, “Incremental Updates,” for details.




FIGURE 3.2 *Initial structure of a PDF file*

As a matter of convention, the tokens in a PDF file are arranged into lines; see Section 3.1, “Lexical Conventions.” Each line is terminated by an end-of-line (EOL) marker, which may be a carriage return (character code 13), a line feed (character code 10), or both. PDF files with binary data may have arbitrarily long lines. However, to increase compatibility with other applications that process PDF files, lines that are not part of stream object data are limited to no more than 255 characters, with one exception: beginning with PDF 1.3, an exception is made to the restriction on line length in the case of the **Contents** string of a signature dictionary (see “Signature Fields” on page 547). See also implementation note 11 in Appendix H.

The rules described here are sufficient to produce a well-formed PDF file. However, there are some additional rules for organizing a PDF file to enable efficient incremental access to a document’s components in a network environment. This form of organization, called *Linearized PDF*, is described in Appendix F.

3.4.1 File Header

The first line of a PDF file is a *header* identifying the version of the PDF specification to which the file conforms. For a file conforming to PDF version 1.4, the header should be

```
%PDF-1.4
```

However, since any file conforming to an earlier version of PDF also conforms to version 1.4, an application that processes PDF 1.4 can also accept files with any of the following headers:

```
%PDF-1.0  
%PDF-1.1  
%PDF-1.2  
%PDF-1.3
```

(See also implementation notes 12 and 13 in Appendix H.)

In PDF 1.4, the version in the file header can be overridden by the **Version** entry in the document's catalog dictionary (located via the **Root** entry in the file's trailer, as described in Section 3.4.4, "File Trailer"). This enables a PDF producer application to update the version using an incremental update (see Section 3.4.5, "Incremental Updates").

Under some conditions, a viewer application may be able to process PDF files conforming to a later version than it was designed to accept. New PDF features are often introduced in such a way that they can safely be ignored by a viewer that does not understand them (see Section H.1, "PDF Version Numbers").

Note: *If a PDF file contains binary data, as most do (see Section 3.1, "Lexical Conventions"), it is recommended that the header line be immediately followed by a comment line containing at least four binary characters—that is, characters whose codes are 128 or greater. This will ensure proper behavior of file transfer applications that inspect data near the beginning of a file to determine whether to treat the file's contents as text or as binary.*

3.4.2 File Body

The *body* of a PDF file consists of a sequence of indirect objects representing the contents of a document. The objects, which are of the basic types described in Section 3.2, “Objects,” represent components of the document such as fonts, pages, and sampled images.

3.4.3 Cross-Reference Table

The *cross-reference table* contains information that permits random access to indirect objects within the file, so that the entire file need not be read to locate any particular object. The table contains a one-line entry for each indirect object, specifying the location of that object within the body of the file.

The cross-reference table is the only part of a PDF file with a fixed format; this permits entries in the table to be accessed randomly. The table comprises one or more *cross-reference sections*. Initially, the entire table consists of a single section (or two sections if the file is linearized; see Appendix F); one additional section is added each time the file is updated (see Section 3.4.5, “Incremental Updates”).

Each cross-reference section begins with a line containing the keyword **xref**. Following this line are one or more *cross-reference subsections*, which may appear in any order. The subsection structure is useful for incremental updates, since it allows a new cross-reference section to be added to the PDF file, containing entries only for objects that have been added or deleted. For a file that has never been updated, the cross-reference section contains only one subsection, whose object numbering begins at 0.

Each cross-reference subsection contains entries for a contiguous range of object numbers. The subsection begins with a line containing two numbers, separated by a space: the object number of the first object in this subsection and the number of entries in the subsection. For example, the line

```
28 5
```

introduces a subsection containing five objects, numbered consecutively from 28 to 32.

Following this line are the cross-reference entries themselves, one per line. Each entry is exactly 20 bytes long, including the end-of-line marker. There are two

kinds of cross-reference entry: one for objects that are in use and another for objects that have been deleted and so are free. Both types of entry have similar basic formats, distinguished by the keyword **n** (for an in-use entry) or **f** (for a free entry). The format of an in-use entry is as follows:

```
nnnnnnnnnn ggggg n eol
```

where

nnnnnnnnnn is a 10-digit byte offset

ggggg is a 5-digit generation number

n is a literal keyword identifying this as an in-use entry

eol is a 2-character end-of-line sequence

The byte offset is a 10-digit number, padded with leading zeros if necessary, giving the number of bytes from the beginning of the file to the beginning of the object. It is separated from the generation number by a single space. The generation number is a 5-digit number, also padded with leading zeros if necessary. Following the generation number is a single space, the keyword **n**, and then a 2-character end-of-line sequence. If the file's end-of-line marker is a single character (either a carriage return or a line feed), it is preceded by a single space; if the marker is 2 characters (both a carriage return and a line feed), it is not preceded by a space. Thus the overall length of the entry is always exactly 20 bytes.

The cross-reference entry for a free object has essentially the same format, except that the keyword is **f** instead of **n** and the interpretation of the first item is different:

```
nnnnnnnnnn ggggg f eol
```

where

nnnnnnnnnn is the 10-digit object number of the next free object

ggggg is a 5-digit generation number

f is a literal keyword identifying this as a free entry

eol is a 2-character end-of-line sequence

The free entries in the cross-reference table form a linked list, with each free entry containing the object number of the next. The first entry in the table (object

number 0) is always free and has a generation number of 65,535; it is the head of the linked list of free objects. The last free entry (the tail of the linked list) links back to object number 0.

Except for object number 0, all objects in the cross-reference table initially have generation numbers of 0. When an indirect object is deleted, its cross-reference entry is marked free and it is added to the linked list of free entries. The entry's generation number is incremented by 1 to indicate the generation number to be used the next time an object with that object number is created. Thus each time the entry is reused, it is given a new generation number. The maximum generation number is 65,535; when a cross-reference entry reaches this value, it will never be reused.

The cross-reference table (comprising the original cross-reference section and all update sections) must contain one entry for each object number from 0 to the maximum object number used in the file, even if one or more of the object numbers in this range do not actually occur in the file.

Example 3.5 shows a cross-reference section consisting of a single subsection with six entries: four that are in use (objects number 1, 2, 4, and 5) and two that are free (objects number 0 and 3). Object number 3 has been deleted, and the next object created with that object number will be given a generation number of 7.

Example 3.5

```
xref
0 6
000000003 65535 f
000000017 00000 n
000000081 00000 n
000000000 00007 f
000000331 00000 n
000000409 00000 n
```

Example 3.6 shows a cross-reference section with four subsections, containing a total of five entries. The first subsection contains one entry, for object number 0, which is free. The second subsection contains one entry, for object number 3, which is in use. The third subsection contains two entries, for objects number 23 and 24, both of which are in use. Object number 23 has been reused, as can be seen from the fact that it has a generation number of 2. The fourth subsection contains one entry, for object number 30, which is in use.

Example 3.6

```
xref
0 1
0000000000 65535 f
3 1
0000025325 00000 n
23 2
0000025518 00002 n
0000025635 00000 n
30 1
0000025777 00000 n
```

See Section G.6, “Updating Example,” for a more extensive example of the structure of a PDF file that has been updated several times.

3.4.4 File Trailer

The *trailer* of a PDF file enables an application reading the file to quickly find the cross-reference table and certain special objects. Applications should read a PDF file from its end. The last line of the file contains only the end-of-file marker, %%EOF. (See implementation note 14 in Appendix H.) The two preceding lines contain the keyword **startxref** and the byte offset from the beginning of the file to the beginning of the **xref** keyword in the last cross-reference section. The **startxref** line is preceded by the *trailer dictionary*, consisting of the keyword **trailer** followed by a series of key-value pairs enclosed in double angle brackets (<<...>>). Thus the trailer has the following overall structure:

```
trailer
<< key1 value1
   key2 value2
   ...
   keyn valuen
>>
startxref
Byte_offset_of_last_cross-reference_section
%%EOF
```

Table 3.12 lists the contents of the trailer dictionary; Example 3.7 shows an example trailer for a file that has never been updated (as indicated by the absence of a **Prev** entry in the trailer dictionary).

TABLE 3.12 Entries in the file trailer dictionary

KEY	TYPE	VALUE
Size	integer	<i>(Required)</i> The total number of entries in the file’s cross-reference table, as defined by the combination of the original section and all update sections. Equivalently, this value is 1 greater than the highest object number used in the file.
Prev	integer	<i>(Present only if the file has more than one cross-reference section)</i> The byte offset from the beginning of the file to the beginning of the previous cross-reference section.
Root	dictionary	<i>(Required; must be an indirect reference)</i> The catalog dictionary for the PDF document contained in the file (see Section 3.6.1, “Document Catalog”).
Encrypt	dictionary	<i>(Required if document is encrypted; PDF 1.1)</i> The document’s encryption dictionary (see Section 3.5, “Encryption”).
Info	dictionary	<i>(Optional; must be an indirect reference)</i> The document’s information dictionary (see Section 9.2.1, “Document Information Dictionary”).
ID	array	<i>(Optional; PDF 1.1)</i> An array of two strings constituting a file identifier (see Section 9.3, “File Identifiers”) for the file.

Example 3.7

```

trailer
  << /Size 22
    /Root 2 0 R
    /Info 1 0 R
    /ID [ <81b14aafa313db63dbd6f981e49f94f4>
          <81b14aafa313db63dbd6f981e49f94f4>
        ]
  >>
startxref
18799
%%EOF

```

3.4.5 Incremental Updates

The contents of a PDF file can be updated incrementally, without rewriting the entire file. Changes are appended to the end of the file, leaving its original contents intact. The main advantage to updating a file in this way (as discussed in

Section 2.2.7, “Incremental Update”) is that it enables small changes to a large document to be saved quickly. Other advantages include the following:

- In some cases, incremental updating is the only way to save changes to a document. An accepted practice for minimizing the risk of data loss when saving a document is to write it to a new file and then rename the new file to replace the old one; however, in certain contexts, such as when editing a document across an HTTP connection or using OLE embedding (a Windows-specific technology), it is not possible to overwrite the contents of the original file in this manner. Incremental updates can be used to save changes to documents in these contexts.
- Once a document has been signed (see Section 2.2.6, “Security”), all changes made to the document must be saved using incremental updates, since altering any existing bytes in the file will invalidate existing signatures.

In an incremental update, any new or changed objects are appended to the file, a cross-reference section is added, and a new trailer is inserted. The resulting file has the structure shown in Figure 3.3. A complete example of an updated file is shown in Section G.6, “Updating Example.”

The cross-reference section added when a file is updated contains entries only for objects that have been changed, replaced, or deleted, plus the entry for object 0. Deleted objects are left unchanged in the file, but are marked as deleted via their cross-reference entries. The added trailer contains all the entries (perhaps modified) from the previous trailer, as well as a **Prev** entry giving the location of the previous cross-reference section (see Table 3.12 on page 68). As shown in Figure 3.3, a file that has been updated several times contains several trailers; note that each trailer is terminated by its own end-of-file (%%EOF) marker.

Because updates are appended to PDF files, it is possible to end up with several copies of an object with the same object identifier (object number and generation number). This can occur, for example, if a text annotation (see Section 8.4, “Annotations”) is changed several times, with the file being saved between changes. Because the text annotation object is not deleted, it retains the same object number and generation number as before. An updated copy of the object is included in the new update section added to the file; the update’s cross-reference section includes a byte offset to this new copy of the object, overriding the old byte offset contained in the original cross-reference section. When a viewer application reads the file, it must build its cross-reference information in such a way that the most recent copy of each object is the one accessed in the file.




FIGURE 3.3 *Structure of an updated PDF file*

In versions of PDF prior to 1.4, it was not possible to use an incremental update to alter the version of PDF to which the document conforms, since the version was specified only in the header at the beginning of the file (see Section 3.4.1, “File Header”). In PDF 1.4, it is possible for a **Version** entry in the document’s catalog dictionary (see Section 3.6.1, “Document Catalog”) to override the ver-

sion specified in the header; this enables the version to be altered using an incremental update.

3.5 Encryption

A PDF document can be *encrypted* (PDF 1.1) to protect its contents from unauthorized access. Encryption applies to all strings and streams in the document's PDF file, but not to other object types such as integers and boolean values, which are used primarily to convey information about the document's structure rather than its content. Leaving these values unencrypted allows random access to the objects within a document, while encrypting the strings and streams protects the document's substantive contents.

Note: *When a PDF stream object (see Section 3.2.7, "Stream Objects") refers to an external file, the stream's contents are not encrypted, since they are not part of the PDF file itself. However, if the contents of the stream are embedded within the PDF file (see Section 3.10.3, "Embedded File Streams"), they are encrypted like any other stream in the file.*

Encryption is controlled by an *encryption dictionary*, which is the value of the **Encrypt** entry in the document's trailer dictionary (see Table 3.12 on page 68). The absence of this entry from the trailer dictionary means that the document is not encrypted. The entries shown in Table 3.13 are common to all encryption dictionaries.

The encryption dictionary's **Filter** entry identifies the file's *security handler*, a software module that implements various aspects of the encryption process and controls access to the contents of the encrypted document. PDF specifies a standard security handler that all viewer applications are expected to support, but applications may optionally substitute security handlers of their own.

The **V** entry, in specifying which algorithm to use, determines the length or lengths allowed for the encryption key, on which the encryption (and decryption) of data in a PDF file is based. If more than one length is allowed (that is, for **V** values 2 and 3), the **Length** entry specifies the exact length of the encryption key.

The remaining contents of the encryption dictionary are determined by the security handler, and may vary from one handler to another. Those for the standard security handler are described in Section 3.5.2, "Standard Security Handler."

TABLE 3.13 Entries common to all encryption dictionaries

KEY	TYPE	VALUE
Filter	name	<i>(Required)</i> The name of the <i>security handler</i> for this document; see below. Default value: Standard , for the built-in security handler. (Names for other security handlers can be registered using the procedure described in Appendix E.)
V	number	<p><i>(Optional but strongly recommended)</i> A code specifying the algorithm to be used in encrypting and decrypting the document:</p> <ul style="list-style-type: none"> 0 An algorithm that is undocumented and no longer supported, and whose use is strongly discouraged. 1 Algorithm 3.1 on page 73, with an encryption key length of 40 bits; see below. 2 <i>(PDF 1.4)</i> Algorithm 3.1 on page 73, but allowing encryption key lengths greater than 40 bits. 3 <i>(PDF 1.4)</i> An unpublished algorithm allowing encryption key lengths ranging from 40 to 128 bits. (This algorithm is unpublished as an export requirement of the U.S. Department of Commerce.) <p>The default value if this entry is omitted is 0, but a value of 1 or greater is strongly recommended. (See implementation note 15 in Appendix H.)</p>
Length	integer	<i>(Optional; PDF 1.4; only if V is 2 or 3)</i> The length of the encryption key, in bits. The value must be a multiple of 8, in the range 40 to 128. Default value: 40.

Unlike strings within the body of the document, those in the encryption dictionary must be direct objects and are *not* encrypted by the usual methods. The security handler itself is responsible for encrypting and decrypting strings in the encryption dictionary, using whatever encryption algorithm it chooses.

Note: Document creators have two choices if the standard security handler and encryption methods provided by PDF are not sufficient for their needs: they can provide an alternate, more secure security handler or they can encrypt whole PDF documents themselves, not making use of PDF security.

3.5.1 General Encryption Algorithm

PDF's standard encryption methods use the MD5 message-digest algorithm (described in Internet RFC 1321, *The MD5 Message-Digest Algorithm*; see the Bibliography) and a proprietary encryption algorithm known as RC4. RC4 is a

symmetric stream cipher: the same algorithm is used for both encryption and decryption, and the algorithm does not change the length of the data.

Note: *RC4 is a copyrighted, proprietary algorithm of RSA Security, Inc. Adobe Systems has licensed this algorithm for use in its Acrobat products. Independent software vendors may be required to license RC4 in order to develop software that encrypts or decrypts PDF documents. For further information, visit the RSA Web site at <<http://www.rsasecurity.com>> or send e-mail to <products@rsasecurity.com>.*

The encryption of data in a PDF file is based on the use of an *encryption key* computed by the security handler. Different security handlers can compute the encryption key in a variety of ways, more or less cryptographically secure. Regardless of how the key is computed, its use in the encryption of data is always the same (see Algorithm 3.1). Because the RC4 algorithm is symmetric, this same sequence of steps can be used both to encrypt and to decrypt data.

Algorithm 3.1 Encryption of data using an encryption key

1. Obtain the object number and generation number from the object identifier of the string or stream to be encrypted (see Section 3.2.9, “Indirect Objects”). If the string is a direct object, use the identifier of the indirect object containing it.
2. Treating the object number and generation number as binary integers, extend the original n -byte encryption key to $n + 5$ bytes by appending the low-order 3 bytes of the object number and the low-order 2 bytes of the generation number in that order, low-order byte first. (n is 5 unless the value of **V** in the encryption dictionary is greater than 1, in which case n is the value of **Length** divided by 8.)
3. Initialize the MD5 hash function and pass the result of step 2 as input to this function.
4. Use the first $(n + 5)$ bytes, up to a maximum of 16, of the output from the MD5 hash as the key for the RC4 encryption function, along with the string or stream data to be encrypted. The output is the encrypted data to be stored in the PDF file.

Stream data is encrypted after applying all stream encoding filters, and is decrypted before applying any stream decoding filters; the number of bytes to be encrypted or decrypted is given by the **Length** entry in the stream dictionary. Decryption of strings (other than those in the encryption dictionary) is done after escape-sequence processing and hexadecimal decoding as appropriate to the string representation described in Section 3.2.3, “String Objects.”

3.5.2 Standard Security Handler

PDF's standard security handler allows *access permissions* and up to two passwords to be specified for a document: an *owner password* and a *user password*. An application's decision to encrypt a document is based on whether the user creating the document specifies any passwords or access restrictions (for example, in a security settings dialog that the user can invoke before saving the PDF file); if so, the document is encrypted, and the permissions and information required to validate the passwords are stored in the encryption dictionary. (An application may also create an encrypted document without any user interaction, if it has some other source of information about what passwords and permissions to use.)

If a user attempts to open an encrypted document that has a user password, the viewer application should prompt for a password. Correctly supplying either password allows the user to open the document, decrypt it, and display it on the screen. If the document does not have a user password, no password is requested; the viewer application can simply open, decrypt, and display the document. Whether additional operations are allowed on a decrypted document depends on which password (if any) was supplied when the document was opened and on any access restrictions that were specified when the document was created:

- Opening the document with the correct owner password (assuming it is not the same as the user password) allows full (owner) access to the document. This unlimited access includes the ability to change the document's passwords and access permissions.
- Opening the document with the correct user password (or opening a document that does not have a user password) allows additional operations to be performed according to the user access permissions specified in the document's encryption dictionary.

Access permissions are specified in the form of flags corresponding to the various operations, and the set of operations to which they correspond depends in turn on the security handler's revision number (also stored in the encryption dictionary). If the revision number is 2, the operations to which user access can be controlled are as follows:

- Modifying the document's contents
- Copying or otherwise extracting text and graphics from the document, including extraction for accessibility purposes (that is, to make the contents of the

document accessible through assistive technologies such as screen readers or Braille output devices; see Section 9.8, “Accessibility Support”)

- Adding or modifying text annotations (“sticky notes”; see “Text Annotations” on page 499) and interactive form fields (Section 8.6, “Interactive Forms”)
- Printing the document

If the security handler’s revision number is 3, user access to the following operations can be controlled more selectively:

- Filling in forms (that is, filling in existing interactive form fields) and signing the document (which amounts to filling in existing signature fields, a type of interactive form field)
- Assembling the document: inserting, rotating, or deleting pages and creating navigation elements such as bookmarks or thumbnail images (see Section 8.2, “Document-Level Navigation”)
- Printing to a representation from which a faithful digital copy of the PDF content could be generated. Disallowing such printing may result in degradation of output quality (a feature implemented as “Print As Image” in Acrobat)

In addition, revision 3 enables the extraction of text and graphics (in support of accessibility to disabled users or for other purposes) to be controlled separately.

***Note:** Once the document has been opened and decrypted successfully, the viewer application has access to the entire contents of the document. There is nothing inherent in PDF encryption that enforces the document permissions specified in the encryption dictionary. It is up to the implementors of PDF viewer applications to respect the intent of the document creator by restricting user access to an encrypted PDF file according to the permissions contained in the file.*

Standard Encryption Dictionary

Table 3.14 shows the encryption dictionary entries for the standard security handler (in addition to those in Table 3.13 on page 72).

TABLE 3.14 Additional encryption dictionary entries for the standard security handler

KEY	TYPE	VALUE
R	number	<i>(Required)</i> A number specifying which revision of the standard security handler should be used to interpret this dictionary. The revision number should be 2 if the document is encrypted with a V value less than 2 (see Table 3.13) and does not have any of the access permissions set (via the P entry, below) that are designated “Revision 3” in Table 3.15; otherwise (that is, if the document is encrypted with a V value greater than 2 or has any “Revision 3” access permissions set), this value should be 3.
O	string	<i>(Required)</i> A 32-byte string, based on both the owner and user passwords, that is used in computing the encryption key and in determining whether a valid owner password was entered. For more information, see “Encryption Key Algorithm” on page 78 and “Password Algorithms” on page 79.
U	string	<i>(Required)</i> A 32-byte string, based on the user password, that is used in determining whether to prompt the user for a password and, if so, whether a valid user or owner password was entered. For more information, see “Password Algorithms” on page 79.
P	integer	<i>(Required)</i> A set of flags specifying which operations are permitted when the document is opened with user access (see Table 3.15).

The values of the **O** and **U** entries in this dictionary are used to determine whether a password entered when the document is opened is the correct owner password, user password, or neither.

The value of the **P** entry is an unsigned 32-bit integer containing a set of flags specifying which access permissions should be granted when the document is opened with user access. Table 3.15 shows the meanings of these flags. Bit positions within the flag word are numbered from 1 (low-order) to 32 (high-order); a 1 bit in any position enables the corresponding access permission. Which bits are meaningful, and in some cases how they are interpreted, depends on the security handler’s revision number (specified in the encryption dictionary’s **R** entry).

***Note:** PDF integer objects in fact are represented internally in signed two-complement form. Since all the reserved high-order flag bits in the encryption dictionary’s **P** value are required to be 1, the value must be specified as a negative integer. For example, assuming revision 2 of the security handler, the value -44 allows printing and copying but disallows modifying the contents and annotations.*

TABLE 3.15 User access permissions

BIT POSITION	MEANING
1–2	Reserved; must be 0.
3	<i>(Revision 2)</i> Print the document. <i>(Revision 3)</i> Print the document (possibly not at the highest quality level, depending on whether bit 12 is also set).
4	Modify the contents of the document by operations other than those controlled by bits 6, 9, and 11.
5	<i>(Revision 2)</i> Copy or otherwise extract text and graphics from the document, including extracting text and graphics (in support of accessibility to disabled users or for other purposes). <i>(Revision 3)</i> Copy or otherwise extract text and graphics from the document by operations other than that controlled by bit 10.
6	Add or modify text annotations, fill in interactive form fields, and, if bit 4 is also set, create or modify interactive form fields (including signature fields).
7–8	Reserved; must be 1.
9	<i>(Revision 3 only)</i> Fill in existing interactive form fields (including signature fields), even if bit 6 is clear.
10	<i>(Revision 3 only)</i> Extract text and graphics (in support of accessibility to disabled users or for other purposes).
11	<i>(Revision 3 only)</i> Assemble the document (insert, rotate, or delete pages and create bookmarks or thumbnail images), even if bit 4 is clear.
12	<i>(Revision 3 only)</i> Print the document to a representation from which a faithful digital copy of the PDF content could be generated. When this bit is clear (and bit 3 is set), printing is limited to a low-level representation of the appearance, possibly of degraded quality. (See implementation note 16 in Appendix H.)
13–32	<i>(Revision 3 only)</i> Reserved; must be 1.

Encryption Key Algorithm

As noted earlier, one function of a security handler is to generate an encryption key for use in encrypting and decrypting the contents of a document. Given a password string, the standard security handler computes an encryption key as shown in Algorithm 3.2.

Algorithm 3.2 *Computing an encryption key*

1. Pad or truncate the password string to exactly 32 bytes. If the password string is more than 32 bytes long, use only its first 32 bytes; if it is less than 32 bytes long, pad it by appending the required number of additional bytes from the beginning of the following padding string:

```
< 28 BF 4E 5E 4E 75 8A 41 64 00 4E 56 FF FA 01 08  
  2E 2E 00 B6 D0 68 3E 80 2F 0C A9 FE 64 53 69 7A >
```

That is, if the password string is n bytes long, append the first $32 - n$ bytes of the padding string to the end of the password string. If the password string is empty (zero-length), meaning there is no user password, substitute the entire padding string in its place.

2. Initialize the MD5 hash function and pass the result of step 1 as input to this function.
3. Pass the value of the encryption dictionary's **O** entry to the MD5 hash function. (Algorithm 3.3 shows how the **O** value is computed.)
4. Treat the value of the **P** entry as an unsigned 4-byte integer and pass these bytes to the MD5 hash function, low-order byte first.
5. Pass the first element of the file's file identifier array (the value of the **ID** entry in the document's trailer dictionary; see Table 3.12 on page 68) to the MD5 hash function and finish the hash.
6. (*Revision 3 only*) Do the following 50 times: Take the output from the previous MD5 hash and pass it as input into a new MD5 hash.
7. Set the encryption key to the first n bytes of the output from the final MD5 hash, where n is always 5 for revision 2 but for revision 3 depends on the value of the encryption dictionary's **Length** entry.

This algorithm, when applied to the user password string, produces the encryption key used to encrypt or decrypt string and stream data according to Algorithm 3.1 on page 73. Parts of this algorithm are also used in the algorithms described below.

Password Algorithms

In addition to the encryption key, the standard security handler must provide the contents of the encryption dictionary (Tables 3.13 on page 72 and 3.14 on page 76). The values of the **Filter**, **V**, **Length**, **R**, and **P** entries are straightforward, but the computation of the **O** (owner password) and **U** (user password) entries requires further explanation. Algorithms 3.3 through 3.5 show how the values of the owner password and user password entries are computed (with separate versions of the latter for revisions 2 and 3 of the security handler).

Algorithm 3.3 Computing the encryption dictionary's **O** (owner password) value

1. Pad or truncate the owner password string as described in step 1 of Algorithm 3.2. If there is no owner password, use the user password instead. (See implementation note 17 in Appendix H.)
2. Initialize the MD5 hash function and pass the result of step 1 as input to this function.
3. (*Revision 3 only*) Do the following 50 times: Take the output from the previous MD5 hash and pass it as input into a new MD5 hash.
4. Create an RC4 encryption key using the first n bytes of the output from the final MD5 hash, where n is always 5 for revision 2 but for revision 3 depends on the value of the encryption dictionary's **Length** entry.
5. Pad or truncate the user password string as described in step 1 of Algorithm 3.2.
6. Encrypt the result of step 5, using an RC4 encryption function with the encryption key obtained in step 4.
7. (*Revision 3 only*) Do the following 19 times: Take the output from the previous invocation of the RC4 function and pass it as input to a new invocation of the function; use an encryption key generated by taking each byte of the encryption key obtained in step 4 and performing an XOR (exclusive or) operation between that byte and the single-byte value of the iteration counter (from 1 to 19).
8. Store the output from the final invocation of the RC4 function as the value of the **O** entry in the encryption dictionary.

Algorithm 3.4 Computing the encryption dictionary's **U** (user password) value (*Revision 2*)

1. Create an encryption key based on the user password string, as described in Algorithm 3.2.

2. Encrypt the 32-byte padding string shown in step 1 of Algorithm 3.2, using an RC4 encryption function with the encryption key from the preceding step.
3. Store the result of step 2 as the value of the **U** entry in the encryption dictionary.

Algorithm 3.5 Computing the encryption dictionary's U (user password) value (Revision 3)

1. Create an encryption key based on the user password string, as described in Algorithm 3.2.
2. Initialize the MD5 hash function and pass the 32-byte padding string shown in step 1 of Algorithm 3.2 as input to this function.
3. Pass the first element of the file's file identifier array (the value of the **ID** entry in the document's trailer dictionary; see Table 3.12 on page 68) to the hash function and finish the hash.
4. Encrypt the 16-byte result of the hash, using an RC4 encryption function with the encryption key from step 1.
5. Do the following 19 times: Take the output from the previous invocation of the RC4 function and pass it as input to a new invocation of the function; use an encryption key generated by taking each byte of the original encryption key (obtained in step 1) and performing an XOR (exclusive or) operation between that byte and the single-byte value of the iteration counter (from 1 to 19).
6. Append 16 bytes of arbitrary padding to the output from the final invocation of the RC4 function and store the 32-byte result as the value of the **U** entry in the encryption dictionary.

The standard security handler uses Algorithms 3.6 and 3.7 to determine whether a supplied password string is the correct user or owner password. Note too that Algorithm 3.6 can be used to determine whether a document's user password is the empty string, and therefore whether to suppress prompting for a password when the document is opened.

Algorithm 3.6 Authenticating the user password

1. Perform all but the last step of Algorithm 3.4 (*Revision 2*) or Algorithm 3.5 (*Revision 3*) using the supplied password string.
2. If the result of step 1 is equal to the value of the encryption dictionary's **U** entry (comparing on the first 16 bytes in the case of Revision 3), the password supplied is the correct user password. The key obtained in step 1 (that is, in the first step of Algorithm 3.4 or 3.5) can be used to decrypt the document using Algorithm 3.1 on page 73.

Algorithm 3.7 Authenticating the owner password

1. Compute an encryption key from the supplied password string, as described in steps 1 to 4 of Algorithm 3.3.
2. (*Revision 2 only*) Decrypt the value of the encryption dictionary's **O** entry, using an RC4 encryption function with the encryption key computed in step 1.

(*Revision 3 only*) Do the following 20 times: Decrypt the value of the encryption dictionary's **O** entry (first iteration) or the output from the previous iteration (all subsequent iterations), using an RC4 encryption function with a different encryption key at each iteration. The key is generated by taking the original key (obtained in step 1) and performing an XOR (exclusive or) operation between each byte of the key and the single-byte value of the iteration counter (from 19 to 0).

3. The result of step 2 purports to be the user password. Authenticate this user password using Algorithm 3.6. If it is found to be correct, the password supplied is the correct owner password.

3.6 Document Structure

A PDF document can be regarded as a hierarchy of objects contained in the body section of a PDF file. At the root of the hierarchy is the document's *catalog* dictionary (see Section 3.6.1, "Document Catalog"). Most of the objects in the hierarchy are dictionaries. For example, each page of the document is represented by a *page object*—a dictionary that includes references to the page's contents and other attributes, such as its thumbnail image (Section 8.2.3, "Thumbnail Images") and any annotations (Section 8.4, "Annotations") associated with it. The individual page objects are tied together in a structure called the *page tree* (described in Section 3.6.2, "Page Tree"), which in turn is located via an indirect reference in the document catalog. Parent, child, and sibling relationships within the hierarchy are defined by dictionary entries whose values are indirect references to other dictionaries. Figure 3.4 illustrates the structure of the object hierarchy.

Note: *The data structures described in this section, particularly the catalog and page dictionaries, combine entries describing document structure with ones dealing with the detailed semantics of documents and pages. All entries are listed here, but many of their descriptions are deferred to subsequent chapters.*




FIGURE 3.4 Structure of a PDF document

3.6.1 Document Catalog

The root of a document’s object hierarchy is the *catalog* dictionary, located via the **Root** entry in the trailer of the PDF file (see Section 3.4.4, “File Trailer”). The catalog contains references to other objects defining the document’s contents, outline, article threads (*PDF 1.1*), named destinations, and other attributes. In addition, it contains information about how the document should be displayed on the screen, such as whether its outline and thumbnail page images should be displayed automatically and whether some location other than the first page should be shown when the document is opened. Table 3.16 shows the entries in the catalog dictionary.

TABLE 3.16 Entries in the catalog dictionary

KEY	TYPE	VALUE
Type	name	<i>(Required)</i> The type of PDF object that this dictionary describes; must be Catalog for the catalog dictionary.
Version	name	<i>(Optional; PDF 1.4)</i> The version of the PDF specification to which the document conforms (for example, 1.4), if later than the version specified in the file’s header (see Section 3.4.1, “File Header”). If the header specifies a later version, or if this entry is absent, the document conforms to the version specified in the header. This entry enables a PDF producer application to update the version using an incremental update; see Section 3.4.5, “Incremental Updates.” (See implementation note 18 in Appendix H.) <i>Note: The value of this entry is a name object, not a number, and so must be preceded by a slash character (/) when written in the PDF file (for example, /1.4).</i>
Pages	dictionary	<i>(Required; must be an indirect reference)</i> The <i>page tree node</i> that is the root of the document’s <i>page tree</i> (see Section 3.6.2, “Page Tree”).
PageLabels	number tree	<i>(Optional; PDF 1.3)</i> A number tree (see Section 3.8.5, “Number Trees”) defining the page labeling for the document. The keys in this tree are page indices; the corresponding values are <i>page label dictionaries</i> (see Section 8.3.1, “Page Labels”). Each page index denotes the first page in a <i>labeling range</i> to which the specified page label dictionary applies. The tree must include a value for page index 0.
Names	dictionary	<i>(Optional; PDF 1.2)</i> The document’s <i>name dictionary</i> (see Section 3.6.3, “Name Dictionary”).

Dests	dictionary	<i>(Optional; PDF 1.1; must be an indirect reference)</i> A dictionary of names and corresponding <i>destinations</i> (see “Named Destinations” on page 476).								
ViewerPreferences	dictionary	<i>(Optional; PDF 1.2)</i> A <i>viewer preferences dictionary</i> (see Section 8.1, “Viewer Preferences”) specifying the way the document is to be displayed on the screen. If this entry is absent, viewer applications should use their own current user preference settings.								
PageLayout	name	<p><i>(Optional)</i> A name object specifying the page layout to be used when the document is opened:</p> <table border="0" style="margin-left: 2em;"> <tr> <td>SinglePage</td> <td>Display one page at a time.</td> </tr> <tr> <td>OneColumn</td> <td>Display the pages in one column.</td> </tr> <tr> <td>TwoColumnLeft</td> <td>Display the pages in two columns, with odd-numbered pages on the left.</td> </tr> <tr> <td>TwoColumnRight</td> <td>Display the pages in two columns, with odd-numbered pages on the right.</td> </tr> </table> <p>(See implementation note 19 in Appendix H.) Default value: SinglePage.</p>	SinglePage	Display one page at a time.	OneColumn	Display the pages in one column.	TwoColumnLeft	Display the pages in two columns, with odd-numbered pages on the left.	TwoColumnRight	Display the pages in two columns, with odd-numbered pages on the right.
SinglePage	Display one page at a time.									
OneColumn	Display the pages in one column.									
TwoColumnLeft	Display the pages in two columns, with odd-numbered pages on the left.									
TwoColumnRight	Display the pages in two columns, with odd-numbered pages on the right.									
PageMode	name	<p><i>(Optional)</i> A name object specifying how the document should be displayed when opened:</p> <table border="0" style="margin-left: 2em;"> <tr> <td>UseNone</td> <td>Neither document outline nor thumbnail images visible</td> </tr> <tr> <td>UseOutlines</td> <td>Document outline visible</td> </tr> <tr> <td>UseThumbs</td> <td>Thumbnail images visible</td> </tr> <tr> <td>FullScreen</td> <td>Full-screen mode, with no menu bar, window controls, or any other window visible</td> </tr> </table> <p>Default value: UseNone.</p>	UseNone	Neither document outline nor thumbnail images visible	UseOutlines	Document outline visible	UseThumbs	Thumbnail images visible	FullScreen	Full-screen mode, with no menu bar, window controls, or any other window visible
UseNone	Neither document outline nor thumbnail images visible									
UseOutlines	Document outline visible									
UseThumbs	Thumbnail images visible									
FullScreen	Full-screen mode, with no menu bar, window controls, or any other window visible									
Outlines	dictionary	<i>(Optional; must be an indirect reference)</i> The <i>outline dictionary</i> that is the root of the document’s <i>outline hierarchy</i> (see Section 8.2.2, “Document Outline”).								
Threads	array	<i>(Optional; PDF 1.1; must be an indirect reference)</i> An array of <i>thread dictionaries</i> representing the document’s <i>article threads</i> (see Section 8.3.2, “Articles”).								
OpenAction	array or dictionary	<i>(Optional; PDF 1.1)</i> A value specifying a <i>destination</i> to be displayed or an <i>action</i> to be performed when the document is opened. The value is either an array defining a destination (see Section 8.2.1, “Destinations”) or an <i>action dictionary</i> representing an action (Section 8.5, “Actions”). If this entry is absent, the document should be opened to the top of the first page at the default magnification factor.								

AA	dictionary	<i>(Optional; PDF 1.4)</i> An <i>additional-actions dictionary</i> defining the actions to be taken in response to various <i>trigger events</i> affecting the document as a whole (see “Trigger Events” on page 514). (See also implementation note 20 in Appendix H.)
URI	dictionary	<i>(Optional)</i> A <i>URI dictionary</i> containing document-level information for <i>URI (uniform resource identifier) actions</i> (see “URI Actions” on page 523).
AcroForm	dictionary	<i>(Optional; PDF 1.2)</i> The document’s <i>interactive form (AcroForm) dictionary</i> (see Section 8.6.1, “Interactive Form Dictionary”).
Metadata	stream	<i>(Optional; PDF 1.4; must be an indirect reference)</i> A <i>metadata stream</i> containing metadata for the document (see Section 9.2.2, “Metadata Streams”).
StructTreeRoot	dictionary	<i>(Optional; PDF 1.3)</i> The document’s <i>structure tree root dictionary</i> (see Section 9.6.1, “Structure Hierarchy”).
MarkInfo	dictionary	<i>(Optional; PDF 1.4)</i> A <i>mark information dictionary</i> containing information about the document’s usage of Tagged PDF conventions (see Section 9.7.1, “Mark Information Dictionary”).
Lang	text string	<i>(Optional; PDF 1.4)</i> A <i>language identifier</i> specifying the natural language for all text in the document except where overridden by language specifications for structure elements or marked content (see Section 9.8.1, “Natural Language Specification”). If this entry is absent, the language is considered unknown.
SpiderInfo	dictionary	<i>(Optional; PDF 1.3)</i> A <i>Web Capture information dictionary</i> containing state information used by the Acrobat Web Capture (AcroSpider) plug-in extension (see Section 9.9.1, “Web Capture Information Dictionary”).
OutputIntents	array	<i>(Optional; PDF 1.4)</i> An array of <i>output intent dictionaries</i> describing the color characteristics of output devices on which the document might be rendered (see “Output Intents” on page 684).

Example 3.8 shows a sample catalog object.

Example 3.8

```

1 0 obj
  << /Type /Catalog
    /Pages 2 0 R
    /PageMode /UseOutlines
    /Outlines 3 0 R
  >>
endobj

```

3.6.2 Page Tree

The pages of a document are accessed through a structure known as the *page tree*, which defines their ordering within the document. The tree structure allows PDF viewer applications to quickly open a document containing thousands of pages using only limited memory. The tree contains nodes of two types—intermediate nodes, called *page tree nodes*, and leaf nodes, called *page objects*—whose form is described in the sections below. Viewer applications should be prepared to handle any form of tree structure built of such nodes. The simplest structure would consist of a single page tree node that references all of the document’s page objects directly; however, to optimize the performance of viewer applications, the Acrobat Distiller and PDF Writer programs construct trees of a particular form, known as *balanced trees*. Further information on this form of tree can be found in *Data Structures and Algorithms*, by Aho, Hopcroft, and Ullman (see the Bibliography).

Page Tree Nodes

Table 3.17 shows the required entries in a page tree node.

TABLE 3.17 Required entries in a page tree node		
KEY	TYPE	VALUE
Type	name	<i>(Required)</i> The type of PDF object that this dictionary describes; must be Pages for a page tree node.
Parent	dictionary	<i>(Required except in root node; must be an indirect reference)</i> The page tree node that is the immediate parent of this one.
Kids	array	<i>(Required)</i> An array of indirect references to the immediate children of this node. The children may be page objects or other page tree nodes.
Count	integer	<i>(Required)</i> The number of leaf nodes (page objects) that are descendants of this node within the page tree.

Note: *The structure of the page tree is not necessarily related to the logical structure of the document itself; that is, page tree nodes do not represent chapters, sections, and so forth. (Other data structures are defined for that purpose; see Section 9.6, “Logical Structure.”) Applications that consume or produce PDF files are not required to preserve the existing structure of the page tree.*

Example 3.9 illustrates the page tree for a document with three pages. See “Page Objects,” below, for the contents of the individual page objects, and Section G.4, “Page Tree Example,” for a more extended example showing the page tree for a longer document.

Example 3.9

```

2 0 obj
  << /Type /Pages
    /Kids [ 4 0 R
            10 0 R
            24 0 R
          ]
    /Count 3
  >>
endobj

4 0 obj
  << /Type /Page
    ...Additional entries describing the attributes of this page...
  >>
endobj

10 0 obj
  << /Type /Page
    ...Additional entries describing the attributes of this page...
  >>
endobj

24 0 obj
  << /Type /Page
    ...Additional entries describing the attributes of this page...
  >>
endobj

```

In addition to the entries shown in Table 3.17, a page tree node may contain further entries defining *inherited attributes* for the page objects that are its descendants (see “Inheritance of Page Attributes” on page 91).

Page Objects

The leaves of the page tree are *page objects*, each of which is a dictionary specifying the attributes of a single page of the document. Table 3.18 shows the contents of this dictionary (see also implementation note 21 in Appendix H). The table

also identifies which attributes a page may inherit from its ancestor nodes in the page tree, as described under “Inheritance of Page Attributes” on page 91. Attributes that are not explicitly identified in the table as inheritable cannot be inherited.

TABLE 3.18 Entries in a page object

KEY	TYPE	VALUE
Type	name	<i>(Required)</i> The type of PDF object that this dictionary describes; must be Page for a page object.
Parent	dictionary	<i>(Required; must be an indirect reference)</i> The page tree node that is the immediate parent of this page object.
LastModified	date	<i>(Required if PieceInfo is present; optional otherwise; PDF 1.3)</i> The date and time (see Section 3.8.2, “Dates”) when the page’s contents were most recently modified. If a page-piece dictionary (PieceInfo) is present, the modification date is used to ascertain which of the application data dictionaries that it contains correspond to the current content of the page (see Section 9.4, “Page-Piece Dictionaries”).
Resources	dictionary	<i>(Required; inheritable)</i> A dictionary containing any resources required by the page (see Section 3.7.2, “Resource Dictionaries”). If the page requires no resources, the value of this entry should be an empty dictionary; omitting the entry entirely indicates that the resources are to be inherited from an ancestor node in the page tree.
MediaBox	rectangle	<i>(Required; inheritable)</i> A rectangle (see Section 3.8.3, “Rectangles”), expressed in default user space units, defining the boundaries of the physical medium on which the page is intended to be displayed or printed (see Section 9.10.1, “Page Boundaries”).
CropBox	rectangle	<i>(Optional; inheritable)</i> A rectangle, expressed in default user space units, defining the visible region of default user space. When the page is displayed or printed, its contents are to be clipped (cropped) to this rectangle and then imposed on the output medium in some implementation-defined manner (see Section 9.10.1, “Page Boundaries”). Default value: the value of MediaBox .
BleedBox	rectangle	<i>(Optional; PDF 1.3)</i> A rectangle, expressed in default user space units, defining the region to which the contents of the page should be clipped when output in a production environment (see Section 9.10.1, “Page Boundaries”). Default value: the value of CropBox .

TrimBox	rectangle	<i>(Optional; PDF 1.3)</i> A rectangle, expressed in default user space units, defining the intended dimensions of the finished page after trimming (see Section 9.10.1, “Page Boundaries”). Default value: the value of CropBox .
ArtBox	rectangle	<i>(Optional; PDF 1.3)</i> A rectangle, expressed in default user space units, defining the extent of the page’s meaningful content (including potential white space) as intended by the page’s creator (see Section 9.10.1, “Page Boundaries”). Default value: the value of CropBox .
BoxColorInfo	dictionary	<i>(Optional)</i> A <i>box color information dictionary</i> specifying the colors and other visual characteristics to be used in displaying guidelines on the screen for the various page boundaries (see “Display of Page Boundaries” on page 679). If this entry is absent, the viewer application should use its own current default settings.
Contents	stream or array	<p><i>(Optional)</i> A <i>content stream</i> (see Section 3.7.1, “Content Streams”) describing the contents of this page. If this entry is absent, the page is empty.</p> <p>The value may be either a single stream or an array of streams. If it is an array, the effect is as if all of the streams in the array were concatenated, in order, to form a single stream. This allows a program generating a PDF file to create image objects and other resources as they occur, even though they interrupt the content stream. The division between streams may occur only at the boundaries between lexical tokens (see Section 3.1, “Lexical Conventions”), but is unrelated to the page’s logical content or organization. Applications that consume or produce PDF files are not required to preserve the existing structure of the Contents array. (See implementation note 22 in Appendix H.)</p>
Rotate	integer	<i>(Optional; inheritable)</i> The number of degrees by which the page should be rotated clockwise when displayed or printed. The value must be a multiple of 90. Default value: 0.
Group	dictionary	<i>(Optional; PDF 1.4)</i> A <i>group attributes dictionary</i> specifying the attributes of the page’s page group for use in the transparent imaging model (see Sections 7.3.6, “Page Group,” and 7.5.5, “Transparency Group XObjects”).
Thumb	stream	<i>(Optional)</i> A stream object defining the page’s <i>thumbnail image</i> (see Section 8.2.3, “Thumbnail Images”).
B	array	<i>(Optional; PDF 1.1; recommended if the page contains article beads)</i> An array of indirect references to <i>article beads</i> appearing on the page (see Section 8.3.2, “Articles”; see also implementation note 23 in Appendix H). The beads are listed in the array in natural reading order.

Dur	number	(<i>Optional; PDF 1.1</i>) The page's <i>display duration</i> (also called its <i>advance timing</i>): the maximum length of time, in seconds, that the page will be displayed during presentations before the viewer application automatically advances to the next page (see Section 8.3.3, "Presentations"). By default, the viewer does not advance automatically.
Trans	dictionary	(<i>Optional; PDF 1.1</i>) A <i>transition dictionary</i> describing the transition effect to be used when displaying the page during presentations (see Section 8.3.3, "Presentations").
Annots	array	(<i>Optional</i>) An array of <i>annotation dictionaries</i> representing annotations associated with the page (see Section 8.4, "Annotations").
AA	dictionary	(<i>Optional; PDF 1.2</i>) An <i>additional-actions dictionary</i> defining actions to be performed when the page is opened or closed (see Section 8.5.2, "Trigger Events"; see also implementation note 24 in Appendix H).
Metadata	stream	(<i>Optional; PDF 1.4</i>) A <i>metadata stream</i> containing metadata for the page (see Section 9.2.2, "Metadata Streams").
PieceInfo	dictionary	(<i>Optional; PDF 1.3</i>) A <i>page-piece dictionary</i> associated with the page (see Section 9.4, "Page-Piece Dictionaries").
StructParents	integer	(<i>Required if the page contains structural content items; PDF 1.3</i>) The integer key of the page's entry in the <i>structural parent tree</i> (see "Finding Structure Elements from Content Items" on page 600).
ID	string	(<i>Optional; PDF 1.3; indirect reference preferred</i>) The digital identifier of the page's parent <i>Web Capture content set</i> (see Section 9.9.5, "Object Attributes Related to Web Capture").
PZ	number	(<i>Optional; PDF 1.3</i>) The page's preferred <i>zoom (magnification) factor</i> : the factor by which it should be scaled to achieve the "natural" display magnification (see Section 9.9.5, "Object Attributes Related to Web Capture").
SeparationInfo	dictionary	(<i>Optional; PDF 1.3</i>) A <i>separation dictionary</i> containing information needed to generate color separations for the page (see Section 9.10.3, "Separation Dictionaries").

Example 3.10 shows the definition of a page object with a thumbnail image and two annotations. The media box specifies that the page is to be printed on letter-size paper. In addition, the resource dictionary is specified as a direct object and shows that the page makes use of three fonts, named F3, F5, and F7.

Example 3.10

```

3 0 obj
  << /Type /Page
    /Parent 4 0 R
    /MediaBox [0 0 612 792]
    /Resources << /Font << /F3 7 0 R
                  /F5 9 0 R
                  /F7 11 0 R
                >>
            >>
        /ProcSet [/PDF]
    >>
    /Contents 12 0 R
    /Thumb 14 0 R
    /Annots [ 23 0 R
             24 0 R
            ]
  >>
endobj

```

Inheritance of Page Attributes

Some of the page attributes shown in Table 3.18 are designated as *inheritable*. If such an attribute is omitted from a page object, its value is inherited from an ancestor node in the page tree. If the attribute is a required one, a value must be supplied in an ancestor node; if it is optional and no inherited value is specified, the default value is used.

An attribute can thus be defined once for a whole set of pages, by specifying it in an intermediate page tree node and arranging the pages that share the attribute as descendants of that node. For example, a document might specify the same media box for all of its pages by including a **MediaBox** entry in the root node of the page tree. If necessary, an individual page object could then override this inherited value with a **MediaBox** entry of its own.

Note: *In a document conforming to the Linearized PDF organization (see Appendix F), all page attributes must be specified explicitly as entries in the page dictionaries to which they apply; they may not be inherited from an ancestor node.*

Figure 3.5 illustrates the inheritance of attributes. In the page tree shown, pages 1, 2, and 4 are rotated clockwise by 90 degrees, page 3 by 270 degrees, page 6 by 180 degrees, and pages 5 and 7 not at all (0 degrees).

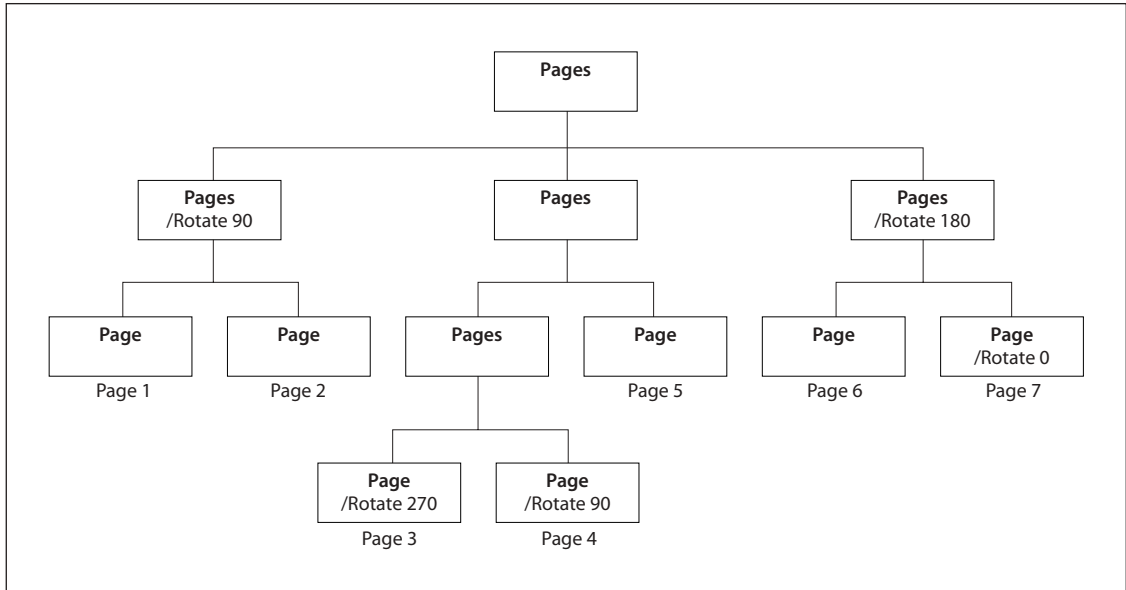


FIGURE 3.5 *Inheritance of attributes*

3.6.3 Name Dictionary

Some categories of objects in a PDF file can be referred to by name rather than by object reference. The correspondence between names and objects is established by the document's *name dictionary* (PDF 1.2), located via the **Names** entry in the document's catalog (see Section 3.6.1, "Document Catalog"). Each entry in this dictionary designates the root of a name tree (Section 3.8.4, "Name Trees") defining names for a particular category of objects. Table 3.19 shows the contents of the name dictionary.

3.7 Content Streams and Resources

Content streams are the primary means for describing the appearance of pages and other graphical elements. A content stream depends on information contained in an associated resource dictionary; in combination, these two objects form a self-contained entity. This section describes these objects.

TABLE 3.19 Entries in the name dictionary

KEY	TYPE	VALUE
Dests	name tree	(<i>Optional; PDF 1.2</i>) A name tree mapping name strings to destinations (see “Named Destinations” on page 476).
AP	name tree	(<i>Optional; PDF 1.3</i>) A name tree mapping name strings to annotation appearance streams (see Section 8.4.4, “Appearance Streams”).
JavaScript	name tree	(<i>Optional; PDF 1.3</i>) A name tree mapping name strings to document-level JavaScript [®] actions (see “JavaScript Actions” on page 556).
Pages	name tree	(<i>Optional; PDF 1.3</i>) A name tree mapping name strings to visible pages for use in interactive forms (see Section 8.6.5, “Named Pages”).
Templates	name tree	(<i>Optional; PDF 1.3</i>) A name tree mapping name strings to invisible (template) pages for use in interactive forms (see Section 8.6.5, “Named Pages”).
IDS	name tree	(<i>Optional; PDF 1.3</i>) A name tree mapping digital identifiers to Web Capture content sets (see Section 9.9.3, “Content Sets”).
URLS	name tree	(<i>Optional; PDF 1.3</i>) A name tree mapping uniform resource locators (URLs) to Web Capture content sets (see Section 9.9.3, “Content Sets”).
EmbeddedFiles	name tree	(<i>Optional; PDF 1.4</i>) A name tree mapping name strings to embedded file streams (see Section 3.10.3, “Embedded File Streams”).

3.7.1 Content Streams

A *content stream* is a PDF stream object whose data consists of a sequence of instructions describing the graphical elements to be painted on a page. The instructions are represented in the form of PDF objects, using the same object syntax as in the rest of the PDF document. However, whereas the document as a whole is a static, random-access data structure, the objects in the content stream are intended to be interpreted and acted upon sequentially.

Each page of a document is represented by one or more content streams. Content streams are also used to package up sequences of instructions as self-contained graphical elements, such as forms (see Section 4.9, “Form XObjects”), patterns (Section 4.6, “Patterns”), certain fonts (Section 5.5.4, “Type 3 Fonts”), and annotation appearances (Section 8.4.4, “Appearance Streams”).

A content stream, after decoding with any specified filters, is interpreted according to the PDF syntax rules described in Section 3.1, “Lexical Conventions.” It consists of PDF objects denoting operands and operators. The operands needed by an operator precede it in the stream. See Example 3.3 on page 43 for an example of a content stream.

An *operand* is a direct object belonging to any of the basic PDF data types except a stream. Dictionaries are permitted as operands only by certain specific operators. Indirect objects and object references are not permitted at all.

An *operator* is a PDF keyword that specifies some action to be performed, such as painting a graphical shape on the page. An operator keyword is distinguished from a name object by the absence of an initial slash character (/). Operators are meaningful only inside a content stream.

Note: This “postfix” notation, in which an operator is preceded by its operands, is superficially the same as in the PostScript language. However, PDF has no concept of an operand stack as PostScript has. In PDF, all of the operands needed by an operator must immediately precede that operator. Operators do not return results, and there may not be operands left over when an operator finishes execution.

Most operators have to do with painting graphical elements on the page or with specifying parameters that affect subsequent painting operations. The individual operators are described in the chapters devoted to their functions:

- Chapter 4 describes operators that paint general graphics, such as filled areas, strokes, and sampled images, and that specify device-independent graphical parameters, such as color.
- Chapter 5 describes operators that paint text using character glyphs defined in fonts.
- Chapter 6 describes operators that specify device-dependent rendering parameters.
- Chapter 9 describes the marked-content operators that associate higher-level logical information with objects in the content stream. These operators do not affect the rendered appearance of the content; rather, they specify information useful to applications that use PDF for document interchange.

Ordinarily, when a viewer application encounters an operator in a content stream that it does not recognize, an error will occur. (See implementation note 25 in Appendix H.) A pair of compatibility operators, **BX** and **EX** (*PDF 1.1*), modify this behavior (see Table 3.20). These operators must occur in pairs and may be nested. They bracket a *compatibility section*, a portion of a content stream within which unrecognized operators are to be ignored without error. This mechanism enables a PDF document to use operators defined in newer versions of PDF without sacrificing compatibility with older viewers; it should be used only in cases where ignoring such newer operators is the appropriate thing to do. The **BX** and **EX** operators are not themselves part of any graphics object (see Section 4.1, “Graphics Objects”) or of the graphics state (Section 4.3, “Graphics State”).

TABLE 3.20 Compatibility operators

OPERANDS	OPERATOR	DESCRIPTION
—	BX	(<i>PDF 1.1</i>) Begin a compatibility section. Unrecognized operators (along with their operands) will be ignored without error until the balancing EX operator is encountered.
—	EX	(<i>PDF 1.1</i>) End a compatibility section begun by a balancing BX operator.

3.7.2 Resource Dictionaries

As stated above, the operands supplied to operators in a content stream may only be direct objects; indirect objects and object references are not permitted. In some cases, an operator needs to refer to a PDF object that is defined outside the content stream, such as a font dictionary or a stream containing image data. This can be accomplished by defining such objects as *named resources* and referring to them by name from within the content stream.

Note: *Named resources are meaningful only in the context of a content stream. The scope of a resource name is local to a particular content stream, and is unrelated to externally known identifiers for objects such as fonts. References from one object to another outside of content streams should be made by means of indirect object references rather than named resources.*

A content stream’s named resources are defined by a *resource dictionary*, which enumerates the named resources needed by the operators in the content stream and the names by which they can be referred to. For example, if a text operator

appearing within the content stream needed a certain font, the content stream's resource dictionary might associate the name F42 with the corresponding font dictionary. The text operator could then use this name to refer to the font.

A resource dictionary is associated with a content stream in one of the following ways:

- For a content stream that is the value of a page's **Contents** entry (or is an element of an array that is the value of that entry), the resource dictionary is designated by the page dictionary's **Resources** entry. (Since a page's **Resources** attribute is inheritable, as described under "Inheritance of Page Attributes" on page 91, it may actually reside in some ancestor node of the page object.)
- For other content streams, the resource dictionary is specified by the **Resources** entry in the stream dictionary of the content stream itself. This applies to content streams that define form XObjects, patterns, Type 3 fonts, and annotation appearances.
- A form XObject or a Type 3 font's glyph description may omit the **Resources** entry, in which case resources will be looked up in the **Resources** entry of the page on which the form or font is used. *This practice is not recommended.*

In the context of a given content stream, the term *current resource dictionary* refers to the resource dictionary associated with the stream in one of the ways described above.

Each key in a resource dictionary is the name of a resource type, as shown in Table 3.21. For most resource types, the corresponding value is a subdictionary whose keys, in turn, are the names of resources of the given type and whose values are the PDF objects representing those resources. (For resource type **ProcSet**, the value is an array of procedure set names instead of a subdictionary.)

Example 3.11 shows a resource dictionary containing procedure sets, fonts, and external objects. The procedure sets are specified by an array, as described in Section 9.1, "Procedure Sets." The fonts are specified with a subdictionary associating the names F5, F6, F7, and F8 with objects 6, 8, 10, and 12, respectively. Likewise, the **XObject** subdictionary associates the names Im1 and Im2 with objects 13 and 15, respectively.

TABLE 3.21 Entries in a resource dictionary

KEY	TYPE	VALUE
ExtGState	dictionary	<i>(Optional)</i> A dictionary mapping resource names to graphics state parameter dictionaries (see Section 4.3.4, “Graphics State Parameter Dictionaries”).
ColorSpace	dictionary	<i>(Optional)</i> A dictionary mapping each resource name to either the name of a device-dependent color space or an array describing a color space (see Section 4.5, “Color Spaces”).
Pattern	dictionary	<i>(Optional)</i> A dictionary mapping resource names to pattern objects (see Section 4.6, “Patterns”).
Shading	dictionary	<i>(Optional; PDF 1.3)</i> A dictionary mapping resource names to shading dictionaries (see “Shading Dictionaries” on page 233).
XObject	stream	<i>(Optional)</i> A dictionary mapping resource names to external objects (see Section 4.7, “External Objects”).
Font	dictionary	<i>(Optional)</i> A dictionary mapping resource names to font dictionaries (see Chapter 5).
ProcSet	array	<i>(Optional)</i> An array of predefined procedure set names (see Section 9.1, “Procedure Sets”).
Properties	dictionary	<i>(Optional; PDF 1.2)</i> A dictionary mapping resource names to property list dictionaries for marked content (see Section 9.5.1, “Property Lists”).

Example 3.11

```

<< /ProcSet [/PDF /ImageB]
  /Font << /F5 6 0 R
           /F6 8 0 R
           /F7 10 0 R
           /F8 12 0 R
        >>
  /XObject << /Im1 13 0 R
             /Im2 15 0 R
            >>
>>

```

3.8 Common Data Structures

As mentioned at the beginning of this chapter, there are some general-purpose data structures that are built from the basic object types described in Section 3.2, “Objects,” and are used in many places throughout PDF. This section describes data structures for text strings, dates, rectangles, name trees, and number trees. The subsequent two sections describe more complex data structures for functions and file specifications.

All of these data structures are meaningful only as part of the document hierarchy; they cannot appear within content streams. In particular, the special conventions for interpreting the values of string objects apply only to strings outside content streams. An entirely different convention is used within content streams for using strings to select sequences of glyphs to be painted on the page (see Chapter 5). Table 3.22 summarizes the basic and higher-level data types that are used throughout this book to describe the values of dictionary entries and other PDF data values.

3.8.1 Text Strings

Certain strings contain information that is intended to be human-readable, such as text annotations, bookmark names, article names, document information, and so forth. Such strings are referred to as *text strings*. Text strings are encoded in either **PDFDocEncoding** or Unicode character encoding. **PDFDocEncoding** is a superset of the ISO Latin 1 encoding and is documented in Appendix D. Unicode is described in the *Unicode Standard* by the Unicode Consortium (see the Bibliography).

For text strings encoded in Unicode, the first two bytes must be 254 followed by 255, representing the Unicode byte order marker, U+FEFF. (This sequence conflicts with the **PDFDocEncoding** character sequence thorn ydieresis, which is unlikely to be a meaningful beginning of a word or phrase.) The remainder of the string consists of Unicode character codes, according to the UTF-16 encoding specified in the Unicode standard, version 2.0. Commonly used Unicode values are represented as 2 bytes per character, with the high-order byte appearing first in the string.

TABLE 3.22 PDF data types

TYPE	DESCRIPTION	SECTION	PAGE
array	Array object	3.2.5	34
boolean	Boolean value	3.2.1	28
date	Date (string)	3.8.2	100
dictionary	Dictionary object	3.2.6	35
file specification	File specification (string or dictionary)	3.10	118
function	Function (dictionary or stream)	3.9	106
integer	Integer number	3.2.2	28
name	Name object	3.2.4	32
name tree	Name tree (dictionary)	3.8.4	101
null	Null object	3.2.8	39
number	Number (integer or real)	3.2.2	28
number tree	Number tree (dictionary)	3.8.5	105
rectangle	Rectangle (array)	3.8.3	101
stream	Stream object	3.2.7	36
string	String object	3.2.3	29
text string	Text string	3.8.1	98

Anywhere in a Unicode text string, an escape sequence may appear to indicate the language in which subsequent text is written; this is useful when the language cannot be determined from the character codes used in the text itself. The escape sequence consists of the following elements, in order:

1. The Unicode value U+001B (that is, the byte sequence 0 followed by 27)
2. A 2-character ISO 639 language code—for example, en for English or ja for Japanese
3. (*Optional*) A 2-character ISO 3166 country code—for example, US for the United States or JP for Japan
4. The Unicode value U+001B

The complete list of codes defined by ISO 639 and ISO 3166 can be obtained from the International Organization for Standardization (see the Bibliography).

3.8.2 Dates

PDF defines a standard date format, which closely follows that of the international standard ASN.1 (Abstract Syntax Notation One), defined in ISO/IEC 8824 (see the Bibliography). A date is a string of the form

(D:YYYYMMDDHHmmSSOHh'mm')

where

YYYY is the year

MM is the month

DD is the day (01–31)

HH is the hour (00–23)

mm is the minute (00–59)

SS is the second (00–59)

O is the relationship of local time to Universal Time (UT), denoted by one of the characters +, -, or Z (see below)

HH followed by ' is the absolute value of the offset from UT in hours (00–23)

mm followed by ' is the absolute value of the offset from UT in minutes (00–59)

The apostrophe character (') after *HH* and *mm* is part of the syntax. All fields after the year are optional. (The prefix D:, although also optional, is strongly recommended.) The default values for *MM* and *DD* are both 01; all other numerical fields default to zero values. A plus sign (+) as the value of the *O* field signifies that local time is later than UT, a minus sign (-) that local time is earlier than UT, and the letter Z that local time is equal to UT. If no UT information is specified, the relationship of the specified time to UT is considered to be unknown. Whether or not the time zone is known, the rest of the date should be specified in local time.

For example, December 23, 1998, at 7:52 PM, U.S. Pacific Standard Time, is represented by the string

D:199812231952-08'00'

3.8.3 Rectangles

Rectangles are used to describe locations on a page and bounding boxes for a variety of objects, such as fonts. A rectangle is written as an array of four numbers giving the coordinates of a pair of diagonally opposite corners. Typically, the array takes the form

$$[ll_x \ ll_y \ ur_x \ ur_y]$$

specifying the lower-left x , lower-left y , upper-right x , and upper-right y coordinates of the rectangle, in that order.

Note: Although rectangles are conventionally specified by their lower-left and upper-right corners, it is acceptable to specify any two diagonally opposite corners. Applications that process PDF should be prepared to normalize such rectangles in situations where specific corners are required.

3.8.4 Name Trees

A *name tree* serves a similar purpose to a dictionary—associating keys and values—but by different means. A name tree differs from a dictionary in the following important ways:

- Unlike the keys in a dictionary, which are name objects, those in a name tree are strings.
- The keys are ordered.
- The values associated with the keys may be objects of any type, but they must always be specified via indirect object references.
- The data structure can represent an arbitrarily large collection of key-value pairs, which can be looked up efficiently without requiring the entire data structure to be read from the PDF file. (In contrast, a dictionary is subject to an implementation limit on the number of entries it can contain.)

A name tree is constructed of *nodes*, each of which is a dictionary object. Table 3.23 shows the entries in a node dictionary. The nodes are of three kinds, depending on the specific entries they contain. The tree always has exactly one *root node*, which contains a single entry: either **Kids** or **Names** but not both. If the root node has a **Names** entry, it is the only node in the tree. If it has a **Kids** entry,

then each of the remaining nodes is either an *intermediate node*, containing a **Limits** entry and a **Kids** entry, or a *leaf node*, containing a **Limits** entry and a **Names** entry.

TABLE 3.23 Entries in a name tree node dictionary

KEY	TYPE	VALUE
Kids	array	<i>(Root and intermediate nodes only; required in intermediate nodes; present in the root node if and only if Names is not present)</i> An array of indirect references to the immediate children of this node. The children may be intermediate or leaf nodes.
Names	array	<i>(Root and leaf nodes only; required in leaf nodes; present in the root node if and only if Kids is not present)</i> An array of the form <p style="text-align: center;">[<i>key</i>₁ <i>value</i>₁ <i>key</i>₂ <i>value</i>₂ ... <i>key</i>_{<i>n</i>} <i>value</i>_{<i>n</i>}]</p> where each <i>key</i> _{<i>i</i>} is a string and the corresponding <i>value</i> _{<i>i</i>} is an indirect reference to the object associated with that key. The keys are sorted in lexical order, as described below.
Limits	array	<i>(Intermediate and leaf nodes only; required)</i> An array of two strings, specifying the (lexically) least and greatest keys included in the Names array of a leaf node or in the Names arrays of any leaf nodes that are descendants of an intermediate node.

The **Kids** entries in the root and intermediate nodes define the tree's structure by identifying the immediate children of each node. The **Names** entries in the leaf (or root) nodes contain the tree's keys and their associated values, arranged in key-value pairs and sorted lexically in ascending order by key. Shorter keys appear before longer ones beginning with the same byte sequence. The encoding of the keys is immaterial as long as it is self-consistent; keys are compared for equality on a simple byte-by-byte basis.

The keys contained within the various nodes' **Names** entries do not overlap; that is, each **Names** entry contains a single contiguous range of all the keys in the tree. In a leaf node, the **Limits** entry specifies the least and greatest keys contained within the node's **Names** entry; in an intermediate node, it specifies the least and greatest keys contained within the **Names** entries of any of that node's descendants. The value associated with a given key can thus be found by walking the tree in order, searching for the leaf node whose **Names** entry contains that key.

Table 3.24 is an abbreviated outline, showing object numbers and nodes, of a name tree that maps the names of all the chemical elements, from actinium to zirconium, to their atomic numbers. Example 3.12 shows the representation of this tree in a PDF file.

TABLE 3.24 Example of a name tree

- 1: Root node
 - 2: Intermediate node: Actinium to Gold
 - 5: Leaf node: Actinium = 25, ..., Astatine = 31
 - 25: Integer: 89
 - ...
 - 31: Integer: 85
 - ...
 - 11: Leaf node: Gadolinium = 56, ..., Gold = 59
 - 56: Integer: 64
 - ...
 - 59: Integer: 79
 - 3: Intermediate node: Hafnium to Protactinium
 - 12: Leaf node: Hafnium = 60, ..., Hydrogen = 65
 - 60: Integer: 72
 - ...
 - 65: Integer: 1
 - ...
 - 19: Leaf node: Palladium = 92, ..., Protactinium = 100
 - 92: Integer: 46
 - ...
 - 100: Integer: 91
 - 4: Intermediate node: Radium to Zirconium
 - 20: Leaf node: Radium = 101, ..., Ruthenium = 107
 - 101: Integer: 89
 - ...
 - 107: Integer: 85
 - ...
 - 24: Leaf node: Xenon = 129, ..., Zirconium = 133
 - 129: Integer: 54
 - ...
 - 133: Integer: 40
-

Example 3.12

```
1 0 obj
  /Kids [ 2 0 R
          3 0 R
          4 0 R
        ]
  >>
endobj

2 0 obj
  << /Limits [(Actinium) (Gold)]
    /Kids [ 5 0 R
            6 0 R
            7 0 R
            8 0 R
            9 0 R
            10 0 R
            11 0 R
          ]
  >>
endobj

3 0 obj
  << /Limits [(Hafnium) (Protactinium)]
    /Kids [ 12 0 R
            13 0 R
            14 0 R
            15 0 R
            16 0 R
            17 0 R
            18 0 R
            19 0 R
          ]
  >>
endobj

4 0 obj
  << /Limits [(Radium) (Zirconium)]
    /Kids [ 20 0 R
            21 0 R
            22 0 R
            23 0 R
            24 0 R
          ]
  >>
endobj
```



```

5 0 obj
  << /Limits [(Actinium) (Astatine)]           % Leaf node
    /Names [ (Actinium) 25 0 R
              (Aluminum) 26 0 R
              (Americium) 27 0 R
              (Antimony) 28 0 R
              (Argon) 29 0 R
              (Arsenic) 30 0 R
              (Astatine) 31 0 R
            ]
  >>
endobj
...
24 0 obj
  << /Limits [(Xenon) (Zirconium)]           % Leaf node
    /Names [ (Xenon) 129 0 R
              (Ytterbium) 130 0 R
              (Yttrium) 131 0 R
              (Zinc) 132 0 R
              (Zirconium) 133 0 R
            ]
  >>
endobj
25 0 obj
  89                                           % Atomic number (Actinium)
endobj
...
133 0 obj
  40                                           % Atomic number (Zirconium)
endobj

```

3.8.5 Number Trees

A *number tree* is similar to a name tree (see Section 3.8.4, “Name Trees”), except that its keys are integers instead of strings, sorted in ascending numerical order. The entries in the leaf (or root) nodes containing the key-value pairs are named **Nums** instead of **Names** as in a name tree. Table 3.25 shows the entries in a number tree’s node dictionaries.

TABLE 3.25 Entries in a number tree node dictionary

KEY	TYPE	VALUE
Kids	array	(<i>Root and intermediate nodes only; required in intermediate nodes; present in the root node if and only if Nums is not present</i>) An array of indirect references to the immediate children of this node. The children may be intermediate or leaf nodes.
Nums	array	(<i>Root and leaf nodes only; required in leaf nodes; present in the root node if and only if Kids is not present</i>) An array of the form $[key_1\ value_1\ key_2\ value_2\ \dots\ key_n\ value_n]$ where each key_i is an integer and the corresponding $value_i$ is an indirect reference to the object associated with that key. The keys are sorted in numerical order, analogously to the arrangement of keys in a name tree as described in Section 3.8.4, “Name Trees.”
Limits	array	(<i>Intermediate and leaf nodes only; required</i>) An array of two integers, specifying the (numerically) least and greatest keys included in the Nums array of a leaf node or in the Nums arrays of any leaf nodes that are descendants of an intermediate node.

3.9 Functions

PDF is not a programming language, and a PDF file is not a program; however, PDF does provide several types of *function object* (PDF 1.2) that represent parameterized classes of functions, including mathematical formulas and sampled representations with arbitrary resolution. Functions are used in various ways in PDF: device-dependent rasterization information for high-quality printing (half-tone spot functions and transfer functions), color transform functions for certain color spaces, and specification of colors as a function of position for smooth shadings.

Functions in PDF represent static, self-contained numerical transformations. A function to add two numbers has two input values and one output value:

$$f(x_0, x_1) = x_0 + x_1$$

Similarly, a function that computes the arithmetic and geometric mean of two numbers could be viewed as a function of two input values and two output values:

$$f(x_0, x_1) = \frac{x_0 + x_1}{2}, \sqrt{x_0 \times x_1}$$

In general, a function can take any number (m) of input values and produce any number (n) of output values:

$$f(x_0, \dots, x_{m-1}) = y_0, \dots, y_{n-1}$$

In PDF functions, all the input values and all the output values are numbers, and functions have no side effects.

Each function definition includes a *domain*, the set of legal values for the input. Some types of functions also define a *range*, the set of legal values for the output. Input values passed to the function are clipped to the domain, and output values produced by the function are clipped to the range. For example, suppose the function

$$f(x) = x + 2$$

is defined with a domain of $[-1 \ 1]$. If the function is called with the input value 6, that value is replaced with the nearest value in the defined domain, 1, before the function is evaluated; the resulting output value is therefore 3. Similarly, if the function

$$f(x_0, x_1) = 3 \times x_0 + x_1$$

is defined with a range of $[0 \ 100]$, and if the input values -6 and 4 are passed to the function (and are within its domain), then the output value produced by the function, -14 , is replaced with 0 , the nearest value in the defined range.

A function object may be a dictionary or a stream, depending on the type of function; the term *function dictionary* will be used generically in this section to refer to either a dictionary object or the dictionary portion of a stream object. A function dictionary specifies the function's representation, the set of attributes that parameterize that representation, and the additional data needed by that representation. Four types of function are available, as indicated by the dictionary's **FunctionType** entry:

- (PDF 1.2) A *sampled function* (type 0) uses a table of *sample values* to define the function. Various techniques are used to interpolate values between the sample values.
- (PDF 1.3) An *exponential interpolation function* (type 2) defines a set of coefficients for an exponential function.

- (PDF 1.3) A *stitching function* (type 3) is a combination of other functions, partitioned across a domain.
- (PDF 1.3) A *PostScript calculator function* (type 4) uses operators from the PostScript language to describe an arithmetic expression.

All function dictionaries share the entries listed in Table 3.26.

TABLE 3.26 Entries common to all function dictionaries

KEY	TYPE	VALUE
FunctionType	integer	(Required) The function type: <ul style="list-style-type: none"> 0 Sampled function 2 Exponential interpolation function 3 Stitching function 4 PostScript calculator function
Domain	array	(Required) An array of $2 \times m$ numbers, where m is the number of input values. For each i from 0 to $m - 1$, Domain _{$2i$} must be less than or equal to Domain _{$2i+1$} , and the i th input value, x_i , must lie in the interval Domain _{$2i$} $\leq x_i \leq$ Domain _{$2i+1$} . Input values outside the declared domain are clipped to the nearest boundary value.
Range	array	(Required for type 0 and type 4 functions, optional otherwise; see below) An array of $2 \times n$ numbers, where n is the number of output values. For each j from 0 to $n - 1$, Range _{$2j$} must be less than or equal to Range _{$2j+1$} , and the j th output value, y_j , must lie in the interval Range _{$2j$} $\leq y_j \leq$ Range _{$2j+1$} . Output values outside the declared range are clipped to the nearest boundary value. If this entry is absent, no clipping is done.

In addition, each type of function dictionary must include entries appropriate to the particular function type. The number of output values can usually be inferred from other attributes of the function; if not (as is always the case for type 0 and type 4 functions), the **Range** entry is required. The dimensionality of the function implied by the **Domain** and **Range** entries must be consistent with that implied by other attributes of the function.

3.9.1 Type 0 (Sampled) Functions

Type 0 functions use a sequence of *sample values* (contained in a stream) to provide an approximation for functions whose domains and ranges are bounded. The samples are organized as an m -dimensional table in which each entry has n components.

Sampled functions are highly general and offer reasonably accurate representations of arbitrary analytic functions at low expense. For example, a 1-input sinusoidal function can be represented over the range [0 180] with an average error of only 1 percent, using just ten samples and linear interpolation. Two-input functions require significantly more samples, but usually not a prohibitive number, so long as the function does not have high frequency variations.

The dimensionality of a sampled function is restricted only by implementation limits. However, the number of samples required to represent functions with high dimensionality multiplies rapidly unless the sampling resolution is very low. Also, the process of multilinear interpolation becomes computationally intensive if the number of inputs m is greater than 2. The multidimensional spline interpolation is even more computationally intensive.

In addition to the entries in Table 3.26, a type 0 function dictionary includes those shown in Table 3.27.

The **Domain**, **Encode**, and **Size** entries determine how the function's input variable values are mapped into the sample table. For example, if **Size** is [21 31], the default **Encode** array is [0 20 0 30], which maps the entire domain into the full set of sample table entries. Other values of **Encode** may be used.

To explain the relationship between **Domain**, **Encode**, **Size**, **Decode**, and **Range**, we use the following notation:

$$\begin{aligned} y &= \text{Interpolate}(x, x_{\min}, x_{\max}, y_{\min}, y_{\max}) \\ &= y_{\min} + \left((x - x_{\min}) \times \frac{y_{\max} - y_{\min}}{x_{\max} - x_{\min}} \right) \end{aligned}$$

For a given value of x , Interpolate calculates the y value on the line defined by the two points (x_{\min}, y_{\min}) and (x_{\max}, y_{\max}) .

TABLE 3.27 Additional entries specific to a type 0 function dictionary

KEY	TYPE	VALUE
Size	array	<i>(Required)</i> An array of m positive integers specifying the number of samples in each input dimension of the sample table.
BitsPerSample	integer	<i>(Required)</i> The number of bits used to represent each sample. (If the function has multiple output values, each one occupies BitsPerSample bits.) Valid values are 1, 2, 4, 8, 12, 16, 24, and 32.
Order	integer	<i>(Optional)</i> The order of interpolation between samples. Valid values are 1 and 3, specifying linear and cubic spline interpolation, respectively. (See implementation note 26 in Appendix H.) Default value: 1.
Encode	array	<i>(Optional)</i> An array of $2 \times m$ numbers specifying the linear mapping of input values into the domain of the function's sample table. Default value: $[0 (\mathbf{Size}_0 - 1) 0 (\mathbf{Size}_1 - 1) \dots]$.
Decode	array	<i>(Optional)</i> An array of $2 \times n$ numbers specifying the linear mapping of sample values into the range appropriate for the function's output values. Default value: same as the value of Range .
<i>other stream attributes</i>	(various)	<i>(Optional)</i> Other attributes of the stream that provides the sample values, as appropriate (see Table 3.4 on page 38).

When a sampled function is called, each input value x_i , for $0 \leq i < m$, is clipped to the domain:

$$x_i' = \min(\max(x_i, \mathbf{Domain}_{2i}), \mathbf{Domain}_{2i+1})$$

That value is encoded:

$$e_i = \text{Interpolate}(x_i', \mathbf{Domain}_{2i}, \mathbf{Domain}_{2i+1}, \mathbf{Encode}_{2i}, \mathbf{Encode}_{2i+1})$$

That value is clipped to the size of the sample table in that dimension:

$$e_i' = \min(\max(e_i, 0), \mathbf{Size}_i - 1)$$

The encoded input values are real numbers, not restricted to integers. Interpolation is then used to determine output values from the nearest surrounding values in the sample table. Each output value r_j , for $0 \leq j < n$, is then decoded:

$$r_j' = \text{Interpolate}(r_j, 0, 2^{\mathbf{BitsPerSample}} - 1, \mathbf{Decode}_{2j}, \mathbf{Decode}_{2j+1})$$

Finally, each decoded value is clipped to the range:

$$y_j = \min(\max(r_j', \mathbf{Range}_{2j}), \mathbf{Range}_{2j+1})$$

Sample data is represented as a stream of unsigned 8-bit bytes (integers in the range 0 to 255). The bytes constitute a continuous bit stream, with the high-order bit of each byte first. Each sample value is represented as a sequence of **BitsPerSample** bits. Successive values are adjacent in the bit stream; there is no padding at byte boundaries.

For a function with multidimensional input (more than one input variable), the sample values in the first dimension vary fastest, and the values in the last dimension vary slowest. For example, for a function $f(a, b, c)$, where a , b , and c vary from 0 to 9 in steps of 1, the sample values would appear in this order: $f(0, 0, 0)$, $f(1, 0, 0)$, ..., $f(9, 0, 0)$, $f(0, 1, 0)$, $f(1, 1, 0)$, ..., $f(9, 1, 0)$, $f(0, 2, 0)$, $f(1, 2, 0)$, ..., $f(9, 2, 0)$, $f(0, 3, 0)$, ..., $f(9, 9, 0)$, $f(0, 0, 1)$, $f(1, 0, 1)$, and so on.

For a function with multidimensional output (more than one output value), the values are stored in the same order as **Range**.

The stream data must be long enough to contain the entire sample array, as indicated by **Size**, **Range**, and **BitsPerSample**; see “Stream Extent” on page 37.

Example 3.13 illustrates a sampled function with 4-bit samples in an array containing 21 columns and 31 rows (651 values). The function takes two arguments, x and y , in the domain $[-1.0 \ 1.0]$, and returns one value, z , in that same range. The x argument is linearly transformed by the encoding to the domain $[0 \ 20]$ and the y argument to the domain $[0 \ 30]$. Using bilinear interpolation between sample points, the function computes a value for z , which (because **BitsPerSample** is 4) will be in the range $[0 \ 15]$, and the decoding transforms z to a number in the range $[-1.0 \ 1.0]$ for the result. The sample array is stored in a stream of 326 bytes, calculated as follows (rounded up):

$$326 \text{ bytes} = 31 \text{ rows} \times 21 \text{ samples/row} \times 4 \text{ bits/sample} \div 8 \text{ bits/byte}$$

The first byte contains the sample for the point $(-1.0, -1.0)$ in the high-order 4 bits and the sample for the point $(-0.9, -1.0)$ in the low-order 4 bits.

Example 3.13

```

14 0 obj
  << /FunctionType 0
    /Domain [-1.0 1.0 -1.0 1.0]
    /Size [21 31]
    /Encode [0 20 0 30]
    /BitsPerSample 4
    /Range [-1.0 1.0]
    /Decode [-1.0 1.0]
    /Length ...
    /Filter ...
  >>
stream
...651 sample values...
endstream
endobj

```

The **Decode** entry can be used creatively to increase the accuracy of encoded samples corresponding to certain values in the range. For example, if the desired range of the function is $[-1.0\ 1.0]$ and **BitsPerSample** is 4, the usual value of **Decode** would be $[-1.0\ 1.0]$ and the sample values would be integers in the interval $[0\ 15]$ (as shown in Figure 3.6). But if these values were used, the midpoint of the range, 0.0, would not be represented exactly by any sample value, since it would fall halfway between 7 and 8. On the other hand, if the **Decode** array were $[-1.0\ +1.1429]$ (1.1429 being approximately equal to $16 \div 14$) and the sample values supplied were in the interval $[0\ 14]$, then the desired effective range of $[-1.0\ 1.0]$ would be achieved, and the range value 0.0 would be represented by the sample value 7.




FIGURE 3.6 Mapping with the **Decode** array

The **Size** value for an input dimension can be 1, in which case all input values in that dimension will be mapped to the single allowed value. If **Size** is less than 4, cubic spline interpolation is not possible and **Order** 3 will be ignored if specified.

3.9.2 Type 2 (Exponential Interpolation) Functions

Type 2 functions (*PDF 1.3*) include a set of parameters that define an *exponential interpolation* of one input value and n output values:

$$f(x) = y_0, \dots, y_{n-1}$$

In addition to the entries in Table 3.26 on page 108, a type 2 function dictionary includes those; listed in Table 3.28. (See implementation note 27 in Appendix H.)

TABLE 3.28 Additional entries specific to a type 2 function dictionary

KEY	TYPE	VALUE
C0	array	(<i>Optional</i>) An array of n numbers defining the function result when $x = 0.0$ (hence the “0” in the name). Default value: [0.0].
C1	array	(<i>Optional</i>) An array of n numbers defining the function result when $x = 1.0$ (hence the “1” in the name). Default value: [1.0].
N	number	(<i>Required</i>) The interpolation exponent. Each input value x will return n values, given by $y_j = \mathbf{C0}_j + x^{\mathbf{N}} \times (\mathbf{C1}_j - \mathbf{C0}_j)$, for $0 \leq j < n$.

Values of **Domain** must constrain x in such a way that if **N** is not an integer, all values of x must be nonnegative, and if **N** is negative, no value of x may be zero. Typically, **Domain** will be declared as [0.0 1.0], and **N** will be a positive number. The **Range** attribute is optional and can be used to clip the output to a desired range. Note that when **N** is 1, the function performs a linear interpolation between **C0** and **C1**; this can also be expressed as a sampled function (type 0).

3.9.3 Type 3 (Stitching) Functions

Type 3 functions (*PDF 1.3*) define a “stitching” of the subdomains of several 1-input functions to produce a single new 1-input function. Since the resulting *stitching function* is a 1-input function, the domain is given by a two-element array, [**Domain**₀ **Domain**₁].

In addition to the entries in Table 3.26 on page 108, a type 3 function dictionary includes those listed in Table 3.29. (See implementation note 28 in Appendix H.)

TABLE 3.29 Additional entries specific to a type 3 function dictionary

KEY	TYPE	VALUE
Functions	array	<i>(Required)</i> An array of k 1-input functions making up the stitching function. The output dimensionality of all functions must be the same, and compatible with the value of Range if Range is present.
Bounds	array	<i>(Required)</i> An array of $k - 1$ numbers that, in combination with Domain , define the intervals to which each function from the Functions array applies. Bounds elements must be in order of increasing value, and each value must be within the domain defined by Domain .
Encode	array	<i>(Required)</i> An array of $2 \times k$ numbers that, taken in pairs, map each subset of the domain defined by Domain and the Bounds array to the domain of the corresponding function.

Domain must be of size 2 (that is, $m = 1$), and **Domain**₀ must be strictly less than **Domain**₁ unless $k = 1$. The domain is partitioned into k subdomains, as indicated by the dictionary's **Bounds** entry, which is an array of $k - 1$ numbers that obey the following relationships (with exceptions as noted below):

$$\mathbf{Domain}_0 < \mathbf{Bounds}_0 < \mathbf{Bounds}_1 < \dots < \mathbf{Bounds}_{k-2} < \mathbf{Domain}_1$$

The **Bounds** array describes a series of half-open intervals, closed on the left and open on the right (except the last, which is closed on the right as well). The value of the **Functions** entry is an array of k functions. The first function applies to x values in the first subdomain, $\mathbf{Domain}_0 \leq x < \mathbf{Bounds}_0$; the second function applies to x values in the second subdomain, $\mathbf{Bounds}_0 \leq x < \mathbf{Bounds}_1$; and so on. The last function applies to x values in the last subdomain, which includes the upper bound: $\mathbf{Bounds}_{k-2} \leq x \leq \mathbf{Domain}_1$. The value of k may be 1, in which case the **Bounds** array is empty and the single item in the **Functions** array applies to all x values, $\mathbf{Domain}_0 \leq x \leq \mathbf{Domain}_1$.

The **Encode** array contains $2 \times k$ numbers. A value x from the i th subdomain is encoded as follows:

$$x' = \text{Interpolate}(x, \mathbf{Bounds}_{i-1}, \mathbf{Bounds}_i, \mathbf{Encode}_{2i}, \mathbf{Encode}_{2i+1})$$

for $0 \leq i < k$. In this equation, **Bounds**₋₁ means **Domain**₀, and **Bounds**_{*k*-1} means **Domain**₁. If the last bound, **Bounds**_{*k*-2}, is equal to **Domain**₁, then x' is defined to be **Encode**_{2*i*}.

The stitching function is designed to make it easy to combine several functions to be used within one shading pattern, over different parts of the shading's domain. (Shading patterns are discussed in Section 4.6.3, "Shading Patterns.") The same effect could be achieved by creating a separate shading dictionary for each of the functions, with adjacent domains. However, since each shading would have similar parameters, and because the overall effect is one shading, it is more convenient to have a single shading with multiple function definitions.

Also, type 3 functions provide a general mechanism for inverting the domains of 1-input functions. For example, consider a function f with a **Domain** of [0.0 1.0], and a stitching function g with a **Domain** of [0.0 1.0], a **Functions** array containing f , and an **Encode** array of [1.0 0.0]. In effect, $g(x) = f(1 - x)$.

3.9.4 Type 4 (PostScript Calculator) Functions

A type 4 function (*PDF 1.3*), also called a *PostScript calculator function*, is represented as a stream containing code written in a small subset of the PostScript language. While any function can be sampled (in a type 0 PDF function) and others can be described with exponential functions (type 2 in PDF), type 4 functions offer greater flexibility and potentially greater accuracy. For example, a tint transformation function for a hexachrome (six-component) **DeviceN** color space with an alternate color space of **DeviceCMYK** (see "DeviceN Color Spaces" on page 205) requires a 6-in, 4-out function. If such a function were sampled with m values for each input variable, the number of samples, $4 \times m^6$, could be prohibitively large. In practice, such functions are often written as short, simple PostScript functions. (See implementation note 29 in Appendix H.)

Type 4 functions also make it possible to include a wide variety of halftone spot functions without the loss of accuracy that comes from sampling, and without adding to the list of predefined spot functions (see Section 6.4.2, "Spot Functions"). All of the predefined spot functions can be written as type 4 functions.

The language that can be used in a type 4 function contains expressions involving integers, real numbers, and boolean values only. There are no composite data structures such as strings or arrays, no procedures, and no variables or names. Table 3.30 lists the operators that can be used in this type of function. (For more

information on these operators, see Appendix B of the *PostScript Language Reference*, Third Edition.) Although the semantics are those of the corresponding PostScript operators, a PostScript interpreter is not required.

TABLE 3.30 Operators in type 4 functions

OPERATOR TYPE	OPERATORS				
Arithmetic operators	abs	cvi	floor	mod	sin
	add	cvr	idiv	mul	sqrt
	atan	div	ln	neg	sub
	ceiling	exp	log	round	truncate
	cos				
Relational, boolean, and bitwise operators	and	false	le	not	true
	bitshift	ge	lt	or	xor
	eq	gt	ne		
Conditional operators	if	ifelse			
Stack operators	copy	exch	pop		
	dup	index	roll		

The operand syntax for type 4 functions follows PDF conventions rather than PostScript conventions. The entire code stream defining the function is enclosed in braces { }. Braces also delimit expressions that are executed conditionally by the **if** and **ifelse** operators:

```
boolean {expression} if
boolean {expression1} {expression2} ifelse
```

Note that this is a purely syntactic construct; unlike in PostScript, no “procedure objects” are involved.

124

A type 4 function dictionary includes the entries in Table 3.26 on page 108, as well as other stream attributes as appropriate (see Table 3.4 on page 38). Example 3.14 shows a type 4 function equivalent to the predefined spot function **Double-Dot** (see Section 6.4.2, “Spot Functions”).

Example 3.14

```
10 0 obj
  << /FunctionType 4
    /Domain [-1.0 1.0 -1.0 1.0]
    /Range [-1.0 1.0]
    /Length 71
  >>
stream
{ 360 mul sin
  2 div
  exch 360 mul sin
  2 div
  add
}
endstream
endobj
```

The **Domain** and **Range** entries are both required. The input variables constitute the initial operand stack; the items remaining on the operand stack after execution of the function are the output variables. It is an error for the number of remaining operands to differ from the number of output variables specified by **Range**, or for any of them to be objects other than numbers.

Implementations of type 4 functions must provide a stack with room for at least 100 entries. No implementation is required to provide a larger stack, and it is an error to overflow the stack.

Although any integers or real numbers that may appear in the stream fall under the same implementation limits (defined in Appendix C) as in other contexts, the *intermediate* results in type 4 function computations do not. An implementation may use a representation that exceeds those limits. Operations on real numbers, for example, might use single-precision or double-precision floating-point numbers. (See implementation note 30 in Appendix H.)

Errors in Type 4 Functions

The code that reads a type 4 function (analogous to the PostScript *scanner*) must detect and report syntax errors. It may also be able to detect some errors that will occur when the function is used, although this is not always possible. Any errors

detected by the scanner are considered to be errors in the PDF file itself and are handled like other errors in the file.

The code that executes a type 4 function (analogous to the PostScript *interpreter*) must detect and report errors. PDF does not define a representation for the errors; those details are provided by the application that processes the PDF file. The following types of errors can occur (among others):

- Stack overflow
- Stack underflow
- A type error (for example, applying **not** to a real number)
- A range error (for example, applying **sqrt** to a negative number)
- An undefined result (for example, dividing by 0)

3.10 File Specifications

A PDF file can refer to the contents of another file by using a *file specification* (PDF 1.1), which can take either of two forms. A *simple* file specification gives just the name of the target file in a standard format, independent of the naming conventions of any particular file system; a *full* file specification includes information related to one or more specific file systems. A simple file specification may take the form of either a string or a dictionary; a full file specification can only be represented as a dictionary.

Although the file designated by a file specification is normally external to the PDF file referring to it, PDF 1.3 permits a copy of the external file to be embedded within the PDF file itself, allowing its contents to be stored or transmitted along with the PDF file. However, embedding a file does not change the presumption that it is external to the PDF file. Consequently, in order for the PDF file to be processed correctly, it may be necessary to copy the embedded files it contains back into a local file system.

3.10.1 File Specification Strings

The standard format for representing a simple file specification in string form divides the string into component substrings separated by the slash character (/). The slash is a generic component separator that is mapped to the appropriate

platform-specific separator when generating a platform-dependent file name. Any of the components may be empty. If a component contains one or more literal slashes, each must be preceded by a backslash (\), which in turn must be preceded by another backslash to indicate that it is part of the string and not an escape character. For example, the string

```
(in\\out)
```

represents the file name

```
in/out
```

The backslashes are removed in processing the string; they are needed only to distinguish the component values from the component separators. The component substrings are stored as bytes and are passed to the operating system without interpretation or conversion of any sort.

Absolute and Relative File Specifications

A simple file specification that begins with a slash is an *absolute* file specification. The last component is the file name; the preceding components specify its context. In some file specifications, the file name may be empty; for example, URL (uniform resource locator) specifications can specify directories instead of files. A file specification that does not begin with a slash is a *relative* file specification giving the location of the file relative to that of the PDF file containing it.

In the case of a URL-based file system, the rules of Internet RFC 1808, *Relative Uniform Resource Locators* (see the Bibliography), are used to compute an absolute URL from a relative file specification and the specification of the PDF file. Prior to this process, the relative file specification is converted to a relative URL by using the escape mechanism of RFC 1738, *Uniform Resource Locators*, to represent any bytes that would be either “unsafe” according to RFC 1738 or not representable in 7-bit U.S. ASCII. In addition, such URL-based relative file specifications are limited to paths as defined in RFC 1808; the scheme, network location/login, fragment identifier, query information, and parameter sections are not allowed.

In the case of other file systems, a relative file specification is converted to an absolute file specification by removing the file name component from the specifica-

tion of the containing PDF file and appending the relative file specification in its place. For example, the relative file specification

```
ArtFiles/Figure1.pdf
```

appearing in a PDF file whose specification is

```
/HardDisk/PDFDocuments/AnnualReport/Summary.pdf
```

yields the absolute specification

```
/HardDisk/PDFDocuments/AnnualReport/ArtFiles/Figure1.pdf
```

The special component `..` (two periods) can be used in a relative file specification to move up a level in the file system hierarchy. When the component immediately preceding `..` is not another `..`, the two cancel each other; both are eliminated from the file specification and the process is repeated. Thus in the example above, the relative file specification

```
../../ArtFiles/Figure1.pdf
```

would yield the absolute specification

```
/HardDisk/ArtFiles/Figure1.pdf
```

Conversion to Platform-Dependent File Names

The conversion of a file specification into a platform-dependent file name depends on the specific file naming conventions of each platform. For example:

- For DOS, the initial component is either a physical or logical drive identifier or a network resource name as returned by the Microsoft Windows function `WNetGetConnection`, and is followed by a colon (`:`). A network resource name is constructed from the first two components; the first component is the server name and the second is the share name (volume name). All components are then separated by backslashes. It is possible to specify an absolute DOS path without a drive by making the first component empty. (Empty components are ignored by other platforms.)
- For Mac OS, all components are separated by colons (`:`).

- For UNIX, all components are separated by slashes (/). An initial slash, if present, is preserved.

Strings used to specify a file name are interpreted in the standard encoding for the platform on which the document is being viewed. Table 3.31 shows examples of file specifications on the most common platforms.

TABLE 3.31 Examples of file specifications

SYSTEM	SYSTEM-DEPENDENT PATHS	WRITTEN FORM
DOS	\pdfdocs\spec.pdf (no drive) r:\pdfdocs\spec.pdf pplib/eng:\pdfdocs\spec.pdf	(//pdfdocs/spec.pdf) (/r/pdfdocs/spec.pdf) (/pplib/eng/pdfdocs/spec.pdf)
Mac OS	Mac HD:PDFDocs:spec.pdf	(/Mac HD/PDFDocs/spec.pdf)
UNIX	/user/fred/pdfdocs/spec.pdf pdfdocs/spec.pdf (relative)	(/user/fred/pdfdocs/spec.pdf) (pdfdocs/spec.pdf)

When creating documents that are to be viewed on multiple platforms, care must be taken to ensure file name compatibility. Only a subset of the U.S. ASCII character set should be used in file specifications: the uppercase alphabetic characters (A–Z), the numeric characters (0–9), and the underscore (_). The period (.) has special meaning in DOS and Windows file names, and as the first character in a Mac OS pathname. In file specifications, the period should be used only to separate a base file name from a file extension.

Some file systems are case-insensitive, so names within a directory should remain distinguishable if lowercase letters are changed to uppercase or vice versa. On DOS and Windows 3.1 systems and on some CD-ROM file systems, file names are limited to 8 characters plus a 3-character extension. File system software typically converts long names to short names by retaining the first 6 or 7 characters of the file name and the first 3 characters after the last period, if any. Since characters beyond the sixth or seventh are often converted to other values unrelated to the original value, file names must be distinguishable from the first 6 characters.

Multiple-Byte Strings in File Specifications

In PDF 1.2 or higher, a file specification may contain multiple-byte character codes, represented in hexadecimal form between angle brackets (< and >). Since the slash character <2F> is used as a component delimiter and the backslash <5C> is used as an escape character, any occurrence of either of these bytes in a multiple-byte character must be preceded by the ASCII code for the backslash character. For example, a file name containing the 2-byte character code <89 5C> must write it as <89 5C 5C>. When the viewer application encounters this sequence of bytes in a file name, it replaces the sequence with the original 2-byte code.

3.10.2 File Specification Dictionaries

The dictionary form of file specification provides more flexibility than the string form, allowing different files to be specified for different file systems or platforms, or for file systems other than the standard ones (DOS/Windows, Mac OS, and UNIX). Table 3.32 shows the entries in a file specification dictionary. Viewer applications running on a particular platform should use the appropriate platform-specific entry (**DOS**, **Mac**, or **Unix**) if available. If the required platform-specific entry is not present and there is no file system entry (**FS**), the generic **F** entry should be used as a simple file specification.

TABLE 3.32 Entries in a file specification dictionary

KEY	TYPE	VALUE
Type	name	<i>(Required if an EF or RF entry is present; recommended always)</i> The type of PDF object that this dictionary describes; must be Filespec for a file specification dictionary.
FS	name	<i>(Optional)</i> The name of the file system to be used to interpret this file specification. If this entry is present, all other entries in the dictionary are interpreted by the designated file system. PDF defines only one standard file system, URL (see Section 3.10.4, “URL Specifications”); a viewer application or plug-in extension can register a different one (see Appendix E). Note that this entry is independent of the F , DOS , Mac , and Unix entries.
F	string	<i>(Required if the DOS, Mac, and Unix entries are all absent)</i> A file specification string of the form described in Section 3.10.1, “File Specification Strings,” or (if the file system is URL) a uniform resource locator, as described in Section 3.10.4, “URL Specifications.”

DOS	string	<i>(Optional)</i> A file specification string (see Section 3.10.1, “File Specification Strings”) representing a DOS file name.
Mac	string	<i>(Optional)</i> A file specification string (see Section 3.10.1, “File Specification Strings”) representing a Mac OS file name.
Unix	string	<i>(Optional)</i> A file specification string (see Section 3.10.1, “File Specification Strings”) representing a UNIX file name.
ID	array	<i>(Optional)</i> An array of two strings constituting a file identifier (see Section 9.3, “File Identifiers”) that is also included in the referenced file. The use of this entry improves a viewer application’s chances of finding the intended file and allows it to warn the user if the file has changed since the link was made.
V	boolean	<i>(Optional; PDF 1.2)</i> A flag indicating whether the file referenced by the file specification is <i>volatile</i> (changes frequently with time). If the value is true , viewer applications should never cache a copy of the file. For example, a movie annotation referencing a URL to a live video camera could set this flag to true , notifying the application that it should reacquire the movie each time it is played. Default value: false .
EF	dictionary	<i>(Required if RF is present; PDF 1.3)</i> A dictionary containing a subset of the keys F , DOS , Mac , and Unix , corresponding to the entries by those names in the file specification dictionary. The value of each such key is an embedded file stream (see Section 3.10.3, “Embedded File Streams”) containing the corresponding file. If this entry is present, the Type entry is required and the file specification dictionary must be indirectly referenced.
RF	dictionary	<i>(Optional; PDF 1.3)</i> A dictionary with the same structure as the EF dictionary, which must also be present. Each key in the RF dictionary must also be present in the EF dictionary. Each value is a related files array (see “Related Files Arrays” on page 125) identifying files that are related to the corresponding file in the EF dictionary. If this entry is present, the Type entry is required and the file specification dictionary must be indirectly referenced.

3.10.3 Embedded File Streams

File specifications ordinarily refer to files external to the PDF file in which they occur. To preserve the integrity of the PDF file, this requires that all external files it refers to must accompany it when it is archived or transmitted. *Embedded file streams (PDF 1.3)* address this problem by allowing the contents of the referenced files to be embedded directly within the body of the PDF file itself. For example, if the file contains OPI (Open Prepress Interface) dictionaries that refer to externally stored high-resolution images (see Section 9.10.6, “Open Prepress Interface

(OPI”)), the image data can be incorporated into the PDF file with embedded file streams. This makes the PDF file a self-contained unit that can be stored or transmitted as a single entity. (The embedded files are included purely for convenience, and need not be directly processed by any PDF consumer application.)

One way of including an embedded file stream in a PDF document is through the **EF** entry in the file specification dictionary. This method still requires a file specification to be provided, associating a location in the file system with the stream data. It is also possible to embed data in a PDF document independently of any file system: the data’s presence in the document is all that is required. This can be done using an embedded file stream referenced by a private entry in a dictionary. However, PDF 1.4 defines a standard location for the data: the **EmbeddedFiles** entry in the PDF document’s name dictionary can contain a name tree that maps name strings to embedded file streams (see Section 3.6.3, “Name Dictionary”).

***Note:** The relationship between the name string provided in the name dictionary and the associated embedded file stream is an artificial one for data management purposes. This allows embedded files to be easily identified by the author of the document, in much the same way that the JavaScript name tree associates name strings with document-level JavaScript actions (see “JavaScript Actions” on page 556).*

The stream dictionary describing an embedded file contains the standard entries for any stream, such as **Length** and **Filter** (see Table 3.4 on page 38), as well as the additional entries shown in Table 3.33.

TABLE 3.33 Additional entries in an embedded file stream dictionary

KEY	TYPE	VALUE
Type	name	<i>(Optional)</i> The type of PDF object that this dictionary describes; if present, must be EmbeddedFile for an embedded file stream.
Subtype	name	<i>(Optional)</i> The subtype of the embedded file. The value of this entry must be a first-class name, as defined in Appendix E. Names without a registered prefix must conform to the MIME media type names defined in Internet RFC 2046, <i>Multipurpose Internet Mail Extensions (MIME), Part Two: Media Types</i> (see the Bibliography), with the provision that characters not allowed in names must use the 2-character hexadecimal code format described in Section 3.2.4, “Name Objects.”
Params	dictionary	<i>(Optional)</i> An <i>embedded file parameter dictionary</i> containing additional, file-specific information (see Table 3.34).

TABLE 3.34 Entries in an embedded file parameter dictionary

KEY	TYPE	VALUE
Size	integer	<i>(Optional)</i> The size of the embedded file, in bytes.
CreationDate	date	<i>(Optional)</i> The date and time when the embedded file was created.
ModDate	date	<i>(Optional)</i> The date and time when the embedded file was last modified.
Mac	dictionary	<i>(Optional)</i> A subdictionary containing additional information specific to Mac OS files (see Table 3.35).
Checksum	string	<i>(Optional)</i> A 16-byte string that is the checksum of the bytes of the uncompressed embedded file. The checksum is calculated by applying the standard MD5 message-digest algorithm (described in Internet RFC 1321, <i>The MD5 Message-Digest Algorithm</i> ; see the Bibliography) to the bytes of the embedded file stream.

For Mac OS files, the **Mac** entry in the embedded file parameter dictionary holds a further subdictionary containing Mac OS–specific file information. Table 3.35 shows the contents of this subdictionary.

TABLE 3.35 Entries in a Mac OS file information dictionary

KEY	TYPE	VALUE
Subtype	string	<i>(Optional)</i> The embedded file’s file type.
Creator	string	<i>(Optional)</i> The embedded file’s creator signature.
ResFork	stream	<i>(Optional)</i> The binary contents of the embedded file’s resource fork.

Related Files Arrays

In some circumstances, a PDF file can refer to a group of related files, such as the set of five files that make up a DCS 1.0 color-separated image. The file specification explicitly names only one of the files; the rest are identified by some systematic variation of that file name (such as by altering the extension). When such a file is to be embedded in a PDF file, the related files must be embedded as well. This is accomplished by including a *related files array* (PDF 1.3) as the value of the

RF entry in the file specification dictionary. The array has $2 \times n$ elements, which are paired in the form

```
[ string1 stream1
  string2 stream2
  ...
  stringn streamn
]
```

The first element of each pair is a string giving the name of one of the related files; the second element is an embedded file stream holding the file's contents.

In Example 3.15, objects 21, 31, and 41 are embedded file streams containing the DOS file SUNSET.EPS, the Mac OS file Sunset.eps, and the UNIX file Sunset.eps, respectively. The file specification dictionary's **RF** entry specifies an array, object 30, identifying a set of embedded files related to the Mac OS file, forming a DCS 1.0 set. The example shows only the first two embedded file streams in the set; an actual PDF file would of course include all of them.

Example 3.15

```
10 0 obj                                % File specification dictionary
  << /Type /Filespec
    /DOS (SUNSET.EPS)
    /Mac (Sunset.eps)                    % Name of Mac OS file
    /Unix (Sunset.eps)
    /EF << /DOS 21 0 R
          /Mac 31 0 R                    % Embedded Mac OS file
          /Unix 41 0 R
    >>
    /RF << /Mac 30 0 R >>                % Related files array for Mac OS file
  >>
endobj

30 0 obj                                % Related files array for Mac OS file
  [ (Sunset.eps) 31 0 R                  % Includes file Sunset.eps itself
    (Sunset.C) 32 0 R
    (Sunset.M) 33 0 R
    (Sunset.Y) 34 0 R
    (Sunset.K) 35 0 R
  ]
endobj
```

```
31 0 obj                                % Embedded file stream for Mac OS file
  << /Type /EmbeddedFile                %  Sunset.eps
    /Length ...
    /Filter ...
  >>
stream
...Data for Sunset.eps...
endstream
endobj

32 0 obj                                % Embedded file stream for related file
  << /Type /EmbeddedFile                %  Sunset.C
    /Length ...
    /Filter ...
  >>
stream
...Data for Sunset.C...
endstream
endobj
```

3.10.4 URL Specifications

When the **FS** entry in a file specification dictionary has the value **URL**, the value of the **F** entry in that dictionary is not a file specification string, but a uniform resource locator (URL) of the form defined in Internet RFC 1738, *Uniform Resource Locators* (see the Bibliography). Example 3.16 shows a URL specification.

Example 3.16

```
<< /FS /URL
  /F (ftp://www.beatles.com/Movies/AbbeyRoad.mov)
>>
```

The URL must adhere to the character-encoding requirements specified in RFC 1738. Because 7-bit U.S. ASCII is a strict subset of **PDFDocEncoding**, this value may also be considered to be in that encoding.

3.10.5 Maintenance of File Specifications

The techniques described in this section can be used to maintain the integrity of the file specifications within a PDF file during operations such as the following:

- Updating the relevant file specification when a referenced file is renamed
- Determining the complete collection of files that must be copied to a mirror site
- When creating new links to external files, discovering existing file specifications that refer to the same files and sharing them
- Finding the file specifications associated with embedded files to be packed or unpacked

It is not possible, in general, to find all file specification strings in a PDF file, because there is no way to determine whether a given string is a file specification string. It is possible, however, to find all file specification *dictionaries*, provided that they meet the following conditions:

- They are indirect objects.
- They contain a **Type** entry whose value is the name **Filespec**.

An application can then locate all of the file specification dictionaries by traversing the PDF file's cross-reference table (see Section 3.4.3, "Cross-Reference Table") and finding all dictionaries with **Type** keys whose value is **Filespec**. For this reason, it is highly recommended that all file specifications be expressed in dictionary form and meet the conditions stated above. Note that any file specification dictionary specifying embedded files (that is, one that contains an **EF** entry) *must* satisfy these conditions (see Table 3.32 on page 122).

***Note:** It may not be possible to locate file specification dictionaries that are direct objects, since they are neither self-typed nor necessarily reachable via any standard path of object references.*

Files may be embedded in a PDF file either directly, using the **EF** entry in a file specification dictionary, or indirectly, using related files arrays specified in the **RF** entry. If a file is embedded indirectly, its name is given by the string that precedes the embedded file stream in the related files array; if it is embedded directly, its name is obtained from the value of the corresponding entry in the file specifica-

tion dictionary. In Example 3.15 on page 126, for instance, the **EF** dictionary has a **DOS** entry identifying object number 21 as an embedded file stream; the name of the embedded DOS file, `SUNSET.EPS`, is given by the **DOS** entry in the file specification dictionary.

A given external file may be referenced from more than one file specification. Therefore, when embedding a file with a given name, it is necessary to check for other occurrences of the same name as the value associated with the corresponding key in other file specification dictionaries. This requires finding all embeddable file specifications and, for each matching key, checking for both of the following conditions:

- The string value associated with the key matches the name of the file being embedded.
- A value has not already been embedded for the file specification. (If there is already a corresponding key in the **EF** dictionary, then a file has already been embedded for that use of the file name.)

Note that there is no requirement that the files associated with a given file name be unique. The same file name, such as `readme.txt`, may be associated with different embedded files in distinct file specifications.

CHAPTER 4

Graphics

THE GRAPHICS OPERATORS used in PDF content streams describe the appearance of pages that are to be reproduced on a raster output device. The facilities described in this chapter are intended for both printer and display applications.

The graphics operators form six main groups:

- *Graphics state operators* manipulate the data structure called the *graphics state*, the global framework within which the other graphics operators execute. The graphics state includes the *current transformation matrix* (CTM), which maps user space coordinates used within a PDF content stream into output device coordinates. It also includes the *current color*, the *current clipping path*, and many other parameters that are implicit operands of the painting operators.
- *Path construction operators* specify *paths*, which define shapes, line trajectories, and regions of various sorts. They include operators for beginning a new path, adding line segments and curves to it, and closing it.
- *Path-painting operators* fill a path with a color, paint a stroke along it, or use it as a clipping boundary.
- *Other painting operators* paint certain self-describing graphics objects. These include sampled images, geometrically defined shadings, and entire content streams that in turn contain sequences of graphics operators.
- *Text operators* select and show *character glyphs* from *fonts* (descriptions of typefaces for representing text characters). Because PDF treats glyphs as general graphical shapes, many of the text operators could be grouped with the graphics state or painting operators. However, the data structures and mechanisms for dealing with glyph and font descriptions are sufficiently specialized that Chapter 5 focuses on them.

- *Marked-content operators* associate higher-level logical information with objects in the content stream. This information does not affect the rendered appearance of the content; it is useful to applications that use PDF for document interchange. Marked content is described in Section 9.5, “Marked Content.”

This chapter presents general information about device-independent graphics in PDF: how a PDF content stream describes the abstract appearance of a page. *Rendering*—the device-dependent part of graphics—is covered in Chapter 6. The Bibliography lists a number of books that give details of these computer graphics concepts and their implementation.

4.1 Graphics Objects

As discussed in Section 3.7.1, “Content Streams,” the data in a content stream is interpreted as a sequence of *operators* and their *operands*, expressed as basic data objects according to standard PDF syntax. A content stream can describe the appearance of a page, or it can be treated as a graphical element in certain other contexts.

The operands and operators are written sequentially using postfix notation. Although this notation resembles the sequential execution model of the PostScript language, a PDF content stream is not a program to be interpreted; rather, it is a static description of a sequence of *graphics objects*. There are specific rules, described below, for writing the operands and operators that describe a graphics object.

PDF provides five types of graphics object:

- A *path object* is an arbitrary shape made up of straight lines, rectangles, and cubic Bézier curves. A path may intersect itself and may have disconnected sections and holes. A path object ends with one or more painting operators that specify whether the path is stroked, filled, used as a clipping boundary, or some combination of these operations.
- A *text object* consists of one or more character strings that identify sequences of glyphs to be painted. Like a path, text can be stroked, filled, or used as a clipping boundary.
- An *external object* (*XObject*) is an object defined outside the content stream and referenced as a named resource (see Section 3.7.2, “Resource Dictionaries”). The interpretation of an *XObject* depends on its type. An *image XObject* defines

a rectangular array of color samples to be painted; a *form XObject* is an entire content stream to be treated as a single graphics object. Specialized types of form XObject are used to import content from one PDF file into another (*reference XObjects*) and to group graphical elements together as a unit for various purposes (*group XObjects*). In particular, the latter are used to define *transparency groups* for use in the transparent imaging model (*transparency group XObjects*, discussed in detail in Chapter 7). There is also a *PostScript XObject*, whose use is discouraged.

- An *inline image object* is a means of expressing the data for a small image directly within the content stream, using a special syntax.
- A *shading object* describes a geometric shape whose color is an arbitrary function of position within the shape. (A shading can also be treated as a color when painting other graphics objects; it is not considered to be a separate graphics object in that case.)

PDF 1.3 and earlier versions use an *opaque imaging model* in which each graphics object is painted in sequence, completely obscuring any previous marks it may overlay on the page. PDF 1.4 introduces a new *transparent imaging model* in which objects can be less than fully opaque, allowing previously painted marks to show through. Each object is painted on the page with a specified *opacity*, which may be constant at every point within the object's shape or may vary from point to point. The previously existing contents of the page form a *backdrop* with which the new object is *composited*, producing results that combine the colors of the object and backdrop according to their respective opacity characteristics. The objects at any given point on the page can be thought of as forming a *transparency stack*, where the stacking order is defined to be the order in which the objects are specified, bottommost object first. All objects in the stack can potentially contribute to the result, depending on their colors, shapes, and opacities.

PDF's graphics parameters are so arranged that objects are painted by default with full opacity, reducing the behavior of the transparent imaging model to that of the opaque model. Accordingly, the material in this chapter applies to both the opaque and transparent models except where explicitly stated otherwise; the transparent model is described in its full generality in Chapter 7.

Although the painting behavior described above is often attributed to individual operators making up an object, it is always the object as a whole that is painted. Figure 4.1 shows the ordering rules for the operations that define graphics objects. Some operations are permitted only in certain types of graphics object or

in the intervals between graphics objects (called the *page description level* in the figure). Every content stream begins at the page description level, where changes can be made to the graphics state, such as colors and text attributes, as discussed in the following sections.

In the figure, arrows indicate the operators that mark the beginning or end of each type of graphics object. Some operators are identified individually, others by general category. Table 4.1 summarizes these categories for all PDF operators. For example, the path construction operators **m** and **re** signal the beginning of a path object. Inside the path object, additional path construction operators are permitted, as are the clipping path operators **W** and **W***, but not general graphics state operators such as **w** or **J**. A path-painting operator, such as **S** or **f**, ends the path object and returns to the page description level.

TABLE 4.1 Operator categories

CATEGORY	OPERATORS	TABLE	PAGE
General graphics state	w, J, j, M, d, ri, i, gs	4.7	156
Special graphics state	q, Q, cm	4.7	156
Path construction	m, l, c, v, y, h, re	4.9	163
Path painting	S, s, f, F, f*, B, B*, b, b*, n	4.10	167
Clipping paths	W, W*	4.11	172
Text objects	BT, ET	5.4	308
Text state	Tc, Tw, Tz, TL, Tf, Tr, Ts	5.2	302
Text positioning	Td, TD, Tm, T*	5.5	310
Text showing	Tj, TJ, ', "	5.6	311
Type 3 fonts	d0, d1	5.10	326
Color	CS, cs, SC, SCN, sc, scn, G, g, RG, rg, K, k	4.21	216
Shading patterns	sh	4.24	232
Inline images	BI, ID, EI	4.38	278
XObjects	Do	4.34	261
Marked content	MP, DP, BMC, BDC, EMC	9.8	584
Compatibility	BX, EX	3.20	95

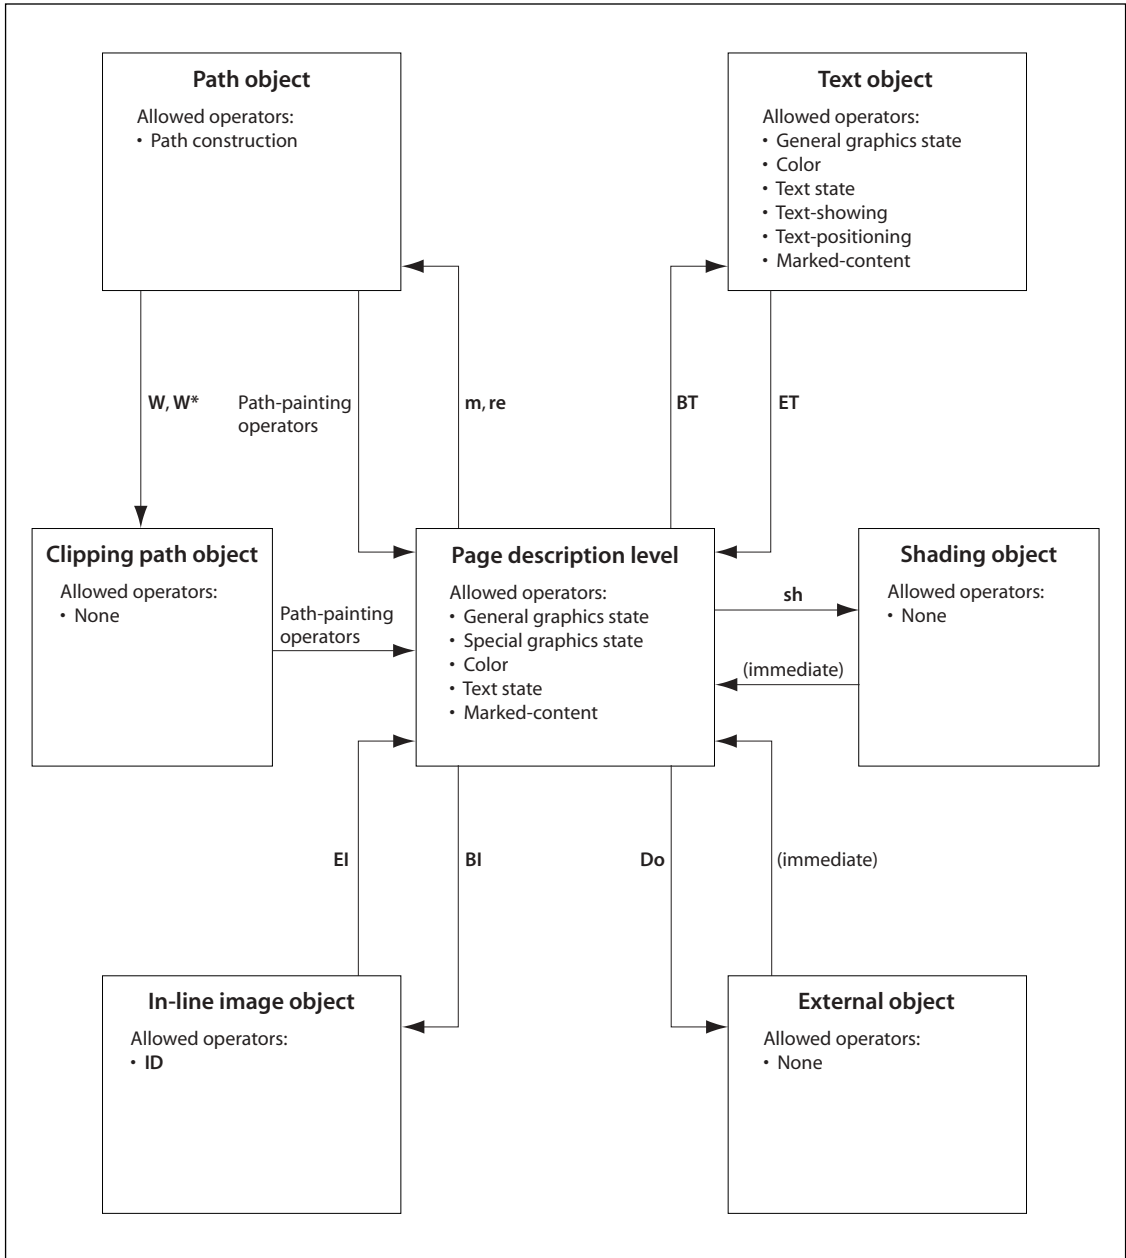


FIGURE 4.1 Graphics objects

Note: A content stream whose operations violate these rules for describing graphics objects can produce unpredictable behavior, even though it may display and print correctly. Applications that attempt to extract graphics objects for editing or other purposes depend on the objects' being well formed. The rules for graphics objects are also important for the proper interpretation of marked content (see Section 9.5, "Marked Content").

A graphics object also implicitly includes all graphics state parameters that affect its behavior. For instance, a path object depends on the value of the current color parameter at the moment the path object is defined. The effect is as if this parameter were specified as part of the definition of the path object. However, the operators that are invoked at the page description level to set graphics state parameters are *not* considered to belong to any particular graphics object. Graphics state parameters need to be specified only when they change. A graphics object may depend on parameters that were defined much earlier.

Similarly, the individual character strings within a text object implicitly include the graphics state parameters on which they depend. Most of these parameters may be set either inside or outside the text object. The effect is as if they were separately specified for each text string.

The important point is that there is no semantic significance to the exact arrangement of graphics state operators. An application that reads and writes a PDF content stream is not required to preserve this arrangement, but is free to change it to any other arrangement that achieves the same values of the relevant graphics state parameters for each graphics object. An application should not infer any higher-level logical semantics from the arrangement of tokens constituting a graphics object. A separate mechanism, *marked content* (see Section 9.5, "Marked Content"), allows such higher-level information to be explicitly associated with the graphics objects.

4.2 Coordinate Systems

Coordinate systems define the canvas on which all painting occurs. They determine the position, orientation, and size of the text, graphics, and images that appear on a page. This section describes each of the coordinate systems used in PDF, how they are related, and how transformations among them are specified.

4.2.1 Coordinate Spaces

Paths and positions are defined in terms of pairs of *coordinates* on the Cartesian plane. A coordinate pair is a pair of real numbers x and y that locate a point horizontally and vertically within a two-dimensional *coordinate space*. A coordinate space is determined by the following properties with respect to the current page:

- The location of the origin
- The orientation of the x and y axes
- The lengths of the units along each axis

PDF defines several coordinate spaces in which the coordinates specifying graphics objects are interpreted. The following sections describe these spaces and the relationships among them.

Transformations among coordinate spaces are defined by *transformation matrices*, which can specify any linear mapping of two-dimensional coordinates, including translation, scaling, rotation, reflection, and skewing. Transformation matrices are discussed in Sections 4.2.2, “Common Transformations,” and 4.2.3, “Transformation Matrices.”

Device Space

The contents of a page ultimately appear on a raster output device such as a display or a printer. Such devices vary greatly in the built-in coordinate systems they use to address pixels within their imageable areas. A particular device’s coordinate system is called its *device space*. The origin of the device space on different devices can fall in different places on the output page; on displays, the origin can vary depending on the window system. Because the paper or other output medium moves through different printers and imagesetters in different directions, the axes of their device spaces may be oriented differently; for instance, vertical (y) coordinates may increase from the top of the page to the bottom on some devices and from bottom to top on others. Finally, different devices have different resolutions; some even have resolutions that differ in the horizontal and vertical directions.

If coordinates in a PDF file were specified in device space, the file would be device-dependent and would appear differently on different devices. For example, images specified in the typical device spaces of a 72-pixel-per-inch display and a 600-dot-per-inch printer would differ in size by more than a factor of 8; an 8-inch line segment on the display would appear less than 1 inch long on the printer. Figure 4.2 shows how the same graphics object, specified in device space, can appear drastically different when rendered on different output devices.

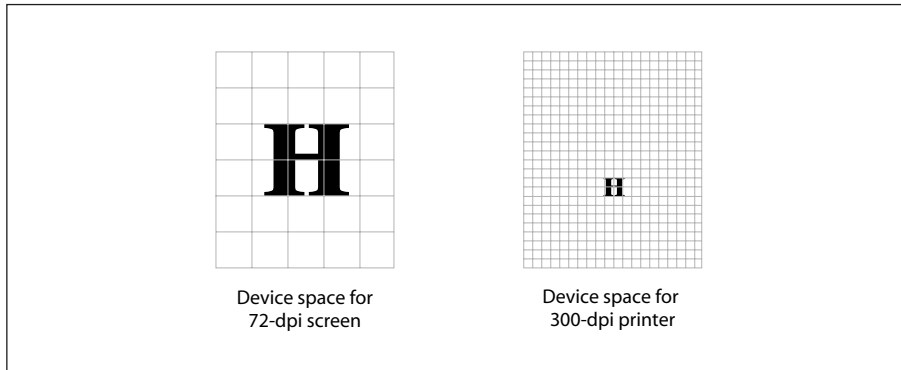


FIGURE 4.2 *Device space*

User Space

To avoid the device-dependent effects of specifying objects in device space, PDF defines a device-independent coordinate system that always bears the same relationship to the current page, regardless of the output device on which printing or displaying will occur. This device-independent coordinate system is called *user space*.

The user space coordinate system is initialized to a default state for each page of a document. The **CropBox** entry in the page dictionary specifies the rectangle of user space corresponding to the visible area of the intended output medium (display window or printed page). The positive x axis extends horizontally to the right and the positive y axis vertically upward, as in standard mathematical practice (subject to alteration by the **Rotate** entry in the page dictionary). The length of a unit along both the x and y axes is $1/72$ inch. This coordinate system is called *default user space*.

***Note:** In PostScript, the origin of default user space always corresponds to the lower-left corner of the output medium. While this convention is common in PDF documents as well, it is not required; the page dictionary’s **CropBox** entry can specify any rectangle of default user space to be made visible on the medium.*

***Note:** The unit size in default user space (1/72 inch) is approximately the same as a point, a unit widely used in the printing industry. It is not exactly the same, however; there is no universal definition of a point.*

Conceptually, user space is an infinite plane. Only a small portion of this plane corresponds to the imageable area of the output device: a rectangular region defined by the **CropBox** entry in the page dictionary. The region of default user space that is viewed or printed can be different for each page, and is described in Section 9.10.1, “Page Boundaries.”

***Note:** Because coordinates in user space (as in any other coordinate space) may be specified as either integers or real numbers, the unit size in default user space does not constrain positions to any arbitrary grid. The resolution of coordinates in user space is not related in any way to the resolution of pixels in device space.*

The transformation from user space to device space is defined by the *current transformation matrix* (CTM), an element of the PDF graphics state (see Section 4.3, “Graphics State”). A PDF viewer application can adjust the CTM for the native resolution of a particular output device, maintaining the device-independence of the PDF page description itself. Figure 4.3 shows how this allows an object specified in user space to appear the same regardless of the device on which it is rendered.

The default user space provides a consistent, dependable starting place for PDF page descriptions regardless of the output device used. If necessary, a PDF content stream may then modify user space to be more suitable to its needs by applying the *coordinate transformation operator*, **cm** (see Section 4.3.3, “Graphics State Operators”). Thus what may appear to be absolute coordinates in a content stream are not absolute with respect to the current page, because they are expressed in a coordinate system that may slide around and shrink or expand. Coordinate system transformation not only enhances device-independence but is a useful tool in its own right. For example, a content stream originally composed to occupy an entire page can be incorporated without change as an element of another page by shrinking the coordinate system in which it is drawn.




FIGURE 4.3 *User space*

Other Coordinate Spaces

In addition to device space and user space, PDF uses a variety of other coordinate spaces for specialized purposes:

- The coordinates of text are specified in *text space*. The transformation from text space to user space is defined by a *text matrix* in combination with several text-related parameters in the graphics state (see Section 5.3.1, “Text-Positioning Operators”).
- Character glyphs in a font are defined in *glyph space* (see Section 5.1.3, “Glyph Positioning and Metrics”). The transformation from glyph space to text space is defined by the *font matrix*. For most types of font, this matrix is predefined to map 1000 units of glyph space to 1 unit of text space; for Type 3 fonts, the font matrix is given explicitly in the font dictionary (see Section 5.5.4, “Type 3 Fonts”).

- All sampled images are defined in *image space*. The transformation from image space to user space is predefined and cannot be changed. All images are 1 unit wide by 1 unit high in user space, regardless of the number of samples in the image. To be painted, an image must be mapped to the desired region of the page by temporarily altering the current transformation matrix (CTM).

Note: In *PostScript*, unlike *PDF*, the relationship between image space and user space can be specified explicitly. The fixed transformation prescribed in *PDF* corresponds to the convention that is recommended for use in *PostScript*.

- A form XObject (discussed in Section 4.9, “Form XObjects”) is a self-contained content stream that can be treated as a graphical element within another content stream. The space in which it is defined is called *form space*. The transformation from form space to user space is specified by a *form matrix* contained in the form XObject.
- PDF 1.2 defines a type of color known as a *pattern*, discussed in Section 4.6, “Patterns.” A pattern is defined either by a content stream that is invoked repeatedly to tile an area or by a shading whose color is a function of position. The space in which a pattern is defined is called *pattern space*. The transformation from pattern space to user space is specified by a *pattern matrix* contained in the pattern.

Relationships among Coordinate Spaces

Figure 4.4 shows the relationships among the coordinate spaces described above. Each arrow in the figure represents a transformation from one coordinate space to another. PDF allows modifications to many of these transformations.

Because PDF coordinate spaces are defined relative to one another, changes made to one transformation can affect the appearance of objects defined in several coordinate spaces. For example, a change in the CTM, which defines the transformation from user space to device space, will affect forms, text, images, and patterns, since they are all “upstream” from user space.

4.2.2 Common Transformations

A *transformation matrix* specifies the relationship between two coordinate spaces. By modifying a transformation matrix, objects can be scaled, rotated, translated, or transformed in other ways.




FIGURE 4.4 Relationships among coordinate systems

A transformation matrix in PDF is specified by six numbers, usually in the form of an array containing six elements. In its most general form, this array is denoted $[a\ b\ c\ d\ e\ f]$; it can represent any linear transformation from one coordinate system to another. This section lists the arrays that specify the most common transformations; Section 4.2.3, “Transformation Matrices,” discusses more mathematical details of transformations, including information on specifying transformations that are combinations of those listed here.

- Translations are specified as $[1\ 0\ 0\ 1\ t_x\ t_y]$, where t_x and t_y are the distances to translate the origin of the coordinate system in the horizontal and vertical dimensions, respectively.
- Scaling is obtained by $[s_x\ 0\ 0\ s_y\ 0\ 0]$. This scales the coordinates so that 1 unit in the horizontal and vertical dimensions of the new coordinate system is the same size as s_x and s_y units, respectively, in the previous coordinate system.
- Rotations are produced by $[\cos\ \theta\ \sin\ \theta\ -\sin\ \theta\ \cos\ \theta\ 0\ 0]$, which has the effect of rotating the coordinate system axes by an angle θ counterclockwise.
- Skew is specified by $[1\ \tan\ \alpha\ \tan\ \beta\ 1\ 0\ 0]$, which skews the x axis by an angle α and the y axis by an angle β .

Figure 4.5 shows examples of each transformation. The directions of translation, rotation, and skew shown in the figure correspond to positive values of the array elements.




FIGURE 4.5 *Effects of coordinate transformations*

If several transformations are combined, the order in which they are applied is significant. For example, first scaling and then translating the x axis is not the same as first translating and then scaling it. In general, to obtain the expected results, transformations should be done in the following order:

1. Translate
2. Rotate
3. Scale or skew

Figure 4.6 shows the effect of the order in which transformations are applied. The figure shows two sequences of transformations applied to a coordinate system. After each successive transformation, an outline of the letter n is drawn.




FIGURE 4.6 *Effect of transformation order*

The transformations shown in the figure are as follows:

- A translation of 10 units in the x direction and 20 units in the y direction
- A rotation of 30 degrees
- A scaling by a factor of 3 in the x direction

In the figure, the axes are shown with a dash pattern having a 2-unit dash and a 2-unit gap. In addition, the original (untransformed) axes are shown in a lighter color for reference. Notice that the scale-rotate-translate ordering results in a distortion of the coordinate system, leaving the x and y axes no longer perpendicular, while the recommended translate-rotate-scale ordering does not.

4.2.3 Transformation Matrices

This section discusses the mathematics of transformation matrices. It is not necessary to read this section in order to use the transformations described previ-

ously; the information is presented for the benefit of readers who want to gain a deeper understanding of the theoretical basis of coordinate transformations.

To understand the mathematics of coordinate transformations in PDF, it is vital to remember two points:

- *Transformations alter coordinate systems, not graphics objects.* All objects painted before a transformation is applied are unaffected by the transformation. Objects painted after the transformation is applied will be interpreted in the transformed coordinate system.
- *Transformation matrices specify the transformation from the new (transformed) coordinate system to the original (untransformed) coordinate system.* All coordinates used after the transformation are expressed in the transformed coordinate system. PDF applies the transformation matrix to find the equivalent coordinates in the untransformed coordinate system.

Note: *Many computer graphics textbooks consider transformations of graphics objects rather than of coordinate systems. Although either approach is correct and self-consistent, some details of the calculations differ depending on which point of view is taken.*

PDF represents coordinates in a two-dimensional space. The point (x, y) in such a space can be expressed in vector form as $[x \ y \ 1]$. The constant third element of this vector (1) is needed so that the vector can be used with 3-by-3 matrices in the calculations described below.

The transformation between two coordinate systems is represented by a 3-by-3 transformation matrix written as

$$\begin{bmatrix} a & b & 0 \\ c & d & 0 \\ e & f & 1 \end{bmatrix}$$

Because a transformation matrix has only six elements that can be changed, it is usually specified in PDF as the six-element array $[a \ b \ c \ d \ e \ f]$.

Coordinate transformations are expressed as matrix multiplications:

$$[x' \ y' \ 1] = [x \ y \ 1] \times \begin{bmatrix} a & b & 0 \\ c & d & 0 \\ e & f & 1 \end{bmatrix}$$

Because PDF transformation matrices specify the conversion from the transformed coordinate system to the original (untransformed) coordinate system, x' and y' in this equation are the coordinates in the untransformed coordinate system, while x and y are the coordinates in the transformed system. Carrying out the multiplication, we have

$$\begin{aligned} x' &= a \times x + c \times y + e \\ y' &= b \times x + d \times y + f \end{aligned}$$

If a series of transformations is carried out, the matrices representing each of the individual transformations can be multiplied together to produce a single equivalent matrix representing the composite transformation.

Matrix multiplication is not commutative—the order in which matrices are multiplied is significant. Consider a sequence of two transformations: a scaling transformation applied to the user space coordinate system, followed by a conversion from the resulting scaled user space to device space. Let M_S be the matrix specifying the scaling and M_C the current transformation matrix, which transforms user space to device space. Recalling that coordinates are always specified in the transformed space, the correct order of transformations must first convert the scaled coordinates to default user space and then the default user space coordinates to device space. This can be expressed as

$$X_D = X_U \times M_C = (X_S \times M_S) \times M_C = X_S \times (M_S \times M_C)$$

where

X_D denotes the coordinates in device space

X_U denotes the coordinates in default user space

X_S denotes the coordinates in scaled user space

This shows that when a new transformation is concatenated with an existing one, the matrix representing it must be multiplied *before* (*premultiplied* with) the existing transformation matrix.

This result is true in general for PDF: when a sequence of transformations is carried out, the matrix representing the combined transformation (M') is calculated by premultiplying the matrix representing the additional transformation (M_T) with the one representing all previously existing transformations (M):

$$M' = M_T \times M$$

Note: When rendering graphics objects, it is sometimes necessary for a viewer application to perform the inverse of a transformation—that is, to find the user space coordinates that correspond to a given pair of device space coordinates. Not all transformations are invertible, however. For example, if a matrix contains a, b, c, and d elements that are all zero, all user coordinates map to the same device coordinates and there is no unique inverse transformation. Such noninvertible transformations are not very useful and generally arise from unintended operations, such as scaling by 0. Use of a noninvertible matrix when painting graphics objects can result in unpredictable behavior.

4.3 Graphics State

A PDF viewer application maintains an internal data structure called the *graphics state* that holds current graphics control parameters. These parameters define the global framework within which the graphics operators execute. For example, the **f** (fill) operator implicitly uses the *current color* parameter, and the **S** (stroke) operator additionally uses the *current line width* parameter from the graphics state.

The graphics state is initialized at the beginning of each page, using the default values specified in Tables 4.2 and 4.3. Table 4.2 lists those graphics state parameters that are device-independent and are appropriate to specify in page descriptions. The parameters listed in Table 4.3 control details of the rendering (scan conversion) process and are device-dependent; a page description that is intended to be device-independent should not modify these parameters.

TABLE 4.2 Device-independent graphics state parameters

PARAMETER	TYPE	VALUE
CTM	array	The <i>current transformation matrix</i> , which maps positions from user coordinates to device coordinates (see Section 4.2, “Coordinate Systems”). This matrix is modified by each application of the coordinate transformation operator, cm . Initial value: a matrix that transforms default user coordinates to device coordinates.
clipping path	(internal)	The <i>current clipping path</i> , which defines the boundary against which all output is to be cropped (see Section 4.4.3, “Clipping Path Operators”). Initial value: the boundary of the entire imageable portion of the output page.
color space	name or array	The <i>current color space</i> in which color values are to be interpreted (see Section 4.5, “Color Spaces”). There are two separate color space parameters: one for stroking and one for all other painting operations. Initial value: DeviceGray .
color	(various)	The <i>current color</i> to be used during painting operations (see Section 4.5, “Color Spaces”). The type and interpretation of this parameter depend on the current color space; for most color spaces, a color value consists of one to four numbers. There are two separate color parameters: one for stroking and one for all other painting operations. Initial value: black.
text state	(various)	A set of nine graphics state parameters that pertain only to the painting of text. These include parameters that select the font, scale the glyphs to an appropriate size, and accomplish other effects. The text state parameters are described in Section 5.2, “Text State Parameters and Operators.”
line width	number	The thickness, in user space units, of paths to be stroked (see “Line Width” on page 152). Initial value: 1.0.
line cap	integer	A code specifying the shape of the endpoints for any open path that is stroked (see “Line Cap Style” on page 153). Initial value: 0, for square butt caps.
line join	integer	A code specifying the shape of joints between connected segments of a stroked path (see “Line Join Style” on page 153). Initial value: 0, for mitered joins.
miter limit	number	The maximum length of mitered line joins for stroked paths (see “Miter Limit” on page 153). This parameter limits the length of “spikes” produced when line segments join at sharp angles. Initial value: 10.0, for a miter cutoff below approximately 11.5 degrees.

dash pattern	array and number	A description of the dash pattern to be used when paths are stroked (see “Line Dash Pattern” on page 155). Initial value: a solid line.
rendering intent	name	The <i>rendering intent</i> to be used when converting CIE-based colors to device colors (see “Rendering Intents” on page 197). Default value: RelativeColorimetric .
stroke adjustment	boolean	(PDF 1.2) A flag specifying whether to compensate for possible rasterization effects when stroking a path with a line width that is small relative to the pixel resolution of the output device (see Section 6.5.4, “Automatic Stroke Adjustment”). Note that this is considered a device-independent parameter, even though the details of its effects are device-dependent. Initial value: false .
blend mode	name or array	(PDF 1.4) The <i>current blend mode</i> to be used in the transparent imaging model (see Sections 7.2.4, “Blend Mode,” and 7.5.2, “Specifying Blending Color Space and Blend Mode”). This parameter is implicitly reset to its initial value at the beginning of execution of a transparency group XObject (see Section 7.5.5, “Transparency Group XObjects”). Initial value: Normal .
soft mask	dictionary or name	(PDF 1.4) A <i>soft-mask dictionary</i> (see “Soft-Mask Dictionaries” on page 445) specifying the mask shape or mask opacity values to be used in the transparent imaging model (see “Source Shape and Opacity” on page 421 and “Mask Shape and Opacity” on page 443), or the name None if no such mask is specified. This parameter is implicitly reset to its initial value at the beginning of execution of a transparency group XObject (see Section 7.5.5, “Transparency Group XObjects”). Initial value: None .
alpha constant	number	(PDF 1.4) The constant shape or constant opacity value to be used in the transparent imaging model (see “Source Shape and Opacity” on page 421 and “Constant Shape and Opacity” on page 444). There are two separate alpha constant parameters: one for stroking and one for all other painting operations. This parameter is implicitly reset to its initial value at the beginning of execution of a transparency group XObject (see Section 7.5.5, “Transparency Group XObjects”). Initial value: 1.0.
alpha source	boolean	(PDF 1.4) A flag specifying whether the current soft mask and alpha constant parameters are to be interpreted as shape values (true) or opacity values (false). This flag also governs the interpretation of the SMask entry, if any, in an image dictionary (see Section 4.8.4, “Image Dictionaries”). Initial value: false .

TABLE 4.3 Device-dependent graphics state parameters

PARAMETER	TYPE	VALUE
overprint	boolean	<i>(PDF 1.2)</i> A flag specifying (on output devices that support the overprint control feature) whether painting in one set of colorants should cause the corresponding areas of other colorants to be erased (false) or left unchanged (true); see Section 4.5.6, “Overprint Control.” In PDF 1.3, there are two separate overprint parameters: one for stroking and one for all other painting operations. Initial value: false .
overprint mode	number	<i>(PDF 1.3)</i> A code specifying whether a color component value of 0 in a DeviceCMYK color space should erase that component (0) or leave it unchanged (1) when overprinting (see Section 4.5.6, “Overprint Control”). Initial value: 0.
black generation	function or name	<i>(PDF 1.2)</i> A function that calculates the level of the black color component to use when converting <i>RGB</i> colors to <i>CMYK</i> (see Section 6.2.3, “Conversion from DeviceRGB to DeviceCMYK”). Initial value: installation-dependent.
undercolor removal	function or name	<i>(PDF 1.2)</i> A function that calculates the reduction in the levels of the cyan, magenta, and yellow color components to compensate for the amount of black added by black generation (see Section 6.2.3, “Conversion from DeviceRGB to DeviceCMYK”). Initial value: installation-dependent.
transfer	function, array, or name	<i>(PDF 1.2)</i> A function that adjusts device gray or color component levels to compensate for nonlinear response in a particular output device (see Section 6.3, “Transfer Functions”). Initial value: installation-dependent.
halftone	dictionary, stream, or name	<i>(PDF 1.2)</i> A halftone screen for gray and color rendering, specified as a halftone dictionary or stream (see Section 6.4, “Halftones”). Initial value: installation-dependent.
flatness	number	The precision with which curves are to be rendered on the output device (see Section 6.5.1, “Flatness Tolerance”). The value of this parameter gives the maximum error tolerance, measured in output device pixels; smaller numbers give smoother curves at the expense of more computation and memory use. Initial value: 1.0.

smoothness	number	(PDF 1.3) The precision with which color gradients are to be rendered on the output device (see Section 6.5.2, “Smoothness Tolerance”). The value of this parameter gives the maximum error tolerance, expressed as a fraction of the range of each color component; smaller numbers give smoother color transitions at the expense of more computation and memory use. Initial value: installation-dependent.
------------	--------	--

Some graphics state parameters are set with specific PDF operators, some are set by including a particular entry in a *graphics state parameter dictionary*, and some can be specified either way. The current line width, for example, can be set either with the **w** operator or (in PDF 1.3) with the **LW** entry in a graphics state parameter dictionary, whereas the current color is set only with specific operators and the current halftone is set only with a graphics state parameter dictionary. It is expected that all future graphics state parameters will be specified with new entries in the graphics state parameter dictionary rather than with new operators.

In general, the operators that set graphics state parameters simply store them unchanged for later use by the painting operators. However, some parameters have special properties or behavior:

- Most parameters must be of the correct type or have values that fall within a certain range.
- Parameters that are numeric values, such as the current color, line width, and miter limit, are forced into valid range, if necessary. However, they are *not* adjusted to reflect capabilities of the raster output device, such as resolution or number of distinguishable colors. Painting operators perform such adjustments, but the adjusted values are not stored back into the graphics state.
- Paths are internal objects that are not directly represented in PDF.

Note: As indicated in Tables 4.2 and 4.3, some of the parameters—*color space*, *color*, and *overprint*—have two values, one used for stroking (of paths and text objects) and one for all other painting operations. The two parameter values can be set independently, allowing for operations such as combined filling and stroking of the same path with different colors. Except where noted, a term such as *current color* should be interpreted to refer to whichever color parameter applies to the operation being performed. When necessary, the individual color parameters are distinguished explicitly as the *stroking color* and the *nonstroking color*.

4.3.1 Graphics State Stack

A well-structured PDF document typically contains many graphical elements that are essentially independent of each other and sometimes nested to multiple levels. The *graphics state stack* allows these elements to make local changes to the graphics state without disturbing the graphics state of the surrounding environment. The stack is a LIFO (last in, first out) data structure in which the contents of the graphics state can be saved and later restored using the following operators:

- The **q** operator pushes a copy of the entire graphics state onto the stack.
- The **Q** operator restores the entire graphics state to its former value by popping it from the stack.

These operators can be used to encapsulate a graphical element so that it can modify parameters of the graphics state and later restore them to their previous values. Occurrences of the **q** and **Q** operators must be balanced within a given content stream (or within the sequence of streams specified in a page dictionary's **Contents** array).

4.3.2 Details of Graphics State Parameters

This section gives details of several of the device-independent graphics state parameters listed in Table 4.2 on page 148.

Line Width




The *line width* parameter specifies the thickness of the line used to stroke a path. It is a nonnegative number expressed in user space units; stroking a path entails painting all points whose perpendicular distance from the path in user space is less than or equal to half the line width. The effect produced in device space depends on the current transformation matrix (CTM) in effect at the time the path is stroked. If the CTM specifies scaling by different factors in the horizontal and vertical dimensions, the thickness of stroked lines in device space will vary according to their orientation. The actual line width achieved can differ from the requested width by as much as 2 device pixels, depending on the positions of lines with respect to the pixel grid. Automatic stroke adjustment can be used to ensure uniform line width; see Section 6.5.4, “Automatic Stroke Adjustment.”

A line width of 0 denotes the thinnest line that can be rendered at device resolution: 1 device pixel wide. However, some devices cannot reproduce 1-pixel lines, and on high-resolution devices, they are nearly invisible. Since the results of rendering such “zero-width” lines are device-dependent, their use is not recommended.

Line Cap Style

The *line cap style* specifies the shape to be used at the ends of open subpaths (and dashes, if any) when they are stroked. Table 4.4 shows the possible values.

TABLE 4.4 Line cap styles

STYLE	APPEARANCE	DESCRIPTION
0		<i>Butt cap.</i> The stroke is squared off at the endpoint of the path. There is no projection beyond the end of the path.
1		<i>Round cap.</i> A semicircular arc with a diameter equal to the line width is drawn around the endpoint and filled in.
2		<i>Projecting square cap.</i> The stroke continues beyond the endpoint of the path for a distance equal to half the line width and is then squared off.




Line Join Style

The *line join style* specifies the shape to be used at the corners of paths that are stroked. Table 4.5 shows the possible values. Join styles are significant only at points where consecutive segments of a path connect at an angle; segments that meet or intersect fortuitously receive no special treatment.

Miter Limit

When two line segments meet at a sharp angle and mitered joins have been specified as the line join style, it is possible for the miter to extend far beyond the thickness of the line stroking the path. The *miter limit* imposes a maximum on the ratio of the miter length to the line width (see Figure 4.7). When the limit is exceeded, the join is converted from a miter to a bevel.

TABLE 4.5 Line join styles

STYLE	APPEARANCE	DESCRIPTION
0		<i>Miter join.</i> The outer edges of the strokes for the two segments are extended until they meet at an angle, as in a picture frame. If the segments meet at too sharp an angle (as defined by the miter limit parameter—see “Miter Limit,” above), a bevel join is used instead.
1		<i>Round join.</i> A circle with a diameter equal to the line width is drawn around the point where the two segments meet and is filled in, producing a rounded corner. <i>Note:</i> If path segments shorter than half the line width meet at a sharp angle, an unintended “wrong side” of the circle may appear.
2		<i>Bevel join.</i> The two segments are finished with butt caps (see “Line Cap Style” on page 153) and the resulting notch beyond the ends of the segments is filled with a triangle.

The ratio of miter length to line width is directly related to the angle φ between the segments in user space by the formula

$$\frac{\text{miterLength}}{\text{lineWidth}} = \frac{1}{\sin\left(\frac{\varphi}{2}\right)}$$

For example, a miter limit of 1.414 converts miters to bevels for φ less than 90 degrees, a limit of 2.0 converts them for φ less than 60 degrees, and a limit of 10.0 converts them for φ less than approximately 11.5 degrees.










FIGURE 4.7 Miter length

Line Dash Pattern

The *line dash pattern* controls the pattern of dashes and gaps used to stroke paths. It is specified by a *dash array* and a *dash phase*. The dash array's elements are numbers that specify the lengths of alternating dashes and gaps; the dash phase specifies the distance into the dash pattern at which to start the dash. The elements of both the dash array and the dash phase are expressed in user space units.

Before beginning to stroke a path, the dash array is cycled through, adding up the lengths of dashes and gaps. When the accumulated length equals the value specified by the dash phase, stroking of the path begins, using the dash array cyclically from that point onward. Table 4.6 shows examples of line dash patterns. As can be seen from the table, an empty dash array and zero phase can be used to restore the dash pattern to a solid line.

TABLE 4.6 Examples of line dash patterns

DASH ARRAY AND PHASE	APPEARANCE	DESCRIPTION
[] 0		No dash; solid, unbroken lines
[3] 0		3 units on, 3 units off, ...
[2] 1		1 on, 2 off, 2 on, 2 off, ...
[2 1] 0		2 on, 1 off, 2 on, 1 off, ...
[3 5] 6		2 off, 3 on, 5 off, 3 on, 5 off, ...
[2 3] 11		1 on, 3 off, 2 on, 3 off, 2 on, ...

Dashed lines wrap around curves and corners just as solid stroked lines do. The ends of each dash are treated with the current line cap style, and corners within dashes are treated with the current line join style. A stroking operation takes no measures to coordinate the dash pattern with features of the path; it simply dispenses dashes and gaps along the path in the pattern defined by the dash array.

When a path consisting of several subpaths is stroked, each subpath is treated independently—that is, the dash pattern is restarted and the dash phase is reapplied to it at the beginning of each subpath.

4.3.3 Graphics State Operators

Table 4.7 shows the operators that set the values of parameters in the graphics state. (See also the color operators listed in Table 4.21 on page 216 and the text state operators in Table 5.2 on page 302.)

TABLE 4.7 Graphics state operators

OPERANDS	OPERATOR	DESCRIPTION
—	q	Save the current graphics state on the graphics state stack (see “Graphics State Stack” on page 152).
—	Q	Restore the graphics state by removing the most recently saved state from the stack and making it the current state (see “Graphics State Stack” on page 152).
<i>a b c d e f</i>	cm	Modify the current transformation matrix (CTM) by concatenating the specified matrix (see Section 4.2.1, “Coordinate Spaces”). Although the operands specify a matrix, they are written as six separate numbers, not as an array.
<i>lineWidth</i>	w	Set the line width in the graphics state (see “Line Width” on page 152).
<i>lineCap</i>	J	Set the line cap style in the graphics state (see “Line Cap Style” on page 153).
<i>lineJoin</i>	j	Set the line join style in the graphics state (see “Line Join Style” on page 153).
<i>miterLimit</i>	M	Set the miter limit in the graphics state (see “Miter Limit” on page 153).
<i>dashArray dashPhase</i>	d	Set the line dash pattern in the graphics state (see “Line Dash Pattern” on page 155).
<i>intent</i>	ri	(PDF 1.1) Set the color rendering intent in the graphics state (see “Rendering Intents” on page 197).
<i>flatness</i>	i	Set the flatness tolerance in the graphics state (see Section 6.5.1, “Flatness Tolerance”). <i>flatness</i> is a number in the range 0 to 100; a value of 0 specifies the output device’s default flatness tolerance.
<i>dictName</i>	gs	(PDF 1.2) Set the specified parameters in the graphics state. <i>dictName</i> is the name of a graphics state parameter dictionary in the ExtGState sub-dictionary of the current resource dictionary (see the next section).

4.3.4 Graphics State Parameter Dictionaries

While some parameters in the graphics state can be set with individual operators, as shown in Table 4.7, others cannot. The latter can only be set with the generic graphics state operator **gs** (PDF 1.2). The operand supplied to this operator is the name of a *graphics state parameter dictionary* whose contents specify the values of one or more graphics state parameters. This name is looked up in the **ExtGState** subdictionary of the current resource dictionary. (The name **ExtGState**, for “extended graphics state,” is a vestige of earlier versions of PDF.)

Note: The *graphics state parameter dictionary* is also used by type 2 patterns, which do not have a content stream in which the graphics state operators could be invoked (see Section 4.6.3, “Shading Patterns”).

Each entry in the parameter dictionary specifies the value of an individual graphics state parameter, as shown in Table 4.8. It is not necessary for all entries to be present for every invocation of the **gs** operator; the parameter dictionary supplied may include any desired combination of parameter entries. The results of **gs** are cumulative; parameter values established in previous invocations will persist until explicitly overridden. Note that some parameters appear in both Tables 4.7 and 4.8; these parameters can be set either with individual graphics state operators or with **gs**. It is expected that any future extensions to the graphics state will be implemented by adding new entries to the graphics state parameter dictionary, rather than by introducing new graphics state operators.

TABLE 4.8 Entries in a graphics state parameter dictionary

KEY	TYPE	DESCRIPTION
Type	name	(Optional) The type of PDF object that this dictionary describes; must be ExtGState for a graphics state parameter dictionary.
LW	number	(Optional; PDF 1.3) The line width (see “Line Width” on page 152).
LC	integer	(Optional; PDF 1.3) The line cap style (see “Line Cap Style” on page 153).
LJ	integer	(Optional; PDF 1.3) The line join style (see “Line Join Style” on page 153).
ML	number	(Optional; PDF 1.3) The miter limit (see “Miter Limit” on page 153).
D	array	(Optional; PDF 1.3) The line dash pattern, expressed as an array of the form [<i>dashArray dashPhase</i>], where <i>dashArray</i> is itself an array and <i>dashPhase</i> is an integer (see “Line Dash Pattern” on page 155).

RI	name	<i>(Optional; PDF 1.3)</i> The name of the rendering intent (see “Rendering Intents” on page 197).
OP	boolean	<i>(Optional)</i> A flag specifying whether to apply overprint (see Section 4.5.6, “Overprint Control”). In PDF 1.2 and earlier, there is a single overprint parameter that applies to all painting operations. Beginning with PDF 1.3, there are two separate overprint parameters: one for stroking and one for all other painting operations. Specifying an OP entry sets both parameters unless there is also an op entry in the same graphics state parameter dictionary, in which case the OP entry sets only the overprint parameter for stroking.
op	boolean	<i>(Optional; PDF 1.3)</i> A flag specifying whether to apply overprint (see Section 4.5.6, “Overprint Control”) for painting operations other than stroking. If this entry is absent, the OP entry, if any, sets this parameter.
OPM	integer	<i>(Optional; PDF 1.3)</i> The overprint mode (see Section 4.5.6, “Overprint Control”).
Font	array	<i>(Optional; PDF 1.3)</i> An array of the form [<i>font size</i>], where <i>font</i> is an indirect reference to a font dictionary and <i>size</i> is a number expressed in text space units. These two objects correspond to the operands of the Tf operator (see Section 5.2, “Text State Parameters and Operators”); however, the first operand is an indirect object reference instead of a resource name.
BG	function	<i>(Optional)</i> The black-generation function, which maps the interval [0.0 1.0] to the interval [0.0 1.0] (see Section 6.2.3, “Conversion from DeviceRGB to DeviceCMYK”).
BG2	function or name	<i>(Optional; PDF 1.3)</i> Same as BG except that the value may also be the name Default , denoting the black-generation function that was in effect at the start of the page. If both BG and BG2 are present in the same graphics state parameter dictionary, BG2 takes precedence.
UCR	function	<i>(Optional)</i> The undercolor-removal function, which maps the interval [0.0 1.0] to the interval [-1.0 1.0] (see Section 6.2.3, “Conversion from DeviceRGB to DeviceCMYK”).
UCR2	function or name	<i>(Optional; PDF 1.3)</i> Same as UCR except that the value may also be the name Default , denoting the undercolor-removal function that was in effect at the start of the page. If both UCR and UCR2 are present in the same graphics state parameter dictionary, UCR2 takes precedence.
TR	function, array, or name	<i>(Optional)</i> The transfer function, which maps the interval [0.0 1.0] to the interval [0.0 1.0] (see Section 6.3, “Transfer Functions”). The value is either a single function (which applies to all process colorants) or an array of four functions (which apply to the process colorants individually). The name Identity may be used to represent the identity function.

TR2	function, array, or name	<i>(Optional; PDF 1.3)</i> Same as TR except that the value may also be the name Default, denoting the transfer function that was in effect at the start of the page. If both TR and TR2 are present in the same graphics state parameter dictionary, TR2 takes precedence.
HT	dictionary, stream, or name	<i>(Optional)</i> The halftone dictionary or stream (see Section 6.4, “Halftones”) or the name Default, denoting the halftone that was in effect at the start of the page.
FL	number	<i>(Optional; PDF 1.3)</i> The flatness tolerance (see Section 6.5.1, “Flatness Tolerance”).
SM	number	<i>(Optional; PDF 1.3)</i> The smoothness tolerance (see Section 6.5.2, “Smoothness Tolerance”).
SA	boolean	<i>(Optional)</i> A flag specifying whether to apply automatic stroke adjustment (see Section 6.5.4, “Automatic Stroke Adjustment”).
BM	name or array	<i>(Optional; PDF 1.4)</i> The current blend mode to be used in the transparent imaging model (see Sections 7.2.4, “Blend Mode,” and 7.5.2, “Specifying Blending Color Space and Blend Mode”).
SMask	dictionary or name	<p><i>(Optional; PDF 1.4)</i> The current soft mask, specifying the mask shape or mask opacity values to be used in the transparent imaging model (see “Source Shape and Opacity” on page 421 and “Mask Shape and Opacity” on page 443).</p> <p><i>Note: Although the current soft mask is sometimes referred to as a “soft clip,” altering it with the gs operator completely replaces the old value with the new one, rather than intersecting the two as is done with the current clipping path parameter (see Section 4.4.3, “Clipping Path Operators”).</i></p>
CA	number	<i>(Optional; PDF 1.4)</i> The current stroking alpha constant, specifying the constant shape or constant opacity value to be used for stroking operations in the transparent imaging model (see “Source Shape and Opacity” on page 421 and “Constant Shape and Opacity” on page 444).
ca	number	<i>(Optional; PDF 1.4)</i> Same as CA , but for nonstroking operations.
AIS	boolean	<i>(Optional; PDF 1.4)</i> The alpha source flag (“alpha is shape”), specifying whether the current soft mask and alpha constant are to be interpreted as shape values (true) or opacity values (false).
TK	boolean	<i>(Optional; PDF 1.4)</i> The text knockout flag, which determines the behavior of overlapping glyphs within a text object in the transparent imaging model (see Section 5.2.7, “Text Knockout”).

Example 4.1 shows two graphics state parameter dictionaries. In the first, automatic stroke adjustment is turned on, and the dictionary includes a transfer function that inverts its value, $f(x) = 1 - x$. In the second, overprint is turned off, and the dictionary includes a parabolic transfer function, $f(x) = (2x - 1)^2$, with a sample of 21 values. The domain of the transfer function, $[0.0 \ 1.0]$, is mapped to $[0 \ 20]$, and the range of the sample values, $[0 \ 255]$, is mapped to the range of the transfer function, $[0.0 \ 1.0]$.

Example 4.1

```

10 0 obj                                % Page object
  << /Type /Page
    /Parent 5 0 R
    /Resources 20 0 R
    /Contents 40 0 R
  >>
endobj

20 0 obj                                % Resource dictionary for page
  << /ProcSet [/PDF /Text]
    /Font << /F1 25 0 R >>
    /ExtGState << /GS1 30 0 R
                /GS2 35 0 R
    >>
  >>
endobj

30 0 obj                                % First graphics state parameter dictionary
  << /Type /ExtGState
    /SA true
    /TR 31 0 R
  >>
endobj

31 0 obj                                % First transfer function
  << /FunctionType 0
    /Domain [0.0 1.0]
    /Range [0.0 1.0]
    /Size 2
    /BitsPerSample 8
    /Length 7
    /Filter /ASCIIHexDecode
  >>

```



```

stream
01 00 >
endstream
endobj

35 0 obj                                % Second graphics state parameter dictionary
  << /Type /ExtGState
    /OP false
    /TR 36 0 R
  >>
endobj

36 0 obj                                % Second transfer function
  << /FunctionType 0
    /Domain [0.0 1.0]
    /Range [0.0 1.0]
    /Size 21
    /BitsPerSample 8
    /Length 63
    /Filter /ASCIIHexDecode
  >>
stream
FF CE A3 7C 5B 3F 28 16 0A 02 00 02 0A 16 28 3F 5B 7C A3 CE FF >
endstream
endobj

```

4.4 Path Construction and Painting

Paths define shapes, trajectories, and regions of all sorts. They are used to draw lines, define the shapes of filled areas, and specify boundaries for clipping other graphics. The graphics state includes a *current clipping path* that defines the clipping boundary for the current page. At the beginning of each page, the clipping path is initialized to include the entire page.

A path is composed of straight and curved line segments, which may connect to one another or may be disconnected. A pair of segments are said to *connect* only if they are defined consecutively, with the second segment starting where the first one ends. Thus the order in which the segments of a path are defined is significant. Nonconsecutive segments that meet or intersect fortuitously are not considered to connect.

A path is made up of one or more disconnected *subpaths*, each comprising a sequence of connected segments. The topology of the path is unrestricted: it may be concave or convex, may contain multiple subpaths representing disjoint areas, and may intersect itself in arbitrary ways. There is an operator, **h**, that explicitly connects the end of a subpath back to its starting point; such a subpath is said to be *closed*. A subpath that has not been explicitly closed is *open*.

As discussed in Section 4.1, “Graphics Objects,” a path object is defined by a sequence of operators to construct the path, followed by one or more operators to paint the path or to use it as a clipping boundary. PDF path operators fall into three categories:

- *Path construction operators* (Section 4.4.1) define the geometry of a path. A path is constructed by sequentially applying one or more of these operators.
- *Path-painting operators* (Section 4.4.2) end a path object, usually causing the object to be painted on the current page in any of a variety of ways.
- *Clipping path operators* (Section 4.4.3), invoked immediately prior to a path-painting operator, cause the path object also to be used for clipping of subsequent graphics objects.

4.4.1 Path Construction Operators

A page description begins with an empty path and builds up its definition by invoking one or more path construction operators to add segments to it. The path construction operators may be invoked in any sequence, but the first one invoked must be **m** or **re** to begin a new subpath. The path definition concludes with the application of a path-painting operator such as **S**, **f**, or **b** (see Section 4.4.2, “Path-Painting Operators”); this may optionally be preceded by one of the clipping path operators **W** or **W*** (Section 4.4.3, “Clipping Path Operators”). Note that the path construction operators in themselves do not place any marks on the page; only the painting operators do that. A path definition is not complete until a path-painting operator has been applied to it.

The path currently under construction is called the *current path*. In PDF (unlike PostScript), the current path is *not* part of the graphics state and is *not* saved and restored along with the other graphics state parameters. PDF paths are strictly internal objects with no explicit representation. Once a path has been painted, it is

no longer defined; there is then no current path until a new one is begun with the **m** or **re** operator.

The trailing endpoint of the segment most recently added to the current path is referred to as the *current point*. If the current path is empty, the current point is undefined. Most operators that add a segment to the current path start at the current point; if the current point is undefined, they generate an error.

Table 4.9 shows the path construction operators. All operands are numbers denoting coordinates in user space.

TABLE 4.9 Path construction operators

OPERANDS	OPERATOR	DESCRIPTION
$x\ y$	m	Begin a new subpath by moving the current point to coordinates (x, y) , omitting any connecting line segment. If the previous path construction operator in the current path was also m , the new m overrides it; no vestige of the previous m operation remains in the path.
$x\ y$	l (lowercase L)	Append a straight line segment from the current point to the point (x, y) . The new current point is (x, y) .
$x_1\ y_1\ x_2\ y_2\ x_3\ y_3$	c	Append a cubic Bézier curve to the current path. The curve extends from the current point to the point (x_3, y_3) , using (x_1, y_1) and (x_2, y_2) as the Bézier control points (see “Cubic Bézier Curves,” below). The new current point is (x_3, y_3) .
$x_2\ y_2\ x_3\ y_3$	v	Append a cubic Bézier curve to the current path. The curve extends from the current point to the point (x_3, y_3) , using the current point and (x_2, y_2) as the Bézier control points (see “Cubic Bézier Curves,” below). The new current point is (x_3, y_3) .
$x_1\ y_1\ x_3\ y_3$	y	Append a cubic Bézier curve to the current path. The curve extends from the current point to the point (x_3, y_3) , using (x_1, y_1) and (x_3, y_3) as the Bézier control points (see “Cubic Bézier Curves,” below). The new current point is (x_3, y_3) .
—	h	Close the current subpath by appending a straight line segment from the current point to the starting point of the subpath. This operator terminates the current subpath; appending another segment to the current path will begin a new subpath, even if the new segment begins at the endpoint reached by the h operation. If the current subpath is already closed, h does nothing.

`x y width height re` Append a rectangle to the current path as a complete subpath, with lower-left corner (x, y) and dimensions *width* and *height* in user space. The operation

`x y width height re`

is equivalent to

```
x y m
(x + width) y l
(x + width) (y + height) l
x (y + height) l
h
```

Cubic Bézier Curves

Curved path segments are specified as *cubic Bézier curves*. Such curves are defined by four points: the two endpoints (the current point P_0 and the final point P_3) and two *control points* P_1 and P_2 . Given the coordinates of the four points, the curve is generated by varying the parameter t from 0.0 to 1.0 in the following equation:

$$R(t) = (1 - t)^3 P_0 + 3t(1 - t)^2 P_1 + 3t^2(1 - t) P_2 + t^3 P_3$$

When $t = 0.0$, the value of the function $R(t)$ coincides with the current point P_0 ; when $t = 1.0$, $R(t)$ coincides with the final point P_3 . Intermediate values of t generate intermediate points along the curve. The curve does not, in general, pass through the two control points P_1 and P_2 .

Cubic Bézier curves have two desirable properties:

- The curve can be very quickly split into smaller pieces for rapid rendering.
- The curve is contained within the convex hull of the four points defining the curve, most easily visualized as the polygon obtained by stretching a rubber band around the outside of the four points. This property allows rapid testing of whether the curve lies completely outside the visible region, and hence does not have to be rendered.

The Bibliography lists several books that describe cubic Bézier curves in more depth.

The most general PDF operator for constructing curved path segments is the **c** operator, which specifies the coordinates of points P_1 , P_2 , and P_3 explicitly, as shown in Figure 4.8. (The starting point, P_0 , is defined implicitly by the current point.)




FIGURE 4.8 Cubic Bézier curve generated by the **c** operator

Two more operators, **v** and **y**, each specify one of the two control points implicitly (see Figure 4.9). In both of these cases, one control point and the final point of the curve are supplied as operands; the other control point is implied, as follows:

- For the **v** operator, the first control point coincides with initial point of the curve.
- For the **y** operator, the second control point coincides with final point of the curve.




FIGURE 4.9 Cubic Bézier curves generated by the **v** and **y** operators

4.4.2 Path-Painting Operators

The path-painting operators end a path object, causing it to be painted on the current page in the manner that the operator specifies. The principal path-painting operators are **S** (for *stroking*) and **f** (for *filling*). Variants of these operators combine stroking and filling in a single operation or apply different rules for determining the area to be filled. Table 4.10 lists all the path-painting operators.

Stroking

The **S** operator paints a line along the current path. The stroked line follows each straight or curved segment in the path, centered on the segment with sides parallel to it. Each of the path's subpaths is treated separately.

The results of the **S** operator depend on the current settings of various parameters in the graphics state. See Section 4.3, "Graphics State," for further information on these parameters.

TABLE 4.10 Path-painting operators

OPERANDS	OPERATOR	DESCRIPTION
—	S	Stroke the path.
—	s	Close and stroke the path. This operator has the same effect as the sequence <code>h S</code> .
—	f	Fill the path, using the nonzero winding number rule to determine the region to fill (see “Nonzero Winding Number Rule” on page 169).
—	F	Equivalent to f ; included only for compatibility. Although applications that read PDF files must be able to accept this operator, those that generate PDF files should use f instead.
—	f*	Fill the path, using the even-odd rule to determine the region to fill (see “Even-Odd Rule” on page 170).
—	B	Fill and then stroke the path, using the nonzero winding number rule to determine the region to fill. This produces the same result as constructing two identical path objects, painting the first with f and the second with S . Note, however, that the filling and stroking portions of the operation consult different values of several graphics state parameters, such as the current color. See also “Special Path-Painting Considerations” on page 462.
—	B*	Fill and then stroke the path, using the even-odd rule to determine the region to fill. This operator produces the same result as B , except that the path is filled as if with f* instead of f . See also “Special Path-Painting Considerations” on page 462.
—	b	Close, fill, and then stroke the path, using the nonzero winding number rule to determine the region to fill. This operator has the same effect as the sequence <code>h B</code> . See also “Special Path-Painting Considerations” on page 462.
—	b*	Close, fill, and then stroke the path, using the even-odd rule to determine the region to fill. This operator has the same effect as the sequence <code>h B*</code> . See also “Special Path-Painting Considerations” on page 462.
—	n	End the path object without filling or stroking it. This operator is a “path-painting no-op,” used primarily for the side effect of changing the current clipping path (see Section 4.4.3, “Clipping Path Operators”).

- The width of the stroked line is determined by the current line width parameter (“Line Width” on page 152).
- The color or pattern of the line is determined by the current color and color space for stroking operations.
- The line can be painted either solid or with a dash pattern, as specified by the current line dash pattern (“Line Dash Pattern” on page 155).
- If a subpath is open, the unconnected ends are treated according to the current line cap style, which may be butt, rounded, or square (“Line Cap Style” on page 153).
- Wherever two consecutive segments are connected, the joint between them is treated according to the current line join style, which may be mitered, rounded, or beveled (“Line Join Style” on page 153). Mitered joins are also subject to the current miter limit (“Miter Limit” on page 153).

***Note:** Points at which unconnected segments happen to meet or intersect receive no special treatment. In particular, “closing” a subpath with an explicit *l* operator rather than with *h* may result in a messy corner, because line caps will be applied instead of a line join.*

- The stroke adjustment parameter (*PDF 1.2*) specifies that coordinates and line widths be adjusted automatically to produce strokes of uniform thickness despite rasterization effects (Section 6.5.4, “Automatic Stroke Adjustment”).

If a subpath is degenerate (consists of a single-point closed path or of two or more points at the same coordinates), the **S** operator paints it only if round line caps have been specified, producing a filled circle centered at the single point. If butt or projecting square line caps have been specified, **S** produces no output, because the orientation of the caps would be indeterminate. (Note that this rule applies only to zero-length subpaths of the path being stroked, and not to zero-length dashes in a dash pattern. In the latter case, the line caps are always painted, since their orientation is determined by the direction of the underlying path.) A single-point open subpath (specified by a trailing **m** operator) produces no output.

Filling

The **f** operator uses the current nonstroking color to paint the entire region enclosed by the current path. If the path consists of several disconnected subpaths, **f**

paints the insides of all subpaths, considered together. Any subpaths that are open are implicitly closed before being filled.

If a subpath is degenerate (consists of a single-point closed path or of two or more points at the same coordinates), **f** paints the single device pixel lying under that point; the result is device-dependent and not generally useful. A single-point open subpath (specified by a trailing **m** operator) produces no output.

For a simple path, it is intuitively clear what region lies inside. However, for a more complex path—for example, a path that intersects itself or has one subpath that encloses another—the interpretation of “inside” is not always obvious. The path machinery uses one of two rules for determining which points lie inside a path: the *nonzero winding number rule* and the *even-odd rule*, both discussed in detail below.

The nonzero winding number rule is more versatile than the even-odd rule and is the standard rule the **f** operator uses. Similarly, the **W** operator uses this rule to determine the inside of the current clipping path. The even-odd rule is occasionally useful for special effects or for compatibility with other graphics systems; the **f*** and **W*** operators invoke this rule.

Nonzero Winding Number Rule

The *nonzero winding number rule* determines whether a given point is inside a path by conceptually drawing a ray from that point to infinity in any direction and then examining the places where a segment of the path crosses the ray. Starting with a count of 0, the rule adds 1 each time a path segment crosses the ray from left to right and subtracts 1 each time a segment crosses from right to left. After counting all the crossings, if the result is 0 then the point is outside the path; otherwise it is inside.

Note: *The method just described does not specify what to do if a path segment coincides with or is tangent to the chosen ray. Since the direction of the ray is arbitrary, the rule simply chooses a ray that does not encounter such problem intersections.*

For simple convex paths, the nonzero winding number rule defines the inside and outside as one would intuitively expect. The more interesting cases are those involving complex or self-intersecting paths like the ones shown in Figure 4.10. For a path consisting of a five-pointed star, drawn with five connected straight

line segments intersecting each other, the rule considers the inside to be the entire area enclosed by the star, including the pentagon in the center. For a path composed of two concentric circles, the areas enclosed by both circles are considered to be inside, *provided that both are drawn in the same direction*. If the circles are drawn in opposite directions, only the “doughnut” shape between them is inside, according to the rule; the “doughnut hole” is outside.





FIGURE 4.10 *Nonzero winding number rule*

Even-Odd Rule

An alternative to the nonzero winding number rule is the *even-odd rule*. This rule determines the “insideness” of a point by drawing a ray from that point in any direction and simply counting the number of path segments that cross the ray, regardless of direction. If this number is odd, the point is inside; if even, the point is outside. This yields the same results as the nonzero winding number rule for paths with simple shapes, but produces different results for more complex shapes.

Figure 4.11 shows the effects of applying the even-odd rule to complex paths. For the five-pointed star, the rule considers the triangular points to be inside the path, but not the pentagon in the center. For the two concentric circles, only the “doughnut” shape between the two circles is considered inside, regardless of the directions in which the circles are drawn.

**FIGURE 4.11** Even-odd rule

4.4.3 Clipping Path Operators

The graphics state contains a *current clipping path* that limits the regions of the page affected by painting operators. The closed subpaths of this path define the area that can be painted. Marks falling inside this area will be applied to the page; those falling outside it will not. (Precisely what is considered to be “inside” a path is discussed under “Filling,” above.)

Note: In the context of the transparent imaging model (PDF 1.4), the *current clipping path* constrains an object’s shape (see Section 7.1, “Overview of Transparency”). The *effective shape* is the intersection of the object’s intrinsic shape with the clipping path; the source shape value is 0.0 outside this intersection. Similarly, the shape of a transparency group (defined as the union of the shapes of its constituent objects) is influenced both by the clipping path in effect when each of the objects is painted and by the one in effect at the time the group’s results are painted onto its backdrop.

The initial clipping path includes the entire page. A clipping path operator (**W** or **W***, shown in Table 4.11) may appear after the last path construction operator and before the path-painting operator that terminates a path object. Although the clipping path operator appears before the painting operator, it does not alter the clipping path at the point where it appears. Rather, it modifies the effect of the succeeding painting operator. *After* the path has been painted, the clipping path in the graphics state is set to the intersection of the current clipping path and the newly constructed path.

TABLE 4.11 Clipping path operators

OPERANDS	OPERATOR	DESCRIPTION
—	W	Modify the current clipping path by intersecting it with the current path, using the nonzero winding number rule to determine which regions lie inside the clipping path.
—	W*	Modify the current clipping path by intersecting it with the current path, using the even-odd rule to determine which regions lie inside the clipping path.

Note: In addition to path objects, text objects can also be used for clipping; see Section 5.2.5, “Text Rendering Mode.”

The **n** operator (see Table 4.10 on page 167) is a “no-op” path-painting operator; it causes no marks to be placed on the page, but can be used with a clipping path operator to establish a new clipping path. That is, after a path has been constructed, the sequence **W n** will intersect that path with the current clipping path to establish a new clipping path.

There is no way to enlarge the current clipping path or to set a new clipping path without reference to the current one. However, since the clipping path is part of the graphics state, its effect can be localized to specific graphics objects by enclosing the modification of the clipping path and the painting of those objects between a pair of **q** and **Q** operators (see Section 4.3.1, “Graphics State Stack”). Execution of the **Q** operator causes the clipping path to revert to the value that was saved by the **q** operator, before the clipping path was modified.

4.5 Color Spaces

PDF includes powerful facilities for specifying the colors of graphics objects to be painted on the current page. The color facilities are divided into two parts:

- *Color specification.* A PDF file can specify abstract colors in a device-independent way. Colors can be described in any of a variety of color systems, or *color spaces*. Some color spaces are related to device color representation (grayscale, *RGB*, *CMYK*), others to human visual perception (CIE-based). Certain special features are also modeled as color spaces: patterns, color mapping, separations, and high-fidelity and multitone color.

- *Color rendering.* The viewer application reproduces colors on the raster output device by a multiple-step process that includes some combination of color conversion, gamma correction, halftoning, and scan conversion. Some aspects of this process use information that is specified in PDF. However, unlike the facilities for color specification, the color rendering facilities are device-dependent and ordinarily should not be included in a page description.

Figures 4.12 and 4.13 on pages 174 and 175 illustrate the division between PDF's (device-independent) color specification and (device-dependent) color rendering facilities. This section describes the color specification features, covering everything that most PDF documents need in order to specify colors. The facilities for controlling color rendering are described in Chapter 6; a PDF document should use these facilities only to configure or calibrate an output device or to achieve special device-dependent effects.

4.5.1 Color Values

As described in Section 4.4.2, “Path-Painting Operators,” marks placed on the page by operators such as **f** and **S** have a color that is determined by the *current color* parameter of the graphics state. A color value consists of one or more *color components*, which are usually numbers. For example, a gray level can be specified by a single number ranging from 0.0 (black) to 1.0 (white). Full color values can be specified in any of several ways; a common method uses three numeric values to specify red, green, and blue components.

Color values are interpreted according to the *current color space*, another parameter of the graphics state. A PDF content stream first selects a color space by invoking the **CS** operator (for the stroking color) or the **cs** operator (for the nonstroking color). It then selects color values within that color space with the **SC** operator (stroking) or the **sc** operator (nonstroking). There are also convenience operators—**G**, **g**, **RG**, **rg**, **K**, and **k**—that select both a color space and a color value within it in a single step. Table 4.21 on page 216 lists all the color-setting operators.

Sampled images (see Section 4.8, “Images”) specify the color values of individual samples with respect to a color space designated by the image object itself. While these values are independent of the current color space and color parameters in the graphics state, all later stages of color processing treat them in exactly the same way as color values specified with the **SC** or **sc** operator.




FIGURE 4.12 Color specification




FIGURE 4.13 Color rendering

4.5.2 Color Space Families

Color spaces can be classified into *color space families*. Spaces within a family share the same general characteristics; they are distinguished by parameter values supplied at the time the space is specified. The families, in turn, fall into three broad categories:

- *Device color spaces* directly specify colors or shades of gray that the output device is to produce. They provide a variety of color specification methods, including grayscale, *RGB* (red-green-blue), and *CMYK* (cyan-magenta-yellow-black), corresponding to the color space families **DeviceGray**, **DeviceRGB**, and **DeviceCMYK**. Since each of these families consists of just a single color space with no parameters, they are often loosely referred to as the **DeviceGray**, **DeviceRGB**, and **DeviceCMYK** color spaces.
- *CIE-based color spaces* are based on an international standard for color specification created by the Commission Internationale de l'Éclairage (International Commission on Illumination). These spaces allow colors to be specified in a way that is independent of the characteristics of any particular output device. Color space families in this category include **CalGray**, **CalRGB**, **Lab**, and **ICCBased**. Individual color spaces within these families are specified by means of dictionaries containing the parameter values needed to define the space.
- *Special color spaces* add features or properties to an underlying color space. They include facilities for patterns, color mapping, separations, and high-fidelity and multitone color. The corresponding color space families are **Pattern**, **Indexed**, **Separation**, and **DeviceN**. Individual color spaces within these families are specified by means of additional parameters.

Table 4.12 summarizes the color space families supported by PDF. (See implementation note 31 in Appendix H.)

TABLE 4.12 Color space families

DEVICE	CIE-BASED	SPECIAL
DeviceGray (<i>PDF 1.1</i>)	CalGray (<i>PDF 1.1</i>)	Indexed (<i>PDF 1.1</i>)
DeviceRGB (<i>PDF 1.1</i>)	CalRGB (<i>PDF 1.1</i>)	Pattern (<i>PDF 1.2</i>)
DeviceCMYK (<i>PDF 1.1</i>)	Lab (<i>PDF 1.1</i>)	Separation (<i>PDF 1.2</i>)
	ICCBased (<i>PDF 1.3</i>)	DeviceN (<i>PDF 1.3</i>)

A color space is defined by an array object whose first element is a name object identifying the color space family. The remaining array elements, if any, are parameters that further characterize the color space; their number and types vary according to the particular family. For families that do not require parameters, the color space can be specified simply by the family name itself instead of an array.

There are two principal ways in which a color space can be specified:

- Within a content stream, the **CS** or **cs** operator establishes the current color space parameter in the graphics state. The operand is always a name object, which either identifies one of the color spaces that need no additional parameters (**DeviceGray**, **DeviceRGB**, **DeviceCMYK**, or some cases of **Pattern**) or is used as a key in the **ColorSpace** subdictionary of the current resource dictionary (see Section 3.7.2, “Resource Dictionaries”). In the latter case, the value of the dictionary entry is in turn a color space array or name. A color space array is never permitted inline within a content stream.
- Outside a content stream, certain objects, such as image XObjects, specify a color space as an explicit parameter, often associated with the key **ColorSpace**. In this case, the color space array or name is always defined directly as a PDF object, not by an entry in the **ColorSpace** resource subdictionary. This convention also applies when color spaces are defined in terms of other color spaces.

The following operators set the current color space and current color parameters in the graphics state:

- **CS** sets the stroking color space; **cs** sets the nonstroking color space.
- **SC** and **SCN** set the stroking color; **sc** and **scn** set the nonstroking color. Depending on the color space, these operators require one or more operands, each specifying one component of the color value.
- **G**, **RG**, and **K** set the stroking color space implicitly and the stroking color as specified by the operands; **g**, **rg**, and **k** do the same for the nonstroking color space and color.

4.5.3 Device Color Spaces

The device color spaces enable a page description to specify color values that are directly related to their representation on an output device. Color values in these spaces map directly (or via simple conversions) to the application of device colorants, such as quantities of ink or intensities of display phosphors. This enables a PDF document to control colors precisely for a particular device, but the results may not be consistent from one device to another.

Output devices form colors either by adding light sources together or by subtracting light from an illuminating source. Computer displays and film recorders typically add colors, while printing inks typically subtract them. These two ways of forming colors give rise to two complementary methods of color specification, called *additive* and *subtractive* color (see Plate 1). The most widely used forms of these two types of color specification are known as *RGB* and *CMYK*, respectively, for the names of the primary colors on which they're based. The corresponding device color spaces are as follows:

- **DeviceGray** controls the intensity of achromatic light, on a scale from black to white.
- **DeviceRGB** controls the intensities of red, green, and blue light, the three additive primary colors used in displays.
- **DeviceCMYK** controls the concentrations of cyan, magenta, yellow, and black inks, the four subtractive process colors used in printing.

Although the notion of explicit color spaces is a PDF 1.1 feature, the operators for specifying colors in the device color spaces—**G**, **g**, **RG**, **rg**, **K**, and **k**—are available in all versions of PDF. Beginning with PDF 1.2, colors specified in device color spaces can optionally be remapped systematically into other color spaces; see “Default Color Spaces” on page 194.

***Note:** In the transparent imaging model (PDF 1.4), the use of device color spaces is subject to special treatment within a transparency group whose group color space is CIE-based (see Sections 7.3, “Transparency Groups,” and 7.5.5, “Transparency Group XObjects”). In particular, the device color space operators should be used only if device color spaces have been remapped to CIE-based spaces by means of the default color space mechanism. Otherwise, the results will be implementation-dependent and unpredictable.*

DeviceGray Color Space

Black, white, and intermediate shades of gray are special cases of full color. A grayscale value is represented by a single number in the range 0.0 to 1.0, where 0.0 corresponds to black, 1.0 to white, and intermediate values to different gray levels. Example 4.2 shows alternative ways to select the **DeviceGray** color space and a specific gray level within that space for stroking operations.

Example 4.2

```
/DeviceGray CS          % Set DeviceGray color space
gray SC                 % Set gray level
gray G                  % Set both in one operation
```

The **CS** and **SC** operators select the current stroking color space and current stroking color separately; **G** sets them in combination. (The **cs**, **sc**, and **g** operators perform the same functions for nonstroking operations.) Setting either current color space to **DeviceGray** initializes the corresponding current color to 0.0.

DeviceRGB Color Space

Colors in the **DeviceRGB** color space are specified according to the additive *RGB* (red-green-blue) color model, in which color values are defined by three components representing the intensities of the additive primary colorants red, green, and blue. Each component is specified by a number in the range 0.0 to 1.0, where 0.0 denotes the complete absence of a primary component and 1.0 denotes maximum intensity. If all three components have equal intensity, the perceived result theoretically is a pure gray on the scale from black to white. If the intensities are not all equal, the result is some color other than a pure gray.

Example 4.3 shows alternative ways to select the **DeviceRGB** color space and a specific color within that space for stroking operations.

Example 4.3

```
/DeviceRGB CS          % Set DeviceRGB color space
red green blue SC     % Set color
red green blue RG     % Set both in one operation
```

The **CS** and **SC** operators select the current stroking color space and current stroking color separately; **RG** sets them in combination. (The **cs**, **sc**, and **rg** operators perform the same functions for nonstroking operations.) Setting either current color space to **DeviceRGB** initializes the red, green, and blue components of the corresponding current color to 0.0.

DeviceCMYK Color Space

The **DeviceCMYK** color space allows colors to be specified according to the subtractive *CMYK* (cyan-magenta-yellow-black) model typical of printers and other paper-based output devices. In theory, each of the three standard *process colorants* used in printing (cyan, magenta, and yellow) absorbs one of the additive primary colors (red, green, and blue, respectively). Black, a fourth standard process colorant, absorbs all of the additive primaries in equal amounts. The four components in a **DeviceCMYK** color value represent the concentrations of these process colorants. Each component is specified by a number in the range 0.0 to 1.0, where 0.0 denotes the complete absence of a process colorant (that is, absorbs none of the corresponding additive primary) and 1.0 denotes maximum concentration (absorbs as much as possible of the additive primary). Note that the sense of these numbers is opposite to that of *RGB* color components.

Example 4.4 shows alternative ways to select the **DeviceCMYK** color space and a specific color within that space for stroking operations.

Example 4.4

```
/DeviceCMYK CS           % Set DeviceCMYK color space
cyan magenta yellow black SC   % Set color
cyan magenta yellow black K    % Set both in one operation
```

The **CS** and **SC** operators select the current stroking color space and current stroking color separately; **K** sets them in combination. (The **cs**, **sc**, and **k** operators perform the same functions for nonstroking operations.) Setting either current color space to **DeviceCMYK** initializes the cyan, magenta, and yellow components of the corresponding current color to 0.0 and the black component to 1.0.

4.5.4 CIE-Based Color Spaces

Calibrated color in PDF is defined in terms of an international standard used in the graphic arts, television, and printing industries. *CIE-based* color spaces enable a page description to specify color values in a way that is related to human visual perception. The goal is for the same color specification to produce consistent results on different output devices, within the limitations of each device; Plate 2 illustrates the kind of variation in color reproduction that can result from the use of uncalibrated color on different devices. PDF 1.1 supports three CIE-based color space families, named **CalGray**, **CalRGB**, and **Lab**; PDF 1.3 adds a fourth, named **ICCBased**.

Note: In PDF 1.1, a color space family named **CalCMYK** was partially defined, with the expectation that its definition would be completed in a future version. However, this is no longer being considered. PDF 1.3 and later versions support calibrated four-component color spaces by means of ICC profiles (see “*ICCBased Color Spaces*” on page 189). PDF consumer applications should ignore **CalCMYK** color space attributes and render colors specified in this family as if they had been specified using **DeviceCMYK**.

The details of the CIE colorimetric system and the theory on which it is based are beyond the scope of this book; see the Bibliography for sources of further information. The semantics of CIE-based color spaces are defined in terms of the relationship between the space’s components and the tristimulus values X , Y , and Z of the CIE 1931 XYZ space. The **CalRGB** and **Lab** color spaces (PDF 1.1) are special cases of three-component CIE-based color spaces, known as *CIE-based ABC* color spaces. These spaces are defined in terms of a two-stage, nonlinear transformation of the CIE 1931 XYZ space. The formulation of such color spaces models a simple *zone theory* of color vision, consisting of a nonlinear trichromatic first stage combined with a nonlinear opponent-color second stage. This formulation allows colors to be digitized with minimum loss of fidelity, an important consideration in sampled images.

Color values in a CIE-based *ABC* color space have three components, arbitrarily named A , B , and C . The first stage transforms these components by first forcing their values to a specified range, then applying *decoding functions*, and finally multiplying the results by a 3-by-3 matrix, producing three intermediate components arbitrarily named L , M , and N . The second stage transforms these intermediate components in a similar fashion, producing the final X , Y , and Z components of the CIE 1931 XYZ space (see Figure 4.14).

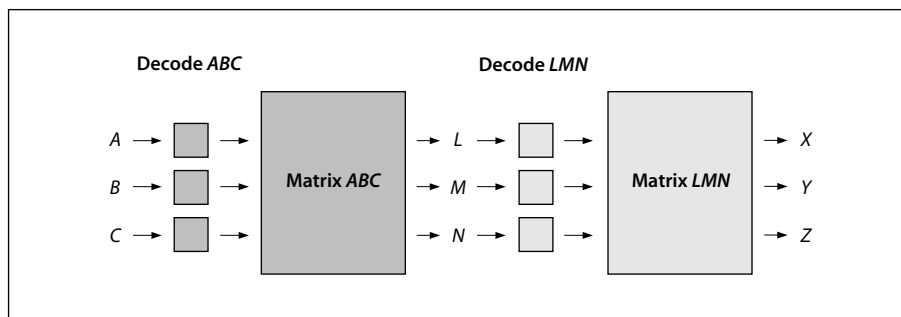


FIGURE 4.14 Component transformations in a CIE-based ABC color space

Color spaces in the CIE-based families are defined by an array

[*name dictionary*]

where *name* is the name of the family and *dictionary* is a dictionary containing parameters that further characterize the space. The entries in this dictionary have specific interpretations that vary depending on the color space; some entries are required and some are optional. See the sections on specific color space families, below, for details.

Setting the current stroking or nonstroking color space to any CIE-based color space initializes all components of the corresponding current color to 0.0 (unless the range of valid values for a given component does not include 0.0, in which case the nearest valid value is substituted.)

Note: *The model and terminology used here—CIE-based ABC (above) and CIE-based A (below)—are derived from the PostScript language, which supports these color space families in their full generality. PDF supports specific useful cases of CIE-based ABC and CIE-based A spaces; most others can be represented as ICCBased spaces.*

CalGray Color Spaces

A **CalGray** color space (*PDF 1.1*) is a special case of a single-component CIE-based color space, known as a *CIE-based A* color space. This type of space is the one-dimensional (and usually achromatic) analog of CIE-based *ABC* spaces. Color values in a CIE-based *A* space have a single component, arbitrarily named *A*. Figure 4.15 illustrates the transformations of the *A* component to *X*, *Y*, and *Z* components of the CIE 1931 *XYZ* space.




FIGURE 4.15 Component transformations in a CIE-based A color space

A **CalGray** color space is a CIE-based A color space with only one transformation stage instead of two. In this type of space, A represents the gray component of a calibrated gray space. This component must be in the range 0.0 to 1.0. The decoding function (denoted by “Decode A” in Figure 4.15) is a gamma function whose coefficient is specified by the **Gamma** entry in the color space dictionary (see Table 4.13). The transformation matrix denoted by “Matrix A” in the figure is derived from the dictionary’s **WhitePoint** entry, as described below. Since there is no second transformation stage, “Decode LMN” and “Matrix LMN” are implicitly taken to be identity transformations.

TABLE 4.13 Entries in a CalGray color space dictionary

KEY	TYPE	VALUE
WhitePoint	array	(Required) An array of three numbers $[X_W Y_W Z_W]$ specifying the tristimulus value, in the CIE 1931 XYZ space, of the diffuse white point; see “CalRGB Color Spaces,” below, for further discussion. The numbers X_W and Z_W must be positive, and Y_W must be equal to 1.0.
BlackPoint	array	(Optional) An array of three numbers $[X_B Y_B Z_B]$ specifying the tristimulus value, in the CIE 1931 XYZ space, of the diffuse black point; see “CalRGB Color Spaces,” below, for further discussion. All three of these numbers must be nonnegative. Default value: [0.0 0.0 0.0].
Gamma	number	(Optional) A number G defining the gamma for the gray (A) component. G must be positive and will generally be greater than or equal to 1. Default value: 1.

The transformation defined by the **Gamma** and **WhitePoint** entries is

$$\begin{aligned} X &= L = X_W \times A^G \\ Y &= M = Y_W \times A^G \\ Z &= N = Z_W \times A^G \end{aligned}$$

In other words, the A component is first decoded by the gamma function, and the result is multiplied by the components of the white point to obtain the L , M , and N components of the intermediate representation. Since there is no second stage, these are also the X , Y , and Z components of the final representation.

The following examples illustrate interesting and useful special cases of **CalGray** spaces. Example 4.5 establishes a space consisting of the Y dimension of the CIE 1931 XYZ space with the CCIR XA/11–recommended D65 white point.

Example 4.5

```
[ /CalGray
  << /WhitePoint [0.9505 1.0000 1.0890] >>
]
```

Example 4.6 establishes a calibrated gray space with the CCIR XA/11–recommended D65 white point and opto-electronic transfer function.

Example 4.6

```
[ /CalGray
  << /WhitePoint [0.9505 1.0000 1.0890]
    /Gamma 2.222
  >>
]
```

CalRGB Color Spaces

A **CalRGB** color space is a CIE-based ABC color space with only one transformation stage instead of two. In this type of space, A , B , and C represent calibrated red, green, and blue color values. These three color components must be in the range 0.0 to 1.0; component values falling outside that range will be adjusted to the nearest valid value without error indication. The decoding functions (denoted by “Decode ABC ” in Figure 4.14 on page 182) are gamma functions whose coefficients are specified by the **Gamma** entry in the color space dictionary

(see Table 4.14). The transformation matrix denoted by “Matrix *ABC*” in Figure 4.14 is defined by the dictionary’s **Matrix** entry. Since there is no second transformation stage, “Decode *LMN*” and “Matrix *LMN*” are implicitly taken to be identity transformations.

TABLE 4.14 Entries in a CalRGB color space dictionary

KEY	TYPE	VALUE
WhitePoint	array	<i>(Required)</i> An array of three numbers $[X_W \ Y_W \ Z_W]$ specifying the tristimulus value, in the CIE 1931 <i>XYZ</i> space, of the diffuse white point; see below for further discussion. The numbers X_W and Z_W must be positive, and Y_W must be equal to 1.0.
BlackPoint	array	<i>(Optional)</i> An array of three numbers $[X_B \ Y_B \ Z_B]$ specifying the tristimulus value, in the CIE 1931 <i>XYZ</i> space, of the diffuse black point; see below for further discussion. All three of these numbers must be nonnegative. Default value: [0.0 0.0 0.0].
Gamma	array	<i>(Optional)</i> An array of three numbers $[G_R \ G_G \ G_B]$ specifying the gamma for the red, green, and blue (<i>A</i> , <i>B</i> , and <i>C</i>) components of the color space. Default value: [1.0 1.0 1.0].
Matrix	array	<i>(Optional)</i> An array of nine numbers $[X_A \ Y_A \ Z_A \ X_B \ Y_B \ Z_B \ X_C \ Y_C \ Z_C]$ specifying the linear interpretation of the decoded <i>A</i> , <i>B</i> , and <i>C</i> components of the color space with respect to the final <i>XYZ</i> representation. Default value: the identity matrix [1 0 0 0 1 0 0 0 1].

The **WhitePoint** and **BlackPoint** entries in the color space dictionary control the overall effect of the CIE-based gamut mapping function described in Section 6.1, “CIE-Based Color to Device Color.” Typically, the colors specified by **WhitePoint** and **BlackPoint** are mapped to the nearly lightest and nearly darkest achromatic colors that the output device is capable of rendering in a way that preserves color appearance and visual contrast.

WhitePoint is assumed to represent the diffuse achromatic highlight, not a specular highlight. Specular highlights, achromatic or otherwise, are often reproduced lighter than the diffuse highlight. **BlackPoint** is assumed to represent the diffuse achromatic shadow; its value is typically limited by the dynamic range of the input device. In images produced by a photographic system, the values of **WhitePoint** and **BlackPoint** vary with exposure, system response, and artistic intent; hence, their values are image-dependent.

The transformation defined by the **Gamma** and **Matrix** entries in the **CalRGB** color space dictionary is

$$\begin{aligned} X = L &= X_A \times A^{G_R} + X_B \times B^{G_G} + X_C \times C^{G_B} \\ Y = M &= Y_A \times A^{G_R} + Y_B \times B^{G_G} + Y_C \times C^{G_B} \\ Z = N &= Z_A \times A^{G_R} + Z_B \times B^{G_G} + Z_C \times C^{G_B} \end{aligned}$$

In other words, the *A*, *B*, and *C* components are first decoded individually by the gamma functions. The results are treated as a three-element vector and multiplied by **Matrix** (a 3-by-3 matrix) to obtain the *L*, *M*, and *N* components of the intermediate representation. Since there is no second stage, these are also the *X*, *Y*, and *Z* components of the final representation.

Example 4.7 shows an example of a **CalRGB** color space for the CCIR XA/11–recommended D65 white point with 1.8 gammas and Sony Trinitron[®] phosphor chromaticities.

Example 4.7

```
[ /CalRGB
  << /WhitePoint [0.9505 1.0000 1.0890]
    /Gamma [1.8000 1.8000 1.8000]
    /Matrix [ 0.4497 0.2446 0.0252
              0.3163 0.6720 0.1412
              0.1845 0.0833 0.9227
            ]
  >>
]
```

In some cases, the parameters of a **CalRGB** color space may be specified in terms of the CIE 1931 chromaticity coordinates (x_R, y_R) , (x_G, y_G) , (x_B, y_B) of the red, green, and blue phosphors, respectively, and the chromaticity (x_W, y_W) of the diffuse white point corresponding to some linear *RGB* value (R, G, B) , where usually $R = G = B = 1.0$. Note that standard CIE notation uses lowercase letters to specify chromaticity coordinates and uppercase letters to specify tristimulus values. Given this information, **Matrix** and **WhitePoint** can be found as follows:

$$z = y_W \times ((x_G - x_B) \times y_R - (x_R - x_B) \times y_G + (x_R - x_G) \times y_B)$$

$$Y_A = \frac{y_R}{R} \times \frac{(x_G - x_B) \times y_W - (x_W - x_B) \times y_G + (x_W - x_G) \times y_B}{z}$$

$$X_A = Y_A \times \frac{x_R}{y_R} \quad Z_A = Y_A \times \left(\frac{1 - x_R}{y_R} - 1 \right)$$

$$Y_B = -\frac{y_G}{G} \times \frac{(x_R - x_B) \times y_W - (x_W - x_B) \times y_R + (x_W - x_R) \times y_B}{z}$$

$$X_B = Y_B \times \frac{x_G}{y_G} \quad Z_B = Y_B \times \left(\frac{1 - x_G}{y_G} - 1 \right)$$

$$Y_C = \frac{y_B}{B} \times \frac{(x_R - x_G) \times y_W - (x_W - x_G) \times y_W + (x_W - x_R) \times y_G}{z}$$

$$X_C = Y_C \times \frac{x_B}{y_B} \quad Z_C = Y_C \times \left(\frac{1 - x_B}{y_B} - 1 \right)$$

$$X_W = X_A \times R + X_B \times G + X_C \times B$$

$$Y_W = Y_A \times R + Y_B \times G + Y_C \times B$$

$$Z_W = Z_A \times R + Z_B \times G + Z_C \times B$$

Lab Color Spaces

A **Lab** color space is a CIE-based *ABC* color space with two transformation stages (see Figure 4.14 on page 182). In a this type of space, *A*, *B*, and *C* represent the L^* , a^* , and b^* components of a CIE 1976 $L^*a^*b^*$ space. The range of the first (L^*) component is always 0 to 100; the ranges of the second and third (a^* and b^*) components are defined by the **Range** entry in the color space dictionary (see Table 4.15).

Plate 3 illustrates the coordinates of a typical **Lab** color space; Plate 4 compares the gamuts (ranges of representable colors) for $L^*a^*b^*$, *RGB*, and *CMYK* spaces.

TABLE 4.15 Entries in a Lab color space dictionary

KEY	TYPE	VALUE
WhitePoint	array	<i>(Required)</i> An array of three numbers [$X_W Y_W Z_W$] specifying the tristimulus value, in the CIE 1931 XYZ space, of the diffuse white point; see “CalRGB Color Spaces” on page 184 for further discussion. The numbers X_W and Z_W must be positive, and Y_W must be equal to 1.0.
BlackPoint	array	<i>(Optional)</i> An array of three numbers [$X_B Y_B Z_B$] specifying the tristimulus value, in the CIE 1931 XYZ space, of the diffuse black point; see “CalRGB Color Spaces” on page 184 for further discussion. All three of these numbers must be nonnegative. Default value: [0.0 0.0 0.0].
Range	array	<p><i>(Optional)</i> An array of four numbers [$a_{\min} a_{\max} b_{\min} b_{\max}$] specifying the range of valid values for the a^* and b^* (B and C) components of the color space—that is,</p> $a_{\min} \leq a^* \leq a_{\max}$ <p>and</p> $b_{\min} \leq b^* \leq b_{\max}$ <p>Component values falling outside the specified range will be adjusted to the nearest valid value without error indication. Default value: [-100 100 -100 100].</p>

A **Lab** color space does not specify explicit decoding functions or matrix coefficients for either stage of the transformation from $L^*a^*b^*$ space to XYZ space (denoted by “Decode ABC,” “Matrix ABC,” “Decode LMN,” and “Matrix LMN” in Figure 4.14 on page 182). Instead, these parameters have constant implicit values. The first transformation stage is defined by the equations

$$L = \frac{L^* + 16}{116} + \frac{a^*}{500}$$

$$M = \frac{L^* + 16}{116}$$

$$N = \frac{L^* + 16}{116} - \frac{b^*}{200}$$

The second transformation stage is given by

$$X = X_W \times g(L)$$

$$Y = Y_W \times g(M)$$

$$Z = Z_W \times g(N)$$

where the function $g(x)$ is defined as

$$g(x) = x^3 \quad \text{if } x \geq \frac{6}{29}$$

$$g(x) = \frac{108}{841} \times \left(x - \frac{4}{29} \right) \quad \text{otherwise}$$

Example 4.8 defines the CIE 1976 $L^*a^*b^*$ space with the CCIR XA/11–recommended D65 white point. The a^* and b^* components, although theoretically unbounded, are defined to lie in the useful range -128 to $+127$.

Example 4.8

```
[ /Lab
  << /WhitePoint [0.9505 1.0000 1.0890]
    /Range [-128 127 -128 127]
  >>
]
```

ICCBased Color Spaces

ICCBased color spaces (*PDF 1.3*) are based on a cross-platform *color profile* as defined by the International Color Consortium (ICC). Unlike the **CalGray**, **CalRGB**, and **Lab** color spaces, which are characterized by entries in the color space dictionary, an **ICCBased** color space is characterized by a sequence of bytes in a standard format. Details of the profile format can be found in the ICC specification (see the Bibliography).

An **ICCBased** color space is specified as an array:

```
[/ICCBased stream]
```

The stream contains the ICC profile. Besides the usual entries common to all streams (see Table 3.4 on page 38), the profile stream has the additional entries listed in Table 4.16.

TABLE 4.16 Additional entries specific to an ICC profile stream dictionary

KEY	TYPE	VALUE
N	integer	<i>(Required)</i> The number of color components in the color space described by the ICC profile data. This number must match the number of components actually in the ICC profile. As of PDF 1.4, N must be 1, 3, or 4.
Alternate	array or name	<i>(Optional)</i> An alternate color space to be used in case the one specified in the stream data is not supported (for example, by viewer applications designed for earlier versions of PDF). The alternate space may be any valid color space (except a Pattern color space) that has the number of components specified by N . If this entry is omitted and the viewer application does not understand the ICC profile data, the color space used will be DeviceGray , DeviceRGB , or DeviceCMYK , depending on whether the value of N is 1, 3, or 4, respectively. <i>Note: Note that there is no conversion of source color values, such as a tint transformation, when using the alternate color space. Color values that are within the range of the ICCBased color space might not be within the range of the alternate color space. In this case, the nearest values within the range of the alternate space will be substituted.</i>
Range	array	<i>(Optional)</i> An array of $2 \times N$ numbers [min_0 max_0 min_1 max_1 ...] specifying the minimum and maximum valid values of the corresponding color components. These values must match the information in the ICC profile. Default value: [0.0 1.0 0.0 1.0 ...].
Metadata	stream	<i>(Optional; PDF 1.4)</i> A <i>metadata stream</i> containing metadata for the color space (see Section 9.2.2, “Metadata Streams”).

The ICC specification is an evolving standard. The **ICCBased** color spaces supported in PDF 1.3 are based on ICC specification version 3.3; those in PDF 1.4 are based on the ICC specification ICC.1:1998-09 and its addendum ICC.1A:1999-04. (Earlier versions of the ICC specification are also supported.) This has the following consequences for producers and consumers of PDF:

- A consumer that supports a given PDF version is required to support ICC profiles conforming to the corresponding version (and earlier versions) of the ICC specification, as described above. It may optionally support later ICC versions.

- For the most predictable and consistent results, a producer of a given PDF version should embed only profiles conforming to the corresponding version of the ICC specification.
- A PDF producer may embed profiles conforming to a later ICC version, with the understanding that the results will vary depending on the capabilities of the consumer. The consumer might process the profile while ignoring newer features, or it might fail altogether to process the profile. In light of this, it is recommended that the producer provide an alternate color space (**Alternate** entry in the **ICCBased** color space dictionary) containing a profile that is appropriate for the PDF version.

As of version 1.4, PDF supports only the profile types shown in Table 4.17; other types may be supported in the future. (In particular, note that XYZ and 16-bit $L^*a^*b^*$ profiles are not supported.) Each of the indicated fields must have one of the values listed for that field in the second column of the table. (Profiles must satisfy *both* the criteria shown in the table.) The terminology is taken from the ICC specifications.

TABLE 4.17 ICC profile types

HEADER FIELD	REQUIRED VALUE
deviceClass	icSigInputClass ('scnr') icSigDisplayClass ('mnr') icSigOutputClass ('prtr') icSigColorSpaceClass ('spac')
colorSpace	icSigGrayData ('GRAY') icSigRgbData ('RGB ') icSigCmykData ('CMYK') icSigLabData ('Lab ')

The terminology used in PDF color spaces and ICC color profiles is similar, but sometimes the same terms are used with different meanings. For example, the default value for each component in an **ICCBased** color space is 0. The range of each color component is a function of the color space specified by the profile and is indicated in the ICC specification. The ranges for several ICC color spaces are shown in Table 4.18.

TABLE 4.18 Ranges for typical ICC color spaces

ICC COLOR SPACE	COMPONENT RANGES
Gray	[0.0 1.0]
RGB	[0.0 1.0]
CMYK	[0.0 1.0]
L*a*b*	L*: [0 100]; a* and b*: [-128 127]

Since the **ICCBased** color space is being used as a source color space, only the “to CIE” profile information (*AToB* in ICC terminology) is used; the “from CIE” (*BToA*) information is ignored when present. An ICC profile may also specify a *rendering intent*, but PDF viewer applications ignore this information; the rendering intent is specified in PDF by a separate parameter (see “Rendering Intents” on page 197).

Note: *The requirements stated above apply to an ICCBased color space that is used to specify the source colors of graphics objects. When such a space is used as the blending color space for a transparency group in the transparent imaging model (see Sections 7.2.3, “Blending Color Space”; 7.3, “Transparency Groups”; and 7.5.5, “Transparency Group XObjects”), it must have both “to CIE” (AToB) and “from CIE” (BToA) information. This is because the group color space is used as both the destination for objects being painted within the group and the source for the group’s results. ICC profiles are also used in specifying output intents for matching the color characteristics of a PDF document with those of a target output device or production environment. When used in this context, they are subject to still other constraints on the “to CIE” and “from CIE” information; see Section 9.10.4, “Output Intents,” for details.*

The representations of **ICCBased** color spaces are less compact than **CalGray**, **CalRGB**, and **Lab**, but can represent a wider range of color spaces. In those cases where a given color space can be expressed by more than one of the CIE-based color space families, the resulting colors are expected to be rendered similarly, regardless of the method selected for representation.

One particular color space is the so-called “standard *RGB*” or *sRGB*, defined in the International Electrotechnical Commission (IEC) document *Colour Measurement and Management in Multimedia Systems and Equipment* (see the Bibliogra-

phy). In PDF, the *sRGB* color space can be expressed precisely only as an **ICCBased** space, although it can be approximated by a **CalRGB** space.

Example 4.9 shows an **ICCBased** color space for a typical three-component *RGB* space. The profile's data has been encoded in hexadecimal representation for readability; in actual practice, a lossless decompression filter such as **FlateDecode** should be used.

Example 4.9

```

10 0 obj                                % Color space
  [/ICCBased 15 0 R]
endobj

15 0 obj                                % ICC profile stream
  << /N 3
    /Alternate /DeviceRGB
    /Length 1605
    /Filter /ASCIIHexDecode
  >>
stream
00 00 02 0C 61 70 70 6C 02 00 00 00 6D 6E 74 72
52 47 42 20 58 59 5A 20 07 CB 00 02 00 16 00 0E
00 22 00 2C 61 63 73 70 41 50 50 4C 00 00 00 00
61 70 70 6C 00 00 04 01 00 00 00 00 00 00 00 02
00 00 00 00 00 00 F6 D4 00 01 00 00 00 00 D3 2B
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 09 64 65 73 63 00 00 00 F0 00 00 00 71
72 58 59 5A 00 00 01 64 00 00 00 14 67 58 59 5A
00 00 01 78 00 00 00 14 62 58 59 5A 00 00 01 8C
00 00 00 14 72 54 52 43 00 00 01 A0 00 00 00 0E
67 54 52 43 00 00 01 B0 00 00 00 0E 62 54 52 43
00 00 01 C0 00 00 00 0E 77 74 70 74 00 00 01 D0
00 00 00 14 63 70 72 74 00 00 01 E4 00 00 00 27
64 65 73 63 00 00 00 00 00 00 00 17 41 70 70 6C
65 20 31 33 22 20 52 47 42 20 53 74 61 6E 64 61
72 64 00 00 00 00 00 00 00 00 00 00 17 41 70
70 6C 65 20 31 33 22 20 52 47 42 20 53 74 61 6E
64 61 72 64 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 58 59 5A 58 59 5A 20 00 00 00 00 00 00 63 0A

```

```

00 00 35 0F 00 00 03 30 58 59 5A 20 00 00 00 00
00 00 53 3D 00 00 AE 37 00 00 15 76 58 59 5A 20
00 00 00 00 00 00 40 89 00 00 1C AF 00 00 BA 82
63 75 72 76 00 00 00 00 00 00 01 01 CC 63 75
63 75 72 76 00 00 00 00 00 00 01 01 CC 63 75
63 75 72 76 00 00 00 00 00 00 01 01 CC 58 59
58 59 5A 20 00 00 00 00 00 00 F3 1B 00 01 00 00
00 01 67 E7 74 65 78 74 00 00 00 00 20 43 6F 70
79 72 69 67 68 74 20 41 70 70 6C 65 20 43 6F 6D
70 75 74 65 72 73 20 31 39 39 34 00 >
endstream
endobj

```

Default Color Spaces

Specifying colors in a device color space (**DeviceGray**, **DeviceRGB**, or **DeviceCMYK**) makes them device-dependent. By setting *default color spaces* (PDF 1.1), a PDF document can request that such colors be systematically transformed (*remapped*) into device-independent CIE-based color spaces. This capability can be useful in a variety of circumstances, such as the following:

- A document originally intended for one output device is redirected to a different device.
- A document is intended to be compatible with viewer applications designed for earlier versions of PDF, and thus cannot specify CIE-based colors directly.
- Color corrections or rendering intents need to be applied to device colors (see “Rendering Intents” on page 197).

A color space is selected for painting each graphics object. This is either the current color space parameter in the graphics state or a color space given as an entry in an image XObject, inline image, or shading dictionary. Regardless of how the color space is specified, it may be subject to remapping as described below.

When a device color space is selected, the **ColorSpace** subdictionary of the current resource dictionary (see Section 3.7.2, “Resource Dictionaries”) is checked for the presence of an entry designating a corresponding default color space (**DefaultGray**, **DefaultRGB**, or **DefaultCMYK**, corresponding to **DeviceGray**, **DeviceRGB**, or **DeviceCMYK**, respectively). If such an entry is present, its value is used as the color space for the operation currently being performed. (If the view-

er application does not recognize this color space, no remapping will occur; the original device color space will be used.)

Color values in the original device color space are passed unchanged to the default color space, which must have the same number of components as the original space. The default color space should be chosen to be compatible with the original, taking into account the components' ranges and whether the components are additive or subtractive. If a color value lies outside the range of the default color space, it will be adjusted to the nearest valid value.

***Note:** Any color space other than a **Lab**, **Indexed**, or **Pattern** color space may be used as a default color space, provided that it is compatible with the original device color space as described above.*

If the selected space is a special color space based on an underlying device color space, the default color space will be used in place of the underlying space. This applies to the following:

- The underlying color space of a **Pattern** color space
- The base color space of an **Indexed** color space
- The alternate color space of a **Separation** or **DeviceN** color space (but only if the alternate color space is actually selected)

See Section 4.5.5, “Special Color Spaces,” for details on these color spaces.

***Note:** Note that there is no conversion of color values, such as a tint transformation, when using the default color space. Color values that are within the range of the device color space might not be within the range of the default color space (particularly if the default is an **ICCBased** color space). In this case, the nearest values within the range of the default space will be used. For this reason, a **Lab** color space is not permitted as the **DefaultRGB** color space.*

Implicit Conversion of CIE-Based Color Spaces

In workflows in which PDF documents are intended for rendering on a specific target output device (such as a printing press with particular inks and media), it is often useful to specify the source colors for some or all of a document's objects in a CIE-based color space that matches the calibration of the intended device. The resulting document, while tailored to the specific characteristics of the target

device, remains device-independent and will produce reasonable results if re-targeted to a different output device. However, the expectation is that if the document is printed on the intended target device, source colors that have been specified in a color space matching the calibration of the device will pass through unchanged, without conversion to and from the intermediate CIE 1931 XYZ space as depicted in Figure 4.14 on page 182.

In particular, when colors intended for a *CMYK* output device are specified in an **ICCBased** color space using a matching *CMYK* printing profile, converting such colors from four components to three and back is unnecessary and will result in an undesirable loss of fidelity in the black component. In such cases, PDF viewer applications may provide the ability for the user to specify a particular calibration to use for printing, proofing, or previewing. This calibration is then considered to be that of the native color space of the intended output device (typically **DeviceCMYK**), and colors expressed in a CIE-based source color space matching it can be treated as if they were specified directly in the device's native color space. Note that the conditions under which such implicit conversion is done cannot be specified in PDF itself, since nothing in PDF describes the calibration of the output device (although an output intent dictionary, if present, may suggest such a calibration; see Section 9.10.4, "Output Intents"). The conversion is completely hidden by the viewer application and plays no part in the interpretation of PDF color spaces.

When this type of implicit conversion is done, all of the semantics of the device color space should also apply, even though they do not apply to CIE-based spaces in general. In particular:

- The nonzero overprint mode (see Section 4.5.6, "Overprint Control") determines the interpretation of color component values in the space.
- If the space is used as the blending color space for a transparency group in the transparent imaging model (see Sections 7.2.3, "Blending Color Space"; 7.3, "Transparency Groups"; and 7.5.5, "Transparency Group XObjects"), components of the space, such as **Cyan**, can be selected in a **Separation** or **DeviceN** color space used within the group (see "Separation Color Spaces" on page 201 and "DeviceN Color Spaces" on page 205).
- Likewise, any uses of device color spaces for objects within such a transparency group have well-defined conversions to the group color space.

Note: A source color space can be specified directly (for example, with an **ICCBased** color space) or indirectly using the default color space mechanism (for example, **DefaultCMYK**; see “Default Color Spaces” on page 194). The implicit conversion of a CIE-based color space to a device space should not depend on whether the CIE-based space is specified directly or indirectly.

Rendering Intents

Although CIE-based color specifications are theoretically device-independent, they are subject to practical limitations in the color reproduction capabilities of the output device. Such limitations may sometimes require compromises to be made among various properties of a color specification when rendering colors for a given device. Specifying a *rendering intent* (PDF 1.1) allows a PDF file to set priorities regarding which of these properties to preserve and which to sacrifice. For example, the PDF file might request that colors falling within the output device’s gamut (the range of colors it can reproduce) be rendered exactly while sacrificing the accuracy of out-of-gamut colors, or that a scanned image such as a photograph be rendered in a perceptually “pleasing” manner at the cost of strict colorimetric accuracy.

Rendering intents are specified with the **ri** operator (see Section 4.3.3, “Graphics State Operators”) and with the **Intent** entry in image dictionaries (Section 4.8.4, “Image Dictionaries”). The value is a name identifying the desired rendering intent. Table 4.19 lists the standard rendering intents recognized in the initial release of PDF viewer applications from Adobe Systems; Plate 5 illustrates their effects. These intents have been deliberately chosen to correspond closely to those defined by the International Color Consortium (ICC), an industry organization that has developed standards for device-independent color. Note, however, that the exact set of rendering intents supported may vary from one output device to another; a particular device may not support all possible intents, or may support additional ones beyond those listed in the table. If the viewer application does not recognize the specified name, it uses the **RelativeColorimetric** intent by default.

See Section 7.6.4, “Rendering Parameters and Transparency,” and in particular “Rendering Intent and Color Conversions” on page 468, for further discussion of the role of rendering intents in the transparent imaging model.

TABLE 4.19 Rendering intents

NAME	DESCRIPTION
AbsoluteColorimetric	Colors are represented solely with respect to the light source; no correction is made for the output medium's white point (such as the color of unprinted paper). Thus, for example, a monitor's white point, which is bluish compared to that of a printer's paper, would be reproduced with a blue cast. In-gamut colors are reproduced exactly; out-of-gamut colors are mapped to the nearest value within the reproducible gamut. This style of reproduction has the advantage of providing exact color matches from one output medium to another. It has the disadvantage of causing colors with Y values between the medium's white point and 1.0 to be out of gamut. A typical use might be for logos and solid colors that require exact reproduction across different media.
RelativeColorimetric	Colors are represented with respect to the combination of the light source and the output medium's white point (such as the color of unprinted paper). Thus, for example, a monitor's white point would be reproduced on a printer by simply leaving the paper unmarked, ignoring color differences between the two media. In-gamut colors are reproduced exactly; out-of-gamut colors are mapped to the nearest value within the reproducible gamut. This style of reproduction has the advantage of adapting for the varying white points of different output media. It has the disadvantage of not providing exact color matches from one medium to another. A typical use might be for vector graphics.
Saturation	Colors are represented in a manner that preserves or emphasizes saturation. Reproduction of in-gamut colors may or may not be colorimetrically accurate. A typical use might be for business graphics, where saturation is the most important attribute of the color.
Perceptual	Colors are represented in a manner that provides a pleasing perceptual appearance. This generally means that both in-gamut and out-of-gamut colors are modified from their precise colorimetric values in order to preserve color relationships. A typical use might be for scanned images.

4.5.5 Special Color Spaces

Special color spaces add features or properties to an underlying color space. There are four special color space families: **Pattern**, **Indexed**, **Separation**, and **DeviceN**.

Pattern Color Spaces

A **Pattern** color space (*PDF 1.2*) enables a PDF content stream to paint an area with a “color” defined as a *pattern*, which may be either a *tiling pattern* (type 1) or a *shading pattern* (type 2). Section 4.6, “Patterns,” discusses patterns in detail.

Indexed Color Spaces

An **Indexed** color space allows a PDF content stream to select from a *color map* or *color table* of arbitrary colors in some other space, using small integers as indices. A PDF viewer application treats each sample value as an index into the color table and uses the color value it finds there. This technique can considerably reduce the amount of data required to represent a sampled image—for example, by using 8-bit index values as samples instead of 24-bit *RGB* color values.

An **Indexed** color space is defined by a four-element array, as follows:

```
[/Indexed base hival lookup]
```

The first element is the color space family name **Indexed**. The remaining elements are parameters that an **Indexed** color space requires; their meanings are discussed below. Setting the current stroking or nonstroking color space to an **Indexed** color space initializes the corresponding current color to 0.

The *base* parameter is an array or name that identifies the *base color space* in which the values in the color table are to be interpreted. It can be any device or CIE-based color space or (in PDF 1.3) a **Separation** or **DeviceN** space, but not a **Pattern** space or another **Indexed** space. For example, if the base color space is **DeviceRGB**, the values in the color table are to be interpreted as red, green, and blue components; if the base color space is a CIE-based *ABC* space such as a **CalRGB** or **Lab** space, the values are to be interpreted as *A*, *B*, and *C* components.

Note: Attempting to use a **Separation** or **DeviceN** color space as the base for an **Indexed** color space will generate an error in PDF 1.2.

The *hival* parameter is an integer that specifies the maximum valid index value. In other words, the color table is to be indexed by integers in the range 0 to *hival*. *hival* can be no greater than 255, which is what would be required to index a table with 8-bit index values.

The color table is defined by the *lookup* parameter, which can be either a stream or (in PDF 1.2) a string. It provides the mapping between index values and the corresponding colors in the base color space.

The color table data must be $m \times (hival + 1)$ bytes long, where m is the number of color components in the base color space. Each byte is an unsigned integer in the range 0 to 255 that is scaled to the range of the corresponding color component in the base color space; that is, 0 corresponds to the minimum value in the range for that component, and 255 corresponds to the maximum.

Note: This is different from the interpretation of an **Indexed** color space's color table in *PostScript*. In *PostScript*, the component value is always scaled to the range 0.0 to 1.0, regardless of the range of color values in the base color space.

The color components for each entry in the table appear consecutively in the string or stream. For example, if the base color space is **DeviceRGB** and the indexed color space contains two colors, the order of bytes in the string or stream is $R_0 G_0 B_0 R_1 G_1 B_1$, where letters denote the color component and numeric subscripts denote the table entry.

Example 4.10 illustrates the specification of an **Indexed** color space that maps 8-bit index values to three-component color values in the **DeviceRGB** color space.

Example 4.10

```
[ /Indexed
  /DeviceRGB
  255
  <000000 FF0000 00FF00 0000FF B57342 ...>
]
```

The example shows only the first five color values in the *lookup* string; in all, there should be 256 color values and the string should be 768 bytes long. Having

established this color space, the program can now specify colors using single-component values in the range 0 to 255. For example, a color value of 4 selects an *RGB* color whose components are coded as the hexadecimal integers B5, 73, and 42. Dividing these by 255 and scaling the results to the range 0.0 to 1.0 yields a color with red, green, and blue components of 0.710, 0.451, and 0.259, respectively.

Although an **Indexed** color space is useful mainly for images, index values can also be used with the color selection operators **SC**, **SCN**, **sc**, and **scn**. For example,

```
123 sc
```

selects the same color as does an image sample value of 123. The index value should be an integer in the range 0 to *hival*. If it is a real number, it is rounded to the nearest integer; if it is outside the range 0 to *hival*, it is adjusted to the nearest value within that range.

Separation Color Spaces

Color output devices produce full color by combining *primary* or *process colorants* in varying amounts. On an additive color device such as a display, the primary colorants consist of red, green, and blue phosphors; on a subtractive device such as a printer, they typically consist of cyan, magenta, yellow, and sometimes black inks. In addition, some devices can apply special colorants, often called *spot colorants*, to produce effects that cannot be achieved with the standard process colorants alone. Examples include metallic and fluorescent colors and special textures.

When printing a page, most devices produce a single *composite* page on which all process colorants (and spot colorants, if any) are combined. However, some devices, such as imagesetters, produce a separate, monochromatic rendition of the page, called a *separation*, for each individual colorant. When the separations are later combined—on a printing press, for example—and the proper inks or other colorants are applied to them, a full-color page results.

A **Separation** color space (*PDF 1.2*) provides a means for specifying the use of additional colorants or for isolating the control of individual color components of a device color space for a subtractive device. When such a space is the current color space, the current color is a single-component value, called a *tint*, that controls the application of the given colorant or color components only.

Note: The term separation is often misused as a synonym for an individual device colorant. In the context of this discussion, a printing system that produces separations generates a separate piece of physical medium (generally film) for each colorant. It is these pieces of physical medium that are correctly referred to as separations. A particular colorant properly constitutes a separation only if the device is generating physical separations, one of which corresponds to the given colorant. The **Separation** color space is so named for historical reasons, but it has evolved to the broader purpose of controlling the application of individual colorants in general, whether or not they are actually realized as physical separations.

Note also that the operation of a **Separation** color space itself is independent of the characteristics of any particular output device. Depending on the device, the space may or may not correspond to a true, physical separation or to an actual colorant. For example, a **Separation** color space could be used to control the application of a single process colorant (such as cyan) on a composite device that does not produce physical separations, or could represent a color (such as orange) for which no specific colorant exists on the device. A **Separation** color space provides consistent, predictable behavior, even on devices that cannot directly generate the requested color.

A **Separation** color space is defined as follows:

```
[/Separation name alternateSpace tintTransform]
```

In other words, it is a four-element array whose first element is the color space family name **Separation**. The remaining elements are parameters that a **Separation** color space requires; their meanings are discussed below.

A color value in a **Separation** color space consists of a single tint component in the range 0.0 to 1.0. The value 0.0 represents the minimum amount of colorant that can be applied; 1.0 represents the maximum. Tints are always treated as *subtractive* colors, even if the device produces output for the designated component by an additive method. Thus a tint value of 0.0 denotes the lightest color that can be achieved with the given colorant, and 1.0 the darkest. (Note that this is the same as the convention for **DeviceCMYK** color components, but opposite to the one for **DeviceGray** and **DeviceRGB**.) The **SCN** and **scn** operators respectively set the current stroking and nonstroking color in the graphics state to a tint value; the initial value in either case is 1.0. A sampled image with single-component samples can also be used as a source of tint values.

The *name* parameter in the color space array is a name object specifying the name of the colorant that this **Separation** color space is intended to represent (or

one of the special names **All** or **None**; see below). Such colorant names are arbitrary, and there can be any number of them, subject to implementation limits.

The special colorant name **All** refers collectively to all colorants available on an output device, including those for the standard process colorants. When a **Separation** space with this colorant name is the current color space, painting operators apply tint values to all available colorants at once. This is useful for purposes such as painting registration targets in the same place on every separation. Such marks would typically be painted as the last step in composing a page, to ensure that they are not overwritten by subsequent painting operations.

The special colorant name **None** will never produce any visible output. Painting operations in a **Separation** space with this colorant name have no effect on the current page.

All devices support **Separation** color spaces with the colorant names **All** and **None**, even if they do not support any others. **Separation** spaces with either of these colorant names ignore the *alternateSpace* and *tintTransform* parameters (discussed below), although valid values must still be provided.

At the moment the color space is set to a **Separation** space, the viewer application determines whether the device has an available colorant corresponding to the name of the requested space. If so, the application ignores the *alternateSpace* and *tintTransform* parameters; subsequent painting operations within the space will apply the designated colorant directly, according to the tint values supplied.

*Note: The preceding paragraph applies only to subtractive output devices such as printers and imagesetters. For an additive device such as a computer display, a **Separation** color space never applies a process colorant directly; it always reverts to the alternate color space as described below. This is because the model of applying process colorants independently does not work as intended on an additive device; for instance, painting tints of the **Red** component on a white background produces a result that varies from white to cyan.*

*Note that this exception applies only to colorants for additive devices, not to the specific names **Red**, **Green**, and **Blue**. In contrast, a printer might have a (subtractive) ink named, say, **Red**, which should work as a **Separation** color space just the same as any other supported colorant.*

If the colorant name associated with a **Separation** color space does not correspond to a colorant available on the device, the viewer application arranges instead for subsequent painting operations to be performed in an *alternate color space*. This enables the intended colors to be approximated by colors in some device or CIE-based color space, which are then rendered using the usual primary or process colorants. This works as follows:

- The *alternateSpace* parameter must be an array or name object that identifies the alternate color space. This can be any device or CIE-based color space, but not another special color space (**Pattern**, **Indexed**, **Separation**, or **DeviceN**).
- The *tintTransform* parameter must be a function (see Section 3.9, “Functions”). During subsequent painting operations, a viewer application will call this function to transform a tint value into color component values in the alternate color space. The function is called with the tint value and must return the corresponding color component values. That is, the number of components and the interpretation of their values depend on the alternate color space.

Note: Painting in the alternate color space may produce a good approximation of the intended color when only opaque objects are painted. However, it will not correctly represent the interactions between an object and its backdrop when the object is painted with transparency or when overprinting (see Section 4.5.6, “Overprint Control”) is enabled.

Example 4.11 illustrates the specification of a **Separation** color space (object 5) that is intended to produce a color named LogoGreen. If the output device has no colorant corresponding to this color, **DeviceCMYK** will be used as the alternate color space; the tint transformation function provided (object 12) maps tint values linearly into shades of a *CMYK* color value approximating the “logo green” color.

Example 4.11

```

5 0 obj                                % Color space
  [ /Separation
    /LogoGreen
    /DeviceCMYK
    12 0 R
  ]
endobj

```

```

12 0 obj                                % Tint transformation function
  << /FunctionType 4
    /Domain [0.0 1.0]
    /Range [0.0 1.0 0.0 1.0 0.0 1.0 0.0 1.0]
    /Length 62
  >>
stream
  { dup 0.84 mul
    exch 0.00 exch dup 0.44 mul
    exch 0.21 mul
  }
endstream
endobj

```

See Section 7.6.2, “Spot Colors and Transparency,” for further discussion of the role of **Separation** color spaces in the transparent imaging model.

DeviceN Color Spaces

DeviceN color spaces (*PDF 1.3*) support the use of high-fidelity and multitone color. *High-fidelity* color is the use of more than the standard *CMYK* process colorants to produce an extended *gamut*, or range of colors. A popular example is the PANTONE® Hexachrome™ system, which uses six colorants: the usual cyan, magenta, yellow, and black, plus orange and green.

Multitone color systems use a single-component image to specify multiple color components. In a *duotone*, for example, a single-component image can be used to specify both the black component and a spot color component. The tone reproduction is generally different for the different components; for example, the black component might be painted with the exact sample data from the single-component image, while the spot color component might be generated as a nonlinear function of the image data in a manner that emphasizes the shadows. Plate 6 shows an example using black and magenta color components. In Plate 7, a single-component grayscale image is used to generate a *quadtone* result using four colorants: black and three PANTONE spot colors. See Example 4.17 on page 212 for the code used to generate this image.

DeviceN color spaces allow any subset of the available device colorants to be treated as a device color space with multiple components. This provides greater flexibility than is possible with standard device color spaces such as **DeviceCMYK** or with individual **Separation** color spaces. For example, it is possible to create a

DeviceN color space consisting of only the cyan, magenta, and yellow color components, while excluding the black component. If overprinting is enabled (see Section 4.5.6, “Overprint Control”), painting in this color space will leave the black component unchanged.

A **DeviceN** color space is specified as follows:

```
[/DeviceN names alternateSpace tintTransform]
```

or

```
[/DeviceN names alternateSpace tintTransform attributes]
```

It is a four- or five-element array whose first element is the color space family name **DeviceN**. The remaining elements are parameters that a **DeviceN** color space requires; their meanings are discussed below.

Color values in the **DeviceN** color space are tint components in the range 0.0 to 1.0. The value 0.0 represents the minimum amount of colorant; 1.0 represents the maximum. The **SCN** and **scn** operators respectively set the current stroking and nonstroking color in the graphics state to a set of tint values; the initial value is 1.0 for each tint. A sampled image can also be treated as a source of tint values.

A **DeviceN** color space works almost the same as a **Separation** color space—in fact, a **DeviceN** color space with only one component is exactly equivalent to a **Separation** color space. The following are the only differences between **DeviceN** and **Separation**:

- Color values in a **DeviceN** color space consist of multiple tint components, rather than only one. The number of components is subject to an implementation limit; see Appendix C.
- The *names* parameter in the color space array is an array of name objects specifying the individual colorants. (The special colorant name **All** is not allowed.) The length of the array determines the number of components, and hence the number of operands required by the **SCN** and **scn** operators when this space is the current color space. Operand values supplied to **SCN** or **scn** are interpreted as color component values in the order in which the colors are given in the *names* array.
- At the moment the color space is set to a **DeviceN** space, the viewer application will select the requested set of colorants only if all of them are available on the

device; otherwise, it will select the alternate color space designated by the *alternateSpace* parameter.

- The tint transformation function is called with n tint values and must return the corresponding m color component values, where n is the number of components needed to specify a color in the **DeviceN** color space and m is the number required by the alternate color space.

In a **DeviceN** color space, one or more of the colorant names in the *names* array may be the name **None**. This indicates that the corresponding color component is never painted on the page, as in a **Separation** color space for the **None** colorant. When a **DeviceN** color space is painting the named device colorants directly, color components corresponding to **None** colorants are discarded. However, when the **DeviceN** color space reverts to its alternate color space, those components are passed to the tint transformation function, which may use them in any desired manner.

Note: A **DeviceN** color space whose component colorant names are all **None** always discards its output, just the same as a **Separation** color space for **None**; it never reverts to the alternate color space. Reversion occurs only if at least one color component (other than **None**) is specified and is not available on the device.

The optional *attributes* parameter is a dictionary containing additional information about the color space. At the time of publication, only one entry is defined in this dictionary, as shown in Table 4.20.

TABLE 4.20 Entry in a **DeviceN** color space attributes dictionary

KEY	TYPE	VALUE
Colorants	dictionary	<p>(Optional) A dictionary describing the individual colorants used in the DeviceN color space. For each entry in this dictionary, the key is a colorant name and the value is an array defining a Separation color space for that colorant (see “Separation Color Spaces” on page 201). The key must match the colorant name given in that color space. The dictionary need not list all colorants used in the DeviceN color space and may list additional colorants.</p> <p>This dictionary has no effect on the operation of the DeviceN color space itself or the appearance that it produces. However, it provides information about the individual colorants that may be useful to some applications. In particular, the alternate color space and tint transformation function of a Separation color space describe the appearance of that colorant alone, whereas those of a DeviceN color space describe only the appearance of its colorants in combination.</p>

Example 4.12 shows a **DeviceN** color space consisting of three color components named **Orange**, **Green**, and **None**. In this example, the **DeviceN** color space, object 30, has an attributes dictionary whose **Colorants** entry is an indirect reference to object 45 (which might also be referenced by attributes dictionaries of other **DeviceN** color spaces). *tintTransform1*, whose definition is not shown, maps three color components (tints of the colorants **Orange**, **Green**, and **None**) to four color components in the alternate color space, **DeviceCMYK**. *tintTransform2* maps a single color component (an orange tint) to four components in **DeviceCMYK**. Likewise, *tintTransform3* maps a green tint to **DeviceCMYK**, and *tintTransform4* maps a tint of PANTONE 131 to **DeviceCMYK**.

Example 4.12

```

30 0 obj                                % Color space
  [ /DeviceN
    [/Orange /Green /None]
    /DeviceCMYK
    tintTransform1
    << /Colorants 45 0 R >>
  ]
endobj

45 0 obj                                % Colorants dictionary
  << /Orange [ /Separation
              /Orange
              /DeviceCMYK
              tintTransform2
            ]
    /Green [ /Separation
             /Green
             /DeviceCMYK
             tintTransform3
           ]
    /PANTONE#20131 [ /Separation
                    /PANTONE#20131
                    /DeviceCMYK
                    tintTransform4
                  ]
  >>
endobj

```


See Section 7.6.2, “Spot Colors and Transparency,” for further discussion of the role of **DeviceN** color spaces in the transparent imaging model.

Multitone Examples

The following examples illustrate various interesting and useful special cases of the use of **Indexed** and **DeviceN** color spaces in combination to produce multi-tone colors.

Examples 4.13 and 4.14 illustrate the use of **DeviceN** to create duotone color spaces. In Example 4.13, an **Indexed** color space maps index values in the range 0 to 255 to a duotone **DeviceN** space in cyan and black. In effect, the index values are treated as if they were tints of the duotone space, which are then mapped into tints of the two underlying colorants. Only the beginning of the lookup table string for the **Indexed** color space is shown; the full table would contain 256 two-byte entries, each specifying a tint value for cyan and black, for a total of 512 bytes. If the alternate color space of the **DeviceN** space is selected, the tint transformation function (object 15 in the example) maps the two tint components for cyan and black to the four components for a **DeviceCMYK** color space by supplying zero values for the other two components. Example 4.14 shows the definition of another duotone color space, this time using black and gold colorants (where gold is a spot colorant) and using a **CalRGB** space as the alternate color space. This could be defined in the same way as in the preceding example, with a tint transformation function that converts from the two tint components to colors in the alternate **CalRGB** color space.

Example 4.13

```

10 0 obj                                % Color space
  [ /Indexed
    [ /DeviceN
      [/Cyan /Black]
      /DeviceCMYK
      15 0 R
    ]
    255
    <6605 6806 6907 6B09 6C0A ...>
  ]
endobj

```

```

15 0 obj                                     % Tint transformation function
  << /FunctionType 4
    /Domain [0.0 1.0 0.0 1.0]
    /Range [0.0 1.0 0.0 1.0 0.0 1.0 0.0 1.0]
    /Length 16
  >>
stream
  {0 0 3 -1 roll}
endstream
endobj

```

Example 4.14

```

30 0 obj                                     % Color space
  [ /Indexed
    [ /DeviceN
      [/Black /Gold]
      [ /CalRGB
        << /WhitePoint [1.0 1.0 1.0]
          /Gamma [2.2 2.2 2.2]
        >>
      ]
      35 0 R                                 % Tint transformation function
    ]
    255
    ... Lookup table ...
  ]
endobj

```

Given a formula for converting any combination of black and gold tints to calibrated *RGB*, a 2-in, 3-out type 4 (PostScript calculator) function could be used for the tint transformation. Alternatively, a type 0 (sampled) function could be used, but this would require a large number of sample points to represent the function accurately; for example, sampling each input variable for 256 tint values between 0.0 and 1.0 would require $256^2 = 65,536$ samples. But since the **DeviceN** color space is being used as the base of an **Indexed** color space, there are actually only 256 possible combinations of black and gold tint values. A more compact way to represent this information is to put the alternate color values directly into the lookup table alongside the **DeviceN** color values, as in Example 4.15.

Example 4.15

```

10 0 obj                                     % Color space
  [ /Indexed
    [ /DeviceN
      [/Black /Gold /None /None /None]
      [ /CalRGB
        << /WhitePoint [1.0 1.0 1.0]
          /Gamma [2.2 2.2 2.2]
        >>
      ]
      20 0 R                                 % Tint transformation function
    ]
    255
    ...Lookup table...
  ]
endobj

```

In this example, each entry in the lookup table has *five* components: two for the black and gold colorants and three more (specified as **None**) for the equivalent **CalRGB** color components. If the black and gold colorants are available on the output device, the **None** components will be ignored; if black and gold are not available, the tint transformation function will be used to convert a five-component color into a three-component equivalent in the alternate **CalRGB** color space. But since, by construction, the third, fourth, and fifth components *are* the **CalRGB** components, the tint transformation function can merely discard the first two components and return the last three. This can be easily expressed with a type 4 (PostScript calculator) function, as shown in Example 4.16.

Example 4.16

```

20 0 obj                                     % Tint transformation function
  << /FunctionType 4
    /Domain [0.0 1.0 0.0 1.0 0.0 1.0 0.0 1.0 0.0 1.0]
    /Range [0.0 1.0 0.0 1.0 0.0 1.0]
    /Length 27
  >>
stream
  {5 3 roll pop pop}
endstream
endobj

```

Finally, Example 4.17 uses an extension of the techniques described above to produce the quadtone (four-component) image shown in Plate 7.

Example 4.17

```

5 0 obj                                % Image XObject
  << /Type /XObject
    /Subtype /Image
    /Width 288
    /Height 288
    /ColorSpace 10 0 R
    /BitsPerComponent 8
    /Length 105278
    /Filter /ASCII85Decode
  >>
stream
...Data for grayscale image...
endstream
endobj

10 0 obj                                % Indexed color space for image
  [ /Indexed
    15 0 R                                % Base color space
    255                                    % Table has 256 entries
    30 0 R                                % Lookup table
  ]
endobj

15 0 obj                                % Base color space (DeviceN) for Indexed space
  [ /DeviceN
    [ /Black                                % Four colorants (black plus three spot colors)
      /PANTONE#20216#20CVC
      /PANTONE#20409#20CVC
      /PANTONE#202985#20CVC
      /None                                  % Three components for alternate space
      /None
      /None
    ]
    16 0 R                                % Alternate color space
    20 0 R                                % Tint transformation function
  ]
endobj

```

```

16 0 obj                                % Alternate color space for DeviceN space
  [ /CalRGB
    << /WhitePoint [1.0 1.0 1.0] >>
  ]
endobj

20 0 obj                                % Tint transformation function for DeviceN space
  << /FunctionType 4
    /Domain [0.0 1.0 0.0 1.0 0.0 1.0 0.0 1.0 0.0 1.0 0.0 1.0]
    /Range [0.0 1.0 0.0 1.0 0.0 1.0]
    /Length 44
  >>
stream
  { 7 3 roll                            % Just discard first four values
    pop pop pop pop
  }
endstream
endobj

30 0 obj                                % Lookup table for Indexed color space
  << /Length 1975
    /Filter [/ASCII85Decode /FlateDecode]
  >>
stream
8;T1BB2"M7*!"psYBt1k\gY1T<D&tO]r*F7Hga*
... Additional data (seven components for each table entry)...
endstream
endobj

```

As in the preceding examples, an **Indexed** color space based on a **DeviceN** space is used to paint the grayscale image shown on the left in the plate with four colorants: black and three PANTONE spot colors. The alternate color space is a simple calibrated *RGB*. Thus the **DeviceN** color space has seven components: the four desired colorants plus the three components of the alternate space. The example shows the image XObject (see Section 4.8.4, “Image Dictionaries”) representing the quadtone image, followed by the color space used to interpret the image data. (See implementation note 32 in Appendix H.)

4.5.6 Overprint Control

The graphics state contains an *overprint parameter*, controlled by the **OP** and **op** entries in a graphics state parameter dictionary. Overprint control is useful mainly on devices that produce true physical separations, but it is available on some

composite devices as well. Although the operation of this parameter is device-dependent, it is described here, rather than in the chapter on color rendering, because it pertains to an aspect of painting in device color spaces that is important to many applications.

Any painting operation marks some specific set of device colorants, depending on the color space in which the painting takes place. In a **Separation** or **DeviceN** color space, the colorants to be marked are specified explicitly; in a device or CIE-based color space, they are implied by the process color model of the output device (see Chapter 6). The overprint parameter is a boolean flag that determines how painting operations affect colorants other than those explicitly or implicitly specified by the current color space.

If the overprint parameter is **false** (the default value), painting a color in any color space causes the corresponding areas of unspecified colorants to be erased (painted with a tint value of 0.0). The effect is that the color at any position on the page is whatever was painted there last; this is consistent with the normal painting behavior of the opaque imaging model.

If the overprint parameter is **true** and the output device supports overprinting, no such erasing actions are performed; anything previously painted in other colorants is left undisturbed. Consequently, the color at a given position on the page may be a combined result of several painting operations in different colorants. The effect produced by such overprinting is device-dependent and is not defined by the PDF language.

***Note:** Not all devices support overprinting. Furthermore, many PostScript printers support it only when separations are being produced, and not for composite output. If overprinting is not supported, the value of the overprint parameter is ignored.*

An additional graphics state parameter, the *overprint mode* (PDF 1.3), affects the interpretation of a tint value of 0.0 for a color component in a **DeviceCMYK** color space when overprinting is enabled. This parameter is controlled by the **OPM** entry in a graphics state parameter dictionary; it has an effect only when the overprint parameter is **true**, as described above.

When colors are specified in a **DeviceCMYK** color space and the native color space of the output device is also **DeviceCMYK**, each of the source color components controls the corresponding device colorant directly. Ordinarily, each source color component value replaces the value previously painted for the corresponding de-

vice colorant, no matter what the new value is; this is the default behavior, specified by overprint mode 0.

When the overprint mode is 1 (also called *nonzero overprint mode*), a tint value of 0.0 for a source color component leaves the corresponding component of the previously painted color unchanged. The effect is equivalent to painting in a **DeviceN** color space that includes only those components whose values are nonzero. For example, if the overprint parameter is **true** and the overprint mode is 1, the operation

```
0.2 0.3 0.0 1.0 k
```

is equivalent to

```
0.2 0.3 1.0 scn
```

in the color space shown in Example 4.18.

Example 4.18

```
10 0 obj                                % Color space
  [ /DeviceN
    [/Cyan /Magenta /Black]
    /DeviceCMYK
    15 0 R
  ]
endobj

15 0 obj                                % Tint transformation function
  << /FunctionType 4
    /Domain [0.0 1.0 0.0 1.0 0.0 1.0]
    /Range [0.0 1.0 0.0 1.0 0.0 1.0 0.0 1.0]
    /Length 13
  >>
stream
  {0 exch}
endstream
endobj
```

Nonzero overprint mode applies only to painting operations that use the current color in the graphics state when the current color space is **DeviceCMYK**. It does not apply to the painting of images or to any colors that are the result of a computation, such as those in a shading pattern or conversions from some other

color space. It also does not apply if the device’s native color space is not **DeviceCMYK**; in that case, source colors must be converted to the device’s native color space, and all components participate in the conversion, whatever their values. (This is shown explicitly in the alternate color space and tint transformation function of the **DeviceN** color space in Example 4.18.)

See Section 7.6.3, “Overprinting and Transparency,” for further discussion of the role of overprinting in the transparent imaging model.

4.5.7 Color Operators

Table 4.21 lists the PDF operators that control color spaces and color values. (Also color-related is the graphics state operator **ri**, listed in Table 4.7 on page 156 and discussed under “Rendering Intents” on page 197.) Color operators may appear at the page description level or inside text objects (see Figure 4.1 on page 135).

TABLE 4.21 Color operators

OPERANDS	OPERATOR	DESCRIPTION
<i>name</i>	CS	<p>(PDF 1.1) Set the current color space to use for stroking operations. The operand <i>name</i> must be a name object. If the color space is one that can be specified by a name and no additional parameters (DeviceGray, DeviceRGB, DeviceCMYK, and certain cases of Pattern), the name may be specified directly. Otherwise, it must be a name defined in the ColorSpace subdictionary of the current resource dictionary (see Section 3.7.2, “Resource Dictionaries”); the associated value is an array describing the color space (see Section 4.5.2, “Color Space Families”).</p> <p><i>Note: The names DeviceGray, DeviceRGB, DeviceCMYK, and Pattern always identify the corresponding color spaces directly; they never refer to resources in the ColorSpace subdictionary.</i></p> <p>The CS operator also sets the current stroking color to its initial value, which depends on the color space:</p> <ul style="list-style-type: none"> • In a DeviceGray, DeviceRGB, CalGray, or CalRGB color space, the initial color has all components equal to 0.0. • In a DeviceCMYK color space, the initial color is [0.0 0.0 0.0 1.0]. • In a Lab or ICCBased color space, the initial color has all components equal to 0.0 unless that falls outside the intervals specified by the space’s Range entry, in which case the nearest valid value is substituted. • In an Indexed color space, the initial color value is 0.

		<ul style="list-style-type: none"> • In a Separation or DeviceN color space, the initial tint value is 1.0 for all colorants. • In a Pattern color space, the initial color is a pattern object that causes nothing to be painted.
<i>name</i>	cs	(PDF 1.1) Same as CS , but for nonstroking operations.
$c_1 \dots c_n$	SC	(PDF 1.1) Set the color to use for stroking operations in a device, CIE-based (other than ICCBased), or Indexed color space. The number of operands required and their interpretation depends on the current stroking color space: <ul style="list-style-type: none"> • For DeviceGray, CalGray, and Indexed color spaces, one operand is required ($n = 1$). • For DeviceRGB, CalRGB, and Lab color spaces, three operands are required ($n = 3$). • For DeviceCMYK, four operands are required ($n = 4$).
$c_1 \dots c_n$	SCN	(PDF 1.2) Same as SC , but also supports Pattern , Separation , DeviceN , and
$c_1 \dots c_n$ <i>name</i>	SCN	ICCBased color spaces.
		If the current stroking color space is a Separation , DeviceN , or ICCBased color space, the operands $c_1 \dots c_n$ are numbers. The number of operands and their interpretation depends on the color space.
		If the current stroking color space is a Pattern color space, <i>name</i> is the name of an entry in the Pattern subdictionary of the current resource dictionary (see Section 3.7.2, “Resource Dictionaries”). For an uncolored tiling pattern (PatternType = 1 and PaintType = 2), $c_1 \dots c_n$ are component values specifying a color in the pattern’s underlying color space. For other types of pattern, these operands must not be specified.
$c_1 \dots c_n$	sc	(PDF 1.1) Same as SC , but for nonstroking operations.
$c_1 \dots c_n$	scn	(PDF 1.2) Same as SCN , but for nonstroking operations.
$c_1 \dots c_n$ <i>name</i>	scn	
<i>gray</i>	G	Set the stroking color space to DeviceGray (or the DefaultGray color space; see “Default Color Spaces” on page 194) and set the gray level to use for stroking operations. <i>gray</i> is a number between 0.0 (black) and 1.0 (white).
<i>gray</i>	g	Same as G , but for nonstroking operations.
<i>r g b</i>	RG	Set the stroking color space to DeviceRGB (or the DefaultRGB color space; see “Default Color Spaces” on page 194) and set the color to use for stroking operations. Each operand must be a number between 0.0 (minimum intensity) and 1.0 (maximum intensity).
<i>r g b</i>	rg	Same as RG , but for nonstroking operations.

<i>c m y k</i>	K	Set the stroking color space to DeviceCMYK (or the DefaultCMYK color space; see “Default Color Spaces” on page 194) and set the color to use for stroking operations. Each operand must be a number between 0.0 (zero concentration) and 1.0 (maximum concentration). The behavior of this operator is affected by the overprint mode (see Section 4.5.6, “Overprint Control”).
<i>c m y k</i>	k	Same as K , but for nonstroking operations.

In certain circumstances, invoking operators that specify colors or other color-related parameters in the graphics state is not allowed. This restriction occurs when defining graphical figures whose colors are to be specified separately each time they are used. Specifically, the restriction applies:

- In any glyph description that uses the **d1** operator (see Section 5.5.4, “Type 3 Fonts”)
- In the content stream of an uncolored tiling pattern (see “Uncolored Tiling Patterns” on page 227)

In these circumstances, the following will cause an error:

- Invoking any of the following operators:

CS	scn	K
cs	G	k
SC	g	ri
SCN	RG	sh
sc	rg	

- Invoking the **gs** operator with any of the following entries in the graphics state parameter dictionary:

TR	BG	UCR
TR2	BG2	UCR2
HT		

- Painting an image. However, painting an *image mask* (see “Stencil Masking” on page 276) is permitted, because it does not specify colors, but rather designates places where the current color is to be painted.

4.6 Patterns

When operators such as **S** (stroke), **f** (fill), and **Tj** (show text) paint an area of the page with the current color, they ordinarily apply a single color that covers the area uniformly. However, it is also possible to apply “paint” that consists of a repeating graphical figure or a smoothly varying color gradient instead of a simple color. Such a repeating figure or smooth gradient is called a *pattern*. Patterns are quite general, and have many uses; for example, they can be used to create various graphical textures, such as weaves, brick walls, sunbursts, and similar geometrical and chromatic effects. (See implementation note 33 in Appendix H.)

Patterns come in two varieties:

- *Tiling patterns* consist of a small graphical figure (called a *pattern cell*) that is replicated at fixed horizontal and vertical intervals to fill the area to be painted. The graphics objects to use for tiling are described by a content stream.
- *Shading patterns* define a *gradient fill* that produces a smooth transition between colors across the area. The color to use is specified as a function of position using any of a variety of methods.

Note: *The ability to paint with patterns is a feature of PDF 1.2 (tiling patterns) and PDF 1.3 (shading patterns). With some effort, it is possible to achieve a limited form of tiling patterns in PDF 1.1 by defining them as character glyphs in a special font and painting them repeatedly with the Tj operator. Another technique, defining patterns as halftone screens, is not recommended, because the effects produced are device-dependent.*

Patterns are specified in a special family of color spaces named **Pattern**, whose “color values” are *pattern objects* instead of the numeric component values used with other spaces. A pattern object may be a dictionary or a stream, depending on the type of pattern; the term *pattern dictionary* is used generically throughout this section to refer to either a dictionary object or the dictionary portion of a stream object. (Those pattern objects that are streams are specifically identified as such in the descriptions of particular pattern types; unless otherwise stated, they are understood to be simple dictionaries instead.) This section describes **Pattern** color spaces and the specification of color values within them; see Section 4.5, “Color Spaces,” for information about color spaces and color values in general, and Section 7.5.6, “Patterns and Transparency,” for further discussion of the treatment of patterns in the transparent imaging model.

4.6.1 General Properties of Patterns

A pattern dictionary contains descriptive information defining the appearance and properties of a pattern. All pattern dictionaries contain an entry named **PatternType**, whose value identifies the kind of pattern the dictionary describes: type 1 for a tiling pattern or type 2 for a shading pattern. The remaining contents of the dictionary depend on the pattern type, and are detailed below in the sections on individual pattern types.

All patterns are treated as colors; a **Pattern** color space is established with the **CS** or **cs** operator just like other color spaces, and a particular pattern is installed as the current color with the **SCN** or **scn** operator (see Table 4.21 on page 216).

A pattern's appearance is described with respect to its own internal coordinate system. Every pattern has a *pattern matrix*, a transformation matrix that maps the pattern's internal coordinate system to the default coordinate system of the pattern's *parent content stream* (the content stream in which the pattern is defined as a resource). The concatenation of the pattern matrix with that of the parent content stream establishes the *pattern coordinate space*, within which all graphics objects in the pattern are interpreted.

For example, if a pattern is used on a page, the pattern will appear in the **Pattern** subdictionary of that page's resource dictionary, and the pattern matrix maps pattern space to the default (initial) coordinate space of the page. Changes to the page's transformation matrix that occur within the page's content stream, such as rotation and scaling, have no effect on the pattern; it maintains its original relationship to the page no matter where on the page it is used. Similarly, if a pattern is used within a form XObject (see Section 4.9, "Form XObjects"), the pattern matrix maps pattern space to the form's default user space (that is, the form coordinate space at the time the form is painted with the **Do** operator). Finally, a pattern may be used within another pattern; the inner pattern's matrix defines its relationship to the pattern space of the outer pattern.

***Note:** PostScript allows a pattern to be defined in one context but used in another. For example, a pattern might be defined on a page (that is, its pattern matrix maps the pattern coordinate space to the user space of the page) but be used in a form on that page, so that its relationship to the page is independent of each individual placement of the form. PDF does not support this feature; in PDF, all patterns are local to the context in which they are defined.*

4.6.2 Tiling Patterns

A *tiling pattern* consists of a small graphical figure called a *pattern cell*. Painting with the pattern replicates the cell at fixed horizontal and vertical intervals to fill an area. The effect is as if the figure were painted on the surface of a clear glass tile, identical copies of which were then laid down in an array covering the area and trimmed to its boundaries. This is called *tiling* the area.

The pattern cell can include graphical elements such as filled areas, text, and sampled images. Its shape need not be rectangular, and the spacing of tiles can differ from the dimensions of the cell itself. When performing painting operations such as **S** (stroke) or **f** (fill), the viewer application paints the cell on the current page as many times as necessary to fill an area. The order in which individual tiles (instances of the cell) are painted is unspecified and unpredictable; it is inadvisable for the figures on adjacent tiles to overlap.

The appearance of the pattern cell is defined by a content stream containing the painting operators needed to paint one instance of the cell. Besides the usual entries common to all streams (see Table 3.4 on page 38), this stream's dictionary has the additional entries listed in Table 4.22.

TABLE 4.22 Additional entries specific to a type 1 pattern dictionary

KEY	TYPE	VALUE
Type	name	<i>(Optional)</i> The type of PDF object that this dictionary describes; if present, must be Pattern for a pattern dictionary.
PatternType	integer	<i>(Required)</i> A code identifying the type of pattern that this dictionary describes; must be 1 for a tiling pattern.
PaintType	integer	<i>(Required)</i> A code that determines how the color of the pattern cell is to be specified: <ol style="list-style-type: none"> 1 <i>Colored tiling pattern.</i> The pattern's content stream itself specifies the colors used to paint the pattern cell. When the content stream begins execution, the current color is the one that was initially in effect in the pattern's parent content stream. (This is similar to the definition of the pattern matrix; see Section 4.6.1, "General Properties of Patterns.") 2 <i>Uncolored tiling pattern.</i> The pattern's content stream does not specify any color information. Instead, the entire pattern cell is painted with a separately specified color each time the pattern is used. Essentially, the content stream describes a <i>stencil</i> through which the cur-

rent color is to be poured. The content stream must not invoke operators that specify colors or other color-related parameters in the graphics state; otherwise, an error will occur (see Section 4.5.7, “Color Operators”). The content stream may paint an image mask, however, since it does not specify any color information (see “Stencil Masking” on page 276).

TilingType	integer	<p>(Required) A code that controls adjustments to the spacing of tiles relative to the device pixel grid:</p> <ol style="list-style-type: none"> 1 <i>Constant spacing.</i> Pattern cells are spaced consistently—that is, by a multiple of a device pixel. To achieve this, the viewer application may need to distort the pattern cell slightly by making small adjustments to XStep, YStep, and the transformation matrix. The amount of distortion does not exceed 1 device pixel. 2 <i>No distortion.</i> The pattern cell is not distorted, but the spacing between pattern cells may vary by as much as 1 device pixel, both horizontally and vertically, when the pattern is painted. This achieves the spacing requested by XStep and YStep <i>on average</i>, but not necessarily for each individual pattern cell. 3 <i>Constant spacing and faster tiling.</i> Pattern cells are spaced consistently as in tiling type 1, but with additional distortion permitted to enable a more efficient implementation.
BBox	rectangle	(Required) An array of four numbers in the pattern coordinate system giving the coordinates of the left, bottom, right, and top edges, respectively, of the pattern cell’s bounding box. These boundaries are used to clip the pattern cell.
XStep	number	(Required) The desired horizontal spacing between pattern cells, measured in the pattern coordinate system.
YStep	number	(Required) The desired vertical spacing between pattern cells, measured in the pattern coordinate system. Note that XStep and YStep may differ from the dimensions of the pattern cell implied by the BBox entry. This allows tiling with irregularly shaped figures. XStep and YStep may be either positive or negative, but not zero.
Resources	dictionary	(Required) A resource dictionary containing all of the named resources required by the pattern’s content stream (see Section 3.7.2, “Resource Dictionaries”).
Matrix	array	(Optional) An array of six numbers specifying the pattern matrix (see Section 4.6.1, “General Properties of Patterns”). Default value: the identity matrix [1 0 0 1 0 0].

The pattern dictionary's **BBox**, **XStep**, and **YStep** values are interpreted in the pattern coordinate system, and the graphics objects in the pattern's content stream are defined with respect to that coordinate system. The placement of pattern cells in the tiling is based on the location of one *key pattern cell*, which is then displaced by multiples of **XStep** and **YStep** to replicate the pattern. The origin of the key pattern cell coincides with the origin of the pattern coordinate system; the phase of the tiling can be controlled by the translation components of the **Matrix** entry in the pattern dictionary.

The first step in painting with a tiling pattern is to establish the pattern as the current color in the graphics state. Subsequent painting operations will tile the painted areas with the pattern cell described by the pattern's content stream. Whenever it needs to obtain the pattern cell, the viewer application does the following:

1. Saves the current graphics state (as if by invoking the **q** operator)
2. Installs the graphics state that was in effect at the beginning of the pattern's parent content stream, with the current transformation matrix altered by the pattern matrix as described in Section 4.6.1, "General Properties of Patterns"
3. Paints the graphics objects specified in the pattern's content stream
4. Restores the saved graphics state (as if by invoking the **Q** operator)

Note: The pattern's content stream should not set any of the device-dependent parameters in the graphics state (see Table 4.3 on page 150). Doing so may result in incorrect output.

Colored Tiling Patterns

A *colored tiling pattern* is one whose color is self-contained. In the course of painting the pattern cell, the pattern's content stream explicitly sets the color of each graphical element it paints. A single pattern cell can contain elements that are painted different colors; it can also contain sampled grayscale or color images. This type of pattern is identified by a pattern type of 1 and a paint type of 1 in the pattern dictionary.

When the current color space is a **Pattern** space, a colored tiling pattern can be selected as the current color by supplying its name as the single operand to the **SCN** or **scn** operator. This name must be the key of an entry in the **Pattern** sub-dictionary of the current resource dictionary (see Section 3.7.2, “Resource Dictionaries”), whose value is the stream object representing the pattern. Since the pattern defines its own color information, no additional operands representing color components are specified to **SCN** or **scn**. For example, if **P1** is the name of a pattern resource in the current resource dictionary, the following code establishes it as the current nonstroking color:

```
/Pattern cs
/P1 scn
```

Subsequent executions of nonstroking painting operators, such as **f** (fill), **Tj** (show text), or **Do** (paint external object) with an image mask, will use the designated pattern to tile the areas to be painted.

Example 4.19 defines a page (object 5) that paints three circles and a triangle using a colored tiling pattern (object 15) over a yellow background. The pattern consists of the symbols for the four suits of playing cards (spades, hearts, diamonds, and clubs), which are character glyphs taken from the ZapfDingbats font (see Section D.4, “ZapfDingbats Set and Encoding”); the pattern’s content stream specifies the color of each glyph. Plate 8 shows the results.

Example 4.19

```
5 0 obj                                % Page object
  << /Type /Page
    /Parent 2 0 R
    /Resources 10 0 R
    /Contents 30 0 R
    /CropBox [0 0 225 225]
  >>
endobj

10 0 obj                                % Resource dictionary for page
  << /Pattern << /P1 15 0 R >>
  >>
endobj
```



```

15 0 obj                                % Pattern definition
  << /Type /Pattern
    /PatternType 1                       % Tiling pattern
    /PaintType 1                         % Colored
    /TilingType 2
    /BBox [0 0 100 100]
    /XStep 100
    /YStep 100
    /Resources 16 0 R
    /Matrix [0.4 0.0 0.0 0.4 0.0 0.0]
    /Length 183
  >>
stream
BT                                       % Begin text object
  /F1 1 Tf                               % Set text font and size
  64 0 0 64 7.1771 2.4414 Tm            % Set text matrix
  0 Tc                                    % Set character spacing
  0 Tw                                    % Set word spacing
  1.0 0.0 0.0 rg                        % Set nonstroking color to red
  (\001) Tj                              % Show spade glyph
  0.7478 -0.007 TD                      % Move text position
  0.0 1.0 0.0 rg                        % Set nonstroking color to green
  (\002) Tj                              % Show heart glyph
  -0.7323 0.7813 TD                    % Move text position
  0.0 0.0 1.0 rg                        % Set nonstroking color to blue
  (\003) Tj                              % Show diamond glyph
  0.6913 0.007 TD                      % Move text position
  0.0 0.0 0.0 rg                        % Set nonstroking color to black
  (\004) Tj                              % Show club glyph
ET                                       % End text object
endstream
endobj

16 0 obj                                % Resource dictionary for pattern
  << /Font << /F1 20 0 R >>
  >>
endobj

20 0 obj                                % Font for pattern
  << /Type /Font
    /Subtype /Type1
    /Encoding 21 0 R
    /BaseFont /ZapfDingbats
  >>
endobj

```

```

21 0 obj                                % Font encoding
  << /Type /Encoding
    /Differences [1 /a109 /a110 /a111 /a112]
  >>
endobj

30 0 obj                                % Contents of page
  << /Length 1252 >>
stream
  0.0 G                                  % Set stroking color to black
  1.0 1.0 0.0 rg                         % Set nonstroking color to yellow
  25 175 175 -150 re                      % Construct rectangular path
  f                                        % Fill path
  /Pattern cs                             % Set pattern color space
  /P1 scn                                 % Set pattern as nonstroking color
  99.92 49.92 m                           % Start new path
  99.92 77.52 77.52 99.92 49.92 99.92 c  % Construct lower-left circle
  22.32 99.92 -0.08 77.52 -0.08 49.92 c
  -0.08 22.32 22.32 -0.08 49.92 -0.08 c
  77.52 -0.08 99.92 22.32 99.92 49.92 c
  B                                        % Fill and stroke path
  224.96 49.92 m                          % Start new path
  224.96 77.52 202.56 99.92 174.96 99.92 c % Construct lower-right circle
  147.36 99.92 124.96 77.52 124.96 49.92 c
  124.96 22.32 147.36 -0.08 174.96 -0.08 c
  202.56 -0.08 224.96 22.32 224.96 49.92 c
  B                                        % Fill and stroke path
  87.56 201.70 m                          % Start new path
  63.66 187.90 55.46 157.32 69.26 133.40 c % Construct upper circle
  83.06 109.50 113.66 101.30 137.56 115.10 c
  161.46 128.90 169.66 159.50 155.86 183.40 c
  142.06 207.30 111.46 215.50 87.56 201.70 c
  B                                        % Fill and stroke path
  50 50 m                                  % Start new path
  175 50 l                                 % Construct triangular path
  112.5 158.253 l
  b                                        % Close, fill, and stroke path
endstream
endobj

```

Several features of Example 4.19 are noteworthy:

- The three circles and the triangle are painted with the same pattern. The pattern cells align, even though the circles and triangle are not aligned with respect to the pattern cell. For example, the position of the blue diamonds varies relative to the three circles.
- The pattern cell does not completely cover the tile: it leaves the spaces between the glyphs unpainted. When the tiling pattern is used as a color, the existing background (the yellow rectangle) shows through these unpainted areas.

Uncolored Tiling Patterns

An *uncolored tiling pattern* is one that has no inherent color: the color must be specified separately whenever the pattern is used. This type of pattern is identified by a pattern type of 1 and a paint type of 2 in the pattern dictionary. The pattern's content stream does not explicitly specify any colors; it can paint an image mask (see “Stencil Masking” on page 276), but no other kind of image. This provides a way to tile different regions of the page with pattern cells having the same shape but different colors.

A **Pattern** color space representing an uncolored tiling pattern requires a parameter: an object identifying the *underlying color space* in which the actual color of the pattern is to be specified. The underlying color space is given as the second element of the array that defines the **Pattern** color space. For example, the array

```
[/Pattern /DeviceRGB]
```

defines a **Pattern** color space with **DeviceRGB** as its underlying color space.

Note: *The underlying color space cannot be another **Pattern** color space.*

Operands supplied to the **SCN** or **scn** operator in such a color space must include a color value in the underlying color space, specified by one or more numeric color components, as well as the name of a pattern object representing an uncolored tiling pattern. For example, if the current resource dictionary (see Section 3.7.2, “Resource Dictionaries”) defines Cs3 as the name of a **ColorSpace**

resource whose value is the **Pattern** color space shown above, and P2 as a **Pattern** resource denoting an uncolored tiling pattern, then the code

```
/Cs3 cs
0.30 0.75 0.21 /P2 scn
```

establishes Cs3 as the current nonstroking color space and P2 as the current nonstroking color, to be painted in the color represented by the specified components in the **DeviceRGB** color space. Subsequent executions of nonstroking painting operators, such as **f** (fill), **Tj** (show text), and **Do** (paint external object) with an image mask, will use the designated pattern and color to tile the areas to be painted. The same pattern can be used repeatedly with a different color each time.

Example 4.20 is similar to Example 4.19 on page 224, except that it uses an uncolored tiling pattern to paint the three circles and the triangle, each in a different color (see Plate 9). To do so, it supplies four operands each time it invokes the **scn** operator: three numbers denoting the components of the desired color in the underlying **DeviceRGB** color space, along with the name of the pattern.

Example 4.20

```
5 0 obj                                % Page object
  << /Type /Page
    /Parent 2 0 R
    /Resources 10 0 R
    /Contents 30 0 R
    /CropBox [0 0 225 225]
  >>
endobj

10 0 obj                                 % Resource dictionary for page
  << /ColorSpace << /Cs12 12 0 R >>
    /Pattern << /P1 15 0 R >>
  >>
endobj

12 0 obj                                 % Color space
  [/Pattern /DeviceRGB]
endobj
```

```

15 0 obj                                % Pattern definition
  << /Type /Pattern
    /PatternType 1                       % Tiling pattern
    /PaintType 2                         % Uncolored
    /TilingType 2
    /BBox [0 0 100 100]
    /XStep 100
    /YStep 100
    /Resources 16 0 R
    /Matrix [0.4 0.0 0.0 0.4 0.0 0.0]
    /Length 127
  >>
stream
  BT                                     % Begin text object
    /F1 1 Tf                             % Set text font and size
    64 0 0 64 7.1771 2.4414 Tm          % Set text matrix
    0 Tc                                 % Set character spacing
    0 Tw                                 % Set word spacing
    (\001) Tj                           % Show spade glyph
    0.7478 -0.007 TD                   % Move text position
    (\002) Tj                           % Show heart glyph
    -0.7323 0.7813 TD                 % Move text position
    (\003) Tj                           % Show diamond glyph
    0.6913 0.007 TD                   % Move text position
    (\004) Tj                           % Show club glyph
  ET                                     % End text object
endstream
endobj

16 0 obj                                % Resource dictionary for pattern
  << /Font << /F1 20 0 R >>
  >>
endobj

20 0 obj                                % Font for pattern
  << /Type /Font
    /Subtype /Type1
    /Encoding 21 0 R
    /BaseFont /ZapfDingbats
  >>
endobj

```

```

21 0 obj                                % Font encoding
  << /Type /Encoding
    /Differences [1 /a109 /a110 /a111 /a112]
  >>
endobj

30 0 obj                                % Contents of page
  << /Length 1316 >>
stream
  0.0 G                                  % Set stroking color to black
  1.0 1.0 0.0 rg                         % Set nonstroking color to yellow
  25 175 175 -150 re                     % Construct rectangular path
  f                                       % Fill path
  /Cs12 cs                                % Set pattern color space
  0.77 0.20 0.00 /P1 scn                 % Set nonstroking color and pattern
  99.92 49.92 m                           % Start new path
  99.92 77.52 77.52 99.92 49.92 99.92 c % Construct lower-left circle
  22.32 99.92 -0.08 77.52 -0.08 49.92 c
  -0.08 22.32 22.32 -0.08 49.92 -0.08 c
  77.52 -0.08 99.92 22.32 99.92 49.92 c
  B                                       % Fill and stroke path
  0.2 0.8 0.4 /P1 scn                     % Change nonstroking color
  224.96 49.92 m                           % Start new path
  224.96 77.52 202.56 99.92 174.96 99.92 c % Construct lower-right circle
  147.36 99.92 124.96 77.52 124.96 49.92 c
  124.96 22.32 147.36 -0.08 174.96 -0.08 c
  202.56 -0.08 224.96 22.32 224.96 49.92 c
  B                                       % Fill and stroke path
  0.3 0.7 1.0 /P1 scn                     % Change nonstroking color
  87.56 201.70 m                           % Start new path
  63.66 187.90 55.46 157.30 69.26 133.40 c % Construct upper circle
  83.06 109.50 113.66 101.30 137.56 115.10 c
  161.46 128.90 169.66 159.50 155.86 183.40 c
  142.06 207.30 111.46 215.50 87.56 201.70 c
  B                                       % Fill and stroke path
  0.5 0.2 1.0 /P1 scn                     % Change nonstroking color
  50 50 m                                  % Start new path
  175 50 l                                 % Construct triangular path
  112.5 158.253 l
  b                                       % Close, fill, and stroke path
endstream
endobj

```

4.6.3 Shading Patterns

Shading patterns (PDF 1.3) provide a smooth transition between colors across an area to be painted, independent of the resolution of any particular output device and without specifying the number of steps in the color transition. Patterns of this type are described by pattern dictionaries with a pattern type of 2. Table 4.23 shows the contents of this type of dictionary.

TABLE 4.23 Entries in a type 2 pattern dictionary

KEY	TYPE	VALUE
Type	integer	(<i>Optional</i>) The type of PDF object that this dictionary describes; if present, must be Pattern for a pattern dictionary.
PatternType	integer	(<i>Required</i>) A code identifying the type of pattern that this dictionary describes; must be 2 for a shading pattern.
Shading	dictionary or stream	(<i>Required</i>) A shading object (see below) defining the shading pattern's gradient fill. The contents of the dictionary consist of the entries in Table 4.25 on page 234, plus those in one of Tables 4.26 to 4.31 on pages 237 to 253.
Matrix	array	(<i>Optional</i>) An array of six numbers specifying the pattern matrix (see Section 4.6.1, "General Properties of Patterns"). Default value: the identity matrix [1 0 0 1 0 0].
ExtGState	dictionary	(<i>Optional</i>) A graphics state parameter dictionary (see Section 4.3.4, "Graphics State Parameter Dictionaries") containing graphics state parameters to be put into effect temporarily while the shading pattern is painted. Any parameters that are not so specified are inherited from the graphics state that was in effect at the beginning of the content stream in which the pattern is defined as a resource.

The most significant entry is **Shading**, whose value is a *shading object* defining the properties of the shading pattern's *gradient fill*. This is a complex "paint" that determines the type of color transition the shading pattern produces when painted across an area. A shading object may be a dictionary or a stream, depending on the type of shading; the term *shading dictionary* is used generically throughout this section to refer to either a dictionary object or the dictionary portion of a stream object. (Those shading objects that are streams are specifically identified as such in the descriptions of particular shading types; unless otherwise stated, they are understood to be simple dictionaries instead.)

By setting a shading pattern as the current color in the graphics state, a PDF content stream can use it with painting operators such as **f** (fill), **S** (stroke), **Tj** (show text), or **Do** (paint external object) with an image mask to paint a path, character glyph, or mask with a smooth color transition. When a shading is used in this way, the geometry of the gradient fill is independent of that of the object being painted.

Shading Operator

When the area to be painted is a relatively simple shape whose geometry is the same as that of the gradient fill itself, the **sh** operator can be used instead of the usual painting operators. **sh** accepts a shading dictionary as an operand and applies the corresponding gradient fill directly to current user space. This operator does not require the creation of a pattern dictionary or a path and works without reference to the current color in the graphics state. Table 4.24 describes the **sh** operator.

***Note:** Patterns defined by type 2 pattern dictionaries do not tile. To create a tiling pattern containing a gradient fill, invoke the **sh** operator from within the content stream of a type 1 (tiling) pattern.*

TABLE 4.24 Shading operator

OPERANDS	OPERATOR	DESCRIPTION
<i>name</i>	sh	<p>(PDF 1.3) Paint the shape and color shading described by a shading dictionary, subject to the current clipping path. The current color in the graphics state is neither used nor altered. The effect is different from that of painting a path using a shading pattern as the current color.</p> <p><i>name</i> is the name of a shading dictionary resource in the Shading subdictionary of the current resource dictionary (see Section 3.7.2, “Resource Dictionaries”). All coordinates in the shading dictionary are interpreted relative to the current user space. (By contrast, when a shading dictionary is used in a type 2 pattern, the coordinates are expressed in pattern space.) All colors are interpreted in the color space identified by the shading dictionary’s ColorSpace entry (see Table 4.25 on page 234). The Background entry, if present, is ignored.</p> <p>This operator should be applied only to bounded or geometrically defined shadings. If applied to an unbounded shading, it will paint the shading’s gradient fill across the entire clipping region, which may be time-consuming.</p>

Shading Dictionaries

A shading dictionary specifies details of a particular gradient fill, including the type of shading to be used, the geometry of the area to be shaded, and the geometry of the gradient fill itself. Various shading types are available, depending on the value of the dictionary's **ShadingType** entry:

- *Function-based shadings* (type 1) define the color of every point in the domain using a mathematical function (not necessarily smooth or continuous).
- *Axial shadings* (type 2) define a color blend along a line between two points, optionally extended beyond the boundary points by continuing the boundary colors.
- *Radial shadings* (type 3) define a blend between two circles, optionally extended beyond the boundary circles by continuing the boundary colors. This type of shading is commonly used to represent three-dimensional spheres and cones.
- *Free-form Gouraud-shaded triangle meshes* (type 4) define a common construct used by many three-dimensional applications to represent complex colored and shaded shapes. Vertices are specified in free-form geometry.
- *Lattice-form Gouraud-shaded triangle meshes* (type 5) are based on the same geometrical construct as type 4, but with vertices specified as a pseudo-rectangular lattice.
- *Coons patch meshes* (type 6) construct a shading from one or more color patches, each bounded by four cubic Bézier curves.
- *Tensor-product patch meshes* (type 7) are similar to type 6, but with additional control points in each patch, affording greater control over color mapping.

Table 4.25 shows the entries that all shading dictionaries share in common; entries specific to particular shading types are described in the relevant sections below.

Note: *The term target coordinate space, used in many of the following descriptions, refers to the coordinate space into which a shading is painted. For shadings used with a type 2 pattern dictionary, this is the pattern coordinate space, discussed in Section 4.6.1, “General Properties of Patterns.” For shadings used directly with the **sh** operator, it is the current user space.*

TABLE 4.25 Entries common to all shading dictionaries

KEY	TYPE	VALUE
ShadingType	integer	<p><i>(Required)</i> The shading type:</p> <ol style="list-style-type: none"> 1 Function-based shading 2 Axial shading 3 Radial shading 4 Free-form Gouraud-shaded triangle mesh 5 Lattice-form Gouraud-shaded triangle mesh 6 Coons patch mesh 7 Tensor-product patch mesh
ColorSpace	name or array	<p><i>(Required)</i> The color space in which color values are expressed. This may be any device, CIE-based, or special color space except a Pattern space. See “Color Space: Special Considerations,” below, for further information.</p>
Background	array	<p><i>(Optional)</i> An array of color components appropriate to the color space, specifying a single background color value. If present, this color is used before any painting operation involving the shading, to fill those portions of the area to be painted that lie outside the bounds of the shading object itself. In the opaque imaging model, the effect is as if the painting operation were performed twice: first with the background color and then again with the shading.</p> <p>Note: <i>The background color is applied only when the shading is used as part of a shading pattern, not when it is painted directly with the sh operator.</i></p>
BBox	rectangle	<p><i>(Optional)</i> An array of four numbers giving the left, bottom, right, and top coordinates, respectively, of the shading’s bounding box. The coordinates are interpreted in the shading’s target coordinate space. If present, this bounding box is applied as a temporary clipping boundary when the shading is painted, in addition to the current clipping path and any other clipping boundaries in effect at that time.</p>
AntiAlias	boolean	<p><i>(Optional)</i> A flag indicating whether to filter the shading function to prevent <i>aliasing</i> artifacts. The shading operators sample shading functions at a rate determined by the resolution of the output device. Aliasing can occur if the function is not smooth—that is, if it has a high spatial frequency relative to the sampling rate. Anti-aliasing can be computationally expensive and is usually unnecessary, since most shading functions are smooth enough, or are sampled at a high enough frequency, to avoid aliasing effects. Anti-aliasing may not be implemented on some output devices, in which case this flag is ignored. Default value: false.</p>

Shading types 4 to 7 are defined by a stream containing descriptive data characterizing the shading's gradient fill. In these cases, the shading dictionary is also a stream dictionary and can contain any of the standard entries common to all streams (see Table 3.4 on page 38). In particular, it will always include a **Length** entry, which is required for all streams.

In addition, some shading dictionaries also include a **Function** entry whose value is a function object (dictionary or stream) defining how colors vary across the area to be shaded. In such cases, the shading dictionary usually defines the geometry of the shading, while the function defines the color transitions across that geometry. The function is required for some types of shading and optional for others. Functions are described in detail in Section 3.9, "Functions."

Note: Discontinuous color transitions, or those with high spatial frequency, may exhibit aliasing effects when painted at low effective resolutions.

Color Space: Special Considerations

Conceptually, a shading determines a color value for each individual point within the area to be painted. In practice, however, the shading may actually be used to compute color values only for some subset of the points in the target area, with the colors of the intervening points determined by interpolation between the ones computed. Viewer applications are free to use this strategy as long as the interpolated color values approximate those defined by the shading to within the smoothness tolerance specified in the graphics state (see Section 6.5.2, "Smoothness Tolerance"). The **ColorSpace** entry common to all shading dictionaries not only defines the color space in which the shading specifies its color values, but also determines the color space in which color interpolation is performed.

Note: Some shading types (4 to 7) perform interpolation on a parametric value supplied as input to the shading's color function, as described in the relevant sections below. This form of interpolation is conceptually distinct from the interpolation described here, which operates on the output color values produced by the color function and takes place within the shading's target color space.

Gradient fills between colors defined by most shadings are implemented using a variety of interpolation algorithms, and these algorithms are sensitive to the characteristics of the color space. Linear interpolation, for example, may have observably different results when applied in a **DeviceCMYK** color space than in a **Lab** color space, even if the starting and ending colors are perceptually identical. The

difference arises because the two color spaces are not linear relative to each other. Shadings are rendered according to the following rules:

- If **ColorSpace** is a device color space different from the native color space of the output device, color values in the shading will be converted to the native color space using the standard conversion formulas described in Section 6.2, “Conversions among Device Color Spaces.” To optimize performance, these conversions may take place at any time (either before or after any interpolation on the color values in the shading). Thus, shadings defined with device color spaces may have color gradient fills that are less accurate and somewhat device-dependent. (This does not apply to axial and radial shadings—shading types 2 and 3—because those shading types perform gradient fill calculations on a single variable and then convert to parametric colors.)
- If **ColorSpace** is a CIE-based color space, all gradient fill calculations will be performed in that space. Conversion to device colors will occur only after all interpolation calculations have been performed. Thus, the color gradients will be device-independent for the colors generated at each point.
- If **ColorSpace** is a **Separation** or **DeviceN** color space and the specified colorants are supported, no color conversion calculations are needed. If the specified colorants are not supported (so that the space’s alternate color space must be used), gradient fill calculations will be performed in the designated **Separation** or **DeviceN** color space before conversion to the alternate space. Thus, non-linear tint transformation functions will be accommodated for the best possible representation of the shading.
- If **ColorSpace** is an **Indexed** color space, all color values specified in the shading will be immediately converted to the base color space. Depending on whether the base color space is a device or CIE-based space, gradient fill calculations will be performed as stated above. Interpolation never occurs in an **Indexed** color space, which is quantized and therefore inappropriate for calculations that assume a continuous range of colors. For similar reasons, an **Indexed** color space is not allowed in any shading whose color values are generated by a function; this applies to any shading dictionary that contains a **Function** entry.

Shading Types

In addition to the entries listed in Table 4.25 on page 234, all shading dictionaries have entries specific to the type of shading they represent, as indicated by the

value of their **ShadingType** entry. The following sections describe the available shading types and the dictionary entries specific to each.

Type 1 (Function-Based) Shadings

In type 1 (function-based) shadings, the color at every point in the domain is defined by a specified mathematical function. The function need not be smooth or continuous. This is the most general of the available shading types, and is useful for shadings that cannot be adequately described with any of the other types. Table 4.26 shows the shading dictionary entries specific to this type of shading, in addition to those common to all shading dictionaries (Table 4.25 on page 234).

*Note: This type of shading may not be used with an **Indexed** color space.*

TABLE 4.26 Additional entries specific to a type 1 shading dictionary

KEY	TYPE	VALUE
Domain	array	<i>(Optional)</i> An array of four numbers [x_{\min} x_{\max} y_{\min} y_{\max}] specifying the rectangular domain of coordinates over which the color function(s) are defined. Default value: [0.0 1.0 0.0 1.0].
Matrix	array	<i>(Optional)</i> An array of six numbers specifying a transformation matrix mapping the coordinate space specified by the Domain entry into the shading's target coordinate space. For example, to map the domain rectangle [0.0 1.0 0.0 1.0] to a 1-inch square with lower-left corner at coordinates (100, 100) in default user space, the Matrix value would be [72 0 0 72 100 100]. Default value: the identity matrix [1 0 0 1 0 0].
Function	function	<i>(Required)</i> A 2-in, n -out function or an array of n 2-in, 1-out functions (where n is the number of color components in the shading dictionary's color space). Each function's domain must be a superset of that of the shading dictionary. If the value returned by the function for a given color component is out of range, it will be adjusted to the nearest valid value.

The domain rectangle (**Domain**) establishes an internal coordinate space for the shading that is independent of the target coordinate space in which it is to be painted. The color function(s) (**Function**) specify the color of the shading at each point within this domain rectangle. The transformation matrix (**Matrix**) then maps the domain rectangle into a corresponding rectangle or parallelogram in the target coordinate space. Points within the shading's bounding box (**BBox**)

that fall outside this transformed domain rectangle will be painted with the shading's background color (**Background**); if the shading dictionary has no **Background** entry, such points will be left unpainted. If the function is undefined at any point within the declared domain rectangle, an error may occur, even if the corresponding transformed point falls outside the shading's bounding box.

Type 2 (Axial) Shadings

Type 2 (axial) shadings define a color blend that varies along a linear axis between two endpoints and extends indefinitely perpendicular to that axis. The shading may optionally be extended beyond either or both endpoints by continuing the boundary colors indefinitely. Table 4.27 shows the shading dictionary entries specific to this type of shading, in addition to those common to all shading dictionaries (Table 4.25 on page 234).

Note: This type of shading may not be used with an **Indexed** color space.

TABLE 4.27 Additional entries specific to a type 2 shading dictionary

KEY	TYPE	VALUE
Coords	array	<i>(Required)</i> An array of four numbers $[x_0 \ y_0 \ x_1 \ y_1]$ specifying the starting and ending coordinates of the axis, expressed in the shading's target coordinate space.
Domain	array	<i>(Optional)</i> An array of two numbers $[t_0 \ t_1]$ specifying the limiting values of a parametric variable t . The variable is considered to vary linearly between these two values as the color gradient varies between the starting and ending points of the axis. The variable t becomes the input argument to the color function(s). Default value: <code>[0.0 1.0]</code> .
Function	function	<i>(Required)</i> A 1-in, n -out function or an array of n 1-in, 1-out functions (where n is the number of color components in the shading dictionary's color space). The function(s) are called with values of the parametric variable t in the domain defined by the Domain entry. Each function's domain must be a superset of that of the shading dictionary. If the value returned by the function for a given color component is out of range, it will be adjusted to the nearest valid value.
Extend	array	<i>(Optional)</i> An array of two boolean values specifying whether to extend the shading beyond the starting and ending points of the axis, respectively. Default value: <code>[false false]</code> .

The color blend is accomplished by linearly mapping each point (x, y) along the axis between the endpoints (x_0, y_0) and (x_1, y_1) to a corresponding point in the domain specified by the shading dictionary's **Domain** entry. The points $(0, 0)$ and $(1, 0)$ in the domain correspond respectively to (x_0, y_0) and (x_1, y_1) on the axis. Since all points along a line in domain space perpendicular to the line from $(0, 0)$ to $(1, 0)$ will have the same color, only the new value of x needs to be computed:

$$x' = \frac{(x_1 - x_0) \times (x - x_0) + (y_1 - y_0) \times (y - y_0)}{(x_1 - x_0)^2 + (y_1 - y_0)^2}$$

The value of the parametric variable t is then determined from x' as follows:

- For $0 \leq x' \leq 1$, $t = t_0 + (t_1 - t_0) \times x'$.
- For $x' < 0$, if the first element of the **Extend** array is **true**, then $t = t_0$; otherwise, t is undefined and the point is left unpainted.
- For $x' > 1$, if the second element of the **Extend** array is **true**, then $t = t_1$; otherwise, t is undefined and the point is left unpainted.

The resulting value of t is then passed as input to the function(s) defined by the shading dictionary's **Function** entry, yielding the component values of the color with which to paint the point (x, y) .

Plate 10 shows three examples of the use of an axial shading to fill a rectangle and display text. The area to be filled extends beyond the shading's bounding box. The shading is the same in all three cases, except for the values of the **Background** and **Extend** entries in the shading dictionary. In the first example, the shading is not extended at either end and no background color is specified, so the shading is clipped to its bounding box at both ends. The second example still has no background color specified, but the shading is extended at both ends; the result is to fill the remaining portions of the filled area with the colors defined at the ends of the shading. In the third example, the shading is extended at both ends and a background color is specified, so the background color is used for the portions of the filled area beyond the ends of the shading.

Type 3 (Radial) Shadings

Type 3 (radial) shadings define a color blend that varies between two circles. Shadings of this type are commonly used to depict three-dimensional spheres and cones. Table 4.28 shows the shading dictionary entries specific to this type of

shading, in addition to those common to all shading dictionaries (Table 4.25 on page 234).

Note: *This type of shading may not be used with an **Indexed** color space.*

TABLE 4.28 Additional entries specific to a type 3 shading dictionary

KEY	TYPE	VALUE
Coords	array	<i>(Required)</i> An array of six numbers $[x_0 y_0 r_0 x_1 y_1 r_1]$ specifying the centers and radii of the starting and ending circles, expressed in the shading's target coordinate space. The radii r_0 and r_1 must both be greater than or equal to 0. If one radius is 0, the corresponding circle is treated as a point; if both are 0, nothing is painted.
Domain	array	<i>(Optional)</i> An array of two numbers $[t_0 t_1]$ specifying the limiting values of a parametric variable t . The variable is considered to vary linearly between these two values as the color gradient varies between the starting and ending circles. The variable t becomes the input argument to the color function(s). Default value: $[0.0 1.0]$.
Function	function	<i>(Required)</i> A 1-in, n -out function or an array of n 1-in, 1-out functions (where n is the number of color components in the shading dictionary's color space). The function(s) are called with values of the parametric variable t in the domain defined by the shading dictionary's Domain entry. Each function's domain must be a superset of that of the shading dictionary. If the value returned by the function for a given color component is out of range, it will be adjusted to the nearest valid value.
Extend	array	<i>(Optional)</i> An array of two boolean values specifying whether to extend the shading beyond the starting and ending circles, respectively. Default value: $[\text{false } \text{false}]$.

The color blend is based on a family of *blend circles* interpolated between the starting and ending circles that are defined by the shading dictionary's **Coords** entry. The blend circles are defined in terms of a subsidiary parametric variable

$$s = \frac{t - t_0}{t_1 - t_0}$$

which varies linearly between 0.0 and 1.0 as t varies across the domain from t_0 to t_1 , as specified by the dictionary's **Domain** entry. The center and radius of each blend circle are given by the parametric equations

$$\begin{aligned}x_c(s) &= x_0 + s \times (x_1 - x_0) \\y_c(s) &= y_0 + s \times (y_1 - y_0) \\r(s) &= r_0 + s \times (r_1 - r_0)\end{aligned}$$

Each value of s between 0.0 and 1.0 determines a corresponding value of t , which is then passed as the input argument to the function(s) defined by the shading dictionary's **Function** entry. This yields the component values of the color with which to fill the corresponding blend circle. For values of s not lying between 0.0 and 1.0, the boolean elements of the shading dictionary's **Extend** array determine whether and how the shading will be extended. If the first of the two elements is **true**, the shading is extended beyond the defined starting circle to values of s less than 0.0; if the second element is **true**, the shading is extended beyond the defined ending circle to s values greater than 1.0.

Note that either of the starting and ending circles may be larger than the other. If the shading is extended at the smaller end, the family of blend circles continues as far as that value of s for which the radius of the blend circle $r(s) = 0$; if the shading is extended at the larger end, the blend circles continue as far as that s value for which $r(s)$ is large enough to encompass the shading's entire bounding box (**BBox**). Extending the shading can thus cause painting to extend beyond the areas defined by the two circles themselves. The two examples in the rightmost column of Plate 11 depict the results of extending the shading at the smaller and larger ends, respectively.

Conceptually, all of the blend circles are painted in order of increasing values of s , from smallest to largest. Blend circles extending beyond the starting circle are painted in the same color defined by the shading dictionary's **Function** entry for the starting circle ($t = t_0, s = 0.0$); those extending beyond the ending circle are painted in the color defined for the ending circle ($t = t_1, s = 1.0$). The painting is opaque, with the color of each circle completely overlaying those preceding it; thus if a point lies within more than one blend circle, its final color will be that of the last of the enclosing circles to be painted, corresponding to the greatest value of s . Note the following points:

- If one of the starting and ending circles entirely contains the other, the shading will depict a sphere, as in Plates 12 and 13. In Plate 12, the inner circle has zero

radius; it is the starting circle in the figure on the left and the ending circle in the one on the right. Neither shading is extended at either the smaller or larger end. In Plate 13, the inner circle in both figures has a nonzero radius and the shading is extended at the larger end. In each plate, a background color is specified for the figure on the right, but not for the one on the left.

- If neither circle contains the other, the shading will depict a cone. If the starting circle is larger, the cone will appear to point out of the page; if the ending circle is larger, the cone will appear to point into the page (see Plate 11).

Example 4.21 paints the leaf-covered branch shown in Plate 14. Each leaf is filled with the same radial shading (object number 5). The color function (object 10) is a stitching function (described in Section 3.9.3, “Type 3 (Stitching) Functions”) whose two subfunctions (objects 11 and 12) are both exponential interpolation functions (see Section 3.9.2, “Type 2 (Exponential Interpolation) Functions”). Each leaf is drawn as a path and then filled with the shading, using code such as that shown in Example 4.22 (where the name `Sh1` is associated with object 5 by the **Shading** subdictionary of the current resource dictionary; see Section 3.7.2, “Resource Dictionaries”).

Example 4.21

```

5 0 obj                                     % Shading dictionary
  << /ShadingType 3
    /ColorSpace /DeviceCMYK
    /Coords [0.0 0.0 0.096 0.0 0.0 1.000]   % Concentric circles
    /Function 10 0 R
    /Extend [true true]
  >>
endobj

10 0 obj                                     % Color function
  << /FunctionType 3
    /Domain [0.0 1.0]
    /Functions [11 0 R 12 0 R]
    /Bounds [0.708]
    /Encode [1.0 0.0 0.0 1.0]
  >>
endobj

```

```

11 0 obj                                     % First subfunction
  << /FunctionType 2
    /Domain [0.0 1.0]
    /C0 [0.929 0.357 1.000 0.298]
    /C1 [0.631 0.278 1.000 0.027]
    /N 1.048
  >>
endobj

12 0 obj                                     % Second subfunction
  << /FunctionType 2
    /Domain [0.0 1.0]
    /C0 [0.929 0.357 1.000 0.298]
    /C1 [0.941 0.400 1.000 0.102]
    /N 1.374
  >>
endobj

```

Example 4.22

```

316.789 140.311 m                           % Move to start of leaf
303.222 146.388 282.966 136.518 279.122 121.983 c % Curved segment
277.322 120.182 l                           % Straight line
285.125 122.688 291.441 121.716 298.156 119.386 c % Curved segment
336.448 119.386 l                           % Straight line
331.072 128.643 323.346 137.376 316.789 140.311 c % Curved segment
W n                                           % Set clipping path
q                                             % Save graphics state
  27.7843 0.0000 0.0000 -27.7843 310.2461 121.1521 cm % Set matrix
  /Sh1 sh                                     % Paint shading
Q                                             % Restore graphics state

```

Type 4 Shadings (Free-Form Gouraud-Shaded Triangle Meshes)

Type 4 shadings (free-form Gouraud-shaded triangle meshes) are commonly used to represent complex colored and shaded three-dimensional shapes. The area to be shaded is defined by a path composed entirely of triangles. The color at each vertex of the triangles is specified, and a technique known as *Gouraud interpolation* is used to color the interiors. The interpolation functions defining the shading may be linear or nonlinear. Table 4.29 shows the entries specific to this type of shading dictionary, in addition to those common to all shading dictionaries (Table 4.25 on page 234) and stream dictionaries (Table 3.4 on page 38).

TABLE 4.29 Additional entries specific to a type 4 shading dictionary

KEY	TYPE	VALUE
BitsPerCoordinate	integer	<i>(Required)</i> The number of bits used to represent each vertex coordinate. Valid values are 1, 2, 4, 8, 12, 16, 24, and 32.
BitsPerComponent	integer	<i>(Required)</i> The number of bits used to represent each color component. Valid values are 1, 2, 4, 8, 12, and 16.
BitsPerFlag	integer	<i>(Required)</i> The number of bits used to represent the edge flag for each vertex (see below). Valid values of BitsPerFlag are 2, 4, and 8, but only the least significant 2 bits in each flag value are used. Valid values for the edge flag itself are 0, 1, and 2.
Decode	rectangle	<p><i>(Required)</i> An array of numbers specifying how to map vertex coordinates and color components into the appropriate ranges of values. The decoding method is similar to that used in image dictionaries (see “Decode Arrays” on page 271). The ranges are specified as follows:</p> $[x_{\min} \ x_{\max} \ y_{\min} \ y_{\max} \ c_{1,\min} \ c_{1,\max} \ \cdots \ c_{n,\min} \ c_{n,\max}]$ <p>Note that only one pair of c values should be specified if a Function entry is present.</p>
Function	function	<p><i>(Optional)</i> A 1-in, n-out function or an array of n 1-in, 1-out functions (where n is the number of color components in the shading dictionary’s color space). If this entry is present, the color data for each vertex must be specified by a single parametric variable rather than by n separate color components; the designated function(s) will be called with each interpolated value of the parametric variable to determine the actual color at each point. Each input value will be forced into the range interval specified for the corresponding color component in the shading dictionary’s Decode array. Each function’s domain must be a superset of that interval. If the value returned by the function for a given color component is out of range, it will be adjusted to the nearest valid value.</p> <p>This entry may not be used with an Indexed color space.</p>

Unlike shading types 1 to 3, types 4 to 7 are represented as streams. Each stream contains a sequence of vertex coordinates and color data that defines the triangle mesh. In a type 4 shading, each vertex is specified by the following values, in the order shown:

$$f \ x \ y \ c_1 \ \dots \ c_n$$

where

f is the vertex's edge flag (discussed below)

x and y are its horizontal and vertical coordinates

$c_1 \dots c_n$ are its color components

All vertex coordinates are expressed in the shading's target coordinate space. If the shading dictionary includes a **Function** entry, then only a single parametric value, t , is permitted for each vertex in place of the color components $c_1 \dots c_n$.

The *edge flag* associated with each vertex determines the way it connects to the other vertices of the triangle mesh. A vertex v_a with an edge flag value $f_a = 0$ begins a new triangle, unconnected to any other. At least two more vertices (v_b and v_c) must be provided, but their edge flags will be ignored. These three vertices define a triangle (v_a, v_b, v_c), as shown in Figure 4.16.

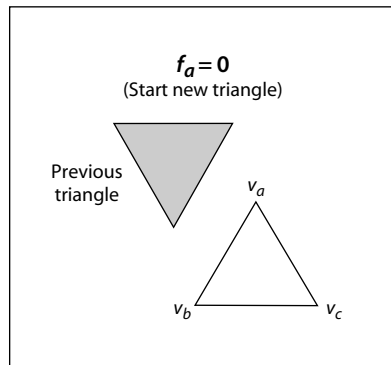


FIGURE 4.16 Starting a new triangle in a free-form Gouraud-shaded triangle mesh

Subsequent triangles are defined by a single new vertex combined with two vertices of the preceding triangle. Given triangle (v_a, v_b, v_c), where vertex v_a precedes vertex v_b in the data stream and v_b precedes v_c , a new vertex v_d can form a new triangle on side v_{bc} or side v_{ac} , as shown in Figure 4.17. (Side v_{ab} is assumed to be shared with a preceding triangle and so is not available for continuing the mesh.) If the edge flag is $f_d = 1$ (side v_{bc}), the next vertex forms the triangle (v_b, v_c, v_d); if the edge flag is $f_d = 2$ (side v_{ac}), the next vertex forms the triangle (v_a, v_c, v_d). An edge flag of $f_d = 0$ would start a new triangle, as described above.




FIGURE 4.17 Connecting triangles in a free-form Gouraud-shaded triangle mesh

Complex shapes can be created by using the edge flags to control the edge on which subsequent triangles are formed. Figure 4.18 shows two simple examples. Mesh 1 begins with triangle 1 and uses the following edge flags to draw each succeeding triangle:

- | | |
|-----------------------------|------------------|
| 1 ($f_a = f_b = f_c = 0$) | 7 ($f_i = 2$) |
| 2 ($f_d = 1$) | 8 ($f_j = 2$) |
| 3 ($f_e = 1$) | 9 ($f_k = 2$) |
| 4 ($f_f = 1$) | 10 ($f_l = 1$) |
| 5 ($f_g = 1$) | 11 ($f_m = 1$) |
| 6 ($f_h = 1$) | |

Mesh 2 again begins with triangle 1 and uses the edge flags

- | | |
|-----------------------------|-----------------|
| 1 ($f_a = f_b = f_c = 0$) | 4 ($f_f = 2$) |
| 2 ($f_d = 1$) | 5 ($f_g = 2$) |
| 3 ($f_e = 2$) | 6 ($f_h = 2$) |

The stream must provide vertex data for a whole number of triangles with appropriate edge flags; otherwise, an error will occur.




FIGURE 4.18 Varying the value of the edge flag to create different shapes

The data for each vertex consists of the following items, reading in sequence from higher-order to lower-order bit positions:

- An edge flag, expressed in **BitsPerFlag** bits
- A pair of horizontal and vertical coordinates, expressed in **BitsPerCoordinate** bits each
- A set of n color components (where n is the number of components in the shading's color space), expressed in **BitsPerComponent** bits each, in the order expected by the **sc** operator

Each set of vertex data must occupy a whole number of bytes; if the total number of bits required is not divisible by 8, the last data byte for each vertex is padded at the end with extra bits, which are ignored. The coordinates and color values are decoded according to the **Decode** array in the same way as in an image dictionary (see “Decode Arrays” on page 271).

If the shading dictionary contains a **Function** entry, the color data for each vertex must be specified by a single parametric value t , rather than by n separate color components. All linear interpolation within the triangle mesh is done using the t values; after interpolation, the results are passed to the function(s) specified in the **Function** entry to determine the color at each point.

Type 5 Shadings (Lattice-Form Gouraud-Shaded Triangle Meshes)

Type 5 shadings (lattice-form Gouraud-shaded triangle meshes) are similar to type 4, but instead of using free-form geometry, their vertices are arranged in a *pseudorectangular lattice*, which is topologically equivalent to a rectangular grid. The vertices are organized into rows, which need not be geometrically linear (see Figure 4.19).




FIGURE 4.19 Lattice-form triangle meshes

Table 4.30 shows the shading dictionary entries specific to this type of shading, in addition to those common to all shading dictionaries (Table 4.25 on page 234) and stream dictionaries (Table 3.4 on page 38).

The data stream for a type 5 shading has the same format as for type 4, except that it does not use edge flags to define the geometry of the triangle mesh. The data for each vertex thus consists of the following values, in the order shown:

$$x \ y \ c_1 \dots c_n$$

where

x and y are the vertex's horizontal and vertical coordinates

$c_1 \dots c_n$ are its color components

TABLE 4.30 Additional entries specific to a type 5 shading dictionary

KEY	TYPE	VALUE
BitsPerCoordinate	integer	<i>(Required)</i> The number of bits used to represent each vertex coordinate. Valid values are 1, 2, 4, 8, 12, 16, 24, and 32.
BitsPerComponent	integer	<i>(Required)</i> The number of bits used to represent each color component. Valid values are 1, 2, 4, 8, 12, and 16.
VerticesPerRow	integer	<i>(Required)</i> The number of vertices in each row of the lattice; the value must be greater than or equal to 2. The number of rows need not be specified.
Decode	array	<p><i>(Required)</i> An array of numbers specifying how to map vertex coordinates and color components into the appropriate ranges of values. The decoding method is similar to that used in image dictionaries (see “Decode Arrays” on page 271). The ranges are specified as follows:</p> $[x_{\min} \ x_{\max} \ y_{\min} \ y_{\max} \ c_{1,\min} \ c_{1,\max} \ \cdots \ c_{n,\min} \ c_{n,\max}]$ <p>Note that only one pair of c values should be specified if a Function entry is present.</p>
Function	function	<p><i>(Optional)</i> A 1-in, n-out function or an array of n 1-in, 1-out functions (where n is the number of color components in the shading dictionary’s color space). If this entry is present, the color data for each vertex must be specified by a single parametric variable rather than by n separate color components; the designated function(s) will be called with each interpolated value of the parametric variable to determine the actual color at each point. Each input value will be forced into the range interval specified for the corresponding color component in the shading dictionary’s Decode array. Each function’s domain must be a superset of that interval. If the value returned by the function for a given color component is out of range, it will be adjusted to the nearest valid value.</p> <p>This entry may not be used with an Indexed color space.</p>

All vertex coordinates are expressed in the shading’s target coordinate space. If the shading dictionary includes a **Function** entry, then only a single parametric value, t , is permitted for each vertex in place of the color components $c_1 \dots c_n$.

The **VerticesPerRow** entry in the shading dictionary gives the number of vertices in each row of the lattice. All of the vertices in a row are specified sequentially, followed by those for the next row. Given m rows of k vertices each, the triangles of

the mesh are constructed using the following triplets of vertices, as shown in Figure 4.19:

$$\begin{aligned} & (V_{i,j}, V_{i,j+1}, V_{i+1,j}) && \text{for } 0 \leq i \leq m-2, 0 \leq j \leq k-2 \\ & (V_{i,j+1}, V_{i+1,j}, V_{i+1,j+1}) \end{aligned}$$

See “Type 4 Shadings (Free-Form Gouraud-Shaded Triangle Meshes)” on page 243 for further details on the format of the vertex data.

Type 6 Shadings (Coons Patch Meshes)

Type 6 shadings (Coons patch meshes) are constructed from one or more *color patches*, each bounded by four cubic Bézier curves. Degenerate Bézier curves are allowed and are useful for certain graphical effects. At least one complete patch must be specified.

A Coons patch generally has two independent aspects:

- Colors are specified for each corner of the unit square, and bilinear interpolation is used to fill in colors over the entire unit square (see the upper figure in Plate 15).
- Coordinates are mapped from the unit square into a four-sided patch whose sides are not necessarily linear (see the lower figure in Plate 15). The mapping is continuous: the corners of the unit square map to corners of the patch and the sides of the unit square map to sides of the patch, as shown in Figure 4.20.

The sides of the patch are given by four cubic Bézier curves, C_1 , C_2 , D_1 , and D_2 , defined over a pair of parametric variables u and v that vary horizontally and vertically across the unit square. The four corners of the Coons patch satisfy the equations

$$\begin{aligned} C_1(0) &= D_1(0) \\ C_1(1) &= D_2(0) \\ C_2(0) &= D_1(1) \\ C_2(1) &= D_2(1) \end{aligned}$$




FIGURE 4.20 Coordinate mapping from a unit square to a four-sided Coons patch

Two surfaces can be described that are linear interpolations between the boundary curves. Along the u axis, the surface S_C is defined by

$$S_C(u, v) = (1 - v) \times C_1(u) + v \times C_2(u)$$

Along the v axis, the surface S_D is given by

$$S_D(u, v) = (1 - u) \times D_1(v) + u \times D_2(v)$$

A third surface is the bilinear interpolation of the four corners:

$$S_B(u, v) = (1 - v) \times [(1 - u) \times C_1(0) + u \times C_1(1)] \\ + v \times [(1 - u) \times C_2(0) + u \times C_2(1)]$$

The coordinate mapping for the shading is given by the surface S , defined as

$$S = S_C + S_D - S_B$$

This defines the geometry of each patch. A patch mesh is constructed from a sequence of one or more such colored patches.

Patches can sometimes appear to fold over on themselves—for example, if a boundary curve intersects itself. As the value of parameter u or v increases in parameter space, the location of the corresponding pixels in device space may

change direction, so that new pixels are mapped onto previous pixels already mapped. If more than one point (u, v) in parameter space is mapped to the same point in device space, the point selected will be the one with the largest value of v ; if multiple points have the same v , the one with the largest value of u will be selected. If one patch overlaps another, the patch that appears later in the data stream paints over the earlier one.

Note also that the patch is a control surface, rather than a painting geometry. The outline of a projected square (that is, the painted area) might not be the same as the patch boundary if, for example, the patch folds over on itself, as shown in Figure 4.21.




FIGURE 4.21 *Painted area and boundary of a Coons patch*

Table 4.31 shows the shading dictionary entries specific to this type of shading, in addition to those common to all shading dictionaries (Table 4.25 on page 234) and stream dictionaries (Table 3.4 on page 38).

TABLE 4.31 Additional entries specific to a type 6 shading dictionary

KEY	TYPE	VALUE
BitsPerCoordinate	integer	<i>(Required)</i> The number of bits used to represent each geometric coordinate. Valid values are 1, 2, 4, 8, 12, 16, 24, and 32.
BitsPerComponent	integer	<i>(Required)</i> The number of bits used to represent each color component. Valid values are 1, 2, 4, 8, 12, and 16.
BitsPerFlag	integer	<i>(Required)</i> The number of bits used to represent the edge flag for each patch (see below). Valid values of BitsPerFlag are 2, 4, and 8, but only the least significant 2 bits in each flag value are used. Valid values for the edge flag itself are 0, 1, 2, and 3.
Decode	array	<p><i>(Required)</i> An array of numbers specifying how to map coordinates and color components into the appropriate ranges of values. The decoding method is similar to that used in image dictionaries (see “Decode Arrays” on page 271). The ranges are specified as follows:</p> $[x_{\min} \ x_{\max} \ y_{\min} \ y_{\max} \ c_{1,\min} \ c_{1,\max} \ \cdots \ c_{n,\min} \ c_{n,\max}]$ <p>Note that only one pair of c values should be specified if a Function entry is present.</p>
Function	function	<p><i>(Optional)</i> A 1-in, n-out function or an array of n 1-in, 1-out functions (where n is the number of color components in the shading dictionary’s color space). If this entry is present, the color data for each vertex must be specified by a single parametric variable rather than by n separate color components; the designated function(s) will be called with each interpolated value of the parametric variable to determine the actual color at each point. Each input value will be forced into the range interval specified for the corresponding color component in the shading dictionary’s Decode array. Each function’s domain must be a superset of that interval. If the value returned by the function for a given color component is out of range, it will be adjusted to the nearest valid value.</p> <p>This entry may not be used with an Indexed color space.</p>

The data stream provides a sequence of Bézier control points and color values that define the shape and colors of each patch. All of a patch’s control points are given first, followed by the color values for its corners. Note that this differs from a triangle mesh (shading types 4 and 5), in which the coordinates and color of each vertex are given together. All control point coordinates are expressed in the shading’s target coordinate space.

As in free-form triangle meshes (type 4), each patch has an *edge flag* that tells which edge, if any, it shares with the previous patch. An edge flag of 0 begins a new patch, unconnected to any other. This must be followed by 12 pairs of coordinates, $x_1 y_1 x_2 y_2 \dots x_{12} y_{12}$, which specify the Bézier control points that define the four boundary curves. Figure 4.22 shows how these control points correspond to the cubic Bézier curves C_1, C_2, D_1 , and D_2 identified in Figure 4.20 on page 251. Color values are then given for the four corners of the patch, in the same order as the control points corresponding to the corners. Thus, c_1 is the color at coordinates (x_1, y_1) , c_2 at (x_4, y_4) , c_3 at (x_7, y_7) , and c_4 at (x_{10}, y_{10}) , as shown in the figure.




FIGURE 4.22 Color values and edge flags in Coons patch meshes

Figure 4.22 also shows how nonzero values of the edge flag ($f = 1, 2$, or 3) connect a new patch to one of the edges of the previous patch. In this case, some of the previous patch's control points serve implicitly as control points for the new patch as well (see Figure 4.23), and so are not explicitly repeated in the data stream. Table 4.32 summarizes the required data values for various values of the edge flag.




FIGURE 4.23 Edge connections in a Coons patch mesh

If the shading dictionary contains a **Function** entry, the color data for each corner of a patch must be specified by a single parametric value t , rather than by n separate color components $c_1 \dots c_n$. All linear interpolation within the mesh is done using the t values; after interpolation, the results are passed to the function(s) specified in the **Function** entry to determine the color at each point.

TABLE 4.32 Data values in a Coons patch mesh

EDGE FLAG	NEXT SET OF DATA VALUES
$f = 0$	$x_1 y_1 x_2 y_2 x_3 y_3 x_4 y_4 x_5 y_5 x_6 y_6$ $x_7 y_7 x_8 y_8 x_9 y_9 x_{10} y_{10} x_{11} y_{11} x_{12} y_{12}$ $c_1 c_2 c_3 c_4$ New patch; no implicit values
$f = 1$	$x_5 y_5 x_6 y_6 x_7 y_7 x_8 y_8 x_9 y_9 x_{10} y_{10} x_{11} y_{11} x_{12} y_{12}$ $c_3 c_4$ Implicit values: $(x_1, y_1) = (x_4, y_4)$ previous $c_1 = c_2$ previous $(x_2, y_2) = (x_5, y_5)$ previous $c_2 = c_3$ previous $(x_3, y_3) = (x_6, y_6)$ previous $(x_4, y_4) = (x_7, y_7)$ previous
$f = 2$	$x_5 y_5 x_6 y_6 x_7 y_7 x_8 y_8 x_9 y_9 x_{10} y_{10} x_{11} y_{11} x_{12} y_{12}$ $c_3 c_4$ Implicit values: $(x_1, y_1) = (x_7, y_7)$ previous $c_1 = c_3$ previous $(x_2, y_2) = (x_8, y_8)$ previous $c_2 = c_4$ previous $(x_3, y_3) = (x_9, y_9)$ previous $(x_4, y_4) = (x_{10}, y_{10})$ previous
$f = 3$	$x_5 y_5 x_6 y_6 x_7 y_7 x_8 y_8 x_9 y_9 x_{10} y_{10} x_{11} y_{11} x_{12} y_{12}$ $c_3 c_4$ Implicit values: $(x_1, y_1) = (x_{10}, y_{10})$ previous $c_1 = c_4$ previous $(x_2, y_2) = (x_{11}, y_{11})$ previous $c_2 = c_1$ previous $(x_3, y_3) = (x_{12}, y_{12})$ previous $(x_4, y_4) = (x_1, y_1)$ previous

Type 7 Shadings (Tensor-Product Patch Meshes)

Type 7 shadings (tensor-product patch meshes) are identical to type 6, except that they are based on a bicubic tensor-product patch defined by 16 control points, instead of the 12 control points that define a Coons patch. The shading

dictionaries representing the two patch types differ only in the value of the **ShadingType** entry and in the number of control points specified for each patch in the data stream. Although the Coons patch is more concise and easier to use, the tensor-product patch affords greater control over color mapping.

Note: The data format for type 7 shadings (as for types 4 through 6) is the same in PDF as it is in PostScript. However, the numbering and order of control points was described incorrectly in the first printing of the PostScript Language Reference, Third Edition. That description has been corrected here.

Like the Coons patch mapping, the tensor-product patch mapping is controlled by the location and shape of four cubic Bézier curves marking the boundaries of the patch. However, the tensor-product patch has four additional, “internal” control points to adjust the mapping. The 16 control points can be arranged in a 4-by-4 array indexed by row and column, as follows (see Figure 4.24):

P_{03}	P_{13}	P_{23}	P_{33}
P_{02}	P_{12}	P_{22}	P_{32}
P_{01}	P_{11}	P_{21}	P_{31}
P_{00}	P_{10}	P_{20}	P_{30}




FIGURE 4.24 Control points in a tensor-product patch

As in a Coons patch mesh, the geometry of the tensor-product patch is described by a surface defined over a pair of parametric variables, u and v , which vary horizontally and vertically across the unit square. The surface is defined by the equation

$$S(u, v) = \sum_{i=0}^3 \sum_{j=0}^3 p_{ij} \times B_i(u) \times B_j(v)$$

where p_{ij} is the control point in column i and row j of the tensor, and B_i and B_j are the *Bernstein polynomials*

$$B_0(t) = (1 - t)^3$$

$$B_1(t) = 3t \times (1 - t)^2$$

$$B_2(t) = 3t^2 \times (1 - t)$$

$$B_3(t) = t^3$$

Since each point p_{ij} is actually a pair of coordinates (x_{ij}, y_{ij}) , the surface can also be expressed as

$$x(u, v) = \sum_{i=0}^3 \sum_{j=0}^3 x_{ij} \times B_i(u) \times B_j(v)$$

$$y(u, v) = \sum_{i=0}^3 \sum_{j=0}^3 y_{ij} \times B_i(u) \times B_j(v)$$

The geometry of the tensor-product patch can be visualized in terms of a cubic Bézier curve moving from the bottom boundary of the patch to the top. At the bottom and top, the control points of this curve coincide with those of the patch's bottom ($p_{00} \dots p_{30}$) and top ($p_{03} \dots p_{33}$) boundary curves, respectively. As the curve moves from the bottom edge of the patch to the top, each of its four control points follows a trajectory that is in turn a cubic Bézier curve defined by the four control points in the corresponding column of the array. That is, the starting point of the moving curve follows the trajectory defined by control points $p_{00} \dots p_{03}$, the trajectory of the ending point is defined by points $p_{30} \dots p_{33}$, and those of the two intermediate control points by $p_{10} \dots p_{13}$ and $p_{20} \dots p_{23}$. Equivalently, the patch can be considered to be traced by a cubic Bézier curve moving

from the left edge to the right, with its control points following the trajectories defined by the rows of the coordinate array instead of the columns.

The Coons patch (type 6) is actually a special case of the tensor-product patch (type 7) in which the four internal control points ($p_{11}, p_{12}, p_{21}, p_{22}$) are implicitly defined by the boundary curves. The values of the internal control points are given by the equations

$$p_{11} = S(1/3, 1/3)$$

$$p_{12} = S(1/3, 2/3)$$

$$p_{21} = S(2/3, 1/3)$$

$$p_{22} = S(2/3, 2/3)$$

where S is the Coons surface equation

$$S = S_C + S_D - S_B$$

discussed above under “Type 6 Shadings (Coons Patch Meshes)” on page 250. In the more general tensor-product patch, the values of these four points are unrestricted.

The coordinates of the control points in a tensor-product patch are actually specified in the shading’s data stream in the following order:

4	5	6	7
3	14	15	8
2	13	16	9
1	12	11	10

All control point coordinates are expressed in the shading’s target coordinate space. These are followed by the color values for the four corners of the patch, in the same order as the corners themselves. If the patch’s edge flag f is 0, all 16 control points and four corner colors must be explicitly specified in the data stream; if f is 1, 2, or 3, the control points and colors for the patch’s shared edge are implicitly understood to be the same as those along the specified edge of the previous patch, and are not repeated in the data stream. Table 4.33 summarizes the data values for various values of the edge flag f , expressed in terms of the row and column indices used in Figure 4.24 above.

TABLE 4.33 Data values in a tensor-product patch mesh

EDGE FLAG	NEXT SET OF DATA VALUES
$f = 0$	$x_{00} y_{00} x_{01} y_{01} x_{02} y_{02} x_{03} y_{03} x_{13} y_{13} x_{23} y_{23} x_{33} y_{33} x_{32} y_{32}$ $x_{31} y_{31} x_{30} y_{30} x_{20} y_{20} x_{10} y_{10} x_{11} y_{11} x_{12} y_{12} x_{22} y_{22} x_{21} y_{21}$ $c_{00} c_{03} c_{33} c_{30}$ New patch; no implicit values
$f = 1$	$x_{13} y_{13} x_{23} y_{23} x_{33} y_{33} x_{32} y_{32} x_{31} y_{31} x_{30} y_{30}$ $x_{20} y_{20} x_{10} y_{10} x_{11} y_{11} x_{12} y_{12} x_{22} y_{22} x_{21} y_{21}$ $c_{33} c_{30}$ Implicit values: $(x_{00}, y_{00}) = (x_{03}, y_{03})$ previous $c_{00} = c_{03}$ previous $(x_{01}, y_{01}) = (x_{13}, y_{13})$ previous $c_{03} = c_{33}$ previous $(x_{02}, y_{02}) = (x_{23}, y_{23})$ previous $(x_{03}, y_{03}) = (x_{33}, y_{33})$ previous
$f = 2$	$x_{13} y_{13} x_{23} y_{23} x_{33} y_{33} x_{32} y_{32} x_{31} y_{31} x_{30} y_{30}$ $x_{20} y_{20} x_{10} y_{10} x_{11} y_{11} x_{12} y_{12} x_{22} y_{22} x_{21} y_{21}$ $c_{33} c_{30}$ Implicit values: $(x_{00}, y_{00}) = (x_{33}, y_{33})$ previous $c_{00} = c_{33}$ previous $(x_{01}, y_{01}) = (x_{32}, y_{32})$ previous $c_{03} = c_{30}$ previous $(x_{02}, y_{02}) = (x_{31}, y_{31})$ previous $(x_{03}, y_{03}) = (x_{30}, y_{30})$ previous
$f = 3$	$x_{13} y_{13} x_{23} y_{23} x_{33} y_{33} x_{32} y_{32} x_{31} y_{31} x_{30} y_{30}$ $x_{20} y_{20} x_{10} y_{10} x_{11} y_{11} x_{12} y_{12} x_{22} y_{22} x_{21} y_{21}$ $c_{33} c_{30}$ Implicit values: $(x_{00}, y_{00}) = (x_{30}, y_{30})$ previous $c_{00} = c_{30}$ previous $(x_{01}, y_{01}) = (x_{20}, y_{20})$ previous $c_{03} = c_{00}$ previous $(x_{02}, y_{02}) = (x_{10}, y_{10})$ previous $(x_{03}, y_{03}) = (x_{00}, y_{00})$ previous

4.7 External Objects

An *external object* (commonly called an *XObject*) is a graphics object whose contents are defined by a self-contained content stream, separate from the content stream in which it is used. There are three types of external object:

- An *image XObject* (Section 4.8.4, “Image Dictionaries”) represents a sampled visual image such as a photograph.
- A *form XObject* (Section 4.9, “Form XObjects”) is a self-contained description of an arbitrary sequence of graphics objects.
- A *PostScript XObject* (Section 4.10, “PostScript XObjects”) contains a fragment of code expressed in the PostScript page description language.

Two further categories of external objects, *group XObjects* and *reference XObjects* (both *PDF 1.4*), are actually specialized types of form XObject with additional properties; see Sections 4.9.2, “Group XObjects,” and 4.9.3, “Reference XObjects,” for additional information.

Any XObject can be painted as part of another content stream by means of the **Do** operator (see Table 4.34). This operator applies to any type of XObject—image, form, or PostScript. The syntax is the same in all cases, although details of the operator’s behavior differ depending on the type. (See implementation note 34 in Appendix H.)

TABLE 4.34 XObject operator

OPERANDS	OPERATOR	DESCRIPTION
<i>name</i>	Do	Paint the specified XObject. The operand <i>name</i> must appear as a key in the XObject subdictionary of the current resource dictionary (see Section 3.7.2, “Resource Dictionaries”); the associated value must be a stream whose Type entry, if present, is XObject . The effect of Do depends on the value of the XObject’s Subtype entry, which may be Image (see Section 4.8.4, “Image Dictionaries”), Form (Section 4.9, “Form XObjects”), or PS (Section 4.10, “PostScript XObjects”).

4.8 Images

PDF's painting operators include general facilities for dealing with sampled images. A *sampled image* (or just *image* for short) is a rectangular array of *sample values*, each representing a color. The image may approximate the appearance of some natural scene obtained through an input scanner or a video camera, or it may be generated synthetically.



FIGURE 4.25 Typical sampled image

An image is defined by a sequence of samples obtained by scanning the image array in row or column order. Each sample in the array consists of as many color components as are needed for the color space in which they are specified—for example, one component for **DeviceGray**, three for **DeviceRGB**, four for **DeviceCMYK**, or whatever number is required by a particular **DeviceN** space. Each component is a 1-, 2-, 4-, or 8-bit integer, permitting the representation of 2, 4, 16, or 256 distinct values for each component.

PDF provides two means for specifying images:

- An *image XObject* (described in Section 4.8.4, “Image Dictionaries”) is a stream object whose dictionary specifies attributes of the image and whose data contains the image samples. Like all external objects, it is painted on the page by invoking the **Do** operator in a content stream (see Section 4.7, “External Objects”). Image XObjects have other uses as well, such as for alternate images (see “Alternate Images” on page 273), image masks (Section 4.8.5,

“Masked Images”), and thumbnail images (Section 8.2.3, “Thumbnail Images”).

- An *inline image* is a small image that is completely defined—both attributes and data—directly inline within a content stream. The kinds of image that can be represented in this way are limited; see Section 4.8.6, “Inline Images,” for details.

4.8.1 Image Parameters

The properties of an image—resolution, orientation, scanning order, and so forth—are entirely independent of the characteristics of the raster output device on which the image is to be rendered. A PDF viewer application usually renders images by a sampling technique that attempts to approximate the color values of the source as accurately as possible. The actual accuracy achieved depends on the resolution and other properties of the output device.

To paint an image, four interrelated items must be specified:

- The format of the image: number of columns (width), number of rows (height), number of color components per sample, and number of bits per color component
- The sample data constituting the image’s visual content
- The correspondence between coordinates in user space and those in the image’s own internal coordinate space, defining the region of user space that will receive the image
- The mapping from color component values in the image data to component values in the image’s color space

All of these items are specified explicitly or implicitly by an image XObject or an inline image.

Note: For convenience, the following sections refer consistently to the object defining an image as an image dictionary. Although this term properly refers only to the dictionary portion of the stream object representing an image XObject, it should be understood to apply equally to the stream’s data portion or to the parameters and data of an inline image.

4.8.2 Sample Representation

The source format for an image can be described by four parameters:

- The width of the image in samples
- The height of the image in samples
- The number of color components per sample
- The number of bits per color component

The image dictionary specifies the width, the height, and the number of bits per component explicitly; the number of color components can be inferred from the color space specified in the dictionary.

Sample data is represented as a stream of bytes, interpreted as 8-bit unsigned integers in the range 0 to 255. The bytes constitute a continuous bit stream, with the high-order bit of each byte first. This bit stream, in turn, is divided into units of n bits each, where n is the number of bits per component. Byte boundaries are ignored, except that each row of sample data must begin on a byte boundary. If the number of data bits per row is not a multiple of 8, the end of the row is padded with extra bits to fill out the last byte. A PDF viewer application ignores these padding bits.

Each n -bit unit within the bit stream is interpreted as an unsigned integer in the range 0 to $2^n - 1$, with the high-order bit first; the image dictionary's **Decode** entry maps this to a color component value, equivalent to what could be used with color operators such as **sc** or **g**. Color components are interleaved sample by sample; for example, in a three-component *RGB* image, the red, green, and blue components for one sample are followed by the red, green, and blue components for the next.

Normally, the color samples in an image are interpreted according to the color space specified in the image dictionary (see Section 4.5, “Color Spaces”), without reference to the color parameters in the graphics state. However, if the image dictionary's **ImageMask** entry is **true**, the sample data is interpreted as a *stencil mask* for applying the graphics state's nonstroking color parameters (see “Stencil Masking” on page 276).

4.8.3 Image Coordinate System

Each image has its own internal coordinate system, or *image space*. The image occupies a rectangle in image space w units wide and h units high, where w and h are the width and height of image in samples. Each sample occupies one square unit. The coordinate origin $(0, 0)$ is at the upper-left corner of the image, with coordinates ranging from 0 to w horizontally and 0 to h vertically.

The image's sample data is ordered by row, with the horizontal coordinate varying most rapidly. This is shown in Figure 4.26, where the numbers inside the squares indicate the order of the samples, counting from 0. The upper-left corner of the first sample is at coordinates $(0, 0)$, the second at $(1, 0)$, and so on through the last sample of the first row, whose upper-left corner is at $(w - 1, 0)$ and whose upper-right corner is at $(w, 0)$. The next samples after that are at coordinates $(0, 1)$, $(1, 1)$, and so on, until the final sample of the image, whose upper-left corner is at $(w - 1, h - 1)$ and whose lower-right corner is at (w, h) .




FIGURE 4.26 Source image coordinate system

Note: The image coordinate system and scanning order imposed by PDF do not preclude using different conventions in the actual image. Coordinate transformations can be used to map from other conventions to the PDF convention.

The correspondence between image space and user space is constant: the unit square of user space, bounded by user coordinates $(0, 0)$ and $(1, 1)$, corresponds to the boundary of the image in image space (see Figure 4.27). Following the normal convention for user space, the coordinate $(0, 0)$ is at the *lower-left* corner of this square, corresponding to coordinates $(0, h)$ in image space. The transformation from image space to user space could be described by the matrix $[1/w \ 0 \ 0 \ -1/h \ 0 \ 1]$.




FIGURE 4.27 Mapping the source image

An image can be placed on the output page in any desired position, orientation, and size by using the `cm` operator to modify the current transformation matrix (CTM) so as to map the unit square of user space to the rectangle or parallelogram in which the image is to be painted. Typically, this is done within a pair of `q` and `Q` operators to isolate the effect of the transformation, which can include translation, rotation, reflection, and skew (see Section 4.2, “Coordinate Systems”). For example, if the `XObject` subdictionary of the current resource dictionary defines the name `Image1` to denote an image `XObject`, the code shown in Example 4.23 paints the image in a rectangle whose lower-left corner is at coordinates $(100, 200)$, that is rotated 45 degrees counterclockwise, and that is 150 units wide and 80 units high.

Example 4.23

```

q                                     % Save graphics state
  1 0 0 1 100 200 cm                 % Translate
  0.7071 0.7071 -0.7071 0.7071 0 0 cm % Rotate
  150 0 0 80 0 0 cm                  % Scale
/Image1 Do                             % Paint image
Q                                       % Restore graphics state

```

(As discussed in Section 4.2.3, “Transformation Matrices,” these three transformations could be combined into one.) Of course, if the aspect ratio (width to height) of the original image in this example is different from 150:80, the result will be distorted.

4.8.4 Image Dictionaries

Table 4.35 lists the entries in an image dictionary—that is, in the dictionary portion of a stream representing an image XObject—in addition to the usual entries common to all streams (see Table 3.4 on page 38). There are many relationships among these entries, and the current color space may limit the choices for some of them. Attempting to use an image dictionary whose entries are inconsistent with each other or with the current color space will cause an error.

Note: The entries described here are those that are appropriate for a base image—one that is invoked directly with the **Do** operator. Some of these entries are not relevant for images used in other ways, such as for alternate images (see “Alternate Images” on page 273), image masks (Section 4.8.5, “Masked Images”), or thumbnail images (Section 8.2.3, “Thumbnail Images”). Except as noted, such irrelevant entries are simply ignored.

TABLE 4.35 Additional entries specific to an image dictionary

KEY	TYPE	VALUE
Type	name	<i>(Optional)</i> The type of PDF object that this dictionary describes; if present, must be XObject for an image XObject.
Subtype	name	<i>(Required)</i> The type of XObject that this dictionary describes; must be Image for an image XObject.
Width	integer	<i>(Required)</i> The width of the image, in samples.
Height	integer	<i>(Required)</i> The height of the image, in samples.

ColorSpace	name or array	<i>(Required except for image masks; not allowed for image masks)</i> The color space in which image samples are specified. This may be any type of color space except Pattern .
BitsPerComponent	integer	<i>(Required except for image masks; optional for image masks)</i> The number of bits used to represent each color component. Only a single value may be specified; the number of bits is the same for all color components. Valid values are 1, 2, 4, and 8. If ImageMask is true , this entry is optional, and if specified, its value must be 1. If the image stream uses a filter, the value of BitsPerComponent must be consistent with the size of the data samples that the filter delivers. In particular, a CCITTFaxDecode or JBIG2Decode filter always delivers 1-bit samples, a RunLengthDecode or DCTDecode filter delivers 8-bit samples, and an LZWDecode or FlateDecode filter delivers samples of a specified size if a predictor function is used.
Intent	name	<i>(Optional; PDF 1.1)</i> The name of a color rendering intent to be used in rendering the image (see “Rendering Intents” on page 197). Default value: the current rendering intent in the graphics state.
ImageMask	boolean	<i>(Optional)</i> A flag indicating whether the image is to be treated as an image mask (see Section 4.8.5, “Masked Images”). If this flag is true , the value of BitsPerComponent must be 1 and Mask and ColorSpace should not be specified; unmasked areas will be painted using the current nonstroking color. Default value: false .
Mask	stream or array	<i>(Optional except for image masks; not allowed for image masks; PDF 1.3)</i> An image XObject defining an image mask to be applied to this image (see “Explicit Masking” on page 277), or an array specifying a range of colors to be applied to it as a color key mask (see “Color Key Masking” on page 277). If ImageMask is true , this entry must not be present. (See implementation note 35 in Appendix H.)
SMask	stream	<i>(Optional; PDF 1.4)</i> A subsidiary image XObject defining a <i>soft-mask image</i> (see “Soft-Mask Images” on page 447) to be used as a source of mask shape or mask opacity values in the transparent imaging model. The alpha source parameter in the graphics state determines whether the mask values are interpreted as shape or opacity. If present, this entry overrides the current soft mask in the graphics state, as well as the image’s Mask entry, if any. (However, the other transparency-related graphics state parameters—blend mode and alpha constant—remain in effect.) If SMask is absent, the image has no associated soft mask (although the current soft mask in the graphics state may still apply).

Decode	array	<i>(Optional)</i> An array of numbers describing how to map image samples into the range of values appropriate for the image’s color space (see “Decode Arrays” on page 271). If ImageMask is true , the array must be either [0 1] or [1 0]; otherwise, its length must be twice the number of color components required by ColorSpace . Default value: see “Decode Arrays” on page 271.
Interpolate	boolean	<i>(Optional)</i> A flag indicating whether image interpolation is to be performed (see “Image Interpolation” on page 273). Default value: false .
Alternates	array	<i>(Optional; PDF 1.3)</i> An array of alternate image dictionaries for this image (see “Alternate Images” on page 273). The order of elements within the array has no significance. This entry may not be present in an image XObject that is itself an alternate image.
Name	name	<i>(Required in PDF 1.0; optional otherwise)</i> The name by which this image XObject is referenced in the XObject subdictionary of the current resource dictionary (see Section 3.7.2, “Resource Dictionaries”). <i>Note: This entry is obsolescent and its use is no longer recommended. (See implementation note 36 in Appendix H.)</i>
StructParent	integer	<i>(Required if the image is a structural content item; PDF 1.3)</i> The integer key of the image’s entry in the structural parent tree (see “Finding Structure Elements from Content Items” on page 600).
ID	string	<i>(Optional; PDF 1.3; indirect reference preferred)</i> The digital identifier of the image’s parent Web Capture content set (see Section 9.9.5, “Object Attributes Related to Web Capture”).
OPI	dictionary	<i>(Optional; PDF 1.2)</i> An OPI version dictionary for the image (see Section 9.10.6, “Open Prepress Interface (OPI)”). If ImageMask is true , this entry is ignored.
Metadata	stream	<i>(Optional; PDF 1.4)</i> A <i>metadata stream</i> containing metadata for the image (see Section 9.2.2, “Metadata Streams”).

Example 4.24 defines an image 256 samples wide by 256 high, with 8 bits per sample in the **DeviceGray** color space. It paints the image on a page with its lower-left corner positioned at coordinates (45, 140) in current user space and scaled to a width and height of 132 user space units.

Example 4.24

```

20 0 obj                                % Page object
  << /Type /Page
    /Parent 1 0 R
    /Resources 21 0 R
    /MediaBox [0 0 612 792]
    /Contents 23 0 R
  >>
endobj

21 0 obj                                % Resource dictionary for page
  << /ProcSet [/PDF /ImageB]
    /XObject << /Im1 22 0 R >>
  >>
endobj

22 0 obj                                % Image XObject
  << /Type /XObject
    /Subtype /Image
    /Width 256
    /Height 256
    /ColorSpace /DeviceGray
    /BitsPerComponent 8
    /Length 83183
    /Filter /ASCII85Decode
  >>
stream
9LhZI9h\GY9i+bb;p;e;G9SP92/)X9MJ>^:f14d;;U(X8P;cO;G9e];c$=k9Mn\
... Image data representing 65,536 samples ...
8P;cO;G9e];c$=k9Mn\]~>
endstream
endobj

23 0 obj                                % Contents of page
  << /Length 56 >>
stream
  q                                        % Save graphics state
    132 0 0 132 45 140 cm                % Translate to (45,140) and scale by 132
    /Im1 Do                               % Paint image
  Q                                        % Restore graphics state
endstream
endobj

```

Decode Arrays

An image's data stream is initially decomposed into integers in the domain 0 to $2^n - 1$, where n is the value of the image dictionary's **BitsPerComponent** entry. The image's **Decode** array specifies a linear mapping of each integer component value to a number that would be appropriate as a component value in the image's color space.

Each pair of numbers in a **Decode** array specifies the lower and upper values to which the domain of sample values in the image is mapped. A **Decode** array contains one pair of numbers for each component in the color space specified by the image's **ColorSpace** entry. The mapping for each color component is a linear transformation. That is, it uses the following formula for linear interpolation:

$$\begin{aligned} y &= \text{Interpolate}(x, x_{\min}, x_{\max}, y_{\min}, y_{\max}) \\ &= y_{\min} + \left((x - x_{\min}) \times \frac{y_{\max} - y_{\min}}{x_{\max} - x_{\min}} \right) \end{aligned}$$

Generally, this is used to convert a value x between x_{\min} and x_{\max} to a corresponding value y between y_{\min} and y_{\max} , projecting along the line defined by the points (x_{\min}, y_{\min}) and (x_{\max}, y_{\max}) . While this formula applies to values outside the domain x_{\min} to x_{\max} and does not require that $x_{\min} < x_{\max}$, note that interpolation used for color conversion, such as the **Decode** array, does require that $x_{\min} < x_{\max}$ and “clips” x values to this domain, so that $y = y_{\min}$ for all $x \leq x_{\min}$, and $y = y_{\max}$ for all $x \geq x_{\max}$.

For a **Decode** array of the form $[D_{\min} \ D_{\max}]$, this can be written as

$$\begin{aligned} y &= \text{Interpolate}(x, 0, 2^n - 1, D_{\min}, D_{\max}) \\ &= D_{\min} + \left(x \times \frac{D_{\max} - D_{\min}}{2^n - 1} \right) \end{aligned}$$

where

n is the value of **BitsPerComponent**

x is the input value, in the domain 0 to $2^n - 1$

D_{\min} and D_{\max} are the values specified in the **Decode** array

y is the output value, to be interpreted in the image's color space

Samples with a value of 0 are mapped to D_{\min} , those with a value of $2^n - 1$ are mapped to D_{\max} , and those with intermediate values are mapped linearly between D_{\min} and D_{\max} . Table 4.36 lists the default **Decode** arrays for use with the various color spaces. For most color spaces, the **Decode** arrays listed in the table map into the full range of allowed component values. For an **Indexed** color space, the default **Decode** array ensures that component values that index a color table are passed through unchanged.

TABLE 4.36 Default Decode arrays

COLOR SPACE	Decode ARRAY
DeviceGray	[0.0 1.0]
DeviceRGB	[0.0 1.0 0.0 1.0 0.0 1.0]
DeviceCMYK	[0.0 1.0 0.0 1.0 0.0 1.0 0.0 1.0]
CalGray	[0.0 1.0]
CalRGB	[0.0 1.0 0.0 1.0 0.0 1.0]
Lab	[0 100 a_{\min} a_{\max} b_{\min} b_{\max}] where a_{\min} , a_{\max} , b_{\min} , and b_{\max} correspond to the values in the Range array of the image's color space
ICCBased	Same as the value of Range in the ICC profile of the image's color space
Indexed	[0 N], where $N = 2^n - 1$
Pattern	(Not permitted with images)
Separation	[0.0 1.0]
DeviceN	[0.0 1.0 0.0 1.0 ... 0.0 1.0] (one pair of elements for each color component)

It is possible to specify a mapping that *inverts* sample color intensities, by specifying a D_{\min} value greater than D_{\max} . For example, if the image's color space is **DeviceGray** and the **Decode** array is [1.0 0.0], an input value of 0 will be mapped to 1.0 (white), while an input value of $2^n - 1$ will be mapped to 0.0 (black).

The D_{\min} and D_{\max} parameters for a color component are not required to fall within the range of values allowed for that component. For instance, if an application uses 6-bit numbers as its native image sample format, it can represent

those samples in PDF in 8-bit form, setting the two unused high-order bits of each sample to 0. The image dictionary should then specify a **Decode** array of [0.00000 4.04762], which maps input values from 0 to 63 into the range 0.0 to 1.0 (4.04762 being approximately equal to $255 \div 63$). If an output value falls outside the range allowed for a component, it will automatically be adjusted to the nearest allowed value.

Image Interpolation

When the resolution of a source image is significantly lower than that of the output device, each source sample covers many device pixels. This can cause images to appear “jaggy” or “blocky.” These visual artifacts can be reduced by applying an *image interpolation* algorithm during rendering. Instead of painting all pixels covered by a source sample with the same color, image interpolation attempts to produce a smooth transition between adjacent sample values. Because it may increase the time required to render the image, image interpolation is disabled by default; setting the **Interpolate** entry in the image dictionary to **true** enables it.

***Note:** The interpolation algorithm is implementation-dependent and is not specified by PDF. Image interpolation may not always be performed for some classes of images or on some output devices.*

Alternate Images

Alternate images (PDF 1.3) provide a straightforward and backward-compatible way to include multiple versions of an image in a PDF file for different purposes. These variant representations of the image may differ, for example, in resolution or in color space. The primary goal is to reduce the need to maintain separate versions of a PDF document for low-resolution on-screen viewing and high-resolution printing.

In PDF 1.3, a *base image* (that is, the image XObject referred to in a resource dictionary) can contain an **Alternates** entry. The value of this entry is an array of *alternate image dictionaries* specifying variant representations of the base image. Each alternate image dictionary contains an image XObject for one variant and specifies its properties. Table 4.37 shows the contents of an alternate image dictionary.

TABLE 4.37 Entries in an alternate image dictionary

KEY	TYPE	VALUE
Image	stream	<i>(Required)</i> The image XObject for the alternate image.
DefaultForPrinting	boolean	<i>(Optional)</i> A flag indicating whether this alternate image is the default version to be used for printing. At most one alternate for a given base image may be so designated. If no alternate has this entry set to true , the base image itself is used for printing.

Example 4.25 shows an image with a single alternate. The base image is a gray-scale image, and the alternate is a high-resolution *RGB* image stored on a Web server.

Example 4.25

```

10 0 obj                                % Image XObject
  << /Type /XObject
    /Subtype /Image
    /Width 100
    /Height 200
    /ColorSpace /DeviceGray
    /BitsPerComponent 8
    /Alternates 15 0 R
    /Length 2167
    /Filter /DCTDecode
  >>
stream
... Image data ...
endstream
endobj

15 0 obj                                % Alternate images array
  [ << /Image 16 0 R
    /DefaultForPrinting true
  >>
  ]
endobj

```

```
16 0 obj                                % Alternate image
  << /Type /XObject
    /Subtype /Image
    /Width 1000
    /Height 2000
    /ColorSpace /DeviceRGB
    /BitsPerComponent 8
    /Length 0                            % This is an external stream
    /F << /FS /URL
      /F (http://www.myserver.mycorp.com/images/exttest.jpg)
    >>
  /FFilter /DCTDecode
  >>
stream
endstream
endobj
```

4.8.5 Masked Images

Ordinarily, in the opaque imaging model, images mark all areas they occupy on the page as if with opaque paint. All portions of the image, whether black, white, gray, or color, completely obscure any marks that may previously have existed in the same place on the page. In the graphic arts industry and page layout applications, however, it is common to crop or “mask out” the background of an image and then place the masked image on a different background, allowing the existing background to show through the masked areas. A number of PDF features are available for achieving such masking effects (see implementation note 37 in Appendix H):

- The **ImageMask** entry in the image dictionary, available in all versions of PDF, specifies that the image data is to be used as a *stencil mask* for painting in the current color.
- The **Mask** entry in the image dictionary (*PDF 1.3*) may specify a separate image XObject to be used as an *explicit mask* specifying which areas of the image to paint and which to mask out.
- Alternatively, the **Mask** entry (*PDF 1.3*) may specify a range of colors to be masked out wherever they occur within the image; this technique is known as *color key masking*.

***Note:** Although the **Mask** entry is a PDF 1.3 feature, its effects are commonly simulated in earlier versions of PDF by defining a clipping path enclosing only those of an image's samples that are to be painted. However, implementation limits can cause errors if the clipping path is very complex (or if there is more than one clipping path). An alternative way to achieve the effect of an explicit mask in PDF 1.2 is to define the image being clipped as a pattern, make it the current color, and then paint the explicit mask as an image whose **ImageMask** entry is **true**. In any case, the PDF 1.3 features allow masked images to be placed on the page without regard to the complexity of the clipping path.*

In the transparent imaging model, a fourth type of masking effect, *soft masking*, is available via the **SMask** entry in the image dictionary (PDF 1.4); see Section 7.5.4, “Specifying Soft Masks,” for further discussion.

Stencil Masking

An *image mask* (an image XObject whose **ImageMask** entry is **true**) is a monochrome image, in which each sample is specified by a single bit. However, instead of being painted in opaque black and white, the image mask is treated as a *stencil mask* that is partly opaque and partly transparent. Sample values in the image do not represent black and white pixels; rather, they designate places on the page that should either be marked with the current color or masked out (not marked at all). Areas that are masked out retain their former contents. The effect is like applying paint in the current color through a cut-out stencil, which allows the paint to reach the page in some places and masks it out in others.

An image mask differs from an ordinary image in the following significant ways:

- The image dictionary does not contain a **ColorSpace** entry, because sample values represent masking properties (1 bit per sample) rather than colors.
- The value of the **BitsPerComponent** entry must be 1.
- The **Decode** entry determines how the source samples are to be interpreted. If the **Decode** array is [0 1] (the default for an image mask), a sample value of 0 marks the page with the current color, while a 1 leaves the previous contents unchanged; if the **Decode** array is [1 0], these meanings are reversed.

One of the most important uses of stencil masking is for painting character glyphs represented as bitmaps. Using such a glyph as a stencil mask transfers only its “black” bits to the page, while leaving the “white” bits (which are really just

background) unchanged. For reasons discussed in Section 5.5.4, “Type 3 Fonts,” an image mask, rather than an image, should almost always be used to paint glyph bitmaps.

Note: If image interpolation (see “Image Interpolation” on page 273) is requested during stencil masking, the effect is to smooth the edges of the mask, not to interpolate the painted color values. This can minimize the “jaggy” appearance of a low-resolution stencil mask.

Explicit Masking

In PDF 1.3, the **Mask** entry in an image dictionary may be an image mask, as described above under “Stencil Masking,” which serves as an *explicit mask* for the primary (base) image. The base image and the image mask need not have the same resolution (**Width** and **Height** values), but since all images are defined on the unit square in user space, their boundaries on the page will coincide; that is, they will overlay each other. The image mask indicates which places on the page are to be painted and which are to be masked out (left unchanged). Unmasked areas are painted with the corresponding portions of the base image; masked areas are not.

Color Key Masking

In PDF 1.3, the **Mask** entry in an image dictionary may alternatively be an array specifying a range of colors to be masked out. Samples in the image that fall within this range are not painted, allowing the existing background to show through. The effect is similar to that of the video technique known as *chroma-key*.

For color key masking, the value of the **Mask** entry is an array of $2 \times n$ integers, $[min_1 max_1 \dots min_n max_n]$, where n is the number of color components in the image’s color space. Each integer must be in the range 0 to $2^{\text{BitsPerComponent}} - 1$, representing color values *before* decoding with the **Decode** array. An image sample is masked (not painted) if all of its color components before decoding, $c_1 \dots c_n$, fall within the specified ranges (that is, if $min_i \leq c_i \leq max_i$ for all $1 \leq i \leq n$).

Note: When color key masking is specified, the use of a **DCTDecode** filter for the stream is not recommended. **DCTDecode** is a lossy filter, meaning that the output is only an approximation of the original input data. This can lead to slight changes in the color values of image samples, possibly causing samples that were intended to be masked to be unexpectedly painted instead, in colors slightly different from the mask color.

4.8.6 Inline Images

As an alternative to the image XObjects described in Section 4.8.4, “Image Dictionaries,” a sampled image may be specified in the form of an *inline image*. This type of image is defined directly within the content stream in which it will be painted, rather than as a separate object. Because the inline format gives the viewer application less flexibility in managing the image data, it should be used only for small images (4 KB or less).

An inline image object is delimited in the content stream by the operators **BI** (begin image), **ID** (image data), and **EI** (end image); these operators are summarized in Table 4.38. **BI** and **ID** bracket a series of key-value pairs specifying the characteristics of the image, such as its dimensions and color space; the image data follows between the **ID** and **EI** operators. The format is thus analogous to that of a stream object such as an image XObject:

```
BI
...Key-value pairs...
ID
...Image data...
EI
```

TABLE 4.38 Inline image operators

OPERANDS	OPERATOR	DESCRIPTION
—	BI	Begin an inline image object.
—	ID	Begin the image data for an inline image object.
—	EI	End an inline image object.

Inline image objects may not be nested; that is, two **BI** operators may not appear without an intervening **EI** to close the first object. Similarly, an **ID** operator may appear only between a **BI** and its balancing **EI**. Unless the image uses **ASCIIHexDecode** or **ASCII85Decode** as one of its filters, the **ID** operator should be followed by a single white-space character; the next character after that is interpreted as the first byte of image data.

The key-value pairs appearing between the **BI** and **ID** operators are analogous to those in the dictionary portion of an image XObject (though the syntax is different). Table 4.39 shows the entries that are valid for an inline image, all of which have the same meanings as in a stream dictionary (Table 3.4 on page 38) or an image dictionary (Table 4.35 on page 267). Entries other than those listed will be ignored; in particular, the **Type**, **Subtype**, and **Length** entries normally found in a stream or image dictionary are unnecessary. For convenience, the abbreviations shown in the table may be used in place of the fully spelled-out keys; Table 4.40 shows additional abbreviations that can be used for the names of color spaces and filters. Note, however, that these abbreviations are valid only in inline images; they may *not* be used in image XObjects. Also note that **JBIG2Decode** is not listed in Table 4.40, because that filter can be applied only to image XObjects.

TABLE 4.39 Entries in an inline image object

FULL NAME	ABBREVIATION
BitsPerComponent	BPC
ColorSpace	CS
Decode	D
DecodeParms	DP
Filter	F
Height	H
ImageMask	IM
Intent (<i>PDF 1.1</i>)	No abbreviation
Interpolate	I (uppercase I)
Width	W

TABLE 4.40 Additional abbreviations in an inline image object

FULL NAME	ABBREVIATION
DeviceGray	G
DeviceRGB	RGB
DeviceCMYK	CMYK
Indexed	I (uppercase I)
ASCIHexDecode	AHx
ASCII85Decode	A85
LZWDecode	LZW
FlateDecode (PDF 1.2)	FI (uppercase F, lowercase L)
RunLengthDecode	RL
CCITTFaxDecode	CCF
DCTDecode	DCT

The color space specified by the **ColorSpace** (or **CS**) entry may be any of the standard device color spaces (**DeviceGray**, **DeviceRGB**, or **DeviceCMYK**). It may not be a CIE-based color space or a special color space, with the exception of a limited form of **Indexed** color space whose base color space is a device space and whose color table is specified by a string (see “Indexed Color Spaces” on page 199). Beginning with PDF 1.2, the value of the **ColorSpace** entry may also be the name of a color space in the **ColorSpace** subdictionary of the current resource dictionary (see Section 3.7.2, “Resource Dictionaries”); in this case, the name may designate any color space that can be used with an image XObject.

Note: The names **DeviceGray**, **DeviceRGB**, and **DeviceCMYK** (as well as their abbreviations **G**, **RGB**, and **CMYK**) always identify the corresponding color spaces directly; they never refer to resources in the **ColorSpace** subdictionary.

The image data in an inline image may be encoded using any of the standard PDF filters. The bytes between the **ID** and **EI** operators are treated much the same as a stream object’s data (see Section 3.2.7, “Stream Objects”), even though they do not follow the standard stream syntax. (This is an exception to the usual rule that

the data in a content stream is interpreted according to the standard PDF syntax for objects.)

Example 4.26 shows an inline image 17 samples wide by 17 high with 8 bits per component in the **DeviceRGB** color space. The image has been encoded using LZW and ASCII base-85 encoding. The **cm** operator is used to scale it to a width and height of 17 units in user space and position it at coordinates (298, 388). The **q** and **Q** operators encapsulate the **cm** operation to limit its effect to resizing the image.

Example 4.26

<code>q</code>	% Save graphics state
<code>17 0 0 17 298 388 cm</code>	% Scale and translate coordinate space
<code>BI</code>	% Begin inline image object
<code> /W 17</code>	% Width in samples
<code> /H 17</code>	% Height in samples
<code> /CS /RGB</code>	% Color space
<code> /BPC 8</code>	% Bits per component
<code> /F [/A85 /LZW]</code>	% Filters
<code>ID</code>	% Begin image data
<code>J1/gKA>.]AN&J?]-<HW]aRVcg*bb.\eKAdVV%/PcZ</code>	
<code>... Omitted data...</code>	
<code>R.s(4KE3&d&7hb*7[%Ct2HCqC~></code>	
<code>EI</code>	% End inline image object
<code>Q</code>	% Restore graphics state

4.9 Form XObjects

A *form XObject* is a self-contained description of any sequence of graphics objects (including path objects, text objects, and sampled images), defined as a PDF content stream. It may be painted multiple times—either on several pages or at several locations on the same page—and will produce the same results each time, subject only to the graphics state at the time it is invoked. Not only is this shared definition economical to represent in the PDF file, but under suitable circumstances the PDF viewer application can optimize execution by caching the results of rendering the form XObject for repeated reuse.

Note: *The term form also refers to a completely different kind of object, an interactive form (sometimes called an AcroForm), discussed in Section 8.6, “Interactive Forms.” Whereas the form XObjects described in this section correspond to the notion*

of forms in the PostScript language, interactive forms are the PDF equivalent of the familiar paper instrument. Any unqualified use of the word form is understood to refer to an interactive form; the type of form described here is always referred to explicitly as a form XObject.

Form XObjects have various uses:

- As its name suggests, a form XObject can serve as the template for an entire page. For example, a program that prints filled-in tax forms can first paint the fixed template as a form XObject and then paint the variable information on top of it.
- Any graphical element that is to be used repeatedly, such as a company logo or a standard component in the output from a computer-aided design system, can be defined as a form XObject.
- Certain document elements that are not part of a page's contents, such as annotation appearances (see Section 8.4.4, "Appearance Streams"), are represented as form XObjects.
- A specialized type of form XObject, called a *group XObject (PDF 1.4)*, can be used to group graphical elements together as a unit for various purposes (see Section 4.9.2, "Group XObjects"). In particular, group XObjects are used to define transparency groups and soft masks for use in the transparent imaging model (see "Soft-Mask Dictionaries" on page 445 and Section 7.5.5, "Transparency Group XObjects").
- Another specialized type of form XObject, a *reference XObject (PDF 1.4)*, can be used to import content from one PDF document into another (see Section 4.9.3, "Reference XObjects").

The use of form XObjects requires two steps:

1. *Define the appearance of the form XObject.* A form XObject is a PDF content stream. The dictionary portion of the stream (called the *form dictionary*) contains descriptive information about the form XObject, while the body of the stream describes the graphics objects that produce its appearance. The contents of the form dictionary are described in Section 4.9.1, "Form Dictionaries."
2. *Paint the form XObject.* The **Do** operator (see Section 4.7, "External Objects") paints a form XObject whose name is supplied as an operand. (The name is defined in the **XObject** subdictionary of the current resource dictionary.)

Before invoking this operator, the content stream in which it appears should set appropriate parameters in the graphics state; in particular, it should alter the current transformation matrix to control the position, size, and orientation of the form XObject in user space.

Each form XObject is defined in its own coordinate system, called *form space*. The **BBox** entry in the form dictionary is expressed in form space, as are any coordinates used in the form XObject's content stream, such as path coordinates. The **Matrix** entry in the form dictionary specifies the mapping from form space to the current user space; each time the form XObject is painted by the **Do** operator, this matrix is concatenated with the current transformation matrix to define the mapping from form space to device space. (This differs from the **Matrix** entry in a pattern dictionary, which maps pattern space to the *initial* user space of the content stream in which the pattern is used.)

When the **Do** operator is applied to a form XObject, it does the following:

1. Saves the current graphics state, as if by invoking the **q** operator (see Section 4.3.3, "Graphics State Operators")
2. Concatenates the matrix from the form dictionary's **Matrix** entry with the current transformation matrix (CTM)
3. Clips according to the form dictionary's **BBox** entry
4. Paints the graphics objects specified in the form's content stream
5. Restores the saved graphics state, as if by invoking the **Q** operator (see Section 4.3.3, "Graphics State Operators")

Except as described above, the initial graphics state for the form is inherited from the graphics state that is in effect at the time **Do** is invoked.

4.9.1 Form Dictionaries

Every form XObject has a *form type*, which determines the format and meaning of the entries in its form dictionary. At the time of publication, only one form type, type 1, has been defined. Table 4.41 shows the contents of the form dictionary for this type of form XObject, in addition to the usual entries common to all streams (see Table 3.4 on page 38).

TABLE 4.41 Additional entries specific to a type 1 form dictionary

KEY	TYPE	VALUE
Type	name	<i>(Optional)</i> The type of PDF object that this dictionary describes; if present, must be XObject for a form XObject.
Subtype	name	<i>(Required)</i> The type of XObject that this dictionary describes; must be Form for a form XObject.
FormType	integer	<i>(Optional)</i> A code identifying the type of form XObject that this dictionary describes. The only valid value defined at the time of publication is 1. Default value: 1.
Name	name	<i>(Required in PDF 1.0; optional otherwise)</i> The name by which this form XObject is referenced in the XObject subdictionary of the current resource dictionary (see Section 3.7.2, “Resource Dictionaries”). <i>Note: This entry is obsolescent and its use is no longer recommended. (See implementation note 38 in Appendix H.)</i>
LastModified	date	<i>(Required if PieceInfo is present; optional otherwise; PDF 1.3)</i> The date and time (see Section 3.8.2, “Dates”) when the form XObject’s contents were most recently modified. If a page-piece dictionary (PieceInfo) is present, the modification date is used to ascertain which of the application data dictionaries it contains correspond to the current content of the form (see Section 9.4, “Page-Piece Dictionaries”).
BBox	rectangle	<i>(Required)</i> An array of four numbers in the form coordinate system (see below), giving the coordinates of the left, bottom, right, and top edges, respectively, of the form XObject’s bounding box. These boundaries are used to clip the form XObject and to determine its size for caching.
Matrix	array	<i>(Optional)</i> An array of six numbers specifying the <i>form matrix</i> , which maps form space into user space (see Section 4.2.3, “Transformation Matrices”). Default value: the identity matrix [1 0 0 1 0 0].
Resources	dictionary	<i>(Optional but strongly recommended; PDF 1.2)</i> A dictionary specifying any resources (such as fonts and images) required by the form XObject (see Section 3.7, “Content Streams and Resources”).

In PDF 1.1 and earlier, all named resources used in the form XObject must be included in the resource dictionary of each page object on which the form XObject appears, whether or not they also appear in the resource dictionary of the form XObject itself. It can be useful to specify these resources in the form XObject’s own resource dictionary as well, in order to determine which resources are used inside the form XObject. If a resource is included in both dictionaries, it should have the same name in both locations.

In PDF 1.2 and later versions, form XObjects can be independent of the content streams in which they appear, and this is strongly recommended although not required. In an independent form XObject, the resource dictionary of the form XObject is required and contains all named resources used by the form XObject. These resources are not “promoted” to the outer content stream’s resource dictionary, although that stream’s resource dictionary will refer to the form XObject itself.

Group	dictionary	<p>(Optional; PDF 1.4) A <i>group attributes dictionary</i> indicating that the contents of the form XObject are to be treated as a group and specifying the attributes of that group (see Section 4.9.2, “Group XObjects”).</p> <p>Note: If a Ref entry (see below) is present, the group attributes also apply to the external page imported by that entry. This allows such an imported page to be treated as a group without further modification.</p>
Ref	dictionary	<p>(Optional; PDF 1.4) A reference dictionary identifying a page to be imported from another PDF file, and for which the form XObject serves as a proxy (see Section 4.9.3, “Reference XObjects”).</p>
Metadata	stream	<p>(Optional; PDF 1.4) A <i>metadata stream</i> containing metadata for the form XObject (see Section 9.2.2, “Metadata Streams”).</p>
PieceInfo	dictionary	<p>(Optional; PDF 1.3) A page-piece dictionary associated with the form XObject (see Section 9.4, “Page-Piece Dictionaries”).</p>
StructParent	integer	<p>(Required if the form XObject is a structural content item; PDF 1.3) The integer key of the form XObject’s entry in the structural parent tree (see “Finding Structure Elements from Content Items” on page 600).</p>
StructParents	integer	<p>(Required if the form XObject contains marked-content sequences that are structural content items; PDF 1.3) The integer key of the form XObject’s entry in the structural parent tree (see “Finding Structure Elements from Content Items” on page 600).</p> <p>Note: At most one of the entries StructParent or StructParents may be present. A form XObject can be either a content item in its entirety or a container for marked-content sequences that are content items, but not both.</p>
OPI	dictionary	<p>(Optional; PDF 1.2) An OPI version dictionary for the form XObject (see Section 9.10.6, “Open Prepress Interface (OPI)”).</p>

Example 4.27 shows a simple form XObject that paints a filled square 1000 units on each side.

Example 4.27

```
6 0 obj                                % Form XObject
  << /Type /XObject
    /Subtype /Form
    /FormType 1
    /BBox [0 0 1000 1000]
    /Matrix [1 0 0 1 0 0]
    /Resources << /ProcSet [/PDF] >>
    /Length 58
  >>
stream
  0 0 m
  0 1000 l
  1000 1000 l
  1000 0 l
  f
endstream
endobj
```

4.9.2 Group XObjects

A *group XObject* (PDF 1.4) is a special type of form XObject that can be used to group graphical elements together as a unit for various purposes. It is distinguished by the presence of the optional **Group** entry in the form dictionary (see Section 4.9.1, “Form Dictionaries”). The value of this entry is a subsidiary *group attributes dictionary* describing the properties of the group.

As shown in Table 4.42, every group XObject has a *group subtype* (specified by the **S** entry in the group attributes dictionary) that determines the format and meaning of the dictionary’s remaining entries. Only one such subtype is defined in PDF 1.4, a *transparency group XObject* (subtype **Transparency**) representing a transparency group for use in the transparent imaging model (see Section 7.3, “Transparency Groups”). The remaining contents of this type of dictionary are described in Section 7.5.5, “Transparency Group XObjects.”

TABLE 4.42 Entries common to all group attributes dictionaries

KEY	TYPE	VALUE
Type	name	(<i>Optional</i>) The type of PDF object that this dictionary describes; if present, must be Group for a group attributes dictionary.
S	name	(<i>Required</i>) The <i>group subtype</i> , which identifies the type of group whose attributes this dictionary describes and determines the format and meaning of the dictionary's remaining entries. The only group subtype defined in PDF 1.4 is Transparency ; see Section 7.5.5, "Transparency Group XObjects," for the remaining contents of this type of dictionary. Other group subtypes may be added in the future.

4.9.3 Reference XObjects

Reference XObjects (PDF 1.4) enable one PDF document to import content from another. The document in which the reference occurs is called the *containing document*; the one whose content is being imported is the *target document*. The target document may reside in a file external to the containing document or may be included within it as an embedded file stream (see Section 3.10.3, "Embedded File Streams").

The reference XObject in the containing document is a form XObject containing the optional **Ref** entry in its form dictionary, as described below. This object serves as a *proxy* that can be displayed or printed in place of the imported content. The proxy might consist of a low-resolution image of the imported content, a piece of descriptive text referring to it, a gray box to be displayed in its place, or any other similar placeholder. Viewer applications that do not recognize the **Ref** entry will simply display or print the proxy as an ordinary form XObject; those that do implement reference XObjects can use the proxy in place of the imported content if the latter is unavailable. A viewer application may also provide a user interface to allow editing and updating of imported content links.

The imported content consists of a single, complete PDF page in the target document. It is designated by a *reference dictionary*, which in turn is the value of the **Ref** entry in the reference XObject's form dictionary (see Section 4.9.1, "Form Dictionaries"). It is the presence of the **Ref** entry that distinguishes reference XObjects from other types of form XObject. Table 4.43 shows the contents of the reference dictionary.

TABLE 4.43 Entries in a reference dictionary

KEY	TYPE	VALUE
F	file specification	<i>(Required)</i> The file containing the target document.
Page	integer or text string	<i>(Required)</i> A page index or page label (see Section 8.3.1, “Page Labels”) identifying the page of the target document containing the content to be imported. Note that the reference is a weak one and can be inadvertently invalidated if the referenced page is changed or replaced in the target document after the reference is created.
ID	array	<i>(Optional)</i> An array of two strings constituting a file identifier (see Section 9.3, “File Identifiers”) for the file containing the target document. The use of this entry improves a viewer application’s chances of finding the intended file and allows it to warn the user if the file has changed since the reference was created.

When the imported content replaces the proxy, it is transformed according to the proxy object’s transformation matrix and clipped to the boundaries of its bounding box, as specified by the **Matrix** and **BBox** entries in the proxy’s form dictionary (see Section 4.9.1, “Form Dictionaries”). The combination of the proxy object’s matrix and bounding box thus implicitly defines the bounding box of the imported page. This bounding box will typically coincide with the imported page’s crop box or art box (see Section 9.10.1, “Page Boundaries”), but it is not required to correspond to any of the defined page boundaries. If the proxy object’s form dictionary contains a **Group** entry, the specified group attributes apply to the imported page as well; this allows the imported page to be treated as a group without further modification.

Printing Reference XObjects

When printing a page containing reference XObjects, a viewer application may emit any of the following, depending on the capabilities of the viewer application, the user’s preferences, and the nature of the print job:

- The imported content designated by the reference XObject
- The reference XObject itself, as a proxy for the imported content
- An OPI proxy or substitute image taken from the reference XObject’s OPI dictionary, if any (see Section 9.10.6, “Open Prepress Interface (OPI)”)

The imported content or the reference XObject itself may also be emitted in place of an OPI proxy when generating OPI comments in a PostScript output stream.

Special Considerations

Certain special considerations arise when reference XObjects interact with other PDF features:

- When the page imported by a reference XObject contains annotations (see Section 8.4, “Annotations”), all annotations that contain a printable, unhidden, visible appearance stream (Section 8.4.4, “Appearance Streams”) must be included in the rendering of the imported page. If the proxy is a snapshot image of the imported page, it must also include the annotation appearances. These appearances must therefore be converted into part of the proxy’s content stream, either as subsidiary form XObjects or by “flattening” them directly into the content stream.
- Logical structure information associated with a page (see Section 9.6, “Logical Structure”) should normally be ignored when importing the page into another document with a reference XObject. In a target document with multiple pages, structure elements occurring on the imported page will typically be part of a larger structure pertaining to the document as a whole; such elements cannot meaningfully be incorporated into the structure of the containing document. In a one-page target document or one made up of independent, structurally unrelated pages, the logical structure for the imported page may be wholly self-contained; in this case, it may be possible to incorporate this structure information into that of the containing document. However, PDF provides no mechanism for the logical structure hierarchy of one document to refer indirectly to that of another.

4.10 PostScript XObjects

In PDF 1.1, a content stream can include PostScript language fragments. These fragments are used only when printing to a PostScript output device; they have no effect either when viewing the document on-screen or when printing to a non-PostScript device. In addition, applications that understand PDF are unlikely to be able to interpret the PostScript fragments. Hence, this capability should be used with extreme caution and only if there is no other way to achieve the same result. Inappropriate use of PostScript XObjects can cause PDF files to print incorrectly.

Note: Since PDF 1.4 encompasses all of the Adobe imaging model features of the PostScript language, there is no longer any reason to use PostScript XObjects. This feature is likely to be removed from PDF in a future version.

A PostScript XObject is an XObject stream whose **Subtype** entry has the value **PS**. Table 4.44 shows the contents of a PostScript XObject dictionary, in addition to the usual entries common to all streams (see Table 3.4 on page 38).

TABLE 4.44 Additional entries specific to a PostScript XObject dictionary

KEY	TYPE	VALUE
Type	name	<i>(Optional)</i> The type of PDF object that this dictionary describes; if present, must be XObject for a PostScript XObject.
Subtype	name	<i>(Required)</i> The type of XObject that this dictionary describes; must be PS for a PostScript XObject.
Level1	stream	<i>(Optional)</i> A stream whose contents are to be used in place of the PostScript XObject's stream when the target PostScript interpreter is known to support only LanguageLevel 1.

When a PDF content stream is translated into the PostScript language, any **Do** operation that references a PostScript XObject is replaced by the contents of the XObject stream itself. The stream is copied without interpretation. The PostScript fragment may use Type 1 and TrueType fonts listed in the **Font** subdictionary of the current resource dictionary (see Section 3.7.2, “Resource Dictionaries”), accessing them by their **BaseFont** names using the PostScript **findfont** operator. The fragment may not use other types of font listed in the **Font** subdictionary. It should not reference the PostScript definitions corresponding to PDF procedure sets (see Section 9.1, “Procedure Sets”), which are subject to change.

CHAPTER 5

Text

THIS CHAPTER DESCRIBES the special facilities in PDF for dealing with text—specifically, for representing characters with *glyphs* from *fonts*. A glyph is a graphical shape and is subject to all graphical manipulations, such as coordinate transformation. Because of the importance of text in most page descriptions, PDF provides higher-level facilities that permit an application to describe, select, and render glyphs conveniently and efficiently.

The first section is a general description of how glyphs from fonts are painted on the page. Subsequent sections cover the following topics in detail:

- *Text state.* A subset of the graphics state parameters pertain to text, including parameters that select the font, scale the glyphs to an appropriate size, and accomplish other graphical effects.
- *Text objects and operators.* The text operators specify the glyphs to be painted, represented by string objects whose values are interpreted as sequences of character codes. A text object encloses a sequence of text operators and associated parameters.
- *Font data structures.* Font dictionaries and associated data structures provide information that a viewer application needs to interpret the text and position the glyphs properly. The definitions of the glyphs themselves are contained in *font programs*, which may be embedded in the PDF file, built into the viewer application, or obtained from an external font file.

5.1 Organization and Use of Fonts

A *character* is an abstract symbol, whereas a *glyph* is a specific graphical rendering of a character. For example, the glyphs A, **A**, and *A* are renderings of the abstract “A” character. Historically these two terms have often been used interchangeably in computer typography (as evidenced by the names chosen for some PDF dictionary keys and PostScript operators), but advances in this area have made the distinction more meaningful in recent times. Consequently, this book distinguishes between characters and glyphs, though with some residual names that are inconsistent.

Glyphs are organized into *fonts*. A font defines glyphs for a particular character set; for example, the Helvetica and Times fonts define glyphs for a set of standard Latin characters. A font for use with a PDF viewer application is prepared in the form of a program. Such a *font program* is written in a special-purpose language, such as the *Type 1* or *TrueType* font format, that is understood by a specialized font interpreter.

In PDF, the term *font* refers to a *font dictionary*, a PDF object that identifies the font program and contains additional information about it. There are several different font types, identified by the **Subtype** entry of the font dictionary.

For most font types, the font program itself is defined in a separate *font file*, which may be either embedded in a PDF stream object or obtained from an external source. The font program contains *glyph descriptions* that generate glyphs.

A content stream paints glyphs on the page by specifying a font dictionary and a string object that is interpreted as a sequence of one or more character codes identifying glyphs in the font. This operation is called *showing* the text string. The glyph description consists of a sequence of graphics operators that produce the specific shape for that character in this font. To render a glyph, the viewer application executes the glyph description.

Programmers who have experience with scan conversion of general shapes may be concerned about the amount of computation that this description seems to imply. However, this is only the abstract behavior of glyph descriptions and font programs, not how they are implemented. In fact, an efficient implementation can be achieved through careful caching and reuse of previously rendered glyphs.

5.1.1 Basics of Showing Text

Example 5.1 illustrates the most straightforward use of a font. It places the text ABC 10 inches from the bottom of the page and 4 inches from the left edge, using 12-point Helvetica.

Example 5.1

```
BT
  /F13 12 Tf
  288 720 Td
  (ABC) Tj
ET
```

The five lines of this example perform the following steps:

1. Begin a text object.
2. Set the font and font size to use, installing them as parameters in the text state. (The font resource identified by the name F13 specifies the font externally known as Helvetica.)
3. Specify a starting position on the page, setting parameters in the text object.
4. Paint the glyphs for a string of characters there.
5. End the text object.

The following paragraphs explain these operations in more detail.

To paint glyphs, a content stream must first identify the font to be used. The **Tf** operator specifies the name of a font resource—that is, an entry in the **Font** subdictionary of the current resource dictionary. The value of that entry is a font dictionary. The font dictionary in turn identifies the font’s externally known name, such as Helvetica, and supplies some additional information that the viewer application needs to paint glyphs from that font; it optionally provides the definition of the font program itself.

***Note:** The font resource name presented to the **Tf** operator is arbitrary, as are the names for all kinds of resources. It bears no relationship to an actual font name, such as Helvetica.*

Example 5.2 illustrates an excerpt from the current page’s resource dictionary, defining the font dictionary that is referenced as F13 in Example 5.1.

Example 5.2

```
/Resources
  << /Font << /F13 23 0 R >>
  >>
23 0 obj
  << /Type /Font
      /Subtype /Type1
      /BaseFont /Helvetica
  >>
endobj
```

A font defines the glyphs for one standard size. This standard is so arranged that the nominal height of tightly spaced lines of text is 1 unit. In the default user coordinate system, this means the standard glyph size is 1 unit in user space, or 1/72 inch. The standard-size font must then be scaled to be usable. The scale factor is specified as the second operand of the **Tf** operator, thereby setting the *text font size* parameter in the graphics state. Example 5.1 establishes the Helvetica font with a 12-unit size in the graphics state.

Once the font has been selected and scaled, it can be used to paint glyphs. The **Td** operator adjusts the current text position (actually, the translation components of the text matrix, as described in Section 5.3.1, “Text-Positioning Operators”). When executed for the first time after **BT**, it establishes the text position in the current user coordinate system. This determines the position on the page at which to begin painting glyphs.

The **Tj** operator takes a string operand and paints the corresponding glyphs using the current font and other text-related parameters in the graphics state. In Example 5.1, the **Tj** operator treats each element of the string (an integer in the range 0 to 255) as a character code. Each code selects a glyph description in the font, and the glyph description is executed to paint that glyph on the page. This is the behavior of **Tj** for simple fonts, such as ordinary Latin text fonts; interpretation of the string as a sequence of character codes is more complex for composite fonts, described in Section 5.6, “Composite Fonts.”

Note: What these steps produce on the page is not a 12-point glyph, but rather a 12-unit glyph, where the unit size is that of the text space at the time the glyphs are rendered on the page. The actual size of the glyph is determined by the text matrix (T_m) in the text object, several text state parameters, and the current transformation

matrix (CTM) in the graphics state; see Section 5.3.3, “Text Space Details.” If the text space is later scaled to make the unit size 1 centimeter, painting glyphs from the same 12-unit font will generate results that are 12 centimeters high.

5.1.2 Achieving Special Graphical Effects

Normal uses of **Tj** and other glyph-painting operators cause black-filled glyphs to be painted. Other effects can be obtained by combining font operators with general graphics operators.

The color used for painting glyphs is the current color in the graphics state: either the nonstroking or the stroking color (or both), depending on the text rendering mode (see Section 5.2.5, “Text Rendering Mode”). The default color is black, but other colors can be obtained by executing an appropriate color-setting operator or operators (see Section 4.5.7, “Color Operators”) before painting the glyphs. Example 5.3 uses text rendering mode 0 and the **g** operator to fill glyphs in 50 percent gray, as shown in Figure 5.1.

Example 5.3

```
BT
  /F13 48 Tf
  20 40 Td
  0 Tr
  0.5 g
  (ABC) Tj
ET
```



FIGURE 5.1 *Glyphs painted in 50% gray*

Other graphical effects can be achieved by treating the glyph outline as a path instead of filling it. The *text rendering mode* parameter in the graphics state specifies whether glyph outlines are to be filled, stroked, used as a clipping boundary, or some combination of these effects. (This parameter does not apply to Type 3 fonts.)

Example 5.4 treats glyph outlines as a path to be stroked. The **Tr** operator sets the text rendering mode to 1 (stroke). The **w** operator sets the line width to 2 units in user space. Given those graphics state parameters, the **Tj** operator strokes the glyph outlines with a line 2 points thick (see Figure 5.2).

Example 5.4

```
BT
  /F13 48 Tf
  20 38 Td
  1 Tr
  2 w
  (ABC) Tj
ET
```




FIGURE 5.2 *Glyph outlines treated as a stroked path*

Example 5.5 treats the glyphs' outlines as a clipping boundary. The **Tr** operator sets the text rendering mode to 7 (clip), causing the subsequent **Tj** operator to impose the glyph outlines as the current clipping path. All subsequent painting operations will mark the page only within this path, as illustrated in Figure 5.3. This state persists until some earlier clipping path is reinstated by the **Q** operator.

Example 5.5

```
BT
  /F13 48 Tf
  20 38 Td
  7 Tr
  (ABC) Tj
ET
...Graphics operators to draw a starburst...
```




FIGURE 5.3 Graphics clipped by a glyph path

5.1.3 Glyph Positioning and Metrics

A glyph's *width*—formally, its *horizontal displacement*—is the amount of space it occupies along the baseline of a line of text that is written horizontally. In other words, it is the distance the current text position moves (by translating text space) when the glyph is painted. Note that the width is distinct from the dimensions of the glyph outline.

In some fonts, the width is constant; it does not vary from glyph to glyph. Such fonts are called *fixed-pitch* or *monospaced*. They are used mainly for typewriter-style printing. However, most fonts used for high-quality typography associate a different width with each glyph. Such fonts are called *proportional* or *variable-pitch* fonts. In either case, the **Tj** operator positions the glyphs for consecutive characters of a string according to their widths.

The width information for each glyph is stored both in the font dictionary and in the font program itself. (The two sets of widths must be identical; storing this information in the font dictionary, although redundant, enables a viewer applica-

tion to determine glyph positioning without having to look inside the font program.) The operators for showing text are designed on the assumption that glyphs are ordinarily positioned according to their standard widths. However, means are provided to vary the positioning in certain limited ways. For example, the **TJ** operator enables the text position to be adjusted between any consecutive pair of glyphs corresponding to characters in a text string. There are graphics state parameters to adjust character and word spacing systematically.

In addition to width, a glyph has several other metrics that influence glyph positioning and painting. For most font types, this information is largely internal to the font program and is not specified explicitly in the PDF font dictionary; however, in a Type 3 font, all metrics are specified explicitly (see Section 5.5.4, “Type 3 Fonts”).

The *glyph coordinate system* is the space in which an individual character’s glyph is defined. All path coordinates and metrics are interpreted in glyph space. For all font types except Type 3, the units of glyph space are one-thousandth of a unit of text space; for a Type 3 font, the transformation from glyph space to text space is defined by a *font matrix* specified in an explicit **FontMatrix** entry in the font. Figure 5.4 shows a typical glyph outline and its metrics.




FIGURE 5.4 *Glyph metrics*

The *glyph origin* is the point (0, 0) in the glyph coordinate system. **Tj** and other text-showing operators position the origin of the first glyph to be painted at the origin of text space. For example, the following code adjusts the origin of text

space to (40, 50) in the user coordinate system, and then places the origin of the A glyph at that point:

```
BT
  40 50 Td
  (ABC) Tj
ET
```

The *glyph displacement* is the distance from the glyph's origin to the point at which the origin of the *next* glyph should normally be placed when painting the consecutive glyphs of a line of text. This distance is a vector (called the *displacement vector*) in the glyph coordinate system; it has horizontal and vertical components. (A displacement that is horizontal is usually called a *width*.) Most Western writing systems, including those based on the Latin alphabet, have a positive horizontal displacement and a zero vertical displacement; some Asian writing systems have a nonzero vertical displacement. In all cases, the text-showing operators transform the displacement vector into text space and then translate text space by that amount.

The *glyph bounding box* is the smallest rectangle (oriented with the axes of the glyph coordinate system) that will just enclose the entire glyph shape. The bounding box is expressed in terms of its left, bottom, right, and top coordinates relative to the glyph origin in the glyph coordinate system.

In some writing systems, text is frequently aligned in two different directions. For example, it is common to write Japanese and Chinese glyphs either horizontally or vertically. To handle this, a font can optionally contain a second set of metrics for each glyph. Which set of metrics to use is selected according to a *writing mode*, where 0 specifies horizontal writing and 1 specifies vertical writing. This feature is available only for composite fonts, discussed in Section 5.6, “Composite Fonts.”

When a glyph has two sets of metrics, each set specifies a glyph origin and a displacement vector for that writing mode. In vertical writing, the glyph position is described by a *position vector* from the origin used for horizontal writing (origin 0) to the origin used for vertical writing (origin 1). Figure 5.5 illustrates the metrics for the two writing modes, as follows:

- The left diagram illustrates the glyph metrics associated with writing mode 0, horizontal writing. The coordinates *ll* and *ur* specify the bounding box of the glyph relative to origin 0. *w0* is the displacement vector that specifies how the

text position is changed after the glyph is painted in writing mode 0; its vertical component is always 0.

- The center diagram illustrates writing mode 1, vertical writing. $w1$ is the displacement vector for writing mode 1; its horizontal component is always 0.
- In the right diagram, v is a position vector defining the position of origin 1 relative to origin 0.




FIGURE 5.5 Metrics for horizontal and vertical writing modes

Glyph metric information is also available separately in the form of Adobe font metrics (AFM) and Adobe composite font metrics (ACFM) files. These files are for use by application programs that generate PDF page descriptions and must make formatting decisions based on the widths and other metrics of glyphs. Also available in the AFM and ACFM files is kerning information, which allows an application generating a PDF file to determine spacing adjustments between characters depending on context. Specifications for the AFM and ACFM file formats are available in Adobe Technical Note #5004, *Adobe Font Metrics File Format Specification*; the files themselves can be obtained from the ASN Developer Program Web site (see the Bibliography).

5.2 Text State Parameters and Operators

The *text state* comprises those graphics state parameters that only affect text. There are nine parameters in the text state (see Table 5.1).

TABLE 5.1 Text state parameters

PARAMETER	DESCRIPTION
T_c	Character spacing
T_w	Word spacing
T_h	Horizontal scaling
T_l	Leading
T_f	Text font
T_{fs}	Text font size
T_{mode}	Text rendering mode
T_{rise}	Text rise
T_k	Text knockout

Except for the self-explanatory T_f and T_{fs} , these parameters are discussed further in the following sections. (As described in Section 5.3, “Text Objects,” three additional text-related parameters are defined only within a text object: T_m , the text matrix; T_{lm} , the text line matrix; and T_{rm} , the text rendering matrix.) The values of the text state parameters are consulted when text is positioned and shown (using the operators described in Sections 5.3.1, “Text-Positioning Operators,” and 5.3.2, “Text-Showing Operators”). In particular, the spacing and scaling parameters participate in a computation described in Section 5.3.3, “Text Space Details.” The text state parameters can be set using the operators listed in Table 5.2.

Note: The text knockout parameter, T_k , is set via the **TK** entry in a graphics state parameter dictionary, using the **gs** operator (see Section 4.3.4, “Graphics State Parameter Dictionaries”). There is no specific operator for setting this parameter.

The text state operators can appear outside text objects, and the values they set are retained across text objects in a single content stream. Like other graphics state parameters, these parameters are initialized to their default values at the beginning of each page.

TABLE 5.2 Text state operators

OPERANDS	OPERATOR	DESCRIPTION
<i>charSpace</i>	Tc	Set the character spacing, T_c , to <i>charSpace</i> , which is a number expressed in unscaled text space units. Character spacing is used by the Tj , TJ , and ' operators. Initial value: 0.
<i>wordSpace</i>	Tw	Set the word spacing, T_w , to <i>wordSpace</i> , which is a number expressed in unscaled text space units. Word spacing is used by the Tj , TJ , and ' operators. Initial value: 0.
<i>scale</i>	Tz	Set the horizontal scaling, T_h , to (<i>scale</i> ÷ 100). <i>scale</i> is a number specifying the percentage of the normal width. Initial value: 100 (normal width).
<i>leading</i>	TL	Set the text leading, T_l , to <i>leading</i> , which is a number expressed in unscaled text space units. Text leading is used only by the T* , ', and " operators. Initial value: 0.
<i>font size</i>	Tf	Set the text font, T_f , to <i>font</i> and the text font size, T_{fs} , to <i>size</i> . <i>font</i> is the name of a font resource in the Font subdictionary of the current resource dictionary; <i>size</i> is a number representing a scale factor. There is no initial value for either <i>font</i> or <i>size</i> ; they must be specified explicitly using Tf before any text is shown.
<i>render</i>	Tr	Set the text rendering mode, T_{mode} , to <i>render</i> , which is an integer. Initial value: 0.
<i>rise</i>	Ts	Set the text rise, T_{rise} , to <i>rise</i> , which is a number expressed in unscaled text space units. Initial value: 0.

Note that some of these parameters are expressed in *unscaled* text space units. This means that they are specified in a coordinate system that is defined by the text matrix, T_m , but is not scaled by the font size parameter, T_{fs} .

5.2.1 Character Spacing

The character spacing parameter, T_c , is a number specified in unscaled text space units (although it is subject to scaling by the T_h parameter if the writing mode is horizontal). When the glyph for each character in the string is rendered, T_c is *added* to the horizontal or vertical component of the glyph's displacement, depending on the writing mode. (See Section 5.1.3, "Glyph Positioning and Metrics," for a discussion of glyph displacements.) In the default coordinate system, horizontal coordinates increase from left to right and vertical coordinates from bottom to top. So for horizontal writing, a positive value of T_c has the effect

of expanding the distance between glyphs (see Figure 5.6), whereas for vertical writing, a *negative* value of T_c has this effect.

$T_c = 0$ (default)	Character
$T_c = 0.25$	Character

FIGURE 5.6 Character spacing in horizontal writing

5.2.2 Word Spacing

Word spacing works the same way as character spacing, but applies only to the space character, code 32. The word spacing parameter, T_w , is added to the glyph's horizontal or vertical displacement (depending on the writing mode). For horizontal writing, a positive value for T_w has the effect of increasing the spacing between words. For vertical writing, a positive value for T_w *decreases* the spacing between words (and a negative value increases it), since vertical coordinates increase from bottom to top. Figure 5.7 illustrates the effect of word spacing in horizontal writing.

$T_w = 0$ (default)	Word Space
$T_w = 2.5$	Word Space

FIGURE 5.7 Word spacing in horizontal writing

Note: Word spacing is applied to every occurrence of the single-byte character code 32 in a string. This can occur when using a simple font or a composite font that defines code 32 as a single-byte code. It does not apply to occurrences of the byte value 32 in multiple-byte codes.

5.2.3 Horizontal Scaling

The horizontal scaling parameter, T_h , adjusts the width of glyphs by stretching or compressing them in the horizontal direction. Its value is specified as a percentage of the normal width of the glyphs, with 100 being the normal width. The scaling always applies to the horizontal coordinate in text space, independently of the writing mode. It affects both the glyph's shape and its horizontal displacement (that is, its displacement vector). If the writing mode is horizontal, it also affects the spacing parameters T_c and T_w , as well as any positioning adjustments performed by the TJ operator. Figure 5.8 shows the effect of horizontal scaling.

$T_h = 100$ (default)	Word
$T_h = 50$	WordWord

FIGURE 5.8 *Horizontal scaling*

5.2.4 Leading

The leading parameter, T_l , is measured in unscaled text space units. It specifies the vertical distance between the baselines of adjacent lines of text, as shown in Figure 5.9.




FIGURE 5.9 *Leading*

The leading parameter is used by the **TD**, **T***, **'**, and **"** operators; see Table 5.5 on page 310 for a precise description of its effects. This parameter always applies to the vertical coordinate in text space, independently of the writing mode.

5.2.5 Text Rendering Mode

The text rendering mode, T_{mode} , determines whether showing text causes glyph outlines to be stroked, filled, used as a clipping boundary, or some combination of the three. Stroking, filling, and clipping have the same effects for a text object as they do for a path object (see Sections 4.4.2, “Path-Painting Operators,” and 4.4.3, “Clipping Path Operators”), although they are specified in an entirely different way. The graphics state parameters affecting those operations, such as line width, are interpreted in user space rather than in text space.








Note: The text rendering mode has no effect on text displayed in a Type 3 font (see Section 5.5.4, “Type 3 Fonts”).

The text rendering modes are shown in Table 5.3. In the examples, a stroke color of black and a fill color of light gray are used. For the clipping modes (4 to 7), a series of lines has been drawn through the glyphs to show where the clipping occurs.

If the text rendering mode calls for filling, the current nonstroking color in the graphics state is used; if it calls for stroking, the current stroking color is used. In modes that perform both filling and stroking, the effect is as if each glyph outline were filled and then stroked in separate operations. If any of the glyphs overlap, the result is equivalent to filling and stroking them one at a time, producing the appearance of stacked opaque glyphs, rather than first filling and then stroking them all at once (see implementation note 39 in Appendix H). In the transparent imaging model, these combined filling and stroking modes are subject to further considerations; see “Special Path-Painting Considerations” on page 462.

The behavior of the clipping modes requires further explanation. Glyph outlines begin accumulating if a **BT** operator is executed while the text rendering mode is set to a clipping mode or if it is set to a clipping mode within a text object. Glyphs accumulate until the text object is ended by an **ET** operator; the text rendering mode must not be changed back to a nonclipping mode before that point.

TABLE 5.3 Text rendering modes

MODE	EXAMPLE	DESCRIPTION
0		Fill text.
1		Stroke text.
2		Fill, then stroke text.
3		Neither fill nor stroke text (invisible).
4		Fill text and add to path for clipping (see above).
5		Stroke text and add to path for clipping.
6		Fill, then stroke text and add to path for clipping.
7		Add text to path for clipping.

At the end of the text object, the accumulated glyph outlines, if any, are combined into a single path, treating the individual outlines as subpaths of that path and applying the nonzero winding number rule (see “Nonzero Winding Number Rule” on page 169). The current clipping path in the graphics state is set to the intersection of this path with the previous clipping path. As is the case for path objects, this clipping occurs *after* all filling and stroking operations for the text object have occurred. It remains in effect until some previous clipping path is restored by an invocation of the **Q** operator.

Note: *If no glyphs are shown, or if the only glyphs shown have no outlines (that is, if they are space characters), no clipping occurs.*

5.2.6 Text Rise

Text rise, T_{rise} , specifies the distance, in unscaled text space units, to move the baseline up or down from its default location. Positive values of text rise move the baseline up. Adjustments to the baseline are useful for drawing superscripts or subscripts. The default location of the baseline can be restored by setting the text rise to 0. Figure 5.10 illustrates the effect of the text rise. Text rise always applies to the vertical coordinate in text space, regardless of the writing mode.

(This text is) Tj 5 Ts (superscripted) Tj	This text is ^{superscripted}
(This text is) Tj -5 Ts (subscripted) Tj	This text is _{subscripted}
(This) Tj -5 Ts (text) Tj 5 Ts (moves) Tj 0 Ts (around) Tj	This _{text} ^{moves} _{around}

FIGURE 5.10 *Text rise*

5.2.7 Text Knockout

The text knockout parameter, T_k (*PDF 1.4*), is a boolean flag that determines what text elements are considered elementary objects for purposes of color compositing in the transparent imaging model. Unlike other text state parameters, there is no specific operator for setting this parameter; it can be set only via the **TK** entry in a graphics state parameter dictionary, using the **gs** operator (see Section 4.3.4, “Graphics State Parameter Dictionaries”).

The text knockout parameter applies only to entire text objects; it may not be set between the **BT** and **ET** operators delimiting a text object. If the text knockout flag is **false**, each glyph in a text object is treated as a separate elementary object; when glyphs overlap, they will composite with one another. If the flag is **true**, all glyphs in the text object are treated together as a single elementary object; when glyphs overlap, later glyphs will overwrite (“knock out”) earlier ones in the area of overlap. (The latter behavior is equivalent to treating the entire text object as if it were a knockout transparency group; see Section 7.3.5, “Knockout Groups.”) The initial value is **true**.

5.3 Text Objects

A PDF *text object* consists of operators that can show text strings, move the text position, and set text state and certain other parameters. In addition, there are three parameters that are defined only within a text object and do not persist from one text object to the next:

- T_m , the *text matrix*
- T_{lm} , the *text line matrix*
- T_{rm} , the *text rendering matrix*, actually just an intermediate result that combines the effects of text state parameters, the text matrix (T_m), and the current transformation matrix

A text object begins with the **BT** operator and ends with the **ET** operator, as shown below and described in Table 5.4.

```
BT
...Zero or more text operators or other allowed operators...
ET
```

TABLE 5.4 Text object operators

OPERANDS	OPERATOR	DESCRIPTION
—	BT	Begin a text object, initializing the text matrix, T_m , and the text line matrix, T_{lm} , to the identity matrix. Text objects cannot be nested; a second BT cannot appear before an ET .
—	ET	End a text object, discarding the text matrix.

The specific categories of text-related operators that can appear in a text object are:

- *Text state operators*, described in Section 5.2, “Text State Parameters and Operators.”
- *Text-positioning operators*, described in Section 5.3.1, “Text-Positioning Operators.”
- *Text-showing operators*, described in Section 5.3.2, “Text-Showing Operators.”

The latter two sections also provide further details about the text object parameters described above. The other operators that can appear in a text object are those related to the general graphics state, color, and marked content, as shown in Figure 4.1 on page 135.

Note: *If a content stream does not contain any text, the **Text** procedure set may be omitted (see Section 9.1, “Procedure Sets”). In those circumstances, no text operators (including operators that merely set the text state) may be present in the content stream, since those operators are defined in the same procedure set.*

Note: *Although text objects cannot be statically nested, text might be shown using a Type 3 font whose glyph descriptions include any graphics objects, including another text object. Likewise, the current color might be a tiling pattern whose pattern cell includes a text object.*

5.3.1 Text-Positioning Operators

Text space is the coordinate system in which text is shown. It is defined by the text matrix, T_m , and the text state parameters T_{fs} , T_h , and T_{rise} , which together determine the transformation from text space to user space. Specifically, the origin of the first glyph shown by a text-showing operator will be placed at the origin of text space. If text space has been translated, scaled, or rotated, then the position, size, or orientation of the glyph in user space will be correspondingly altered.

At the beginning of a text object, T_m is the identity matrix, so the origin of text space is initially the same as that of user space. The *text-positioning operators*, described in Table 5.5, alter T_m and thereby control the placement of glyphs that are subsequently painted. Also, the *text-showing operators*, described in Table 5.6 in

the next section, update T_m (by altering its e and f translation components) to take into account the horizontal or vertical displacement of each glyph painted as well as any character or word spacing parameters in the text state.

TABLE 5.5 Text-positioning operators

OPERANDS	OPERATOR	DESCRIPTION
$t_x t_y$	Td	Move to the start of the next line, offset from the start of the current line by (t_x, t_y) . t_x and t_y are numbers expressed in unscaled text space units. More precisely, this operator performs the following assignments: $T_m = T_{lm} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ t_x & t_y & 1 \end{bmatrix} \times T_{lm}$
$t_x t_y$	TD	Move to the start of the next line, offset from the start of the current line by (t_x, t_y) . As a side effect, this operator sets the leading parameter in the text state. This operator has the same effect as the code $\begin{array}{l} -t_y \text{ TL} \\ t_x t_y \text{ Td} \end{array}$
$a b c d e f$	Tm	Set the text matrix, T_m , and the text line matrix, T_{lm} , as follows: $T_m = T_{lm} = \begin{bmatrix} a & b & 0 \\ c & d & 0 \\ e & f & 1 \end{bmatrix}$ <p>The operands are all numbers, and the initial value for T_m and T_{lm} is the identity matrix, $[1 \ 0 \ 0 \ 1 \ 0 \ 0]$. Although the operands specify a matrix, they are passed to Tm as six separate numbers, not as an array.</p> <p>The matrix specified by the operands is not concatenated onto the current text matrix, but replaces it.</p>
—	T*	Move to the start of the next line. This operator has the same effect as the code $0 T_l \text{ Td}$ <p>where T_l is the current leading parameter in the text state.</p>

Additionally, a text object keeps track of a text line matrix, T_{lm} , which captures the value of T_m at the beginning of a line of text. This is convenient for aligning

evenly spaced lines of text. The text-positioning and text-showing operators read and set T_{lm} on specific occasions mentioned in Tables 5.5 and 5.6.

Note: *The text-positioning operators can appear only within text objects.*

5.3.2 Text-Showing Operators

The *text-showing operators* (Table 5.6) show text on the page, repositioning text space as they do so. All of the operators interpret the text string and apply the text state parameters as described below.

TABLE 5.6 Text-showing operators

OPERANDS	OPERATOR	DESCRIPTION
<i>string</i>	Tj	Show a text string.
<i>string</i>	'	Move to the next line and show a text string. This operator has the same effect as the code T^* $string \ Tj$
a_w a_c <i>string</i>	"	Move to the next line and show a text string, using a_w as the word spacing and a_c as the character spacing (setting the corresponding parameters in the text state). a_w and a_c are numbers expressed in unscaled text space units. This operator has the same effect as the code $a_w \ Tw$ $a_c \ Tc$ $string \ '$
<i>array</i>	TJ	Show one or more text strings, allowing individual glyph positioning (see implementation note 40 in Appendix H). Each element of <i>array</i> can be a string or a number. If the element is a string, this operator shows the string. If it is a number, the operator adjusts the text position by that amount; that is, it translates the text matrix, T_m . The number is expressed in thousandths of a unit of text space (see Section 5.3.3, "Text Space Details," and implementation note 41 in Appendix H). This amount is <i>subtracted</i> from the current horizontal or vertical coordinate, depending on the writing mode. In the default coordinate system, a positive adjustment has the effect of moving the next glyph painted either to the left or down by the given amount. Figure 5.11 shows an example of the effect of passing offsets to TJ.

[(AWAY again)] TJ	AWAY again
[(A) 120 (W) 120 (A) 95 (Y again)] TJ	AWAY again

FIGURE 5.11 Operation of the TJ operator in horizontal writing

Note: The text-showing operators can appear only within text objects.

A string operand of a text-showing operator is interpreted as a sequence of character codes identifying the glyphs to be painted. With most font types, each byte of the string is treated as a separate character code. The character code is then looked up in the font’s encoding to select the glyph, as described in Section 5.5.5, “Character Encoding.”

Beginning with PDF 1.2, a string may be shown in a composite font that uses multiple-byte codes to select some of its glyphs. In that case, one or more consecutive bytes of the string are treated as a single character code. The code lengths and the mappings from codes to glyphs are defined in a data structure called a *CMap*, described in Section 5.6, “Composite Fonts.”

The strings must conform to the syntax for string objects. When a string is written by enclosing the data in parentheses, bytes whose values are the same as those of the ASCII characters left parenthesis (40), right parenthesis (41), and backslash (92) must be preceded by a backslash character. All other byte values between 0 and 255 may be used in a string object. These rules apply to each individual byte in a string object, whether the string is interpreted by the text-showing operators as single-byte or multiple-byte character codes.

Strings presented to the text-showing operators may be of any length—even a single character per string—and may be placed on the page in any order. The grouping of glyphs into strings has no significance; showing multiple glyphs with one invocation of a text-showing operator such as **Tj** produces the same results as showing them with a separate invocation for each glyph. However, the perfor-

mance of text searching (and other text extraction operations) is significantly better if the text strings are as long as possible and are shown in natural reading order.

5.3.3 Text Space Details

As stated in Section 5.3.1, “Text-Positioning Operators,” text is shown in *text space*, which is defined by the combination of the text matrix, T_m , and the text state parameters T_{fs} , T_h , and T_{rise} . This determines how text coordinates are transformed into user space. Both the glyph’s shape and its displacement (horizontal or vertical) are interpreted in text space.

Note: *Glyphs are actually defined in glyph space, whose definition varies according to the font type as discussed in Section 5.1.3, “Glyph Positioning and Metrics.” Glyph coordinates are first transformed from glyph space to text space before being subjected to the transformations described below.*

The entire transformation from text space to device space can be represented by a *text rendering matrix*, T_{rm} :

$$T_{rm} = \begin{bmatrix} T_{fs} \times T_h & 0 & 0 \\ 0 & T_{fs} & 0 \\ 0 & T_{rise} & 1 \end{bmatrix} \times T_m \times CTM$$

T_{rm} is a temporary matrix; conceptually, it is recomputed before each glyph is painted during a text-showing operation.

After the glyph is painted, the text matrix is updated according to the glyph displacement and any spacing parameters that apply. First, a combined displacement is computed, denoted by either t_x (in horizontal writing mode) or t_y (in vertical writing mode); the variable corresponding to the other writing mode is set to 0.

$$t_x = \left(\left(w0 - \frac{T_j}{1000} \right) \times T_{fs} + T_c + T_w \right) \times T_h$$

$$t_y = \left(w1 - \frac{T_j}{1000} \right) \times T_{fs} + T_c + T_w$$

where

$w0$ and $w1$ are the glyph's horizontal and vertical displacements

T_j is a position adjustment specified by a number in a **TJ** array, if any

T_{fs} and T_h are the current text font size and horizontal scaling parameters in the graphics state

T_c and T_w are the current character and word spacing parameters in the graphics state, if applicable

The text matrix is then updated as follows:

$$T_m = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ t_x & t_y & 1 \end{bmatrix} \times T_m$$

5.4 Introduction to Font Data Structures

A font is represented in PDF as a dictionary specifying the type of font, its PostScript name, its encoding, and information that can be used to provide a substitute when the font program is not available. Optionally, the font program itself can be embedded as a stream object in the PDF file.

The font types are distinguished by the **Subtype** entry in the font dictionary. Table 5.7 lists the font types defined in PDF. Type 0 fonts are called *composite fonts*; other types of font are called *simple fonts*. In addition to fonts, PDF supports two classes of font-related objects, called *CIDFonts* and *CMaps*, described in Section 5.6.1, “CID-Keyed Fonts Overview.” *CIDFonts* are listed in Table 5.7 because, like fonts, they are collections of glyphs; however, a *CIDFont* is never used directly, but only as a component of a Type 0 font.

For all font types, the term *font dictionary* refers to a PDF dictionary containing information about the font; likewise, a *CIDFont dictionary* contains information about a *CIDFont*. Except for Type 3, this dictionary is distinct from the *font program* that defines the font's glyphs. That font program may be embedded in the PDF file as a stream object or be obtained from some external source.

TABLE 5.7 Font types

TYPE	SUBTYPE VALUE	DESCRIPTION
Type 0	Type0	(PDF 1.2) A <i>composite</i> font—a font composed of other fonts, organized hierarchically (see Section 5.6, “Composite Fonts”)
Type 1	Type1	A font that defines glyph shapes by using a special encoded format (see Section 5.5.1, “Type 1 Fonts”)
	MMType1	A <i>multiple master</i> font—an extension of the Type 1 font that allows the generation of a wide variety of typeface styles from a single font (see “Multiple Master Fonts” on page 319)
Type 3	Type3	A font that defines glyphs with streams of PDF graphics operators (see Section 5.5.4, “Type 3 Fonts”)
TrueType	TrueType	A font based on the TrueType font format (see Section 5.5.2, “TrueType Fonts”)
CIDFont	CIDFontType0	(PDF 1.2) A CIDFont whose glyph descriptions are based on Type 1 font technology (see Section 5.6.3, “CIDFonts”)
	CIDFontType2	(PDF 1.2) A CIDFont whose glyph descriptions are based on TrueType font technology (see Section 5.6.3, “CIDFonts”)

Note: This terminology differs from that used in the PostScript language. In PostScript, a font dictionary is a PostScript data structure that is created as a direct result of interpreting a font program. In PDF, a font program is always treated as if it were a separate file, even if its contents are embedded in the PDF file. The font program is interpreted by a specialized font interpreter when necessary; its contents never materialize as PDF objects.

Most font programs (and related programs, such as CIDFonts and CMaps) conform to external specifications, such as the *Adobe Type 1 Font Format*. This book does not include those specifications. See the Bibliography for more information about the specifications mentioned in this chapter.

The most predictable and dependable results are produced when all font programs used to show text are embedded in the PDF file. The following sections describe precisely how to do so. On the other hand, if a PDF file refers to font programs that are not embedded, the results depend on the availability of fonts in the viewer application’s environment. The following sections specify some con-

ventions for referring to external font programs; however, some details of font naming, font substitution, and glyph selection are implementation-dependent and may vary among different viewer applications and operating system environments.

5.5 Simple Fonts

There are several types of simple font, all of which have the following properties:

- Glyphs in the font are selected by single-byte character codes obtained from a string that is shown by the text-showing operators. Logically, these codes index into a table of 256 glyphs; the mapping from codes to glyphs is called the font's *encoding*. Each font program has a built-in encoding. Under some circumstances, the encoding can be altered by means described in Section 5.5.5, "Character Encoding."
- Each glyph has a single set of metrics, including a *horizontal displacement* or *width*, as described in Section 5.1.3, "Glyph Positioning and Metrics." That is, simple fonts support only horizontal writing mode.
- Except for Type 3 fonts and certain standard Type 1 fonts, every font dictionary contains a subsidiary dictionary, the *font descriptor*, containing fontwide metrics and other attributes of the font; see Section 5.7, "Font Descriptors." Among those attributes is an optional *font file* stream containing the font program itself.

5.5.1 Type 1 Fonts

A Type 1 font program is a stylized PostScript program that describes glyph shapes. It uses a compact encoding for the glyph descriptions, and it includes hint information that enables high-quality rendering even at small sizes and low resolutions. Details on this format are provided in a separate book, *Adobe Type 1 Font Format*. An alternative, more compact but functionally equivalent representation of a Type 1 font program is documented in Adobe Technical Note #5176, *The Compact Font Format Specification*.

Note: Although a Type 1 font program uses PostScript language syntax, using it does not require a full PostScript interpreter; a specialized Type 1 font interpreter suffices.

A Type 1 font dictionary contains the entries listed in Table 5.8. Some entries are optional for the standard 14 fonts listed under “Standard Type 1 Fonts” on page 319, but are required otherwise.

TABLE 5.8 Entries in a Type 1 font dictionary

KEY	TYPE	VALUE
Type	name	<i>(Required)</i> The type of PDF object that this dictionary describes; must be Font for a font dictionary.
Subtype	name	<i>(Required)</i> The type of font; must be Type1 for a Type 1 font.
Name	name	<i>(Required in PDF 1.0; optional otherwise)</i> The name by which this font is referenced in the Font subdictionary of the current resource dictionary. <i>Note: This entry is obsolescent and its use is no longer recommended. (See implementation note 42 in Appendix H.)</i>
BaseFont	name	<i>(Required)</i> The PostScript name of the font. For Type 1 fonts, this is usually the value of the FontName entry in the font program; for more information, see Section 5.2 of the <i>PostScript Language Reference</i> , Third Edition. The PostScript name of the font can be used to find the font’s definition in the viewer application or its environment. It is also the name that will be used when printing to a PostScript output device.
FirstChar	integer	<i>(Required except for the standard 14 fonts)</i> The first character code defined in the font’s Widths array.
LastChar	integer	<i>(Required except for the standard 14 fonts)</i> The last character code defined in the font’s Widths array.
Widths	array	<i>(Required except for the standard 14 fonts; indirect reference preferred)</i> An array of (LastChar – FirstChar + 1) widths, each element being the glyph width for the character whose code is FirstChar plus the array index. For character codes outside the range FirstChar to LastChar , the value of MissingWidth from the FontDescriptor entry for this font is used. The glyph widths are measured in units in which 1000 units corresponds to 1 unit in text space. These widths must be consistent with the actual widths given in the font program itself. (See implementation note 43 in Appendix H.) For more information on glyph widths and other glyph metrics, see Section 5.1.3, “Glyph Positioning and Metrics.”
FontDescriptor	dictionary	<i>(Required except for the standard 14 fonts; must be an indirect reference)</i> A font descriptor describing the font’s metrics other than its glyph widths (see Section 5.7, “Font Descriptors”).

Note: For the standard 14 fonts, the entries **FirstChar**, **LastChar**, **Widths**, and **FontDescriptor** must either all be present or all absent. Ordinarily, they are absent; specifying them enables a standard font to be overridden (see “Standard Type 1 Fonts,” below).

Encoding	name or dictionary	(<i>Optional</i>) A specification of the font’s character encoding, if different from its built-in encoding. The value of Encoding may be either the name of a pre-defined encoding (MacRomanEncoding , MacExpertEncoding , or WinAnsiEncoding , as described in Appendix D) or an encoding dictionary that specifies differences from the font’s built-in encoding or from a specified pre-defined encoding (see Section 5.5.5, “Character Encoding”).
ToUnicode	stream	(<i>Optional; PDF 1.2</i>) A stream containing a CMap file that maps character codes to Unicode values (see Section 5.9, “ToUnicode CMaps”).

Example 5.6 shows the font dictionary for the Adobe Garamond™ Semibold font. The font has an encoding dictionary (object 25), although neither the encoding dictionary nor the font descriptor (object 7) is shown in the example.

Example 5.6

```

14 0 obj
  << /Type /Font
    /Subtype /Type1
    /BaseFont /AGaramond-Semibold
    /FirstChar 0
    /LastChar 255
    /Widths 21 0 R
    /FontDescriptor 7 0 R
    /Encoding 25 0 R
  >>
endobj

21 0 obj
  [ 255 255 255 255 255 255 255 255 255 255 255 255 255 255 255
    255 255 255 255 255 255 255 255 255 255 255 255 255 255 255
    255 280 438 510 510 868 834 248 320 320 420 510 255 320 255 347
    510 510 510 510 510 510 510 510 510 510 255 255 510 510 510 330
    781 627 627 694 784 580 533 743 812 354 354 684 560 921 780 792
    588 792 656 504 682 744 650 968 648 590 638 320 329 320 510 500
    380 420 510 400 513 409 301 464 522 268 259 484 258 798 533 492
    516 503 349 346 321 520 434 684 439 448 390 320 255 320 510 255
    627 627 694 580 780 792 744 420 420 420 420 420 420 402 409 409
    409 409 268 268 268 268 533 492 492 492 492 492 520 520 520 520
  ]

```

```

486 400 510 510 506 398 520 555 800 800 1044 360 380 549 846 792
713 510 549 549 510 522 494 713 823 549 274 354 387 768 615 496
330 280 510 549 510 549 612 421 421 1000 255 627 627 792 1016 730
500 1000 438 438 248 248 510 494 448 590 100 510 256 256 539 539
486 255 248 438 1174 627 580 627 580 580 354 354 354 354 792 792
790 792 744 744 744 268 380 380 380 380 380 380 380 380 380

```

```

]

```

```

endobj

```

Standard Type 1 Fonts

The PostScript names of 14 Type 1 fonts, known as the *standard fonts*, are as follows:

Times–Roman	Helvetica	Courier	Symbol
Times–Bold	Helvetica–Bold	Courier–Bold	ZapfDingbats
Times–Italic	Helvetica–Oblique	Courier–Oblique	
Times–BoldItalic	Helvetica–BoldOblique	Courier–BoldOblique	

These fonts, or their font metrics and suitable substitution fonts, are guaranteed to be available to the viewer application. The character sets and encodings for these fonts are given in Appendix D. The Adobe font metrics (AFM) files for the standard 14 fonts are available from the ASN Developer Program Web site (see the Bibliography). For more information on font metrics, see Adobe Technical Note #5004, *Adobe Font Metrics File Format Specification*.

Ordinarily, a font dictionary that refers to one of the standard fonts should omit the **FirstChar**, **LastChar**, **Widths**, and **FontDescriptor** entries. However, it is permissible to override a standard font by including these entries and embedding the font program in the PDF file. (See implementation note 44 in Appendix H.)

Multiple Master Fonts

The *multiple master* font format is an extension of the Type 1 font format that allows the generation of a wide variety of typeface styles from a single font program. This is accomplished through the presence of various design dimensions in the font. Examples of design dimensions are weight (light to extra-bold) and width (condensed to expanded). Coordinates along these design dimensions

(such as the degree of boldness) are specified by numbers. A particular choice of numbers selects an *instance* of the multiple master font. Adobe Technical Note #5015, *Type 1 Font Format Supplement*, describes multiple master fonts in detail.

The font dictionary for a multiple master font instance has the same entries as a Type 1 font dictionary (Table 5.8 on page 317), except note the following:

- The value of **Subtype** is **MMType1**.
- If the PostScript name of the instance contains spaces, the spaces are replaced by underscores in the value of **BaseFont**. For instance, as illustrated in Example 5.7, the name “MinionMM 366 465 11 ” (which ends with a space character) becomes `/MinionMM_366_465_11_`.

Example 5.7

```

7 0 obj
  << /Type /Font
    /Subtype /MMType1
    /BaseFont /MinionMM_366_465_11_
    /FirstChar 32
    /LastChar 255
    /Widths 19 0 R
    /FontDescriptor 6 0 R
    /Encoding 5 0 R
  >>
endobj

19 0 obj
  [ 187 235 317 430 427 717 607 168 326 326 421 619 219 317 219 282 427
    ...Omitted data...
    569 0 569 607 607 607 239 400 400 400 400 253 400 400 400 400
  ]
endobj

```

This example illustrates a convention for including the numeric values of the design coordinates as part of the instance’s **BaseFont** name. This convention is commonly used for accessing multiple master font instances from an external source in the viewer application’s environment; it is documented in Adobe Technical Note #5088, *Font Naming Issues*. However, this convention is not prescribed as part of the PDF specification. In particular, if the font program for this in-

stance is embedded in the PDF file, it must be an ordinary Type 1 font program, not a multiple master font program. This font program is called a *snapshot* of the multiple master font instance, incorporating the chosen values of the design coordinates.

5.5.2 TrueType Fonts

The *TrueType* font format was developed by Apple Computer, Inc., and has been adopted as a standard font format for the Microsoft Windows operating system. Specifications for the TrueType font file format are available in Apple's *TrueType Reference Manual* and Microsoft's *TrueType 1.0 Font Files Technical Specification*.

Note: *A TrueType font program can be embedded directly in a PDF file as a stream object. The Type 42 font format that is defined for PostScript does not apply to PDF.*

A TrueType font dictionary can contain the same entries as a Type 1 font dictionary (Table 5.8 on page 317), except note the following:

- The value of **Subtype** is **TrueType**.
- The value of **BaseFont** is derived differently, as described below.
- The value of **Encoding** is subject to limitations that are described in Section 5.5.5, “Character Encoding.”

The PostScript name for the value of **BaseFont** is determined in one of two ways:

- Use the PostScript name that is an optional entry in the “name” table of the TrueType font itself.
- In the absence of such an entry in the “name” table, derive a PostScript name from the name by which the font is known in the host operating system: on a Windows system, it is based on the `lfFaceName` field in a LOGFONT structure; in the Mac OS, it is based on the name of the FONDR resource. If the name contains any spaces, the spaces are removed.

If the font in a source document uses a bold or italic style, but there is no font data for that style, the host operating system will synthesize the style. In this case, a comma and the style name (one of Bold, Italic, or BoldItalic) are appended to the font name. For example, for a TrueType font that is a bold variant of the New

York font, the **BaseFont** value is written as `/NewYork,Bold` (as illustrated in Example 5.8).

Example 5.8

```

17 0 obj
  << /Type /Font
    /Subtype /TrueType
    /BaseFont /NewYork,Bold
    /FirstChar 0
    /LastChar 255
    /Widths 23 0 R
    /FontDescriptor 7 0 R
    /Encoding /MacRomanEncoding
  >>
endobj

23 0 obj
  [ 0 333 333 333 333 333 333 333 0 333 333 333 333 333 333 333 333
    ...Omitted data...
    803 790 803 780 780 780 340 636 636 636 636 636 636 636 636
  ]
endobj

```

Note that for CJK (Chinese, Japanese, and Korean) fonts, the host font system’s “font name” is often encoded in the host operating system’s script. For instance, a Japanese font may have a name that is written in Japanese using some (unidentified) Japanese encoding. Thus, TrueType font names may contain multiple-byte character codes, each of which requires multiple characters to represent in a PDF name object (using the # notation to quote special characters as needed).

5.5.3 Font Subsets

PDF 1.1 permits documents to include subsets of Type 1 and TrueType fonts. The font and font descriptor that describe a font subset are slightly different from those of ordinary fonts. These differences allow an application to recognize font subsets and to merge documents containing different subsets of the same font. (For more information on font descriptors, see Section 5.7, “Font Descriptors.”)

For a font subset, the PostScript name of the font—the value of the font’s **BaseFont** entry and the font descriptor’s **FontName** entry—begins with a *tag*

followed by a plus sign (+). The tag consists of exactly six uppercase letters; the choice of letters is arbitrary, but different subsets in the same PDF file must have different tags. For example, EOODIA+Poetica is the name of a subset of Poetica[®], a Type 1 font. (See implementation note 45 in Appendix H.)

5.5.4 Type 3 Fonts

Type 3 fonts differ from the other fonts supported by PDF. A Type 3 font dictionary defines the font itself, while the other font dictionaries simply contain information *about* the font and refer to a separate font program for the actual glyph descriptions. In Type 3 fonts, glyphs are defined by streams of PDF graphics operators. These streams are associated with character names. A separate encoding entry maps character codes to the appropriate character names for the glyphs.

Type 3 fonts are more flexible than Type 1 fonts, because the glyph descriptions may contain arbitrary PDF graphics operators. However, Type 3 fonts have no hinting mechanism for improving output at small sizes or low resolutions. A Type 3 font dictionary contains the entries listed in Table 5.9.

TABLE 5.9 Entries in a Type 3 font dictionary

KEY	TYPE	VALUE
Type	name	<i>(Required)</i> The type of PDF object that this dictionary describes; must be Font for a font dictionary.
Subtype	name	<i>(Required)</i> The type of font; must be Type3 for a Type 3 font.
Name	name	<i>(Required in PDF 1.0; optional otherwise)</i> See Table 5.8 on page 317.
FontBBox	rectangle	<i>(Required)</i> A rectangle (see Section 3.8.3, “Rectangles”), expressed in the glyph coordinate system, specifying the <i>font bounding box</i> . This is the smallest rectangle enclosing the shape that would result if all of the glyphs of the font were placed with their origins coincident and then filled. If all four elements of the rectangle are zero, no assumptions are made based on the font bounding box. If any element is nonzero, it is essential that the font bounding box be accurate; if any glyph’s marks fall outside this bounding box, incorrect behavior may result.
FontMatrix	array	<i>(Required)</i> An array of six numbers specifying the <i>font matrix</i> , mapping glyph space to text space (see Section 5.1.3, “Glyph Positioning and Metrics”). A common practice is to define glyphs in terms of a 1000-unit

		glyph coordinate system, in which case the font matrix is [0.001 0 0 0.001 0 0].
CharProcs	dictionary	<i>(Required)</i> A dictionary in which each key is a character name and the value associated with that key is a content stream that constructs and paints the glyph for that character. The stream must include as its first operator either d0 or d1 . This is followed by operators describing one or more graphics objects, which may include path, text, or image objects. See below for more details about Type 3 glyph descriptions.
Encoding	name or dictionary	<i>(Required)</i> An encoding dictionary whose Differences array specifies the complete character encoding for this font (see Section 5.5.5, “Character Encoding”; also see implementation note 46 in Appendix H).
FirstChar	integer	<i>(Required)</i> The first character code defined in the font’s Widths array.
LastChar	integer	<i>(Required)</i> The last character code defined in the font’s Widths array.
Widths	array	<i>(Required; indirect reference preferred)</i> An array of (LastChar – FirstChar + 1) widths, each element being the glyph width for the character whose code is FirstChar plus the array index. For character codes outside the range FirstChar to LastChar , the width is 0. These widths are interpreted in glyph space as specified by FontMatrix (unlike the widths of a Type 1 font, which are in thousandths of a unit of text space). <i>Note: If FontMatrix specifies a rotation, only the horizontal component of the transformed width is used. That is, the resulting displacement is always horizontal in text space, as is the case for all simple fonts.</i>
Resources	dictionary	<i>(Optional but strongly recommended; PDF 1.2)</i> A list of the named resources, such as fonts and images, required by the glyph descriptions in this font (see Section 3.7.2, “Resource Dictionaries”). If any glyph descriptions refer to named resources but this dictionary is absent, the names are looked up in the resource dictionary of the page on which the font is used. (See implementation note 47 in Appendix H.)
ToUnicode	stream	<i>(Optional; PDF 1.2)</i> A stream containing a CMap file that maps character codes to Unicode values (see Section 5.9, “ToUnicode CMaps”).

For each character shown by a text-showing operator using a Type 3 font, the viewer application does the following:

1. Looks up the character code in the font’s **Encoding** entry, as described in Section 5.5.5, “Character Encoding,” to obtain a character name.

2. Looks up the character name in the font's **CharProcs** dictionary to obtain a stream object containing a glyph description. (If the name is not present as a key in **CharProcs**, no glyph is painted.)
3. Invokes the glyph description, as described below. The graphics state is saved before this invocation and restored afterward, so any changes the glyph description makes to the graphics state do not persist after it finishes.

When the glyph description begins execution, the current transformation matrix (CTM) is the concatenation of the font matrix (**FontMatrix** in the current font dictionary) and the text space that was in effect at the time the text-showing operator was invoked (see Section 5.3.3, “Text Space Details”). This means that shapes described in the glyph coordinate system will be transformed into the user coordinate system and will appear in the appropriate size and orientation on the page. The glyph description should describe the glyph in terms of absolute coordinates in the glyph coordinate system, placing the glyph origin at (0, 0) in this space. It should make no assumptions about the initial text position.

Aside from the CTM, the graphics state is inherited from the environment of the text-showing operator that caused the glyph description to be invoked. To ensure predictable results, the glyph description must initialize any graphics state parameters on which it depends. In particular, if it invokes the **S** (stroke) operator, it should explicitly set the line width, line join, line cap, and dash pattern to appropriate values. Normally, it is unnecessary and undesirable to initialize the current color parameter, because the text-showing operators are designed to paint glyphs with the current color.

The glyph description must execute one of the operators described in Table 5.10 to pass width and bounding box information to the font machinery. This must precede the execution of any path construction or path-painting operators describing the glyph.

***Note:** Type 3 fonts in PDF are very similar to those in PostScript. Some of the information provided in Type 3 font dictionaries and glyph descriptions, while seemingly redundant or unnecessary, is nevertheless required for correct results when a PDF viewer application prints to a PostScript output device. This applies particularly to the operands of the **d0** and **d1** operators, which in PostScript are named **setcharwidth** and **setcachedevice**. For further explanation, see Section 5.7 of the PostScript Language Reference, Third Edition.*

TABLE 5.10 Type 3 font operators

OPERANDS	OPERATOR	DESCRIPTION
w_x w_y	d0	<p>Set width information for the glyph and declare that the glyph description specifies both its shape and its color. (Note that this operator name ends in the digit 0.) w_x specifies the horizontal displacement in the glyph coordinate system; it must be consistent with the corresponding width in the font's Widths array. w_y must be 0 (see Section 5.1.3, "Glyph Positioning and Metrics").</p> <p>This operator is permitted only in a content stream appearing in a Type 3 font's CharProcs dictionary. It is typically used only if the glyph description executes operators to set the color explicitly.</p>
w_x w_y ll_x ll_y ur_x ur_y	d1	<p>Set width and bounding box information for the glyph and declare that the glyph description specifies only shape, not color. (Note that this operator name ends in the digit 1.) w_x specifies the horizontal displacement in the glyph coordinate system; it must be consistent with the corresponding width in the font's Widths array. w_y must be 0 (see Section 5.1.3, "Glyph Positioning and Metrics").</p> <p>ll_x and ll_y are the coordinates of the lower-left corner, and ur_x and ur_y the upper-right corner, of the glyph bounding box. The glyph bounding box is the smallest rectangle, oriented with the axes of the glyph coordinate system, that completely encloses all marks placed on the page as a result of executing the glyph's description. The declared bounding box must be correct—in other words, sufficiently large to enclose the entire glyph. If any marks fall outside this bounding box, the result is unpredictable.</p> <p>A glyph description that begins with the d1 operator should not execute any operators that set the color (or other color-related parameters) in the graphics state; any use of such operators will be ignored. The glyph description is executed solely to determine the glyph's shape; its color is determined by the graphics state in effect each time this glyph is painted by a text-showing operator. For the same reason, the glyph description may not include an image; however, an image mask is acceptable, since it merely defines a region of the page to be painted with the current color.</p> <p>This operator is permitted only in a content stream appearing in a Type 3 font's CharProcs dictionary.</p>

Example of a Type 3 Font

Example 5.9 shows the definition of a Type 3 font with only two glyphs—a filled square and a filled triangle, selected by the characters a and b. Figure 5.12 shows the result of showing the string (ababab) using this font.




FIGURE 5.12 *Output from Example 5.9*

Example 5.9

```
4 0 obj
  << /Type /Font
    /Subtype /Type3
    /FontBBox [0 0 750 750]
    /FontMatrix [0.001 0 0 0.001 0 0]
    /CharProcs 10 0 R
    /Encoding 9 0 R
    /FirstChar 97
    /LastChar 98
    /Widths [1000 1000]
  >>
endobj

9 0 obj
  << /Type /Encoding
    /Differences [97 /square /triangle]
  >>
endobj
```

```
10 0 obj
  << /square 11 0 R
    /triangle 12 0 R
  >>
endobj

11 0 obj
  << /Length 39 >>
stream
  1000 0 0 0 750 750 d1
  0 0 750 750 re
  f
endstream
endobj

12 0 obj
  << /Length 48 >>
stream
  1000 0 0 0 750 750 d1
  0 0 m
  375 750 l
  750 0 l
  f
endstream
endobj
```

5.5.5 Character Encoding

A font's *encoding* is the association between character codes (obtained from text strings that are shown) and glyph descriptions. This section describes the character encoding scheme used with simple PDF fonts. Composite fonts (Type 0) use a different character mapping algorithm, as discussed in Section 5.6, "Composite Fonts."

Except for Type 3 fonts, every font program has a built-in encoding. Under certain circumstances, a PDF font dictionary can change a font's built-in encoding to match the requirements of the application generating the text being shown. This flexibility in character encoding is valuable for two reasons:

- It permits showing text that is encoded according to any of the various existing conventions. For example, the Microsoft Windows and Apple Mac OS oper-

ating systems use different standard encodings for Latin text, and many applications use their own special-purpose encodings.

- It allows applications to specify how characters selected from a large character set are to be encoded. Some character sets consist of more than 256 characters, including ligatures, accented characters, and other symbols required for high-quality typography or non-Latin writing systems. Different encodings can select different subsets of the same character set.

Latin-text font programs produced by Adobe Systems use the *Adobe standard encoding*, often referred to as **StandardEncoding**. The name **StandardEncoding** has no special meaning in PDF, but this encoding does play a role as a default encoding (as shown in Table 5.11 below). The regular encodings used for Latin-text fonts on Windows and Mac OS systems are named **MacRomanEncoding** and **WinAnsiEncoding**, respectively. Additionally, an encoding named **MacExpertEncoding** is used with “expert” fonts that contain additional characters useful for sophisticated typography. Complete details of these encodings and of the characters present in typical fonts are provided in Appendix D.

In PDF, a font is classified as either *nonsymbolic* or *symbolic* according to whether or not all of its characters are members of the Adobe standard Latin character set. This is indicated by flags in the font descriptor; see Section 5.7.1, “Font Descriptor Flags.” Symbolic fonts contain other character sets, to which the encodings mentioned above ordinarily do not apply. Such font programs have built-in encodings that are usually unique to each font. The standard 14 fonts include two symbolic fonts, Symbol and ZapfDingbats, whose encodings and character sets are documented in Appendix D.

A font program’s built-in encoding can be overridden or altered by including an **Encoding** entry in the PDF font dictionary. The possible encoding modifications depend on the font type, as discussed below. The value of the **Encoding** entry is either a named encoding (the name of one of the predefined encodings **MacRomanEncoding**, **MacExpertEncoding**, or **WinAnsiEncoding**) or an *encoding dictionary*. An encoding dictionary contains the entries listed in Table 5.11.

TABLE 5.11 Entries in an encoding dictionary

KEY	TYPE	VALUE
Type	name	<i>(Optional)</i> The type of PDF object that this dictionary describes; if present, must be Encoding for an encoding dictionary.
BaseEncoding	name	<i>(Optional)</i> The <i>base encoding</i> —that is, the encoding from which the Differences entry (if present) describes differences—specified as the name of a predefined encoding MacRomanEncoding , MacExpertEncoding , or WinAnsiEncoding (see Appendix D). If this entry is absent, the Differences entry describes differences from an implicit base encoding. For a font program that is embedded in the PDF file, the implicit base encoding is the font program’s built-in encoding, as described above and further elaborated in the sections on specific font types below. Otherwise, for a nonsymbolic font, it is StandardEncoding , and for a symbolic font, it is the font’s built-in encoding.
Differences	array	<i>(Optional; not recommended with TrueType fonts)</i> An array describing the differences from the encoding specified by BaseEncoding or, if BaseEncoding is absent, from an implicit base encoding. The Differences array is described above.

The value of the **Differences** entry is an array of character codes and character names organized as follows:

```
code1 name1,1 name1,2 ...
code2 name2,1 name2,2 ...
...
coden namen,1 namen,2 ...
```

Each code is the first index in a sequence of characters to be changed. The first character name after the code becomes the name corresponding to that code. Subsequent names replace consecutive code indices until the next code appears in the array or the array ends. These sequences may be specified in any order but should not overlap.

For example, in the encoding dictionary in Example 5.10, the name `quotesingle` (') is associated with character code 39, `Adieresis` (Ä) with code 128, `Aring` (Å) with 129, and `trademark` (™) with 170.

Example 5.10

```

25 0 obj
  << /Type /Encoding
    /Differences
      [ 39 /quotesingle
        96 /grave
        128 /Adieresis /Aring /Ccedilla /Eacute /Ntilde /Odieresis /Udieresis
          /aacute /agrave /acircumflex /adieresis /atilde /aring /ccedilla
          /eacute /egrave /ecircumflex /edieresis /iacute /igrave /icircumflex
          /idieresis /ntilde /oacute /ograve /ocircumflex /odieresis /otilde
          /uacute /ugrave /ucircumflex /udieresis /dagger /degree /cent
          /sterling /section /bullet /paragraph /germandbls /registered
          /copyright /trademark /acute /dieresis
        174 /AE /Oslash
        177 /plusminus
        180 /yen /mu
        187 /ordfeminine /ordmasculine
        190 /ae /oslash /questiondown /exclamdown /logicalnot
        196 /florin
        199 /guillemotleft /guillemotright /ellipsis
        203 /Agrave /Atilde /Otilde /OE /oe /endash /emdash /quotedblleft
          /quotedblright /quoteleft /quoteright /divide
        216 /ydieresis /Ydieresis /fraction /currency /guilsinglleft /guilsinglright
          /fi /fl /daggerdbl /periodcentered /quotesinglbase /quotedblbase
          /perthousand /Acircumflex /Ecircumflex /Aacute /Edieresis /Egrave
          /Iacute /Icircumflex /Idieresis /Igrave /Oacute /Ocircumflex
        241 /Ograve /Uacute /Ucircumflex /Ugrave /dotlessi /circumflex /tilde
          /macron /breve /dotaccent /ring /cedilla /hungarumlaut /ogonek
          /caron
      ]
  >>
endobj

```

By convention, the name `.notdef` can be used to indicate that there is no character name associated with a given character code.

Encodings for Type 1 Fonts

A Type 1 font program's glyph descriptions are keyed by character *names*, not by character *codes*. Character names are ordinary PDF name objects. Descriptions of Latin alphabetic characters are normally associated with names consisting of single letters, such as **A** or **a**. Other characters are associated with names com-

posed of words, such as three, ampersand, or parenleft. A Type 1 font's built-in encoding is defined by an **Encoding** array that is part of the font program itself; this is not to be confused with the **Encoding** entry in the PDF font dictionary.

An **Encoding** entry can alter a Type 1 font's mapping from character codes to character names. The **Differences** array can map a code to the name of any glyph description that exists in the font program, whether or not that glyph is referenced by the font's built-in encoding or by the encoding specified in the **BaseEncoding** entry.

All Type 1 font programs contain an actual glyph for the character named `.notdef`. The effect produced by showing the `.notdef` character is at the discretion of the font designer; in Type 1 font programs produced by Adobe, it is the same as the space character. If an encoding maps to a character name that does not exist in the Type 1 font program, the `.notdef` character is substituted.

Encodings for Type 3 Fonts

A Type 3 font, like Type 1, contains glyph descriptions that are keyed by character names; in this case, they appear as explicit keys in the font's **CharProcs** dictionary. A Type 3 font's mapping from character codes to character names is entirely defined by its **Encoding** entry, which is required in this case.

Encodings for TrueType Fonts

A TrueType font program's built-in encoding maps directly from character codes to glyph descriptions, using an internal data structure called a "cmap" (not to be confused with the CMap described in Section 5.6.4, "CMaps"). A TrueType font program can contain multiple encodings that are intended for use on different platforms (such as Mac OS and Windows). The PDF font dictionary's **Encoding** entry can select among the available encodings; in the absence of this entry, an implementation-dependent encoding is chosen.

There is no standard support for named characters in TrueType, although some font programs have an optional "post" table listing character names for the glyphs. If the viewer application needs to select glyph descriptions by name in a font program lacking a "post" table, it translates from character names to codes in one of the encodings given in the font program's "cmap" table.

Because some aspects of TrueType glyph selection are dependent on the viewer implementation or the operating system, PDF files that use TrueType fonts should follow certain guidelines to ensure predictable behavior across all viewer applications. The font program should be embedded. A nonsymbolic font should specify **MacRomanEncoding** or **WinAnsiEncoding** as the value of its **Encoding** entry, with no **Differences** array. A symbolic font should not specify an **Encoding** entry; its font program’s “cmap” table should contain exactly one encoding. See below for specific guidelines on the contents of this “cmap” table.

***Note:** Some popular TrueType font programs contain incorrect encoding information. Implementations of TrueType font interpreters have evolved heuristics for dealing with such problems; those heuristics are not described here. For maximum portability, only well-formed TrueType font programs should be used in PDF files.*

The following paragraphs describe the treatment of TrueType font encodings beginning with PDF 1.3, as implemented in Acrobat 4.0 and later viewers. This information does not necessarily apply to earlier versions or implementations.

A TrueType font program’s “cmap” table consists of one or more subtables, each identified by the combination of a *platform ID* and a *platform-specific encoding ID*. If a named encoding (**WinAnsiEncoding**, **MacRomanEncoding**, or **MacExpertEncoding**) is specified in a font dictionary’s **Encoding** entry or in an encoding dictionary’s **BaseEncoding** entry, a “cmap” subtable is selected and used as described below.

- If a “cmap” subtable with platform ID 3 and encoding ID 1 (Microsoft Unicode) is present, it is used as follows: A character code is first mapped to a character name as specified by the font’s **Encoding** entry. The character name is then mapped to a Unicode value by consulting the *Adobe Glyph List* (see the Bibliography). Finally, the Unicode value is mapped to a glyph description according to the (3, 1) subtable.
- If a “cmap” subtable with platform ID 1 and encoding 0 (Macintosh Roman) is present, it is used as follows: A character code is first mapped to a character name as specified by the font’s **Encoding** entry. The character name is then mapped back to a character code according to **MacRomanEncoding** (see Appendix D). Finally, the code is mapped to a glyph description according to the (1, 0) subtable.

- In either of the cases above, if the character name cannot be mapped as specified, the character name is looked up in the font program’s “post” table (if one is present) and the associated glyph description is used.

If no **Encoding** entry is specified in the font dictionary, the “cmap” subtable with platform ID 1 and encoding 0 will be used to map directly from character codes to glyph descriptions, without any consideration of character names. This is the normal convention for symbolic fonts.

If a character cannot be mapped in any of the ways described above, the results are implementation-dependent.

5.6 Composite Fonts

A *composite font* is one whose glyphs are obtained from other fonts or from font-like objects called *CIDFonts*, organized hierarchically. In PDF, a composite font is represented by a font dictionary whose **Subtype** value is **Type0**; this is also called a Type 0 font. The Type 0 font at the top level of the hierarchy is the *root font*. Fonts and CIDFonts immediately below a Type 0 font are called its *descendants*. The Type 0 font immediately above a descendant is called its *parent font*.

When the current font is composite, the text-showing operators behave differently than with simple fonts. Whereas for simple fonts each byte of a string to be shown selects one character, for composite fonts a sequence of one or more bytes can be decoded to select a character from any of the descendant fonts or CIDFonts. This facility supports the use of very large character sets, such as those for the Chinese, Japanese, and Korean languages. It also simplifies the organization of fonts that have complex encoding requirements.

PDF 1.2 introduces a general architecture for composite fonts that theoretically allows a Type 0 font to have multiple descendants, which might themselves be Type 0 fonts. However, in versions up to and including PDF 1.4, only a single descendant is allowed, which must be a CIDFont (not a font). This restriction may be relaxed in a future PDF version.

Note: *Composite fonts in PDF are analogous to composite fonts in PostScript, but with some limitations. In particular, PDF requires that the character encoding be defined by a CMap (described below), which is only one of several encoding methods available in PostScript.*

This section first introduces the architecture of *CID-keyed fonts*, which are the only kind of composite font supported in PDF. Then it describes the *CIDFont* and *CMap* dictionaries, which are the PDF objects that represent the correspondingly named components of a CID-keyed font. Finally, it describes the Type 0 font dictionary, which combines a *CIDFont* and a *CMap* to produce a font whose glyphs can be accessed by means of variable-length character codes in a string to be shown.

5.6.1 CID-Keyed Fonts Overview

CID-keyed fonts provide a convenient and efficient method for defining multiple-byte character encodings, fonts with a large number of glyphs, and fonts that incorporate glyphs obtained from other fonts. These capabilities provide great flexibility for representing text in writing systems for languages with large character sets, such as Chinese, Japanese, and Korean (CJK).

The CID-keyed font architecture specifies the external representation of certain font programs, called *CMap* and *CIDFont* files, along with some conventions for combining and using those files. This architecture is independent of PDF; CID-keyed fonts can be used in other environments. For complete documentation on the architecture and the file formats, see Adobe Technical Notes #5092, *CID-Keyed Font Technology Overview*, and #5014, *Adobe CMap and CIDFont Files Specification*. This section describes only the PDF objects that represent these font programs.

The term *CID-keyed font* reflects the fact that *CID* (character identifier) numbers are used to index and access the glyph descriptions in the font. This method is more efficient for large fonts than the method of accessing by character name, as is used for some simple fonts. CIDs range from 0 to a maximum value that is subject to an implementation limit (see Appendix C).

A *character collection* is an ordered set of all characters needed to support one or more popular character sets for a particular language. The order of the characters in the character collection determines the CID number for each character. Each CID-keyed font must explicitly reference the character collection on which its CID numbers are based; see Section 5.6.2, “CIDSystemInfo Dictionaries.”

A *CMap* (character map) file specifies the correspondence between character codes and the CID numbers used to identify characters. It is equivalent to the concept of an encoding in simple fonts. Whereas a simple font allows a maximum

of 256 characters to be encoded and accessible at one time, a CMap can describe a mapping from multiple-byte codes to thousands of characters in a large CID-keyed font. For example, it can describe Shift-JIS, one of several widely used encodings for Japanese, or Unicode, an international standard encoding that covers many languages.

A CMap can reference an entire character collection, a subset, or multiple character collections. It can also reference characters in other fonts by character code or character name. The CMap mapping yields a *font number* and a *character selector* that can be a CID, a character code, or a character name. Furthermore, a CMap can incorporate another CMap by reference, without having to duplicate it. These features enable character collections to be combined or supplemented, and make all the constituent characters accessible to text-showing operations through a single encoding.

Note: As mentioned earlier, PDF versions up to and including PDF 1.4 do not support the entire CID-keyed font architecture. In PDF, a CID-keyed font may have only a single descendant, whose characters must be referenced by CID.

A *CIDFont* file contains the glyph descriptions for a character collection. The glyph descriptions themselves are typically in a format similar to those used in simple fonts, such as Type 1. However, they are identified by CIDs rather than by names, and they are organized differently.

In PDF, the CMap and CIDFont are represented by PDF objects, which are described below. The CMap and CIDFont programs themselves can be either referenced by name or embedded as stream objects in the PDF file. As stated earlier, the external file formats are not documented here, but in Adobe Technical Note #5014, *Adobe CMap and CIDFont Files Specification*.

A CID-keyed font, then, is the combination of a CMap with one or more CIDFonts, simple fonts, or composite fonts containing glyph descriptions. In PDF, a CID-keyed font is represented as a Type 0 font. It contains an **Encoding** entry whose value is a CMap dictionary, and its **DescendantFonts** array references the CIDFont or font dictionaries with which the CMap has been combined.

5.6.2 CIDSystemInfo Dictionaries

CIDFont and CMap dictionaries contain a **CIDSystemInfo** entry specifying the character collection assumed by the CIDFont or by each CIDFont associated with

the CMap—that is, the interpretation of the CID numbers used by the CIDFont. A character collection is uniquely identified by the **Registry**, **Ordering**, and **Supplement** entries in the **CIDSystemInfo** dictionary, as described in Table 5.12. Character collections whose **Registry** and **Ordering** values are the same are compatible.

TABLE 5.12 Entries in a **CIDSystemInfo** dictionary

KEY	TYPE	VALUE
Registry	string	<i>(Required)</i> A string identifying the issuer of the character collection—for example, Adobe. For information about assigning a registry identifier, consult the ASN Developer Program Web site or contact the Adobe Solutions Network (see the Bibliography).
Ordering	string	<i>(Required)</i> A string that uniquely names the character collection within the specified registry—for example, Japan1.
Supplement	integer	<i>(Required)</i> The <i>supplement number</i> of the character collection. An original character collection has a supplement number of 0. Whenever additional CIDs are assigned in a character collection, the supplement number is increased. Supplements do not alter the ordering of existing CIDs in the character collection. This value is not used in determining compatibility between character collections.

In a CIDFont, the **CIDSystemInfo** entry is a dictionary that specifies the CIDFont’s character collection. Note that the CIDFont need not contain glyph descriptions for all the CIDs in a collection; it can contain a subset. In a CMap, the **CIDSystemInfo** entry is either a single dictionary or an array of dictionaries, depending on whether it associates codes with a single character collection or with multiple character collections; see Section 5.6.4, “CMaps.”

For proper behavior, the **CIDSystemInfo** entry of a CMap should be compatible with that of the CIDFont or CIDFonts with which it is used. If they are incompatible, the effects produced will be unpredictable.

5.6.3 CIDFonts

A CIDFont program contains glyph descriptions that are accessed using a CID as the character selector. There are two types of CIDFont. A Type 0 CIDFont contains glyph descriptions based on Adobe’s Type 1 font format, whereas those in a Type 2 CIDFont are based on the TrueType font format.

A CIDFont dictionary is a PDF object that contains information about a CIDFont program. Although its **Type** value is **Font**, a CIDFont is not actually a font. It does not have an **Encoding** entry, it cannot be listed in the **Font** subdictionary of a resource dictionary, and it cannot be used as the operand of the **Tf** operator. It is used only as a descendant of a Type 0 font. The CMap in the Type 0 font is what defines the encoding that maps character codes to CIDs in the CIDFont. Table 5.13 lists the entries in a CIDFont dictionary.

TABLE 5.13 Entries in a CIDFont dictionary

KEY	TYPE	VALUE
Type	name	<i>(Required)</i> The type of PDF object that this dictionary describes; must be Font for a CIDFont dictionary.
Subtype	name	<i>(Required)</i> The type of CIDFont; CIDFontType0 or CIDFontType2 .
BaseFont	name	<i>(Required)</i> The PostScript name of the CIDFont. For Type 0 CIDFonts, this is usually the value of the CIDFontName entry in the CIDFont program. For Type 2 CIDFonts, it is derived the same way as for a simple TrueType font; see Section 5.5.2, “TrueType Fonts.” In either case, the name can have a subset prefix if appropriate; see Section 5.5.3, “Font Subsets.”
CIDSystemInfo	dictionary	<i>(Required)</i> A dictionary containing entries that define the character collection of the CIDFont. See Table 5.12 on page 337.
FontDescriptor	dictionary	<i>(Required; must be an indirect reference)</i> A font descriptor describing the CIDFont’s default metrics other than its glyph widths (see Section 5.7, “Font Descriptors”).
DW	integer	<i>(Optional)</i> The default width for glyphs in the CIDFont (see “Glyph Metrics in CIDFonts” on page 340). Default value: 1000.
W	array	<i>(Optional)</i> A description of the widths for the glyphs in the CIDFont. The array’s elements have a variable format that can specify individual widths for consecutive CIDs or one width for a range of CIDs (see “Glyph Metrics in CIDFonts” on page 340). Default value: none (the DW value is used for all glyphs).
DW2	array	<i>(Optional; applies only to CIDFonts used for vertical writing)</i> An array of two numbers specifying the default metrics for vertical writing (see “Glyph Metrics in CIDFonts” on page 340). Default value: [880 –1000].
W2	array	<i>(Optional; applies only to CIDFonts used for vertical writing)</i> A description of the metrics for vertical writing for the glyphs in the CIDFont (see “Glyph Metrics in CIDFonts” on page 340). Default value: none (the DW2 value is used for all glyphs).

CIDToGIDMap	stream or name	<p><i>(Optional; Type 2 CIDFonts only)</i> A specification of the mapping from CIDs to glyph indices. If the value is a stream, the bytes in the stream contain the mapping from CIDs to glyph indices: the glyph index for a particular CID value c is a 2-byte value stored in bytes $2 \times c$ and $2 \times c + 1$, where the first byte is the high-order byte. If the value of CIDToGIDMap is a name, it must be Identity, indicating that the mapping between CIDs and glyph indices is the identity mapping. Default value: Identity.</p> <p>This entry may appear only in a Type 2 CIDFont whose associated TrueType font program is embedded in the PDF file (see the next section).</p>
--------------------	-------------------	--

Glyph Selection in CIDFonts

Type 0 and Type 2 CIDFonts handle the mapping from CIDs to glyph descriptions in somewhat different ways.

For Type 0, the CIDFont program itself contains glyph descriptions that are identified by CIDs. The CIDFont program identifies the character collection by a **CIDSystemInfo** dictionary, which should simply be copied into the PDF CIDFont dictionary. CIDs are interpreted uniformly in all CIDFont programs supporting a given character collection, whether the program is embedded in the PDF file or obtained from an external source.

For Type 2, the CIDFont program is actually a TrueType font program, which has no native notion of CIDs. In a TrueType font program, glyph descriptions are identified by *glyph index* values. Glyph indices are internal to the font and are not defined consistently from one font to another. Instead, a TrueType font program contains a “cmap” table that provides mappings directly from character codes to glyph indices for one or more predefined encodings.

TrueType font programs are integrated with the CID-keyed font architecture in one of two ways, depending on whether the font program is embedded in the PDF file.

- If the TrueType font program is embedded, the Type 2 CIDFont dictionary must contain a **CIDToGIDMap** entry that maps CIDs to the glyph indices for the appropriate glyph descriptions in that font program.
- If the TrueType font program is not embedded but is referenced by name, the Type 2 CIDFont dictionary must *not* contain a **CIDToGIDMap** entry, since it is not meaningful to refer to glyph indices in an external font program. In this

case, CIDs do not participate in glyph selection, and only predefined CMaps may be used with this CIDFont (see Section 5.6.4, “CMaps”). The viewer application selects glyphs by translating characters from the encoding specified by the predefined CMap to one of the encodings given in the TrueType font’s “cmap” table. The means by which this is accomplished are implementation-dependent.

Even though the CIDs are sometimes not used to select glyphs in a Type 2 CIDFont, they are always used to determine the glyph metrics, as described in the next section.

Every CIDFont must contain a glyph description for CID 0, which is analogous to the .notdef character name in simple fonts (see “Handling Undefined Characters” on page 355).

Glyph Metrics in CIDFonts

As discussed in Section 5.1.3, “Glyph Positioning and Metrics,” the *width* of a glyph refers to the horizontal displacement between the origin of the glyph and the origin of the next glyph when writing in horizontal mode. In this mode, the vertical displacement between origins is always 0. Widths for a CIDFont are defined using the **DW** and **W** entries in the CIDFont dictionary. These widths must be consistent with the actual widths given in the CIDFont program itself. (See implementation note 43 in Appendix H.)

The **DW** entry defines the default width, which is used for all glyphs whose widths are not specified individually. This entry is particularly useful for Chinese, Japanese, and Korean fonts, in which many of the glyphs have the same width.

The **W** array allows the definition of widths for individual CIDs. The elements of the array are organized in groups of two or three, where each group is in one of the following two formats:

$$c [w_1 \ w_2 \ \dots \ w_n]$$

$$c_{first} \ c_{last} \ w$$

In the first format, c is an integer specifying a starting CID value; it is followed by an array of n numbers that specify the widths for n consecutive CIDs, starting with c . The second format defines the same width, w , for all CIDs in the range c_{first} to c_{last} .

The following is an example of a **W** entry:

```
/W [ 120 [400 325 500]
      7080 8032 1000
    ]
```

In this example, the glyphs for the characters having CIDs 120, 121, and 122 are 400, 325, and 500 units wide, respectively. CIDs in the range 7080 through 8032 all have a width of 1000 units.

Glyphs from a CIDFont can be shown in vertical writing mode. (This is selected by the **WMode** entry in the associated CMap dictionary; see Section 5.6.4, “CMaps.”) To be used in this way, the CIDFont must define the vertical displacement for each glyph and the position vector that relates the horizontal and vertical writing origins.

The default position vector and vertical displacement vector are specified by the **DW2** entry in the CIDFont dictionary. **DW2** is an array of two values: the vertical component of the position vector v and the vertical component of the displacement vector wl (see Figure 5.5 on page 300). The horizontal component of the position vector is always half the glyph width, and that of the displacement vector is always 0. For example, if the **DW2** entry is

```
/DW2 [880 -1000]
```

then a glyph’s position vector and vertical displacement vector are

$$v = (w0 \div 2, 880)$$

$$wl = (0, -1000)$$

where $w0$ is the width (horizontal displacement) for the same glyph. Note that a negative value for the vertical component will place the origin of the next glyph *below* the current glyph, because vertical coordinates in a standard coordinate system increase from bottom to top.

The **W2** array allows the definition of vertical metrics for individual CIDs. The elements of the array are organized in groups of two or five, where each group is in one of the following two formats:

```
c [w11y v1x v1y w12y v2x v2y ...]
cfirst clast w11y v1x v1y
```

In the first format, c is a starting CID and is followed by an array containing numbers interpreted in groups of three. Each group consists of the vertical component of the vertical displacement vector w_1 (whose horizontal component is always 0) followed by the horizontal and vertical components for the position vector v . Successive groups define the vertical metrics for consecutive CIDs starting with c . The second format defines a range of CIDs from c_{first} to c_{last} , followed by three numbers that define the vertical metrics for all CIDs in this range. For example:

```
/W2 [ 120 [-1000 250 772]
      7080 8032 -1000 500 900
    ]
```

This **W2** entry defines the vertical displacement vector for the character with CID 120 as (0, -1000) and the position vector as (250, 772). It also defines the displacement vector for CIDs in the range 7080 through 8032 as (0, -1000) and the position vector as (500, 900).

5.6.4 CMaps

As stated earlier, a CMap specifies the mapping from character codes to character selectors (CIDs, character names, or character codes) in one or more associated fonts or CIDFonts. It serves a function analogous to the **Encoding** dictionary for a simple font. The CMap does not refer directly to specific fonts or CIDFonts; instead, it is combined with them as part of a CID-keyed font, represented in PDF as a Type 0 font dictionary (see Section 5.6.5, “Type 0 Font Dictionaries”).

Within the CMap, the character mappings refer to the associated fonts or CIDFonts by *font number*, indexing from 0. All of the mappings for a particular font number must specify the same kind of character selector. If the character selectors are CIDs, the associated dictionary is expected to be a CIDFont. If the character selectors are names or codes, the associated dictionary is expected to be a font.

Note: As mentioned earlier, PDF versions up to and including PDF 1.4 do not support the entire CID-keyed font architecture. In PDF, a CID-keyed font may have only a single descendant, selected by font number 0, whose characters must be referenced by CID.

A CMap also specifies the writing mode—horizontal or vertical—for any CIDFont with which the CMap is combined. This determines which metrics are to be used when glyphs are painted from that font. (Writing mode is specified as part of the CMap because some glyphs have different shapes when written horizontally and vertically. In that case, the horizontal and vertical variants of a CMap specify different CIDs for a given character code.)

A CMap may be specified in two ways:

- As a name object identifying a predefined CMap, whose definition is known to the viewer application.
- As a stream object whose contents are a CMap file. (See implementation note 48 in Appendix H.)

Predefined CMaps

Table 5.14 lists the names of the predefined CMaps. These CMaps map character codes to CIDs in a single descendant CIDFont. CMaps whose names end in H specify horizontal writing mode; those ending in V specify vertical writing mode.

Note: Several of the CMaps define mappings from Unicode (UCS-2) encodings to character collections. Unicode values appearing in a text string are represented in “big-endian” order (high-order byte first).

TABLE 5.14 Predefined CJK CMap names

NAME	DESCRIPTION
<i>Chinese (Simplified)</i>	
GB–EUC–H	Microsoft Code Page 936 (IfCharSet 0x86), GB 2312-80 character set, EUC-CN encoding
GB–EUC–V	Vertical version of GB–EUC–H
GBpc–EUC–H	Mac OS, GB 2312-80 character set, EUC-CN encoding, Script Manager code 19
GBpc–EUC–V	Vertical version of GBpc–EUC–H
GBK–EUC–H	Microsoft Code Page 936 (IfCharSet 0x86), GBK character set, GBK encoding
GBK–EUC–V	Vertical version of GBK–EUC–H
GBKp–EUC–H	Same as GBK–EUC–H, but replaces half-width Latin characters with proportional forms and maps character code 0x24 to a dollar sign (\$) instead of a yuan symbol (¥)
GBKp–EUC–V	Vertical version of GBKp–EUC–H

GBK2K–H	GB 18030-2000 character set, mixed 1-, 2-, and 4-byte encoding
GBK2K–V	Vertical version of GBK2K–H
UniGB–UCS2–H	Unicode (UCS-2) encoding for the Adobe-GB1 character collection
UniGB–UCS2–V	Vertical version of UniGB–UCS2–H

Chinese (Traditional)

B5pc–H	Mac OS, Big Five character set, Big Five encoding, Script Manager code 2
B5pc–V	Vertical version of B5pc–H
HKscs–B5–H	Hong Kong SCS, an extension to the Big Five character set and encoding
HKscs–B5–V	Vertical version of HKscs–B5–H
ETen–B5–H	Microsoft Code Page 950 (IfCharSet 0x88), Big Five character set with ETen extensions
ETen–B5–V	Vertical version of ETen–B5–H
ETenms–B5–H	Same as ETen–B5–H, but replaces half-width Latin characters with proportional forms
ETenms–B5–V	Vertical version of ETenms–B5–H
CNS–EUC–H	CNS 11643-1992 character set, EUC-TW encoding
CNS–EUC–V	Vertical version of CNS–EUC–H
UniCNS–UCS2–H	Unicode (UCS-2) encoding for the Adobe-CNS1 character collection
UniCNS–UCS2–V	Vertical version of UniCNS–UCS2–H

Japanese

83pv–RKSJ–H	Mac OS, JIS X 0208 character set with KanjiTalk6 extensions, Shift-JIS encoding, Script Manager code 1
90ms–RKSJ–H	Microsoft Code Page 932 (IfCharSet 0x80), JIS X 0208 character set with NEC and IBM extensions
90ms–RKSJ–V	Vertical version of 90ms–RKSJ–H
90msp–RKSJ–H	Same as 90ms–RKSJ–H, but replaces half-width Latin characters with proportional forms
90msp–RKSJ–V	Vertical version of 90msp–RKSJ–H
90pv–RKSJ–H	Mac OS, JIS X 0208 character set with KanjiTalk7 extensions, Shift-JIS encoding, Script Manager code 1
Add–RKSJ–H	JIS X 0208 character set with Fujitsu FMR extensions, Shift-JIS encoding
Add–RKSJ–V	Vertical version of Add–RKSJ–H
EUC–H	JIS X 0208 character set, EUC-JP encoding
EUC–V	Vertical version of EUC–H
Ext–RKSJ–H	JIS C 6226 (JIS78) character set with NEC extensions, Shift-JIS encoding
Ext–RKSJ–V	Vertical version of Ext–RKSJ–H

H	JIS X 0208 character set, ISO-2022-JP encoding
V	Vertical version of H
UniJIS–UCS2–H	Unicode (UCS-2) encoding for the Adobe-Japan1 character collection
UniJIS–UCS2–V	Vertical version of UniJIS–UCS2–H
UniJIS–UCS2–HW–H	Same as UniJIS–UCS2–H, but replaces proportional Latin characters with half-width forms
UniJIS–UCS2–HW–V	Vertical version of UniJIS–UCS2–HW–H

Korean

KSC–EUC–H	KS X 1001:1992 character set, EUC-KR encoding
KSC–EUC–V	Vertical version of KSC–EUC–H
KSCms–UHC–H	Microsoft Code Page 949 (IfCharSet 0x81), KS X 1001:1992 character set plus 8822 additional hangul, Unified Hangul Code (UHC) encoding
KSCms–UHC–V	Vertical version of KSCms–UHC–H
KSCms–UHC–HW–H	Same as KSCms–UHC–H, but replaces proportional Latin characters with half-width forms
KSCms–UHC–HW–V	Vertical version of KSCms–UHC–HW–H
KSCpc–EUC–H	Mac OS, KS X 1001:1992 character set with Mac OS KH extensions, Script Manager Code 3
UniKS–UCS2–H	Unicode (UCS-2) encoding for the Adobe-Korea1 character collection
UniKS–UCS2–V	Vertical version of UniKS–UCS2–H

Generic

Identity–H	The horizontal identity mapping for 2-byte CIDs; may be used with CIDFonts using any Registry , Ordering , and Supplement values. It maps 2-byte character codes ranging from 0 to 65,535 to the same 2-byte CID value, interpreted high-order byte first (see below).
Identity–V	Vertical version of Identity–H. The mapping is the same as for Identity–H.

The Identity–H and Identity–V CMaps can be used to refer to characters directly by their CIDs when showing a text string. When the current font is a Type 0 font whose **Encoding** entry is Identity–H or Identity–V, the string to be shown is interpreted as pairs of bytes representing CIDs, high-order byte first. This works with any CIDFont, independently of its character collection. Additionally, when used in conjunction with a Type 2 CIDFont whose **CIDToGIDMap** entry is **Identity**, the 2-byte CIDs values in fact represent glyph indices for the glyph descriptions in the TrueType font program. This works only if the TrueType font program is embedded in the PDF file.

Table 5.15 lists the character collections referenced by the predefined CMaps for the different versions of PDF. A dash (—) indicates that the CMap is not predefined in that PDF version.

TABLE 5.15 Character collections for predefined CMaps, by PDF version

CMAP	PDF 1.2	PDF 1.3	PDF 1.4
<i>Chinese (Simplified)</i>			
GB-EUC-H/V	Adobe-GB1-0	Adobe-GB1-0	Adobe-GB1-0
GBpc-EUC-H	Adobe-GB1-0	Adobe-GB1-0	Adobe-GB1-0
GBpc-EUC-V	—	Adobe-GB1-0	Adobe-GB1-0
GBK-EUC-H/V	—	Adobe-GB1-2	Adobe-GB1-2
GBKp-EUC-H/V	—	—	Adobe-GB1-2
GBK2K-H/V	—	—	Adobe-GB1-4
UniGB-UCS2-H/V	—	Adobe-GB1-2	Adobe-GB1-4
<i>Chinese (Traditional)</i>			
B5pc-H/V	Adobe-CNS1-0	Adobe-CNS1-0	Adobe-CNS1-0
HKscs-B5-H/V	—	—	Adobe-CNS1-3
ETen-B5-H/V	Adobe-CNS1-0	Adobe-CNS1-0	Adobe-CNS1-0
ETenms-B5-H/V	—	Adobe-CNS1-0	Adobe-CNS1-0
CNS-EUC-H/V	Adobe-CNS1-0	Adobe-CNS1-0	Adobe-CNS1-0
UniCNS-UCS2-H/V	—	Adobe-CNS1-0	Adobe-CNS1-3
<i>Japanese</i>			
83pv-RKSJ-H	Adobe-Japan1-1	Adobe-Japan1-1	Adobe-Japan1-1
90ms-RKSJ-H/V	Adobe-Japan1-2	Adobe-Japan1-2	Adobe-Japan1-2
90msp-RKSJ-H/V	—	Adobe-Japan1-2	Adobe-Japan1-2
90pv-RKSJ-H	Adobe-Japan1-1	Adobe-Japan1-1	Adobe-Japan1-1
Add-RKSJ-H/V	Adobe-Japan1-1	Adobe-Japan1-1	Adobe-Japan1-1
EUC-H/V	—	Adobe-Japan1-1	Adobe-Japan1-1

Ext-RKSJ-H/V	Adobe-Japan1-2	Adobe-Japan1-2	Adobe-Japan1-2
H/V	Adobe-Japan1-1	Adobe-Japan1-1	Adobe-Japan1-1
UniJIS-UCS2-H/V	—	Adobe-Japan1-2	Adobe-Japan1-4
UniJIS-UCS2-HW-H/V	—	Adobe-Japan1-2	Adobe-Japan1-4
<i>Korean</i>			
KSC-EUC-H/V	Adobe-Korea1-0	Adobe-Korea1-0	Adobe-Korea1-0
KSCms-UHC-H/V	Adobe-Korea1-1	Adobe-Korea1-1	Adobe-Korea1-1
KSCms-UHC-HW-H/V	—	Adobe-Korea1-1	Adobe-Korea1-1
KSCpc-EUC-H	Adobe-Korea1-0	Adobe-Korea1-0	Adobe-Korea1-0
UniKS-UCS2-H/V	—	Adobe-Korea1-1	Adobe-Korea1-1

As noted in Section 5.6.2, “CIDSystemInfo Dictionaries,” a character collection is identified by registry, ordering, and supplement number, and supplements are cumulative; that is, a higher-numbered supplement includes the CIDs contained in lower-numbered supplements, as well as some additional CIDs. Consequently, text encoded according to the predefined CMaps for a given PDF version will be valid when interpreted by a viewer application supporting the same or a later PDF version. When interpreted by a viewer supporting an earlier PDF version, such text will cause an error if a CMap is encountered that is not predefined for that PDF version. If character codes are encountered that were added in a higher-numbered supplement than the one corresponding to the supported PDF version, no characters will be displayed for those codes; see “Handling Undefined Characters” on page 355.

Note: *If an application producing a PDF file encounters text to be included that uses characters from a higher-numbered supplement than the one corresponding to the PDF version being generated, the application should embed the CMap for the higher-numbered supplement rather than refer to the predefined CMap (see the next section).*

The CMap programs that define the predefined CMaps are available through the ASN Developer Program Web site and are also provided in conjunction with the book *CJKV Information Processing*, by Ken Lunde. Details on the character collec-

tions, including sample glyphs for all the CIDs, can be found in a number of Adobe Technical Notes; for more information about these Notes and the aforementioned book, see the Bibliography.

Embedded CMap Files

For character encodings that are not predefined, the PDF file must contain a stream that defines the CMap. In addition to the standard entries for streams (listed in Table 3.4 on page 38), the CMap stream dictionary contains the entries listed in Table 5.16. The data in the stream defines the mapping from character codes to a font number and a character selector. The data must follow the syntax defined in Adobe Technical Note #5014, *Adobe CMap and CIDFont Files Specification*.

CMap Example and Operator Summary

The following example of a CMap stream object illustrates and partially explains the contents of a CMap file. This is fully documented in Adobe Technical Note #5014, *Adobe CMap and CIDFont Files Specification*. There are several reasons for including this material here:

- It documents some restrictions on the contents of a CMap file that can be embedded in a PDF file.
- It provides background to aid in understanding subsequent material, particularly “CMap Mapping” on page 354.
- It is the basis for a PDF feature, the **ToUnicode** CMap, which is a minor extension of the CMap file format. This extension is described in Section 5.9, “ToUnicode CMaps.”

Example 5.11 is a sample CMap for a Japanese Shift-JIS encoding. Character codes in this encoding can be either 1 or 2 bytes in length. This CMap maps all character codes to CIDs in font number 0. It could be used with a CIDFont that uses the same CID ordering as specified in the **CIDSystemInfo** entry. Note that several of the entries in the stream dictionary are also replicated in the stream data itself.

TABLE 5.16 Additional entries in a CMap dictionary

KEY	TYPE	VALUE
Type	name	<i>(Required)</i> The type of PDF object that this dictionary describes; must be CMap for a CMap dictionary. (Note that although this object is the value of an entry named Encoding in a Type 0 font, its type is CMap .)
CMapName	name	<i>(Required)</i> The PostScript name of the CMap. This should be the same as the value of CMapName in the CMap file itself.
CIDSystemInfo	dictionary or array	<i>(Required)</i> A dictionary or array containing entries that define the character collection for the CIDFont or CIDFonts associated with the CMap. If the CMap selects only font number 0 and specifies character selectors that are CIDs, this entry can be a dictionary identifying the character collection for the associated CIDFont. Otherwise, it is an array indexed by the font number. If the character selectors for a given font number are CIDs, the corresponding array element is a dictionary identifying the character collection for the associated CIDFont. If the character selectors are names or codes (to be used with an associated font, not a CIDFont), the array element should be null . For details of the CIDSystemInfo dictionaries, see Section 5.6.2, “CIDSystemInfo Dictionaries.” <i>Note:</i> In all PDF versions up to and including PDF 1.4, CIDSystemInfo must be either a dictionary or a one-element array containing a dictionary. The value of this entry should be the same as the value of CIDSystemInfo in the CMap file itself.
WMode	integer	<i>(Optional)</i> A code that determines the writing mode for any CIDFont with which this CMap is combined: 0 Horizontal 1 Vertical Default value: 0. The value of this entry should be the same as the value of WMode in the CMap file itself.
UseCMap	name or stream	<i>(Optional)</i> The name of a predefined CMap, or a stream containing a CMap, that is to be used as the base for this CMap. This allows the CMap to be defined differentially, specifying only the character mappings that differ from the base CMap.

Example 5.11

```
22 0 obj
  << /Type /CMap
    /CMapName /90ms-RKSJ-H
    /CIDSystemInfo << /Registry (Adobe)
                      /Ordering (Japan1)
                      /Supplement 2
                    >>
    /WMode 0
    /Length 23 0 R
  >>

stream
%!PS-Adobe-3.0 Resource-CMap
%%DocumentNeededResources: ProcSet (CIDInit)
%%IncludeResource: ProcSet (CIDInit)
%%BeginResource: CMap (90ms-RKSJ-H)
%%Title: (90ms-RKSJ-H Adobe Japan1 2)
%%Version: 10.001
%%Copyright: Copyright 1990-2001 Adobe Systems Inc.
%%Copyright: All Rights Reserved.
%%EndComments

/CIDInit /ProcSet findresource begin
12 dict begin
begincmap
/CIDSystemInfo
3 dict dup begin
/Registry (Adobe) def
/Ordering (Japan1) def
/Supplement 2 def
end def

/CMapName /90ms-RKSJ-H def
/CMapVersion 10.001 def
/CMapType 1 def
/UIOffset 950 def
/XUID [1 10 25343] def
/WMode 0 def

4 begincodespacerange
<00> <80>
<8140> <9FFC>
<A0> <DF>
<E040> <FCFC>
endcodespacerange
```

```

1 beginnotdefrange
<00> <1F> 231
endnotdefrange

100 begincidrange
<20> <7D> 231
<7E> <7E> 631
<8140> <817E> 633
<8180> <81AC> 696
<81B8> <81BF> 741
<81C8> <81CE> 749
...Additional ranges...
<FB40> <FB7E> 8518
<FB80> <FBFC> 8581
<FC40> <FC4B> 8706
endcidrange
endcmap
CMapName currentdict /CMap defineresource pop
end
end
%%EndResource
%%EOF
endstream
endobj

```

As can be seen from this example, a CMap file conforms to PostScript language syntax; however, a full PostScript interpreter is not needed to interpret it. Aside from some required boilerplate, the CMap file consists of one or more occurrences of several special CMap construction operators, invoked in a specific order. The following is a summary of these operators.

- **begincmap** and **endcmap** enclose the CMap definition.
- **usecmap** incorporates the code mappings from another CMap file. In PDF, the other CMap must also be identified in the **UseCMap** entry in the CMap dictionary (see Table 5.16 on page 349).
- **begincodespacerange** and **endcodespacerange** define *codespace ranges*—the valid input character code ranges—by specifying a pair of codes of some particular length giving the lower and upper bounds of each range; see “CMap Mapping” on page 354.

- **usefont** specifies a font number that is an implicit operand of all the character code mapping operations that follow. In PDF versions up to and including PDF 1.4, the font number must be 0; since this is the default, **usefont** typically does not appear at all.
- **beginbfchar** and **endbfchar** define mappings of individual input character codes to character codes or character names in the associated font. **beginbfrange** and **endbfrange** do the same, but for ranges of input codes. In PDF versions up to and including PDF 1.4, these operators may not appear in a CMap that is used as the **Encoding** entry of a Type 0 font; however, they may appear in the definition of a **ToUnicode** CMap (see Section 5.9, “ToUnicode CMaps”).
- **begincidchar** and **endcidchar** define mappings of individual input character codes to CIDs in the associated CIDFont. **begincidrange** and **endcidrange** do the same, but for ranges of input codes.
- **beginnotdefchar**, **endnotdefchar**, **beginnotdefrange**, and **endnotdefrange** define notdef mappings from character codes to CIDs. As described in the section “Handling Undefined Characters” on page 355, a notdef mapping is used if the normal mapping produces a CID for which no glyph is present in the associated CIDFont.

The **beginrearrangedfont**, **endrearrangedfont**, **beginusematrix**, and **endusematrix** operators, described in Adobe Technical Note #5014, *Adobe CMap and CIDFont Files Specification*, cannot be used in CMap files embedded in a PDF file.

5.6.5 Type 0 Font Dictionaries

A Type 0 font dictionary contains the entries listed in Table 5.17.

Example 5.12 shows a Type 0 font that refers to a single CIDFont. The CMap used is one of the predefined CMaps listed in Table 5.14 on page 343, and is referred to by name.

TABLE 5.17 Entries in a Type 0 font dictionary

KEY	TYPE	VALUE
Type	name	<i>(Required)</i> The type of PDF object that this dictionary describes; must be Font for a font dictionary.
Subtype	name	<i>(Required)</i> The type of font; must be Type0 for a Type 0 font.
BaseFont	name	<i>(Required)</i> The PostScript name of the font. In principle, this is an arbitrary name, since there is no font program associated directly with a Type 0 font dictionary. The conventions described here ensure maximum compatibility with existing Acrobat products. If the descendant is a Type 0 CIDFont, this name should be the concatenation of the CIDFont's BaseFont name, a hyphen, and the CMap name given in the Encoding entry (or the CMapName entry in the CMap program itself). If the descendant is a Type 2 CIDFont, this name should be the same as the CIDFont's BaseFont name.
Encoding	name or stream	<i>(Required)</i> The name of a predefined CMap, or a stream containing a CMap program, that maps character codes to font numbers and CIDs. If the descendant is a Type 2 CIDFont whose associated TrueType font program is not embedded in the PDF file, the Encoding entry must be a predefined CMap name (see “Glyph Selection in CIDFonts” on page 339).
DescendantFonts	array	<i>(Required)</i> An array specifying one or more fonts or CIDFonts that are descendants of this composite font. This array is indexed by the font number that is obtained by mapping a character code through the CMap specified in the Encoding entry. <i>Note:</i> In all PDF versions up to and including PDF 1.4, DescendantFonts must be a one-element array containing a CIDFont dictionary.
ToUnicode	stream	<i>(Optional)</i> A stream containing a CMap file that maps character codes to Unicode values (see Section 5.9, “ToUnicode CMaps”).

Example 5.12

```

14 0 obj
  << /Type /Font
    /Subtype /Type0
    /BaseFont /HeiseiMin-W5-90ms-RKSJ-H
    /Encoding /90ms-RKSJ-H
    /DescendantFonts [15 0 R]
  >>
endobj

```

CMap Mapping

When the current font is a Type 0 font, the text-showing operators (such as **Tj**) interpret the bytes in the string to be shown according to the CMap specified as the **Encoding** entry of the Type 0 font dictionary. The following paragraphs describe how the characters in the string are decoded and mapped into character selectors (which in PDF 1.3 must always be CIDs).

The number of bytes extracted from the string for each successive character is determined exclusively by the codespace ranges in the CMap (delimited by **begincodespacerange** and **endcodespacerange**). A codespace range is specified by a pair of codes of some particular length giving the lower and upper bounds of that range. A code is considered to match the range if it is the same length as the bounding codes and the value of each of its bytes lies between the corresponding bytes of the lower and upper bounds. The code length cannot exceed the number of bytes representable in an integer (see Appendix C).

A sequence of one or more bytes is extracted from the string and matched against the codespace ranges in the CMap. That is, the first byte is matched against 1-byte codespace ranges; if no match is found, a second byte is extracted, and the 2-byte code is matched against 2-byte codespace ranges. This continues for successively longer codes until a match is found or all codespace ranges have been tested. There will be at most one match, since codespace ranges do not overlap.

The code extracted from the string is then looked up in the character code mappings for codes of that length. (These are the mappings defined by **beginbfchar**, **endbfchar**, **begincidchar**, **endcidchar**, and corresponding operators for ranges.) Failing that, it is looked up in the notdef mappings, as described in the next section.

The results of the CMap mapping algorithm are a font number and a character selector. In PDF 1.3, the font number must always be 0 and the character selector must always be a CID; this is the only case described here. The font number is used as an index into the Type 0 font's **DescendantFonts** array, selecting a CIDFont. The CID is then used to select a glyph in the CIDFont. If the CIDFont contains no glyph for that CID, the notdef mappings are consulted, as described in the next section.

Handling Undefined Characters

A CMap mapping operation can fail to select a glyph for any of a variety of reasons. This section describes what happens when that occurs.

If a code maps to a CID for which there is no such glyph in the descendant CIDFont, the *notdef mappings* in the CMap are consulted to obtain a substitute character selector. These mappings (so called by analogy with the `.notdef` character mechanism in simple fonts) are delimited by the operators **beginnotdefchar**, **endnotdefchar**, **beginnotdefrange**, and **endnotdefrange**; they always map to a CID. If a matching notdef mapping is found, the CID selects a glyph in the associated descendant, which must be a CIDFont. If there is no glyph for that CID, the glyph for CID 0 (which is required to be present) is substituted.

If the CMap does not contain either a character mapping or a notdef mapping for the code, descendant 0 is selected and the glyph for CID 0 is substituted from the associated CIDFont.

If the code is invalid—that is, the bytes extracted from the string to be shown do not match any codespace range in the CMap—a substitute glyph is chosen as just described. The character mapping algorithm is reset to its original position in the string, and a modified mapping algorithm chooses the best partially matching codespace range, as follows:

1. If the first byte extracted from the string to be shown does not match the first byte of any codespace range, the range having the shortest codes is chosen.
2. Otherwise (that is, if there is a partial match), for each additional byte extracted, the code accumulated so far is matched against the beginnings of all longer codespace ranges until the longest such partial match has been found. If multiple codespace ranges have partial matches of the same length, the one having the shortest codes is chosen.

The length of the codes in the chosen codespace range determines the total number of bytes to consume from the string for the current mapping operation.

5.7 Font Descriptors

A *font descriptor* specifies metrics and other attributes of a simple font or a CIDFont as a whole, as distinct from the metrics of individual glyphs. These font

metrics provide information that enables a viewer application to synthesize a substitute font or select a similar font when the font program is unavailable. The font descriptor may also be used to embed the font program in the PDF file. (Font descriptors are not used with Type 0 or Type 3 fonts.)

A font descriptor is a dictionary whose entries specify various font attributes. The entries common to all font descriptors—for both simple fonts and CIDFonts—are listed in Table 5.18; additional entries in the font descriptor for a CIDFont are described in Section 5.7.2, “Font Descriptors for CIDFonts.” All integer values are units in glyph space. The conversion from glyph space to text space is described in Section 5.1.3, “Glyph Positioning and Metrics.”

TABLE 5.18 Entries common to all font descriptors

KEY	TYPE	VALUE
Type	name	<i>(Required)</i> The type of PDF object that this dictionary describes; must be FontDescriptor for a font descriptor.
FontName	name	<i>(Required)</i> The PostScript name of the font. This should be the same as the value of BaseFont in the font or CIDFont dictionary that refers to this font descriptor.
Flags	integer	<i>(Required)</i> A collection of flags defining various characteristics of the font (see Section 5.7.1, “Font Descriptor Flags”).
FontBBox	rectangle	<i>(Required)</i> A rectangle (see Section 3.8.3, “Rectangles”), expressed in the glyph coordinate system, specifying the <i>font bounding box</i> . This is the smallest rectangle enclosing the shape that would result if all of the glyphs of the font were placed with their origins coincident and then filled.
ItalicAngle	number	<i>(Required)</i> The angle, expressed in degrees counterclockwise from the vertical, of the dominant vertical strokes of the font. (For example, the 9-o’clock position is 90 degrees, and the 3-o’clock position is –90 degrees.) The value is negative for fonts that slope to the right, as almost all italic fonts do.
Ascent	number	<i>(Required)</i> The maximum height above the baseline reached by glyphs in this font, excluding the height of glyphs for accented characters.
Descent	number	<i>(Required)</i> The maximum depth below the baseline reached by glyphs in this font. The value is a negative number.
Leading	number	<i>(Optional)</i> The desired spacing between baselines of consecutive lines of text. Default value: 0.
CapHeight	number	<i>(Required)</i> The vertical coordinate of the top of flat capital letters, measured from the baseline.

XHeight	number	<i>(Optional)</i> The font’s <i>x height</i> : the vertical coordinate of the top of flat non-ascending lowercase letters (like the letter <i>x</i>), measured from the baseline. Default value: 0.
StemV	number	<i>(Required)</i> The thickness, measured horizontally, of the dominant vertical stems of glyphs in the font.
StemH	number	<i>(Optional)</i> The thickness, measured invertically, of the dominant horizontal stems of glyphs in the font. Default value: 0.
AvgWidth	number	<i>(Optional)</i> The average width of glyphs in the font. Default value: 0.
MaxWidth	number	<i>(Optional)</i> The maximum width of glyphs in the font. Default value: 0.
MissingWidth	number	<i>(Optional)</i> The width to use for character codes whose widths are not specified in a font dictionary’s Widths array. This has a predictable effect only if all such codes map to glyphs whose actual widths are the same as the Missing-Width value. Default value: 0.
FontFile	stream	<i>(Optional)</i> A stream containing a Type 1 font program (see Section 5.8, “Embedded Font Programs”).
FontFile2	stream	<i>(Optional; PDF 1.1)</i> A stream containing a TrueType font program (see Section 5.8, “Embedded Font Programs”).
FontFile3	stream	<i>(Optional; PDF 1.2)</i> A stream containing a font program other than Type 1 or TrueType. The format of the font program is specified by the Subtype entry in the stream dictionary (see Section 5.8, “Embedded Font Programs,” and implementation note 49 in Appendix H).
		At most, only one of the FontFile , FontFile2 , and FontFile3 entries may be present.
CharSet	string	<i>(Optional; meaningful only in Type 1 fonts; PDF 1.1)</i> A string listing the character names defined in a font subset. The names in this string must be in PDF syntax—that is, each name preceded by a slash (/). The names can appear in any order. The name <code>.notdef</code> should be omitted; it is assumed to exist in the font subset. If this entry is absent, the only indication of a font subset is the subset tag in the FontName entry (see Section 5.5.3, “Font Subsets”).

5.7.1 Font Descriptor Flags

The value of the **Flags** entry in a font descriptor is an unsigned 32-bit integer containing flags specifying various characteristics of the font. Bit positions within the flag word are numbered from 1 (low-order) to 32 (high-order). Table 5.19 shows the meanings of the flags; all undefined flag bits are reserved and must be set to 0. Figure 5.13 shows examples of fonts with these characteristics.

TABLE 5.19 Font flags

BIT POSITION	NAME	MEANING
1	FixedPitch	All glyphs have the same width (as opposed to proportional or variable-pitch fonts, which have different widths).
2	Serif	Glyphs have serifs, which are short strokes drawn at an angle on the top and bottom of glyph stems (as opposed to <i>sans serif</i> fonts, which do not).
3	Symbolic	Font contains characters outside the Adobe standard Latin character set. This flag and the Nonsymbolic flag cannot both be set or both be clear (see below).
4	Script	Glyphs resemble cursive handwriting.
6	Nonsymbolic	Font uses the Adobe standard Latin character set or a subset of it (see below).
7	Italic	Glyphs have dominant vertical strokes that are slanted.
17	AllCap	Font contains no lowercase letters; typically used for display purposes such as titles or headlines.
18	SmallCap	Font contains both uppercase and lowercase letters. The uppercase letters are similar to ones in the regular version of the same typeface family. The glyphs for the lowercase letters have the same shapes as the corresponding uppercase letters, but they are sized and their proportions adjusted so that they have the same size and stroke weight as lowercase glyphs in the same typeface family.
19	ForceBold	See below.

The Nonsymbolic flag (bit 6 in the **Flags** entry) indicates that the font’s character set is the Adobe standard Latin character set (or a subset of it) and that it uses the standard names for those characters. The characters in this character set are shown in Section D.1, “Latin Character Set and Encodings.” If the font uses any characters outside this set, the Symbolic flag should be set and the Nonsymbolic flag clear; in other words, any font whose character set is not a subset of the Adobe standard character set is considered to be symbolic. This influences the font’s implicit base encoding and may affect a viewer application’s font substitution strategies.

Fixed-pitch font	The quick brown fox jumped.
Serif font	The quick brown fox jumped.
Sans serif font	The quick brown fox jumped.
Symbolic font	* ** □ ◆ * * * * * * □ □ □ ■ * □ * ◆ ○ □ * * * * * \
Script font	<i>The quick brown fox jumped.</i>
Italic font	<i>The quick brown fox jumped.</i>
All-cap font	THE QUICK BROWN FOX JUMPED
Small-cap font	THE QUICK BROWN FOX JUMPED.

FIGURE 5.13 Characteristics represented in the **Flags** entry of a font descriptor

Note: This classification of nonsymbolic and symbolic fonts is peculiar to PDF. A font may contain additional characters that are used in Latin writing systems but are outside the Adobe standard Latin character set; PDF considers such a font to be symbolic. The use of two flags to represent a single binary choice is a historical accident.

The ForceBold flag (bit 19) determines whether bold glyphs are painted with extra pixels even at very small text sizes. Typically, when glyphs are painted at small sizes on very low-resolution devices such as display screens, features of bold glyphs may appear only 1 pixel wide. Because this is the minimum feature width on a pixel-based device, ordinary (nonbold) glyphs also appear with 1-pixel-wide features, and so cannot be distinguished from bold glyphs. If the ForceBold flag is set, features of bold glyphs may be thickened at small text sizes.

Example 5.13 illustrates a font descriptor whose **Flags** entry has the Serif, Nonsymbolic, and ForceBold flags (bits 2, 6, and 19) set.

Example 5.13

```
7 0 obj
  << /Type /FontDescriptor
    /FontName /AGaramond-Semibold
    /Flags 262178 % Bits 2, 6, and 19
    /FontBBox [-177 -269 1123 866]
    /MissingWidth 255
    /StemV 105
    /StemH 45
    /CapHeight 660
    /XHeight 394
    /Ascent 720
    /Descent -270
    /Leading 83
    /MaxWidth 1212
    /AvgWidth 478
    /ItalicAngle 0
  >>
endobj
```

5.7.2 Font Descriptors for CIDFonts

In addition to the entries in Table 5.18 on page 356, the **FontDescriptor** dictionaries of CIDFonts may contain the entries listed in Table 5.20.

Style

The **Style** dictionary contains entries that define style attributes and values for the CIDFont. Currently, only the **Panose** entry is defined. The value of **Panose** is a 12-byte string consisting of the following:

- The font family class and subclass ID bytes, given in the `sFamilyClass` field of the “OS/2” table in a TrueType font. This field is documented in Microsoft’s *TrueType 1.0 Font Files Technical Specification*.
- Ten bytes for the PANOSE™ classification number for the font. The PANOSE classification system is documented in Hewlett-Packard Company’s *PANOSE Classification Metrics Guide*.

See the Bibliography for more information about these documents.

TABLE 5.20 Additional font descriptor entries for CIDFonts

KEY	TYPE	VALUE
Style	dictionary	<i>(Optional)</i> A dictionary containing entries that describe the style of the glyphs in the font (see “Style,” above).
Lang	name	<i>(Optional)</i> A name specifying the language of the font, used for encodings where the language is not implied by the encoding itself. The possible values are the 2-character language codes defined by ISO 639—for example, en for English and ja for Japanese. The complete list of these codes be obtained from the International Organization for Standardization (see the Bibliography).
FD	dictionary	<i>(Optional)</i> A dictionary whose keys identify a class of characters in a CIDFont. Each value is a dictionary containing entries that override the corresponding values in the main font descriptor dictionary for that class of characters (see “FD,” below).
CIDSet	stream	<i>(Optional)</i> A stream identifying which CIDs are present in the CIDFont file. If this entry is present, the CIDFont contains only a subset of the glyphs in the character collection defined by the CIDSystemInfo dictionary. If it is absent, the only indication of a CIDFont subset is the subset tag in the FontName entry (see Section 5.5.3, “Font Subsets”). The stream’s data is organized as a table of bits indexed by CID. The bits should be stored in bytes with the high-order bit first. Each bit corresponds to a CID. The first bit of the first byte corresponds to CID 0, the next bit to CID 1, and so on.

The following is an example of a **Style** entry in the font descriptor:

```
/Style << /Panose <01 05 02 02 03 00 00 00 00 00 00 00> >>
```

FD

A CIDFont may be made up of different classes of characters, each class requiring different sets of the fontwide attributes that appear in font descriptors. Latin characters, for example, may require different attributes than kanji characters. The font descriptor defines a set of default attributes that apply to all characters in the CIDFont; the **FD** entry in the font descriptor contains exceptions to these defaults.

The key for each entry in an **FD** dictionary is the name of a class of characters—that is, a particular subset of the CIDFont’s character collection. The entry’s value is a font descriptor whose contents are to override the fontwide attributes for that class only. This font descriptor should contain entries for metric information only; it should not include **FontFile**, **FontFile2**, **FontFile3**, or any of the entries listed in Table 5.20.

It is strongly recommended that the **FD** dictionary contain at least the metrics for the proportional Latin characters. With the information for these characters, a more accurate substitution font can be created.

The names of the character classes depend on the character collection, as identified by the **Registry**, **Ordering**, and **Supplement** entries in the **CIDSystemInfo** dictionary. Table 5.21 lists the valid keys for the Adobe-GB1, Adobe-CNS1, Adobe-Japan1, Adobe-Japan2, and Adobe-Korea1 character collections.

TABLE 5.21 Character classes in CJK fonts

CHARACTER COLLECTION	CLASS	CHARACTERS IN CLASS
Adobe-GB1	Alphabetic	Full-width Latin, Greek, and Cyrillic characters
	Dingbats	Special symbols
	Generic	Typeface-independent characters, such as line-drawing
	Hanzi	Full-width hanzi (Chinese) characters
	HRoman	Half-width Latin characters
	HRomanRot	Same as HRoman , but rotated for use in vertical writing
	Kana	Japanese kana (katakana and hiragana) characters
	Proportional	Proportional Latin characters
Adobe-CNS1	ProportionalRot	Same as Proportional , but rotated for use in vertical writing
	Alphabetic	Full-width Latin, Greek, and Cyrillic characters
	Dingbats	Special symbols
	Generic	Typeface-independent characters, such as line-drawing
	Hanzi	Full-width hanzi (Chinese) characters
	HRoman	Half-width Latin characters
	HRomanRot	Same as HRoman , but rotated for use in vertical writing
	Kana	Japanese kana (katakana and hiragana) characters
Proportional	Proportional Latin characters	
	ProportionalRot	Same as Proportional , but rotated for use in vertical writing

Adobe-Japan1	Alphabetic	Full-width Latin, Greek, and Cyrillic characters
	AlphaNum	Numeric characters
	Dingbats	Special symbols
	DingbatsRot	Same as Dingbats , but rotated for use in vertical writing
	Generic	Typeface-independent characters, such as line-drawing
	GenericRot	Same as Generic , but rotated for use in vertical writing
	HKana	Half-width kana (katakana and hiragana) characters
	HKanaRot	Same as HKana , but rotated for use in vertical writing
	HRoman	Half-width Latin characters
	HRomanRot	Same as HRoman , but rotated for use in vertical writing
	Kana	Full-width kana (katakana and hiragana) characters
	Kanji	Full-width kanji (Chinese) characters
	Proportional	Proportional Latin characters
	ProportionalRot	Same as Proportional , but rotated for use in vertical writing
Ruby	Characters used for setting ruby (small characters that serve to annotate other glyphs with meanings or readings)	
Adobe-Japan2	Alphabetic	Full-width Latin, Greek, and Cyrillic characters
	Dingbats	Special symbols
	HojoKanji	Full-width kanji characters
Adobe-Korea1	Alphabetic	Full-width Latin, Greek, and Cyrillic characters
	Dingbats	Special symbols
	Generic	Typeface-independent characters, such as line-drawing
	Hangul	Hangul and jamo characters
	Hanja	Full-width hanja (Chinese) characters
	HRoman	Half-width Latin characters
	HRomanRot	Same as HRoman , but rotated for use in vertical writing
	Kana	Japanese kana (katakana and hiragana) characters
	Proportional	Proportional Latin characters
ProportionalRot	Same as Proportional , but rotated for use in vertical writing	

Example 5.14 illustrates an FD dictionary containing two entries.

Example 5.14

```
/FD << /Proportional 25 0 R
      /HKana 26 0 R
>>
```

```
25 0 obj
  << /Type /FontDescriptor
    /FontName /HeiseiMin-W3-Proportional
    /Flags 2
    /AvgWidth 478
    /MaxWidth 1212
    /MissingWidth 250
    /StemV 105
    /StemH 45
    /CapHeight 660
    /XHeight 394
    /Ascent 720
    /Descent -270
    /Leading 83
  >>
endobj

26 0 obj
  << /Type /FontDescriptor
    /FontName /HeiseiMin-W3-HKana
    /Flags 3
    /Style << /Panose <01 05 02 02 03 00 00 00 00 00 00> >>
    /AvgWidth 500
    /MaxWidth 500
    /MissingWidth 500
    /StemV 50
    /StemH 75
    /Ascent 720
    /Descent 0
    /Leading 83
  >>
endobj
```

5.8 Embedded Font Programs

A font program can be embedded in a PDF file as data contained in a PDF stream object. Such a stream object is also called a *font file*, by analogy with font programs that are available from sources external to the viewer application. (See also implementation note 50 in Appendix H.)

Font programs are subject to copyright, and the copyright owner may impose conditions under which a font program can be used. These permissions are re-

corded either in the font program itself or as part of a separate license. One of the conditions may be that the font program cannot be embedded, in which case it should not be incorporated into a PDF file. A font program may allow embedding for the sole purpose of viewing and printing the document, but not for creating new or modified text using the font (in either the same document or other documents); the latter operation would require the user performing the operation to have a licensed copy of the font program, not a copy extracted from the PDF file. In the absence of explicit information to the contrary, a PDF consumer should assume that any embedded font programs are to be used only to view and print the document and not for any other purposes.

Table 5.22 summarizes the ways in which font programs are embedded in a PDF file, depending on the representation of the font program. The key is the name used in the font descriptor to refer to the font file stream; the subtype is the value of the **Subtype** key, if present, in the font file stream dictionary. Further details of specific font program representations are given below.

TABLE 5.22 Embedded font organization for various font types

KEY	SUBTYPE	DESCRIPTION
FontFile	—	Type 1 font program, in the original (noncompact) format described in <i>Adobe Type 1 Font Format</i> . This entry can appear in the font descriptor for a Type1 or MMType1 font dictionary.
FontFile2	—	(PDF 1.1) TrueType font program, as described in the <i>TrueType Reference Manual</i> . This entry can appear in the font descriptor for a TrueType font dictionary or (in PDF 1.3) for a CIDFontType2 CIDFont dictionary.
FontFile3	Type1C	(PDF 1.2) Type 1–equivalent font program represented in the Compact Font Format (CFF), as described in Adobe Technical Note #5176, <i>The Compact Font Format Specification</i> . This entry can appear in the font descriptor for a Type1 or MMType1 font dictionary.
	CIDFontType0C	(PDF 1.3) Type 0 CIDFont program represented in the Compact Font Format (CFF), as described in Adobe Technical Note #5176, <i>The Compact Font Format Specification</i> . This entry can appear in the font descriptor for a CIDFontType0 CIDFont dictionary.
	<i>otherName</i>	Font or CIDFont program represented in some future format, identified by <i>otherName</i> as the font file subtype.

The stream dictionary for a font file contains the normal entries for a stream, such as **Length** and **Filter** (listed in Table 3.4 on page 38), plus the additional entries listed in Table 5.23.

TABLE 5.23 Additional entries in an embedded font stream dictionary

KEY	TYPE	VALUE
Length1	integer	<i>(Required for Type 1 and TrueType fonts)</i> The length in bytes of the clear-text portion of the Type 1 font program (see below), or the entire TrueType font program, after it has been decoded using the filters specified by the stream's Filter entry, if any.
Length2	integer	<i>(Required for Type 1 fonts)</i> The length in bytes of the encrypted portion of the Type 1 font program (see below) after it has been decoded using the filters specified by the stream's Filter entry.
Length3	integer	<i>(Required for Type 1 fonts)</i> The length in bytes of the fixed-content portion of the Type 1 font program (see below), after it has been decoded using the filters specified by the stream's Filter entry. If Length3 is 0, it indicates that the 512 zeros and cleartomark have not been included in the FontFile font program and must be added.
Subtype	name	<i>(Required if referenced from FontFile3; PDF 1.2)</i> A name specifying the format of the embedded font program. The name must be Type1C for Type 1 compact fonts or CID-FontType0C for Type 0 compact CIDFonts. When additional font formats are added to PDF, more values will be defined for Subtype .
Metadata	stream	<i>(Optional; PDF 1.4)</i> A <i>metadata stream</i> containing metadata for the embedded font program (see Section 9.2.2, "Metadata Streams").

A standard Type 1 font program, as described in the *Adobe Type 1 Font Format* specification, consists of three parts: a clear-text portion (written using PostScript syntax), an encrypted portion, and a fixed-content portion. The fixed-content portion contains 512 ASCII zeros followed by a **cleartomark** operator, and perhaps followed by additional data. While the encrypted portion of a standard Type 1 font may be in binary or ASCII hexadecimal format, PDF supports only the binary format; however, the entire font program may be encoded using any filters.

Example 5.15 shows the structure of an embedded standard Type 1 font.

Example 5.15

```

12 0 obj
  << /Filter /ASCII85Decode
    /Length 41116
    /Length1 2526
    /Length2 32393
    /Length3 570
  >>
stream
,p>`rDKJj'E+LaU0eP.@+AH9dBOu$hFD55nC
...Omitted data...
JJQ&Nt')<=^p&mGf(%:%h1%9c//K(/*o=.C>UXkbVGTrr~>
endstream
endobj

```

As noted in Table 5.22, a Type 1–equivalent font program or a Type 0 CIDFont program can be represented in the Compact Font Format (CFF). The **Length1**, **Length2**, and **Length3** entries are not needed in that case. Although CFF enables multiple font or CIDFont programs to be bundled together in a single file, an embedded CFF font file in PDF must consist of exactly one font or CIDFont (as appropriate for the associated font dictionary).

***Note:** According to the Adobe Type 1 Font Format specification, a Type 1 font program may contain a **PaintType** entry specifying whether the glyphs’ outlines are to be filled or stroked. For fonts embedded in a PDF file, this entry is ignored; the decision whether to fill or stroke glyph outlines is entirely determined by the PDF text rendering mode parameter (see Section 5.2.5, “Text Rendering Mode”). This also applies to Type 1 compact fonts and Type 0 compact CIDFonts.*

A TrueType font program may be used as part of either a font or a CIDFont. Although the basic font file format is the same in both cases, there are different requirements for what information must be present in the font program. The following TrueType tables are always required: “head,” “hhea,” “loca,” “maxp,” “cvt_,” “prep,” “glyf,” “hmtx,” and “fpgm.” If used with a simple font dictionary, the font program must additionally contain a “cmap” table defining one or more encodings, as discussed in “Encodings for TrueType Fonts” on page 332. If used with a CIDFont dictionary, the “cmap” table is not needed, since the mapping from character codes to glyph descriptions is provided separately.

Note: The “*vhea*” and “*vmtx*” tables that specify vertical metrics are never used by a PDF viewer application. The only way to specify vertical metrics in PDF is by means of the **DW2** and **W2** entries in a CIDFont dictionary.

As discussed in Section 5.5.3, “Font Subsets,” an embedded font program may contain only the subset of glyphs that are used in the PDF document. This may be indicated by the presence of a **CharSet** or **CIDSet** entry in the font descriptor that refers to the font file, although subset fonts are not always so identified.

5.9 ToUnicode CMaps

The preceding sections describe all the facilities for showing text and causing the corresponding glyphs to be painted on the page. However, a viewer application sometimes needs to determine the information content of text—that is, its meaning according to some standard character identification, as opposed to its rendered appearance. This need arises during operations such as searching, indexing, and exporting of text to other applications.

The Unicode standard defines a system for numbering all of the common characters used in a large number of languages. It is a suitable scheme for representing the information content of text, but not its appearance, since Unicode values identify characters, not glyphs. For information about Unicode, see the *Unicode Standard* by the Unicode Consortium (see the Bibliography).

If a font’s characters are identified according to a standard character set that is known to the viewer application, the viewer can easily convert text to Unicode values with no additional information. This character identification can occur if either the font uses a standard named encoding or the characters in the font are identified by standard character names or CIDs. A viewer application can identify characters if the font with which the text is shown is defined in any of the following ways:

- It uses one of the predefined encodings **MacRomanEncoding**, **MacExpertEncoding**, **WinAnsiEncoding**, or any of the predefined CMaps listed in Table 5.14 on page 343 except **Identity–H** and **Identity–V**.
- It is a Type 1 font whose character names are taken from the Adobe standard Latin character set and the set of named characters in the Symbol font, documented in Appendix D.

- It is a Type 0 font whose descendant CIDFont uses the Adobe-GB1, Adobe-CNS1, Adobe-Japan1, or Adobe-Korea1 character collection (as specified in the **CIDSystemInfo** dictionary) whose supplement number is the one corresponding to the version of PDF supported by the viewer. See Table 5.15 on page 346 for a list of the character collections corresponding to a given PDF version. (Other supplements of these character collections can be used, but if the supplement is higher-numbered than the one corresponding to the supported PDF version, only the CIDs in the latter supplement are considered to be standard CIDs.)

If a font is not defined in one of these ways, the viewer application has no way to identify characters in the text. The glyphs can still be shown, but the characters cannot be converted to Unicode values without additional information. This information can be provided as an optional **ToUnicode** entry (*PDF 1.2*) in the font dictionary, whose value is a stream object containing a special kind of CMap file that maps character codes to Unicode values. This entry is discussed below; the overall algorithm for mapping character codes to Unicode values is described in more detail under “Unicode Mapping” on page 620.

The CMap defined in the **ToUnicode** entry must follow the syntax for CMaps introduced in Section 5.6.4, “CMaps” and fully documented in Adobe Technical Note #5014, *Adobe CMap and CIDFont Files Specification*. Additional guidance regarding the CMap defined in this entry is provided in Adobe Technical Note #5411, *ToUnicode Mapping File Tutorial*. This CMap differs from an ordinary one in the following ways:

- In the CMap stream dictionary, no CMap-specific entries are required. The only pertinent entry from the list in Table 5.16 on page 349 is **UseCMap**, which may be used if the CMap is based on another **ToUnicode** CMap.
- The CMap file must contain **begincodespacerange** and **endcodespacerange** operators that are consistent with the encoding that the font uses. In particular, for a simple font, the codespace must be one byte long.
- It must use the **beginbfchar**, **endbfchar**, **beginbfrange**, and **endbfrange** operators to define the mapping from character codes to 2-byte Unicode values, interpreted high-order byte first. These operators have been extended to handle cases where a single character code maps to one or more Unicode values, as described below.

Example 5.16 illustrates a Type 0 font that uses the Identity–H CMap to map from character codes to CIDs, and whose descendant CIDFont uses the **Identity** mapping from CIDs to TrueType glyph indices. Text strings shown using this font simply use a 2-byte glyph index for each character. In the absence of a **ToUnicode** entry, there would be no information available about what the characters mean.

Example 5.16

```

14 0 obj
  << /Type /Font
    /Subtype /Type0
    /BaseFont /Ryumin-Light
    /Encoding /Identity-H
    /DescendantFonts [15 0 R]
    /ToUnicode 16 0 R
  >>
endobj

15 0 obj
  << /Type /Font
    /Subtype /CIDFontType2
    /BaseFont /Ryumin-Light
    /CIDSystemInfo 17 0 R
    /FontDescriptor 18 0 R
    /CIDToGIDMap /Identity
  >>
endobj

```

The value of the **ToUnicode** entry is a stream object that contains the definition of the CMap, as shown in Example 5.17.

Example 5.17

```

16 0 obj
  << /Length 433 >>
stream
/CIDInit /ProcSet findresource begin
12 dict begin
begincmap
/CIDSystemInfo
<< /Registry (Adobe)
/Ordering (UCS)
/Supplement 0
>> def

```

```

/CMAPName /Adobe-Identity-UCS def
/CMAPType 2 def
1 begincodespacerange
<0000> <FFFF>
endcodespacerange
2 beginbfrange
<0000> <005E> <0020>
<005F> <0061> [<00660066> <00660069> <00660066006C>]
endbfrange
endcmap
CMAPName currentdict /CMAP defineresource pop
end
end
endstream
endobj

```

The **begincodespacerange** and **endcodespacerange** operators in Example 5.17 define the source character code range to be the 2-byte character codes from <00 00> to <FF FF>. The specific mappings for several of the character codes are shown. For example, <00 00> to <00 5E> are mapped to the Unicode values U+0020 to U+007E (where Unicode values are conventionally written as U+ followed by four hexadecimal digits). This is followed by the definition of a mapping used where each character code represents more than one Unicode value:

```
<005F> <0061> [<00660066> <00660069> <00660066006C>]
```

In this case, the original character codes are the glyph indices for the ligatures ff, fi, and ffi. The entry defines the mapping from the character codes <00 5F>, <00 60>, and <00 61> to the string of Unicode values with a code for each character in the ligature: U+0066 U+0066 are the Unicode values for the character sequence ff, U+0066 U+0069 for fi, and U+0066 U+0066 U+006c for ffi.

Example 5.17 illustrates several extensions to the way destination values can be defined. To support mappings from a source code to a string of destination codes, the following extension has been made to the ranges defined after a **beginbfchar** operator:

```

n beginbfchar
srcCode dstString
endbfchar

```

where *dstString* can be a string of up to 512 bytes. Likewise, mappings after the **beginbfrange** operator may be defined as:

```
n beginbfrange  
srcCode1 srcCode2 dstString  
endbfrange
```

In this case, the last byte of the string will be incremented for each consecutive code in the source code range. When defining ranges of this type, care must be taken to ensure that the value of the last byte in the string is less than or equal to $255 - (srcCode_2 - srcCode_1)$. This ensures that the last byte of the string will not be incremented past 255; otherwise the result of mapping is undefined and an error occurs.

To support more compact representations of mappings from a range of source character codes to a discontinuous range of destination codes, the CMaps used for the **ToUnicode** entry may use the following syntax for the mappings following a **beginbfrange** definition:

```
n beginbfrange  
srcCode1 srcCoden [dstString1 dstString2 ... dstStringn]  
endbfrange
```

Consecutive codes starting with *srcCode*₁ and ending with *srcCode*_{*n*} are mapped to the destination strings in the array starting with *dstString*₁ and ending with *dstString*_{*n*}.

CHAPTER 6

Rendering

THE ADOBE IMAGING MODEL separates *graphics* (the specification of shapes and colors) from *rendering* (controlling a raster output device). Figures 4.12 and 4.13 on pages 174 and 175 illustrate this division. Chapter 4 describes the facilities for specifying the appearance of pages in a device-independent way. This chapter describes the facilities for controlling how shapes and colors are rendered on the raster output device. All of the facilities discussed here depend on the specific characteristics of the output device; PDF documents that are intended to be device-independent should limit themselves to the general graphics facilities described in Chapter 4.

Nearly all of the rendering facilities that are under the control of a PDF document have to do with the reproduction of color. Colors are rendered by a multiple-step process outlined below. (Depending on the current color space and on the characteristics of the device, it is not always necessary to perform every step.)

1. If a color has been specified in a CIE-based color space (see Section 4.5.4, “CIE-Based Color Spaces”), it must first be transformed to the *native color space* of the raster output device (also called its *process color model*).
2. If a color has been specified in a device color space that is inappropriate for the output device (for example, *RGB* color with a *CMYK* or grayscale device), a *color conversion function* is invoked.
3. The device color values are now mapped through *transfer functions*, one for each color component. The transfer functions compensate for peculiarities of the output device, such as nonlinear gray-level response. This step is sometimes called *gamma correction*.
4. If the device cannot reproduce continuous tones, but only certain discrete colors such as black and white pixels, a *halftone function* is invoked, which approximates the desired colors by means of patterns of pixels.

5. Finally, *scan conversion* is performed to mark the appropriate pixels of the raster output device with the requested colors.

Once these operations have been performed for all graphics objects on the page, the resulting raster data is used to mark the physical output medium, such as pixels on a display or ink on a printed page. A PDF document specifies very little about the properties of the physical medium on which the output will be produced; that information is obtained from the following sources:

- The media box and a few other entries in the page dictionary (see Section 9.10.1, “Page Boundaries”).
- An interactive dialog conducted when the user requests viewing or printing.
- A *job ticket*, either embedded in the PDF file or provided separately, specifying detailed instructions for imposing PDF pages onto media and for controlling special features of the output device. Various standards exist for the format of job tickets; two of them, JDF (Job Definition Format) and PJTF (Portable Job Ticket Format), are described in the CIP4 document *JDF Specification* and in Adobe Technical Note #5620, *Portable Job Ticket Format* (see the Bibliography).

Some of the rendering facilities described in this chapter are controlled by device-dependent graphics state parameters, listed in Table 4.3 on page 150. These parameters can be changed by invoking the **gs** operator with a parameter dictionary containing entries shown in Table 4.8 on page 157.

6.1 CIE-Based Color to Device Color

To render CIE-based colors on an output device, the viewer application must convert from the specified CIE-based color space to the device’s native color space (typically **DeviceGray**, **DeviceRGB**, or **DeviceCMYK**), taking into account the known properties of the device. As discussed in Section 4.5.4, “CIE-Based Color Spaces,” CIE-based color is based on a model of human color perception. The goal of CIE-based color rendering is to produce output in the device’s native color space that accurately reproduces the requested CIE-based color values as perceived by a human observer. CIE-based color specification and rendering are a feature of PDF 1.1 (**CalGray**, **CalRGB**, and **Lab**) and PDF 1.3 (**ICCBased**).

The conversion from CIE-based color to device color is complex, and the theory on which it is based is beyond the scope of this book; see the Bibliography for sources of further information. The algorithm has many parameters, including

an optional, full three-dimensional color lookup table. The color fidelity of the output depends on having these parameters properly set, usually by a method that includes some form of calibration. The colors that a device can produce are characterized by a *device profile*, which is usually specified by an ICC profile associated with the device (and entirely separate from the profile that is specified in an **ICCBased** color space).

Note: PDF has no equivalent of the PostScript color rendering dictionary. The means by which a device profile is associated with a viewer application's output device are implementation-dependent and cannot be specified in a PDF file. Typically, this is done through a color management system (CMS) that is provided by the operating system. Beginning with PDF 1.4, a PDF document can also specify one or more output intents providing possible profiles that might be used to process the document (see Section 9.10.4, "Output Intents").

Conversion from a CIE-based color value to a device color value requires two main operations:

1. Adjust the CIE-based color value according to a *CIE-based gamut mapping function*. A *gamut* is a subset of all possible colors in some color space. A page description has a *source gamut* consisting of all the colors it uses. An output device has a *device gamut* consisting of all the colors it can reproduce. This step transforms colors from the source gamut to the device gamut in a way that attempts to preserve color appearance, visual contrast, or some other explicitly specified *rendering intent* (see "Rendering Intents" on page 197).
2. Generate a corresponding device color value according to a *CIE-based color mapping function*. For a given CIE-based color value, this function computes a color value in the device's native color space.

The CIE-based gamut and color mapping functions are applied only to color values presented in a CIE-based color space. By definition, color values in device color spaces directly control the device color components (though this can be altered by the **DefaultGray**, **DefaultRGB**, and **DefaultCMYK** color space resources; see "Default Color Spaces" on page 194).

The source gamut is specified by a page description when it selects a CIE-based color space. This specification is device-independent. The corresponding properties of the output device are given in the device profile associated with the device. The gamut mapping and color mapping functions are part of the implementation of the viewer application.

6.2 Conversions among Device Color Spaces

Each raster output device has a *native color space*, which typically is one of the standard device color spaces (**DeviceGray**, **DeviceRGB**, or **DeviceCMYK**). In other words, most devices support reproduction of colors according to a grayscale (monochrome), *RGB* (red-green-blue), or *CMYK* (cyan-magenta-yellow-black) model. If the device supports continuous-tone output, reproduction occurs directly. Otherwise, it is accomplished by means of halftoning.

A device's native color space is also called its *process color model*. Process colors are ones that are produced by combinations of one or more standard *process colorants*. Colors specified in any device or CIE-based color space are rendered as process colors. (A device can also support additional *spot colorants*, which can be painted only by means of **Separation** or **DeviceN** color spaces. They are not involved in the rendering of device or CIE-based color spaces, nor are they subject to the conversions described below.)

Note: Some devices provide a native color space that is not one of the three named above but consists of a different combination of colorants. In that case, conversion from the standard device color spaces to the device's native color space is performed by device-dependent means.

Knowing the native color space and other output capabilities of the device, the viewer application can automatically convert the color values specified in a document to those appropriate for the device's native color space. For example, if a document specifies colors in the **DeviceRGB** color space but the device supports grayscale (such as a monochrome display) or *CMYK* (such as a color printer), the viewer application performs the necessary conversions. If the document specifies colors directly in the device's native color space, no conversions are necessary.

The algorithms used to convert among device color spaces are very simple. As perceived by a human viewer, the conversions produce only crude approximations of the original colors. More sophisticated control over color conversion can be achieved by means of CIE-based color specification and rendering. Additionally, device color spaces can be remapped into CIE-based color spaces (see "Default Color Spaces" on page 194).

6.2.1 Conversion between DeviceGray and DeviceRGB

Black, white, and intermediate shades of gray can be considered special cases of *RGB* color. A grayscale value is described by a single number: 0.0 corresponds to black, 1.0 to white, and intermediate values to different gray levels.

A gray level is equivalent to an *RGB* value with all three components the same. In other words, the *RGB* color value equivalent to a specific gray value is simply

$$\begin{aligned}red &= gray \\green &= gray \\blue &= gray\end{aligned}$$

The gray value for a given *RGB* value is computed according to the NTSC video standard, which determines how a color television signal is rendered on a black-and-white television set:

$$gray = 0.3 \times red + 0.59 \times green + 0.11 \times blue$$

6.2.2 Conversion between DeviceGray and DeviceCMYK

Nominally, a gray level is the complement of the black component of *CMYK*. Therefore, the *CMYK* color value equivalent to a specific gray level is simply

$$\begin{aligned}cyan &= 0.0 \\magenta &= 0.0 \\yellow &= 0.0 \\black &= 1.0 - gray\end{aligned}$$

To obtain the equivalent gray level for a given *CMYK* value, the contributions of all components must be taken into account:

$$gray = 1.0 - \min(1.0, 0.3 \times cyan + 0.59 \times magenta + 0.11 \times yellow + black)$$

The interactions between the black component and the other three are elaborated below.

6.2.3 Conversion from DeviceRGB to DeviceCMYK

Conversion of a color value from *RGB* to *CMYK* is a two-step process. The first step is to convert the red-green-blue value to equivalent cyan, magenta, and yel-

low components. The second step is to generate a black component and alter the other components to produce a better approximation of the original color.

The subtractive color primaries cyan, magenta, and yellow are the complements of the additive primaries red, green, and blue. For example, a cyan ink subtracts the red component of white light. In theory, the conversion is very simple:

$$\begin{aligned} \text{cyan} &= 1.0 - \text{red} \\ \text{magenta} &= 1.0 - \text{green} \\ \text{yellow} &= 1.0 - \text{blue} \end{aligned}$$

For example, a color that is 0.2 red, 0.7 green, and 0.4 blue can also be expressed as $1.0 - 0.2 = 0.8$ cyan, $1.0 - 0.7 = 0.3$ magenta, and $1.0 - 0.4 = 0.6$ yellow.

Logically, only cyan, magenta, and yellow are needed to generate a printing color. An equal level of cyan, magenta, and yellow should create the equivalent level of black. In practice, however, colored printing inks do not mix perfectly; such combinations often form dark brown shades instead of true black. To obtain a truer color rendition on a printer, it is often desirable to substitute true black ink for the mixed-black portion of a color. Most color printers support a black component (the *K* component of *CMYK*). Computing the quantity of this component requires some additional steps:

1. *Black generation* calculates the amount of black to be used when trying to reproduce a particular color.
2. *Undercolor removal* reduces the amounts of the cyan, magenta, and yellow components to compensate for the amount of black that was added by black generation.

The complete conversion from *RGB* to *CMYK* is as follows, where $BG(k)$ and $UCR(k)$ are invocations of the black-generation and undercolor-removal functions, respectively:

$$\begin{aligned} c &= 1.0 - \text{red} \\ m &= 1.0 - \text{green} \\ y &= 1.0 - \text{blue} \\ k &= \min(c, m, y) \\ \text{cyan} &= \min(1.0, \max(0.0, c - UCR(k))) \\ \text{magenta} &= \min(1.0, \max(0.0, m - UCR(k))) \\ \text{yellow} &= \min(1.0, \max(0.0, y - UCR(k))) \\ \text{black} &= \min(1.0, \max(0.0, BG(k))) \end{aligned}$$

In PDF 1.2, the black-generation and undercolor-removal functions are defined as PDF function dictionaries (see Section 3.9, “Functions”) that are parameters in the graphics state. They are specified as the values of the **BG** and **UCR** (or **BG2** and **UCR2**) entries in a graphics state parameter dictionary (see Table 4.8 on page 157). Each function is called with a single numeric operand and is expected to return a single numeric result.

The input of both the black-generation and undercolor-removal functions is k , the minimum of the intermediate c , m , and y values that have been computed by subtracting the original *red*, *green*, and *blue* components from 1.0. Nominally, k is the amount of black that can be removed from the cyan, magenta, and yellow components and substituted as a separate black component.

The black-generation function computes the black component as a function of the nominal k value. It can simply return its k operand unchanged or it can return a larger value for extra black, a smaller value for less black, or 0.0 for no black at all.

The undercolor-removal function computes the amount to subtract from each of the intermediate c , m , and y values to produce the final cyan, magenta, and yellow components. It can simply return its k operand unchanged or it can return 0.0 (so no color is removed), some fraction of the black amount, or even a negative amount, thereby adding to the total amount of colorant.

The final component values that result after applying black generation and undercolor removal are expected to be in the range 0.0 to 1.0. If a value falls outside this range, the nearest valid value is substituted automatically, without error indication. This is indicated explicitly by the *min* and *max* operations in the formulas above.

The correct choice of black-generation and undercolor-removal functions depends on the characteristics of the output device—for example, how inks mix. Each device is configured with default values that are appropriate for that device.

See Section 7.6.4, “Rendering Parameters and Transparency,” and in particular “Rendering Intent and Color Conversions” on page 468, for further discussion of the role of black-generation and undercolor-removal functions in the transparent imaging model.

6.2.4 Conversion from DeviceCMYK to DeviceRGB

Conversion of a color value from *CMYK* to *RGB* is a simple operation that does not involve black generation or undercolor removal:

$$\begin{aligned}red &= 1.0 - \min(1.0, cyan + black) \\green &= 1.0 - \min(1.0, magenta + black) \\blue &= 1.0 - \min(1.0, yellow + black)\end{aligned}$$

In other words, the black component is simply added to each of the other components, which are then converted to their complementary colors by subtracting them each from 1.0.

6.3 Transfer Functions

In PDF 1.2, a *transfer function* adjusts the values of color components to compensate for nonlinear response in an output device and in the human eye. Each component of a device color space—for example, the red component of the **DeviceRGB** space—is intended to represent the perceived lightness or intensity of that color component in proportion to the component's numeric value. Many devices do not actually behave this way, however; the purpose of a transfer function is to compensate for the device's actual behavior. This operation is sometimes called *gamma correction* (not to be confused with the *CIE-based gamut mapping function* performed as part of CIE-based color rendering).

In the sequence of steps for processing colors, the viewer application applies the transfer function *after* performing any needed conversions between color spaces, but *before* applying a halftone function, if necessary. Each color component has its own separate transfer function; there is no interaction between components.

Transfer functions always operate in the native color space of the output device, regardless of the color space in which colors were originally specified. (For example, for a *CMYK* device, the transfer functions apply to the device's cyan, magenta, yellow, and black color components, even if the colors were originally specified in, say, a **DeviceRGB** or **CalRGB** color space.) The transfer function is called with a numeric operand in the range 0.0 to 1.0 and must return a number in the same range. The input is the value of a color component in the device's native color space, either specified directly or produced by conversion from some other color space. The output is the transformed component value to be transmitted to the device (after halftoning, if necessary).

Both the input and the output of a transfer function are always interpreted as if the corresponding color component were additive (red, green, blue, or gray): the greater the numeric value, the lighter the color. If the component is subtractive (cyan, magenta, yellow, black, or a spot color), it is converted to additive form by subtracting it from 1.0 before it is passed to the transfer function. The output of the function is always in additive form, and is passed on to the halftone function in that form.

In PDF 1.2, transfer functions are defined as PDF function objects (see Section 3.9, “Functions”). There are two ways to specify transfer functions:

- The *current transfer function* parameter in the graphics state consists of either a single transfer function or an array of four separate transfer functions, one each for red, green, blue, and gray or their complements cyan, magenta, yellow, and black. (If only a single function is specified, it applies to all components.) An *RGB* device uses the first three; a monochrome device uses the gray transfer function only; and a *CMYK* device uses all four. The current transfer function can be specified as the value of the **TR** or **TR2** entry in a graphics state parameter dictionary; see Table 4.8 on page 157.
- The *current halftone* parameter in the graphics state can specify transfer functions as optional entries in *halftone dictionaries* (see Section 6.4.4, “Halftone Dictionaries”). This is the only way to set transfer functions for nonprimary color components, or for any component in devices whose native color space uses components other than the ones listed above. A transfer function specified in a halftone dictionary overrides the corresponding one specified by the current transfer function parameter in the graphics state.

In addition to their intended use for gamma correction, transfer functions can be used to produce a variety of special, device-dependent effects. For example, on a monochrome device, the PostScript calculator function

```
{1 exch sub}
```

inverts the output colors, producing a negative rendition of the page. In general, this method does not work for color devices; inversion can be more complicated than merely inverting each of the components. Because transfer functions produce device-dependent effects, a page description that is intended to be device-independent should not alter them.

Note: When the current color space is **DeviceGray** and the output device's native color space is **DeviceCMYK**, the interpreter uses only the gray transfer function. The normal conversion from **DeviceGray** to **DeviceCMYK** produces 0.0 for the cyan, magenta, and yellow components. These components are not passed through their respective transfer functions but are rendered directly, producing output containing no colored inks. This special case exists for compatibility with existing applications that use a transfer function to obtain special effects on monochrome devices, and applies only to colors specified in the **DeviceGray** color space.

See Section 7.6.4, “Rendering Parameters and Transparency,” and in particular “Halftone and Transfer Function” on page 467, for further discussion of the role of transfer functions in the transparent imaging model.

6.4 Halftones

Halftoning is a process by which continuous-tone colors are approximated on an output device that can achieve only a limited number of discrete colors. Colors that the device cannot produce directly are simulated by using patterns of pixels in the colors available. Perhaps the most familiar example is the rendering of gray tones with black and white pixels, as in a newspaper photograph.

Some output devices can reproduce continuous-tone colors directly. Halftoning is not required for such devices; after gamma correction by the transfer functions, the color components are transmitted directly to the device. On devices that do require halftoning, it occurs after all color components have been transformed by the applicable transfer functions. The input to the halftone function consists of continuous-tone, gamma-corrected color components in the device's native color space. Its output consists of pixels in colors the device can reproduce.

PDF provides a high degree of control over details of the halftoning process. For example, in color printing, independent halftone screens can be specified for each of several colorants. When rendering on low-resolution displays, fine control over halftone patterns is needed to achieve the best approximations of gray levels or colors and to minimize visual artifacts.

Note: Remember that everything pertaining to halftones is, by definition, device-dependent. In general, when a PDF document provides its own halftone specifications, it sacrifices portability. Associated with every output device is a default halftone definition that is appropriate for most purposes. Only relatively sophisticated documents need to define their own halftones to achieve special effects.

All halftones are defined in device space, unaffected by the current transformation matrix. For correct results, a PDF document that defines a new halftone must make assumptions about the resolution and orientation of device space. The best choice of halftone parameters often depends on specific physical properties of the output device, such as pixel shape, overlap between pixels, and the effects of electronic or mechanical noise.

6.4.1 Halftone Screens

In general, halftoning methods are based on the notion of a *halftone screen*, which divides the array of device pixels into *cells* that can be modified to produce the desired halftone effects. A screen is defined by conceptually laying a uniform rectangular grid over the device pixel array. Each pixel belongs to one cell of the grid; a single cell typically contains many pixels. The screen grid is defined entirely in device space, and is unaffected by modifications to the current transformation matrix. This property is essential to ensure that adjacent areas colored by halftones are properly stitched together without visible “seams.”

On a bilevel (black-and-white) device, each cell of a screen can be made to approximate a shade of gray by painting some of the cell’s pixels black and some white. Numerically, the gray level produced within a cell is the ratio of white pixels to the total number of pixels in the cell. A cell containing n pixels can render $n + 1$ different gray levels, ranging from all pixels black to all pixels white. A desired gray value g in the range 0.0 to 1.0 is produced by making i pixels white, where $i = \text{floor}(g \times n)$.

The foregoing description also applies to color output devices whose pixels consist of primary colors that are either completely on or completely off. Most color printers, but not color displays, work this way. Halftoning is applied to each color component independently, producing shades of that color.

Color components are presented to the halftoning machinery in additive form, regardless of whether they were originally specified additively (*RGB* or gray) or subtractively (*CMYK* or tint). Larger values of a color component represent lighter colors—greater intensity in an additive device such as a display, or less ink in a subtractive device such as a printer. Transfer functions produce color values in additive form; see Section 6.3, “Transfer Functions.”

6.4.2 Spot Functions

A common way of defining a halftone screen is by specifying a *frequency*, *angle*, and *spot function*. The frequency is the number of halftone cells per inch; the angle indicates the orientation of the grid lines relative to the device coordinate system. As a cell's desired gray level varies from black to white, individual pixels within the cell change from black to white in a well-defined sequence: if a particular gray level includes certain white pixels, lighter grays will include the same white pixels along with some additional ones. The order in which pixels change from black to white for increasing gray levels is determined by a *spot function*, which specifies that order in an indirect way that minimizes interactions with the screen frequency and angle.

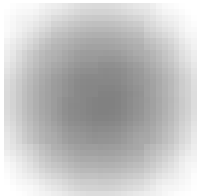
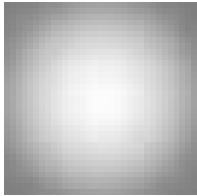
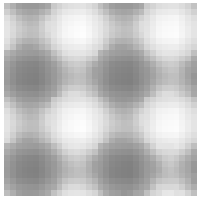
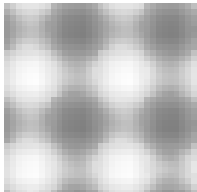
Consider a halftone cell to have its own coordinate system: the center of the cell is the origin and the corners are at coordinates ± 1.0 horizontally and vertically. Each pixel in the cell is centered at horizontal and vertical coordinates that both lie in the range -1.0 to $+1.0$. For each pixel, the spot function is invoked with the pixel's coordinates as input and must return a single number in the range -1.0 to $+1.0$, defining the pixel's position in the whitening order.


The specific values the spot function returns are not significant; all that matters are the *relative* values returned for different pixels. As a cell's gray level varies from black to white, the first pixel whitened is the one for which the spot function returns the lowest value, the next pixel is the one with the next higher spot function value, and so on. If two pixels have the same spot function value, their relative order is chosen arbitrarily.

PDF provides built-in definitions for many of the most commonly used spot functions. A halftone can simply specify any of these predefined spot functions by name instead of giving an explicit function definition. For example, the name **SimpleDot** designates a spot function whose value is inversely related to a pixel's distance from the center of the halftone cell. This produces a "dot screen" in which the black pixels are clustered within a circle whose area is inversely proportional to the gray level. The predefined function **Line** is a spot function whose value is the distance from a given pixel to a line through the center of the cell, producing a "line screen" in which the white pixels grow away from that line.

Table 6.1 shows the predefined spot functions. The table gives the mathematical definition of each function along with the corresponding PostScript language code as it would be defined in a PostScript calculator function (see Section 3.9.4, “Type 4 (PostScript Calculator) Functions”). The image accompanying each function shows how the relative values of the function are distributed over the halftone cell, indicating the approximate order in which pixels are whitened; pixels corresponding to darker points in the image are whitened later than those corresponding to lighter points. (See implementation note 51 in Appendix H.)


TABLE 6.1 Predefined spot functions

NAME	APPEARANCE	DEFINITION
SimpleDot		$1 - (x^2 + y^2)$ <code>{ dup mul exch dup mul add 1 exch sub }</code>
InvertedSimpleDot		$x^2 + y^2 - 1$ <code>{ dup mul exch dup mul add 1 sub }</code>
DoubleDot		$\frac{\sin(360 \times x)}{2} + \frac{\sin(360 \times y)}{2}$ <code>{ 360 mul sin 2 div exch 360 mul sin 2 div add }</code>
InvertedDoubleDot		$-\left(\frac{\sin(360 \times x)}{2} + \frac{\sin(360 \times y)}{2}\right)$ <code>{ 360 mul sin 2 div exch 360 mul sin 2 div add neg }</code>

CosineDot


$$\frac{\cos(180 \times x)}{2} + \frac{\cos(180 \times y)}{2}$$

{ 180 mul cos exch 180 mul cos add 2 div }

Double


$$\frac{\sin\left(360 \times \frac{x}{2}\right)}{2} + \frac{\sin(360 \times y)}{2}$$

{ 360 mul sin 2 div exch 2 div 360 mul sin 2 div add }

InvertedDouble


$$-\left(\frac{\sin\left(360 \times \frac{x}{2}\right)}{2} + \frac{\sin(360 \times y)}{2}\right)$$

{ 360 mul sin 2 div exch 2 div 360 mul sin 2 div add neg }

Line


$$-|y|$$

{ exch pop abs neg }

LineX


$$x$$

{ pop }

LineY


$$y$$

{ exch pop }

Round

if $|x| + |y| \leq 1$ then $1 - (x^2 + y^2)$
 else $(|x| - 1)^2 + (|y| - 1)^2 - 1$


```
{ abs  exch abs
  2 copy add 1 le
    { dup mul  exch dup mul  add 1 exch sub }
    { 1 sub dup mul  exch 1 sub dup mul  add 1 sub }
  ifelse }
```

Ellipse

let $w = (3 \times |x|) + (4 \times |y|) - 3$


$x^2 + \left(\frac{|y|}{0.75}\right)^2$
 if $w < 0$ then $1 - \frac{\quad}{4}$
 else if $w > 1$ then $\frac{(1 - |x|)^2 + \left(\frac{1 - |y|}{0.75}\right)^2}{4} - 1$
 else $0.5 - w$

```
{ abs  exch abs  2 copy 3 mul  exch 4 mul  add 3 sub  dup 0 lt
  { pop dup mul  exch 0.75 div  dup mul  add
    4 div 1 exch sub }
  { dup 1 gt
    { pop 1 exch sub  dup mul
      exch 1 exch sub  0.75 div  dup mul  add
      4 div 1 sub }
    { 0.5 exch sub  exch pop  exch pop }
  ifelse }
ifelse }
```

EllipseA


$1 - (x^2 + 0.9 \times y^2)$

```
{ dup mul 0.9 mul  exch dup mul  add 1 exch sub }
```

InvertedEllipseA


$x^2 + 0.9 \times y^2 - 1$

```
{ dup mul 0.9 mul  exch dup mul  add 1 sub }
```

EllipseB


$$1 - \sqrt{x^2 + \frac{5}{8} \times y^2}$$

```
{ dup 5 mul 8 div mul exch dup mul exch add sqrt
  1 exch sub }
```

EllipseC


$$1 - (0.9 \times x^2 + y^2)$$

```
{ dup mul exch dup mul 0.9 mul add 1 exch sub }
```

InvertedEllipseC


$$0.9 \times x^2 + y^2 - 1$$

```
{ dup mul exch dup mul 0.9 mul add 1 sub }
```

Square


$$-\max(|x|, |y|)$$

```
{ abs exch abs 2 copy lt
  { exch }
  if
  pop neg }
```

Cross

$$-\min(|x|, |y|)$$


```
{ abs exch abs 2 copy gt
  { exch }
  if
  pop neg }
```

Rhomboid

$$\frac{0.9 \times |x| + |y|}{2}$$

```
{ abs exch abs 0.9 mul add 2 div }
```

Diamond



if $|x| + |y| \leq 0.75$ *then* $1 - (x^2 + y^2)$

else if $|x| + |y| \leq 1.23$ *then* $1 - (0.85 \times |x| + |y|)$

else $(|x| - 1)^2 + (|y| - 1)^2 - 1$

{ abs exch abs 2 copy add 0.75 le

{ dup mul exch dup mul add 1 exch sub }

{ 2 copy add 1.23 le

{ 0.85 mul add 1 exch sub }

{ 1 sub dup mul exch 1 sub dup mul add 1 sub }

ifelse }

ifelse }

Figure 6.1 illustrates the effects of some of the predefined spot functions.

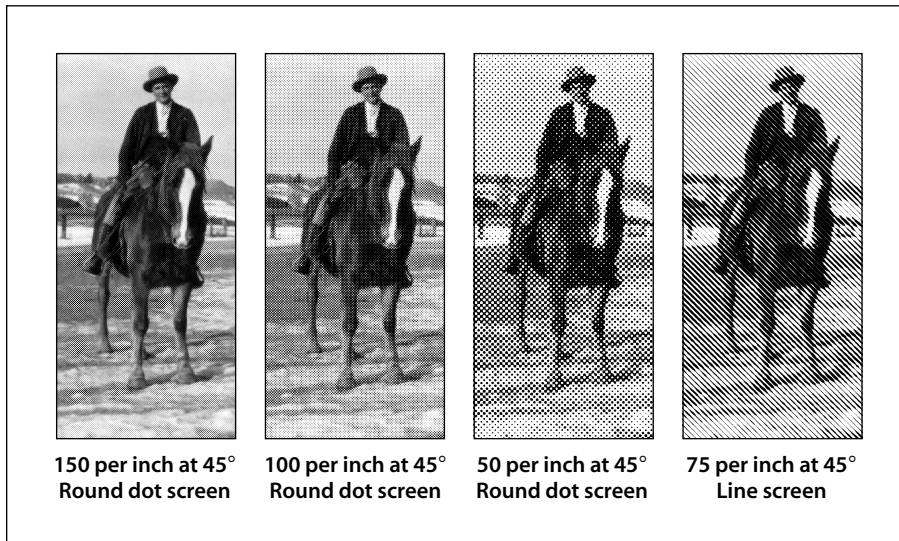


FIGURE 6.1 Various halftoning effects

6.4.3 Threshold Arrays

Another way to define a halftone screen is with a *threshold array* that directly controls individual device pixels in a halftone cell. This technique provides a high

degree of control over halftone rendering. It also permits halftone cells to be arbitrary rectangles, whereas those controlled by a spot function are always square.

A threshold array is much like a sampled image—a rectangular array of pixel values—but is defined entirely in device space. Depending on the halftone type, the threshold values occupy 8 or 16 bits each. Threshold values nominally represent gray levels in the usual way, from 0 for black up to the maximum (255 or 65,535) for white. The threshold array is replicated to tile the entire device space: each pixel in device space is mapped to a particular sample in the threshold array. On a bilevel device, where each pixel is either black or white, halftoning with a threshold array proceeds as follows:

1. For each device pixel that is to be painted with some gray level, consult the corresponding threshold value from the threshold array.
2. If the requested gray level is less than the threshold value, paint the device pixel black; otherwise, paint it white. Gray levels in the range 0.0 to 1.0 correspond to threshold values from 0 to the maximum available (255 or 65,535).

Note: *A threshold value of 0 is treated as if it were 1; therefore, a gray level of 0.0 paints all pixels black, regardless of the values in the threshold array.*

This scheme easily generalizes to monochrome devices with multiple bits per pixel. For example, if there are 2 bits per pixel, each pixel can directly represent one of four different gray levels: black, dark gray, light gray, or white, encoded as 0, 1, 2, and 3, respectively. For any device pixel that is specified with some in-between gray level, the halftoning algorithm consults the corresponding value in the threshold array to determine whether to use the next-lower or next-higher representable gray level. In this situation, the threshold values do not represent absolute gray levels, but rather gradations between any two adjacent representable gray levels.

A halftone defined in this way can also be used with color displays that have a limited number of values for each color component. The red, green, and blue components are simply treated independently as gray levels, applying the appropriate threshold array to each. (This technique also works for a screen defined as a spot function, since the spot function is used to compute a threshold array internally.)

6.4.4 Halftone Dictionaries

In PDF 1.2, the graphics state includes a *current halftone* parameter, which determines the halftoning process to be used by the painting operators. The current halftone can be specified as the value of the **HT** entry in a graphics state parameter dictionary; see Table 4.8 on page 157. It may be defined by either a dictionary or a stream, depending on the type of halftone; the term *halftone dictionary* is used generically throughout this section to refer to either a dictionary object or the dictionary portion of a stream object. (Those halftones that are defined by streams are specifically identified as such in the descriptions of particular halftone types; unless otherwise stated, they are understood to be defined by simple dictionaries instead.)

Every halftone dictionary must have a **HalftoneType** entry whose value is an integer specifying the overall type of halftone definition. The remaining entries in the dictionary are interpreted according to this type. PDF supports the halftone types listed in Table 6.2.

TABLE 6.2 PDF halftone types

TYPE	MEANING
1	Defines a single halftone screen by a <i>frequency</i> , <i>angle</i> , and <i>spot function</i> .
5	Defines an arbitrary number of halftone screens, one for each colorant or color component (including both primary and spot colorants). The keys in this dictionary are names of colorants; the values are halftone dictionaries of other types, each defining the halftone screen for a single colorant.
6	Defines a single halftone screen by a threshold array containing 8-bit sample values.
10	Defines a single halftone screen by a threshold array containing 8-bit sample values, representing a halftone cell that may have a nonzero screen angle.
16	(PDF 1.3) Defines a single halftone screen by a threshold array containing 16-bit sample values, representing a halftone cell that may have a nonzero screen angle.

The dictionaries representing these halftone types contain the same entries as the corresponding PostScript language halftone dictionaries (as described in Section 7.4 of the *PostScript Language Reference*, Third Edition), with the following exceptions:

- The PDF dictionaries may contain a **Type** entry with the value **Halftone**, identifying the type of PDF object that the dictionary describes.
- Spot functions and transfer functions are represented by function objects instead of PostScript procedures.
- Threshold arrays are specified as streams instead of files.
- In type 5 halftone dictionaries, the keys for colorants must be name objects; they may not be strings as they may in PostScript.

Halftone dictionaries have an optional entry, **HalftoneName**, that identifies the desired halftone by name. In PDF 1.3, if this entry is present, all other entries, including **HalftoneType**, are optional. At rendering time, if the output device has a halftone with the specified name, that halftone will be used, overriding any other halftone parameters specified in the dictionary. This provides a way for PDF documents to select the proprietary halftones supplied by some device manufacturers, which would not otherwise be accessible because they are not explicitly defined in PDF. If there is no **HalftoneName** entry, or if the requested halftone name does not exist on the device, the halftone's parameters are defined by the other entries in the dictionary, if any. If no other entries are present, the default halftone is used.

See Section 7.6.4, “Rendering Parameters and Transparency,” and in particular “Halftone and Transfer Function” on page 467, for further discussion of the role of halftones in the transparent imaging model.

Type 1 Halftones

Table 6.3 describes the contents of a halftone dictionary of type 1, which defines a halftone screen in terms of its frequency, angle, and spot function.

TABLE 6.3 Entries in a type 1 halftone dictionary

KEY	TYPE	VALUE
Type	name	<i>(Optional)</i> The type of PDF object that this dictionary describes; if present, must be Halftone for a halftone dictionary.
HalftoneType	integer	<i>(Required)</i> A code identifying the halftone type that this dictionary describes; must be 1 for this type of halftone.
HalftoneName	string	<i>(Optional)</i> The name of the halftone dictionary.
Frequency	number	<i>(Required)</i> The screen frequency, measured in halftone cells per inch in device space.
Angle	number	<i>(Required)</i> The screen angle, in degrees of rotation counterclockwise with respect to the device coordinate system. (Note that most output devices have left-handed device spaces; on such devices, a counterclockwise angle in device space will correspond to a clockwise angle in default user space and on the physical medium.)
SpotFunction	function or name	<i>(Required)</i> A function object defining the order in which device pixels within a screen cell are adjusted for different gray levels, or the name of one of the predefined spot functions (see Table 6.1 on page 385).
AccurateScreens	boolean	<i>(Optional)</i> A flag specifying whether to invoke a special halftone algorithm that is extremely precise, but computationally expensive; see below for further discussion. Default value: false .
TransferFunction	function or name	<i>(Optional)</i> A transfer function, which overrides the current transfer function in the graphics state for the same component. This entry is required if the dictionary is a component of a type 5 halftone (see “Type 5 Halftones” on page 400) and represents either a nonprimary or nonstandard primary color component (see Section 6.3, “Transfer Functions”). The name Identity may be used to specify the identity function.

If the optional entry **AccurateScreens** is present with a boolean value of **true**, a highly precise halftoning algorithm is substituted in place of the standard one; if the **AccurateScreens** entry is **false** or is not present, ordinary halftoning is used. Accurate halftoning achieves the requested screen frequency and angle with very high accuracy, whereas ordinary halftoning adjusts them so that a single screen cell is quantized to device pixels. High accuracy is important mainly for making color separations on high-resolution devices. However, it may be computationally expensive and so is ordinarily disabled.

In principle, PDF permits the use of halftone screens with arbitrarily large cells—in other words, arbitrarily low frequencies. However, cells that are very large relative to the device resolution or that are oriented at unfavorable angles may exceed the capacity of available memory. If this happens, an error will occur. The **AccurateScreens** feature often requires very large amounts of memory to achieve the highest accuracy.

Example 6.1 shows a halftone dictionary for a type 1 halftone.

Example 6.1

```
28 0 obj
  << /Type /Halftone
    /HalftoneType 1
    /Frequency 120
    /Angle 30
    /SpotFunction /CosineDot
    /TransferFunction /Identity
  >>
endobj
```

Type 6 Halftones

A type 6 halftone defines a halftone screen with a threshold array. The halftone is represented as a stream containing the threshold values; the parameters defining the halftone are specified by entries in the stream dictionary. Table 6.4 shows the contents of this dictionary, in addition to the usual entries common to all streams (see Table 3.4 on page 38). The **Width** and **Height** entries specify the dimensions of the threshold array in device pixels; the stream must contain **Width** × **Height** bytes, each representing a single threshold value. Threshold values are defined in device space in the same order as image samples in image space (see Figure 4.26 on page 265), with the first value at device coordinates (0, 0) and horizontal coordinates changing faster than vertical.

Type 10 Halftones

Although type 6 halftones can be used to specify a threshold array with a zero screen angle, they make no provision for other angles. The type 10 halftone removes this restriction and allows the use of threshold arrays for halftones with nonzero screen angles as well.

TABLE 6.4 Additional entries specific to a type 6 halftone dictionary

KEY	TYPE	VALUE
Type	name	<i>(Optional)</i> The type of PDF object that this dictionary describes; if present, must be Halftone for a halftone dictionary.
HalftoneType	integer	<i>(Required)</i> A code identifying the halftone type that this dictionary describes; must be 6 for this type of halftone.
HalftoneName	string	<i>(Optional)</i> The name of the halftone dictionary.
Width	integer	<i>(Required)</i> The width of the threshold array, in device pixels.
Height	integer	<i>(Required)</i> The height of the threshold array, in device pixels.
TransferFunction	function or name	<i>(Optional)</i> A transfer function, which overrides the current transfer function in the graphics state for the same component. This entry is required if the dictionary is a component of a type 5 halftone (see “Type 5 Halftones” on page 400) and represents either a nonprimary or nonstandard primary color component (see Section 6.3, “Transfer Functions”). The name Identity may be used to specify the identity function.

Halftone cells at nonzero angles can be difficult to specify, because they may not line up well with scan lines and because it may be difficult to determine where a given sampled point goes. The type 10 halftone addresses these difficulties by dividing the halftone cell into a pair of squares that line up at zero angles with the output device’s pixel grid. The squares contain the same information as the original cell, but are much easier to store and manipulate. In addition, they can be mapped easily into the internal representation used for all rendering.

Figure 6.2 shows a halftone cell with a frequency of 38.4 cells per inch and an angle of 50.2 degrees, represented graphically in device space at a resolution of 300 dots per inch. Each asterisk in the figure represents a location in device space that is mapped to a specific location in the threshold array.




FIGURE 6.2 *Halftone cell with a nonzero angle*

Figure 6.3 shows how the halftone cell can be divided into two squares. If the squares and the original cell are tiled across device space, the area to the right of the upper square maps exactly into the empty area of the lower square, and vice versa (see Figure 6.4). The last row in the first square is immediately adjacent to the first row in the second and starts in the same column.




FIGURE 6.3 *Angled halftone cell divided into two squares*

Like a type 6 halftone, a type 10 halftone is represented as a stream containing the threshold values, with the parameters defining the halftone specified by entries in the stream dictionary. Table 6.5 shows the contents of this dictionary, in addition to the usual entries common to all streams (see Table 3.4 on page 38); the **Xsquare** and **Ysquare** entries replace the type 6 halftone’s **Width** and **Height** entries.

TABLE 6.5 Additional entries specific to a type 10 halftone dictionary

KEY	TYPE	VALUE
Type	name	<i>(Optional)</i> The type of PDF object that this dictionary describes; if present, must be Halftone for a halftone dictionary.
HalftoneType	integer	<i>(Required)</i> A code identifying the halftone type that this dictionary describes; must be 10 for this type of halftone.
HalftoneName	string	<i>(Optional)</i> The name of the halftone dictionary.
Xsquare	integer	<i>(Required)</i> The side of square <i>X</i> , in device pixels; see below.
Ysquare	integer	<i>(Required)</i> The side of square <i>Y</i> , in device pixels; see below.
TransferFunction	function or name	<i>(Optional)</i> A transfer function, which overrides the current transfer function in the graphics state for the same component. This entry is required if the dictionary is a component of a type 5 halftone (see “Type 5 Halftones” on page 400) and represents either a nonprimary or nonstandard primary color component (see Section 6.3, “Transfer Functions”). The name Identity may be used to specify the identity function.

The **Xsquare** and **Ysquare** entries specify the dimensions of the two squares in device pixels; the stream must contain $\text{Xsquare}^2 + \text{Ysquare}^2$ bytes, each representing a single threshold value. The contents of square *X* are specified first, followed by those of square *Y*. Threshold values within each square are defined in device space in the same order as image samples in image space (see Figure 4.26 on page 265), with the first value at device coordinates (0, 0) and horizontal coordinates changing faster than vertical.

Type 16 Halftones

Like type 10, a type 16 halftone (*PDF 1.3*) defines a halftone screen with a threshold array and allows nonzero screen angles. In type 16, however, each element of the threshold array is 16 bits wide instead of 8. This allows the threshold array to distinguish 65,536 levels of color rather than only 256 levels. The threshold array can consist of either one or two rectangles. If two rectangles are specified, they will tile the device space as shown in Figure 6.5. The last row in the first rectangle is immediately adjacent to the first row in the second and starts in the same column.

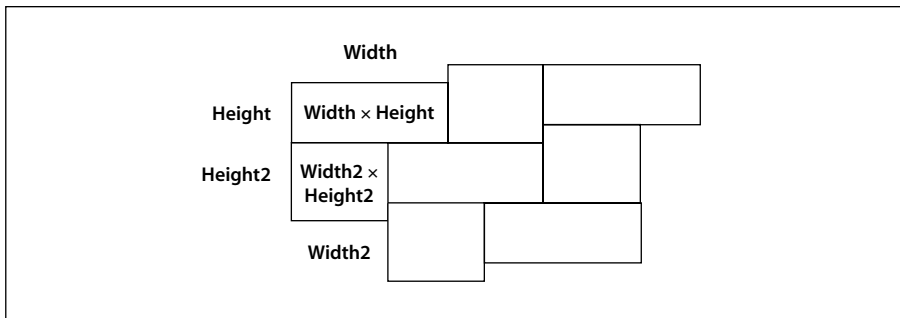


FIGURE 6.5 Tiling of device space in a type 16 halftone

A type 16 halftone, like type 6 and type 10, is represented as a stream containing the threshold values, with the parameters defining the halftone specified by entries in the stream dictionary. Table 6.6 shows the contents of this dictionary, in addition to the usual entries common to all streams (see Table 3.4 on page 38). The dictionary's **Width** and **Height** entries define the dimensions of the first (or only) rectangle; those of the second, optional rectangle are defined by the optional entries **Width2** and **Height2**. Each threshold value is represented as 2 bytes, with the high-order byte first. The stream must thus contain $2 \times \text{Width} \times \text{Height}$ bytes if there is only one rectangle, or $2 \times (\text{Width} \times \text{Height} + \text{Width2} \times \text{Height2})$ bytes if there are two. The contents of the first rectangle are specified first, followed by those of the second rectangle. Threshold values within each rectangle are defined in device space in the same order as image samples in image space (see Figure 4.26 on page 265), with the first value at device coordinates (0, 0) and horizontal coordinates changing faster than vertical.

TABLE 6.6 Additional entries specific to a type 16 halftone dictionary

KEY	TYPE	VALUE
Type	name	<i>(Optional)</i> The type of PDF object that this dictionary describes; if present, must be Halftone for a halftone dictionary.
HalftoneType	integer	<i>(Required)</i> A code identifying the halftone type that this dictionary describes; must be 16 for this type of halftone.
HalftoneName	string	<i>(Optional)</i> The name of the halftone dictionary.
Width	integer	<i>(Required)</i> The width of the first (or only) rectangle in the threshold array, in device pixels.
Height	integer	<i>(Required)</i> The height of the first (or only) rectangle in the threshold array, in device pixels.
Width2	integer	<i>(Optional)</i> The width of the optional second rectangle in the threshold array, in device pixels. If this entry is present, the Height2 entry must be present as well; if this entry is absent, the Height2 entry must also be absent and the threshold array has only one rectangle.
Height2	integer	<i>(Optional)</i> The height of the optional second rectangle in the threshold array, in device pixels.
TransferFunction	function or name	<i>(Optional)</i> A transfer function, which overrides the current transfer function in the graphics state for the same component. This entry is required if the dictionary is a component of a type 5 halftone (see “Type 5 Halftones,” below) and represents either a nonprimary or nonstandard primary color component (see Section 6.3, “Transfer Functions”). The name Identity may be used to specify the identity function.

Type 5 Halftones

Some devices, particularly color printers, require separate halftones for each individual colorant. Also, devices that can produce named separations may require individual halftones for each separation. Halftone dictionaries of type 5 allow individual halftones to be specified for an arbitrary number of colorants or color components.

A type 5 halftone dictionary (Table 6.7) is a composite dictionary containing independent halftone definitions for multiple colorants. Its keys are name objects representing the names of individual colorants or color components. The values associated with these keys are other halftone dictionaries, each defining the halftone screen and transfer function for a single colorant or color component. The component halftone dictionaries may be of any supported type except 5.

TABLE 6.7 Entries in a type 5 halftone dictionary

KEY	TYPE	VALUE
Type	name	<i>(Optional)</i> The type of PDF object that this dictionary describes; if present, must be Halftone for a halftone dictionary.
HalftoneType	number	<i>(Required)</i> A code identifying the halftone type that this dictionary describes; must be 5 for this type of halftone.
HalftoneName	string	<i>(Optional)</i> The name of the halftone dictionary.
<i>any colorant name</i>	dictionary or stream	<i>(Required, one per colorant)</i> The halftone corresponding to the colorant or color component named by the key. The halftone may be of any type other than 5. Note that the key must be a name object; strings are not permitted, as they are in type 5 PostScript halftone dictionaries.
Default	dictionary or stream	<i>(Required)</i> A halftone to be used for any colorant or color component that does not have an entry of its own. The value may not be a type 5 halftone. If there are any nonprimary colorants, the default halftone must have a transfer function.

The colorants or color components represented in a type 5 halftone dictionary fall into two categories:

- Primary color components for the standard native device color spaces (**Red**, **Green**, and **Blue** for **DeviceRGB**; **Cyan**, **Magenta**, **Yellow**, and **Black** for **DeviceCMYK**; **Gray** for **DeviceGray**).
- Nonstandard color components for use as spot colorants in **Separation** and **DeviceN** color spaces. Some of these may also be used as process colorants if the native color space is nonstandard.

The dictionary must also contain an entry whose key is **Default**; the value of this entry is a halftone dictionary to be used for any color component that does not have an entry of its own.

When a halftone dictionary of some other type appears as the value of an entry in a type 5 halftone dictionary, it applies only to the single colorant or color component named by that entry's key. This is in contrast to such a dictionary's being used as the current halftone parameter in the graphics state, which applies to all color components. If nonprimary colorants are requested when the current halftone is defined by any means other than a type 5 halftone dictionary, the gray halftone screen and transfer function are used for all such colorants.

Example 6.2 shows a type 5 halftone dictionary with the primary color components for a *CMYK* device. In this example, the halftone dictionaries for the color components and for the default all use the same spot function.

Example 6.2

```
27 0 obj
  << /Type /Halftone
    /HalftoneType 5
    /Cyan 31 0 R
    /Magenta 32 0 R
    /Yellow 33 0 R
    /Black 34 0 R
    /Default 35 0 R
  >>
endobj

31 0 obj
  << /Type /Halftone
    /HalftoneType 1
    /Frequency 89.827
    /Angle 15
    /SpotFunction /Round
    /AccurateScreens true
  >>
endobj

32 0 obj
  << /Type /Halftone
    /HalftoneType 1
    /Frequency 89.827
    /Angle 75
    /SpotFunction /Round
    /AccurateScreens true
  >>
endobj
```

```
33 0 obj
  << /Type /Halftone
    /HalftoneType 1
    /Frequency 90.714
    /Angle 0
    /SpotFunction /Round
    /AccurateScreens true
  >>
endobj

34 0 obj
  << /Type /Halftone
    /HalftoneType 1
    /Frequency 89.803
    /Angle 45
    /SpotFunction /Round
    /AccurateScreens true
  >>
endobj

35 0 obj
  << /Type /Halftone
    /HalftoneType 1
    /Frequency 90.000
    /Angle 45
    /SpotFunction /Round
    /AccurateScreens true
  >>
endobj
```

6.5 Scan Conversion Details

The final step of rendering is *scan conversion*. As discussed in Section 2.1.4, “Scan Conversion,” the viewer application executes a scan conversion algorithm to paint graphics, text, and images in the raster memory of the output device.

The specifics of the scan conversion algorithm are not defined as part of PDF. Different implementations can perform scan conversion in different ways; techniques that are appropriate for one device may be inappropriate for another. Still, it is useful to have a general understanding of how scan conversion works, particularly when creating PDF documents intended for viewing on a display. At the low resolutions typical of displays, variations of even one pixel’s width can have a noticeable effect on the appearance of painted shapes.

The following sections describe the scan conversion algorithms that are typical of Adobe Acrobat products. (These details also apply to Adobe PostScript products, yielding consistent results when a viewer application prints a document on a PostScript printer.) Most scan conversion details are not under program control, but a few are; the parameters for controlling them are described here.

6.5.1 Flatness Tolerance

The *flatness tolerance* controls the maximum permitted distance in device pixels between the mathematically correct path and an approximation constructed from straight line segments, as shown in Figure 6.6. Flatness can be specified as the operand of the `i` operator (see Table 4.7 on page 156) or as the value of the **FL** entry in a graphics state parameter dictionary (see Table 4.8 on page 157). It must be a positive number; smaller values yield greater precision at the cost of more computation.

Note: Although the figure exaggerates the difference between the curved and flattened paths for the sake of clarity, the purpose of the flatness tolerance is to control the precision of curve rendering, not to draw inscribed polygons. If the parameter's value is large enough to cause visible straight line segments to appear, the result is unpredictable.




FIGURE 6.6 Flatness tolerance

6.5.2 Smoothness Tolerance

The *smoothness tolerance* (PDF 1.3) controls the quality of smooth shading (type 2 patterns and the **sh** operator), and thus indirectly controls the rendering performance. Smoothness is the allowable color error between a shading approximated by piecewise linear interpolation and the true value of a (possibly non-linear) shading function. The error is measured for each color component, and the maximum error is used. The allowable error (or tolerance) is expressed as a fraction of the range of the color component, from 0.0 to 1.0. Thus, a smoothness tolerance of 0.1 represents a tolerance of 10 percent in each color component. Smoothness can be specified as the value of the **SM** entry in a graphics state parameter dictionary (see Table 4.8 on page 157).

Each output device may have internal limits on the maximum and minimum tolerances attainable. For example, setting smoothness to 1.0 may result in an internal smoothness of 0.5 on a high-quality color device, while setting it to 0.0 on the same device may result in an internal smoothness of 0.01 if an error of that magnitude is imperceptible on the device.

The smoothness tolerance may also interact with the accuracy of color conversion. In the case of a color conversion defined by a sampled function, the conversion function is unknown. Thus the error may be sampled at too low a frequency, in which case the accuracy defined by the smoothness tolerance cannot be guaranteed. In most cases, however, where the conversion function is smooth and continuous, the accuracy should be within the specified tolerance.

The effect of the smoothness tolerance is similar to that of the flatness tolerance. Note, however, that flatness is measured in device-dependent units of pixel width, whereas smoothness is measured as a fraction of color component range.

6.5.3 Scan Conversion Rules

The following rules determine which device pixels a painting operation will affect. All references to coordinates and pixels are in device space. A *shape* is a path to be painted with the current color or with an image. Its coordinates are mapped into device space, but not rounded to device pixel boundaries. At this level, curves have been flattened to sequences of straight lines, and all “insideness” computations have been performed.

Pixel boundaries always fall on integer coordinates in device space. A pixel is a square region identified by the location of its corner with minimum horizontal and vertical coordinates. The region is *half-open*, meaning that it includes its lower but not its upper boundaries. More precisely, for any point whose real-number coordinates are (x, y) , let $i = \text{floor}(x)$ and $j = \text{floor}(y)$. The pixel that contains this point is the one identified as (i, j) . The region belonging to that pixel is defined to be the set of points (x', y') such that $i \leq x' < i + 1$ and $j \leq y' < j + 1$. Like pixels, shapes to be painted by filling and stroking operations are also treated as half-open regions that include the boundaries along their “floor” sides, but not along their “ceiling” sides.

A shape is scan-converted by painting any pixel whose square region intersects the shape, no matter how small the intersection is. This ensures that no shape ever disappears as a result of unfavorable placement relative to the device pixel grid, as might happen with other possible scan conversion rules. The area covered by painted pixels is always at least as large as the area of the original shape. This rule applies both to fill operations and to strokes with nonzero width. Zero-width strokes are done in a device-dependent manner that may include fewer pixels than the rule implies.

Note: *Normally, the intersection of two regions is defined as the intersection of their interiors. However, for purposes of scan conversion, a filling region is considered to intersect every pixel through which its boundary passes, even if the interior of the filling region is empty. Thus, for example, a zero-width or zero-height rectangle will paint a line 1 pixel wide.*

The region of device space to be painted by a sampled image is determined similarly to that of a filled shape, though not identically. The viewer application transforms the image’s source rectangle into device space and defines a half-open region, just as for fill operations. However, only those pixels whose *centers* lie within the region are painted. The position of the center of such a pixel—in other words, the point whose coordinate values have fractional parts of one-half—is mapped back into source space to determine how to color the pixel. There is no averaging over the pixel area; if the resolution of the source image is higher than that of device space, some source samples will not be used.

For clipping, the clipping region consists of the set of pixels that would be included by a fill operation. Subsequent painting operations affect a region that is the intersection of the set of pixels defined by the clipping region with the set of pixels for the region to be painted.

Scan conversion of character glyphs is performed by a different algorithm from the one above. That font rendering algorithm uses hints in the glyph descriptions and techniques that are specialized to glyph rasterization.

6.5.4 Automatic Stroke Adjustment

When a stroke is drawn along a path, the scan conversion algorithm may produce lines of nonuniform thickness because of rasterization effects. In general, the line width and the coordinates of the endpoints, transformed into device space, are arbitrary real numbers not quantized to device pixels. A line of a given width can intersect with different numbers of device pixels, depending on where it is positioned. Figure 6.7 illustrates this effect.

For best results, it is important to compensate for the rasterization effects to produce strokes of uniform thickness. This is especially important in low-resolution display applications. To meet this need, PDF 1.2 provides an optional *automatic stroke adjustment* feature. When stroke adjustment is enabled, the line width and the coordinates of a stroke are automatically adjusted as necessary to produce lines of uniform thickness. The thickness is as near as possible to the requested line width—no more than half a pixel different.




FIGURE 6.7 Rasterization without stroke adjustment

***Note:** If stroke adjustment is enabled and the requested line width, transformed into device space, is less than half a pixel, the stroke is rendered as a single-pixel line. This is the thinnest line that can be rendered at device resolution. It is equivalent to the effect produced by setting the line width to 0 (see Section 6.5.3, “Scan Conversion Rules”).*

Because automatic stroke adjustment can have a substantial effect on the appearance of lines, a PDF document must be able to control whether the adjustment is to be performed. This can be specified with the stroke adjustment parameter in the graphics state, set via the **SA** entry in a graphics state parameter dictionary (see Section 4.3.4, “Graphics State Parameter Dictionaries”).

CHAPTER 7

Transparency

PDF 1.4 EXTENDS the Adobe imaging model to include the notion of *transparency*. Transparent objects do not necessarily obey a strict opaque painting model, but can blend (*composite*) in interesting ways with other overlapping objects. This chapter describes the general transparency model, but does not attempt to cover how it is to be implemented. Although implementation-like descriptions are used at various points to describe how things work, this is only for the purpose of elucidating the behavior of the model; the actual implementation will almost certainly be different from what these descriptions might imply.

The chapter is organized as follows:

- Section 7.1, “Overview of Transparency,” introduces the basic concepts of the transparency model and its associated terminology.
- Section 7.2, “Basic Compositing Computations,” describes the mathematics involved in compositing a single object with its backdrop.
- Section 7.3, “Transparency Groups,” introduces the concept of *transparency groups* and describes their properties and behavior.
- Section 7.4, “Soft Masks,” covers the creation and use of masks to specify position-dependent shape and opacity.
- Section 7.5, “Specifying Transparency in PDF,” describes how transparency properties are represented in a PDF document.
- Section 7.6, “Color Space and Rendering Issues,” deals with some specific interactions between transparency and other aspects of color specification and rendering.

7.1 Overview of Transparency

The original Adobe imaging model paints objects (fills, strokes, text, and images), possibly clipped by a path, opaquely onto a page. The color of the page at any point is that of the topmost enclosing object, disregarding any previous objects it may overlap. This effect can be—and often is—realized simply by rendering objects directly to the page in the order in which they are specified, with each object completely overwriting any others that it overlaps.

Under the transparent imaging model, all of the objects on a page can potentially contribute to the result. Objects at a given point can be thought of as forming a *transparency stack* (or just *stack* for short), arranged from bottom to top in the order in which they are specified. The color of the page at each point is determined by combining the colors of all enclosing objects in the stack according to *compositing* rules defined by the transparency model.

Note: *The order in which objects are specified determines the stacking order, but not necessarily the order in which the objects are actually painted onto the page. In particular, the transparency model does not require a viewer application to rasterize objects immediately or to commit to a raster representation at any time before rendering the entire stack onto the page. This is important, since rasterization often causes significant loss of information and precision that is best avoided during intermediate stages of the transparency computation.*

A given object is composited with a *backdrop*. Ordinarily, the backdrop consists of the stack of all objects that have been specified previously; the result of compositing is then treated as the backdrop for the next object. However, within certain kinds of transparency group (see below), a different backdrop is chosen.

When an object is composited with its backdrop, the color at each point is computed using a specified *blend mode*, which is a function of both the object's color and the backdrop color. The blend mode determines how colors interact; different blend modes can be used to achieve a variety of useful effects. A single blend mode is in effect for compositing all of a given object, but different blend modes can be applied to different objects.

Compositing of an object with its backdrop is mediated by two scalar quantities called *shape* and *opacity*. Conceptually, for each object, these quantities are defined at every point in the plane, just as if they were additional color components.

(In actual practice, they are often obtained from auxiliary sources rather than being intrinsic to the object itself.)

Both shape and opacity vary from 0.0 (no contribution) to 1.0 (maximum contribution). At any point where either the shape or the opacity of an object is 0.0, its color is undefined. At points where the shape is 0.0, the opacity is also undefined. The shape and opacity are themselves subject to compositing rules, so that the stack as a whole also has a shape and opacity at each point.

An object's opacity, in combination with the backdrop's opacity, determines the relative contributions of the backdrop color, the object's color, and the blended color to the resulting composite color. The object's shape then determines the degree to which the composite color replaces the backdrop color. Shape values of 0.0 and 1.0 identify points that lie "outside" and "inside" a conventional sharp-edged object; intermediate values are useful in defining soft-edged objects.

Shape and opacity are conceptually very similar. In fact, they can usually be combined into a single value, called *alpha*, which controls both the color compositing computation and the fading between an object and its backdrop. However, there are a few situations in which they must be treated separately; see Section 7.3.5, "Knockout Groups." Moreover, raster-based implementations must maintain a separate shape parameter in order to do anti-aliasing properly; it is therefore convenient to have it be an explicit part of the model.

One or more consecutive objects in a stack can be collected together into a *transparency group* (often referred to hereafter simply as a *group*). The group as a whole can have various properties that modify the compositing behavior of objects within the group and their interactions with its backdrop. An additional blend mode, shape, and opacity can also be associated with the group as a whole and used when compositing it with its backdrop. Groups can be nested within other groups, forming a tree-structured hierarchy.

Note: *The concept of a transparency group is independent of existing notions of "group" or "layer" in applications such as Adobe Illustrator®. Those groupings reflect logical relationships among objects that are meaningful when editing those objects, but they are not part of the imaging model.*

Plate 16 illustrates the effects of transparency grouping. In the upper two figures, three colored circles are painted as independent objects, with no grouping. At the upper left, the three objects are painted opaquely (opacity = 1.0); each object

completely replaces its backdrop (including previously painted objects) with its own color. At the upper right, the same three independent objects are painted with an opacity of 0.5, causing them to composite with each other and with the gray and white backdrop. In the lower two figures, the three objects are combined as a transparency group. At the lower left, the individual objects have an opacity of 1.0 within the group, but the group as a whole is painted in the **Normal** blend mode with an opacity of 0.5. The objects thus completely overwrite each other within the group, but the resulting group then composites transparently with the gray and white backdrop. At the lower right, the objects have an opacity of 0.5 within the group and thus composite with each other; the group as a whole is painted against the backdrop with an opacity of 1.0, but in a different blend mode (**HardLight**), producing a different visual effect.

The color result of compositing a group can be converted to a single-component luminosity value and treated as a *soft mask*. Such a mask can then be used as an additional source of shape or opacity values for subsequent compositing operations. When the mask is used as a shape, this technique is known as *soft clipping*; it is a generalization of the current clipping path in the opaque imaging model (see Section 4.4.3, “Clipping Path Operators”).

The notion of *current page* is generalized to refer to a transparency group consisting of the entire stack of objects placed on the page, composited with a backdrop that is pure white and fully opaque. Logically, this entire stack is then rasterized to determine the actual pixel values to be transmitted to the output device.

Note: *In contexts where a PDF page is treated as a piece of artwork to be placed on some other page—such as an Illustrator artboard or an Encapsulated PostScript (EPS) file—it is treated not as a page but as a group, whose backdrop may be defined differently from that of a page.*

7.2 Basic Compositing Computations

This section describes the basic computations for compositing a single object with its backdrop. These computations will be extended in Section 7.3, “Transparency Groups,” to cover groups consisting of multiple objects.

7.2.1 Basic Notation for Compositing Computations

In general, variable names in this chapter consisting of a lowercase letter denote a scalar quantity, such as an opacity; uppercase letters denote a value with multiple scalar components, such as a color. In the descriptions of the basic color compositing computations, color values are generally denoted by the letter C , with a mnemonic subscript indicating which of several color values is being referred to; for instance, C_s stands for “source color.” Shape and opacity values are denoted respectively by the letters f (for “form factor”) and q (for “opaqueness”)—again with a mnemonic subscript, such as q_s for “source opacity.” The symbol α (alpha) stands for a product of shape and opacity values.

In certain computations, one or more variables may have undefined values; for instance, when opacity is zero, the corresponding color is undefined. A quantity can also be undefined if it results from division by zero. In any formula that uses such an undefined quantity, the quantity has no effect on the ultimate result, because it is subsequently multiplied by zero or otherwise canceled out. The significant point is that while any arbitrary value can be chosen for such an undefined quantity, the computation must not malfunction because of exceptions caused by overflow or division by zero. It is convenient to adopt the further convention that $0 \div 0 = 0$.

7.2.2 Basic Compositing Formula

The primary change in the imaging model to accommodate transparency is in how colors are painted. In the transparent model, the result of painting (the *result color*) is a function of both the color being painted (the *source color*) and the color it is painted over (the *backdrop color*). Both of these colors may vary as a function of position on the page, but for the purposes of this section we will focus our attention on some fixed point on the page and assume a fixed backdrop and source color.

Other parameters in this computation are the *alpha*, which controls the relative contributions of the backdrop and source colors, and the *blend function*, which specifies how they are combined in the painting operation. The resulting *basic*

color compositing formula (or just *basic compositing formula* for short) determines the result color produced by the painting operation:

$$C_r = \left(1 - \frac{\alpha_s}{\alpha_r}\right) \times C_b + \frac{\alpha_s}{\alpha_r} \times [(1 - \alpha_b) \times C_s + \alpha_b \times B(C_b, C_s)]$$

where the variables have the meanings shown in Table 7.1.

TABLE 7.1 Variables used in the basic compositing formula

VARIABLE	MEANING
C_b	Backdrop color
C_s	Source color
C_r	Result color
α_b	Backdrop alpha
α_s	Source alpha
α_r	Result alpha
$B(C_b, C_s)$	Blend function

This is actually a simplified form of the compositing formula in which the shape and opacity values are combined and represented as a single alpha value; the more general form is presented later. This function is based on the **over** operation defined in the article “Compositing Digital Images,” by Porter and Duff (see the Bibliography), extended to include a blend mode in the region of overlapping coverage. The following sections elaborate on the meaning and implications of this formula.

7.2.3 Blending Color Space

Note that the compositing formula shown above is actually a vector function: the colors it operates on are represented in the form of n -element vectors, where n is the number of components required by the color space in which compositing is performed. The i th component of the result color C_r is obtained by applying the compositing formula to the i th components of the constituent colors C_b , C_s , and $B(C_b, C_s)$. The result of the computation thus depends on the color space in

which the colors are represented. For this reason, the color space used for compositing, called the *blending color space*, is explicitly made part of the transparent imaging model. When necessary, backdrop and source colors are converted to the blending color space prior to the compositing computation.

Of the PDF color spaces described in Section 4.5, “Color Spaces,” the following are supported as blending color spaces:

- **DeviceGray**
- **DeviceRGB**
- **DeviceCMYK**
- **CalGray**
- **CalRGB**
- **ICCBased** color spaces equivalent to those above (including calibrated *CMYK*)

The **Lab** space and **ICCBased** spaces that represent lightness and chromaticity separately (such as $L^*a^*b^*$, $L^*u^*v^*$, and *HSV*) are not allowed as blending color spaces, because the compositing computations in such spaces do not give meaningful results when applied separately to each component. In addition, an **ICCBased** space used as a blending color space must be bidirectional; that is, the ICC profile must contain both *AToB* and *BToA* transformations.

The blending color space is consulted only for process colors. Although blending can also be done on individual spot colors specified in a **Separation** or **DeviceN** color space, such colors are never converted to a blending color space (except in the case where they first revert to their alternate color space, as described under “Separation Color Spaces” on page 201 and “DeviceN Color Spaces” on page 205). Instead, the specified color components are blended individually with the corresponding components of the backdrop.

The blend functions for the various blend modes assume that the range for each color component is 0.0 to 1.0 and that the color space is additive. The former condition is true for all of the allowed blending color spaces, but the latter is not. In particular, the **DeviceCMYK**, **Separation**, and **DeviceN** spaces are subtractive. When performing blending operations in subtractive color spaces, it is assumed that the color component values are complemented (subtracted from 1.0) before the blend function is applied and that the results of the function are then complemented back before being used. This adjustment makes the effects of the vari-

ous blend modes numerically consistent across all color spaces. However, the actual visual effect produced by a given blend mode still depends on the color space. Blending in a device color space produces device-dependent results, whereas in a CIE-based space it produces results that are consistent across all devices. See Section 7.6, “Color Space and Rendering Issues,” for additional details concerning color spaces.

7.2.4 Blend Mode

In principle, the blend function $B(C_b, C_s)$, used in the compositing formula to customize the blending operation, could be any function of the backdrop and source colors that yields another color, C_r , for the result. PDF defines a standard set of named blend functions, or *blend modes*, listed in Tables 7.2 and 7.3. Plates 18 and 19 illustrate the resulting visual effects for *RGB* and *CMYK* colors, respectively.

A blend mode is termed *separable* if each component of the result color is completely determined by the corresponding components of the constituent backdrop and source colors—that is, if the blend mode function B is applied separately to each set of corresponding components:

$$c_r = B(c_b, c_s)$$

where the lowercase variables c_r , c_b , and c_s denote corresponding components of the colors C_r , C_b , and C_s , expressed in additive form. (Theoretically, a blend mode could have a different function for each color component and still be separable; however, none of the standard PDF blend modes have this property.) A separable blend mode can be used with any color space, since it applies independently to any number of components. Only separable blend modes can be used for blending spot colors.

Table 7.2 lists the standard separable blend modes available in PDF. Some of them are defined by actual mathematical formulas; the rest are characterized only by a general description of their intended effects.

TABLE 7.2 Standard separable blend modes

NAME	RESULT
Normal	<p>Selects the source color, ignoring the backdrop:</p> $B(c_b, c_s) = c_s$
Multiply	<p>Multiplies the backdrop and source color values:</p> $B(c_b, c_s) = c_b \times c_s$ <p>The result color is always at least as dark as either of the two constituent colors. Multiplying any color with black produces black; multiplying with white leaves the original color unchanged. Painting successive overlapping objects with a color other than black or white produces progressively darker colors.</p>
Screen	<p>Multiplies the complements of the backdrop and source color values, then complements the result:</p> $\begin{aligned} B(c_b, c_s) &= 1 - [(1 - c_b) \times (1 - c_s)] \\ &= c_b + c_s - (c_b \times c_s) \end{aligned}$ <p>The result color is always at least as light as either of the two constituent colors. Screening any color with white produces white; screening with black leaves the original color unchanged. The effect is similar to projecting multiple photographic slides simultaneously onto a single screen.</p>
Overlay	<p>Multiplies or screens the colors, depending on the backdrop color. Source colors overlay the backdrop while preserving its highlights and shadows. The backdrop color is not replaced, but is mixed with the source color to reflect the lightness or darkness of the backdrop.</p>
Darken	<p>Selects the darker of the backdrop and source colors:</p> $B(c_b, c_s) = \min(c_b, c_s)$ <p>The backdrop is replaced with the source where the source is darker; otherwise it is left unchanged.</p>
Lighten	<p>Selects the lighter of the backdrop and source colors:</p> $B(c_b, c_s) = \max(c_b, c_s)$ <p>The backdrop is replaced with the source where the source is lighter; otherwise it is left unchanged.</p>

ColorDodge	Brightens the backdrop color to reflect the source color. Painting with black produces no change.
ColorBurn	Darkens the backdrop color to reflect the source color. Painting with white produces no change.
HardLight	Multiplies or screens the colors, depending on the source color value. If the source color is lighter than 0.5, the backdrop is lightened, as if it were screened; this is useful for adding highlights to a scene. If the source color is darker than 0.5, the backdrop is darkened, as if it were multiplied; this is useful for adding shadows to a scene. The degree of lightening or darkening is proportional to the difference between the source color and 0.5; if it is equal to 0.5, the backdrop is unchanged. Painting with pure black or white produces pure black or white. The effect is similar to shining a harsh spotlight on the backdrop.
SoftLight	Darkens or lightens the colors, depending on the source color value. If the source color is lighter than 0.5, the backdrop is lightened, as if it were dodged; this is useful for adding highlights to a scene. If the source color is darker than 0.5, the backdrop is darkened, as if it were burned in. The degree of lightening or darkening is proportional to the difference between the source color and 0.5; if it is equal to 0.5, the backdrop is unchanged. Painting with pure black or white produces a distinctly darker or lighter area, but does not result in pure black or white. The effect is similar to shining a diffused spotlight on the backdrop.
Difference	Subtracts the darker of the two constituent colors from the lighter: $B(c_b, c_s) = c_b - c_s $ Painting with white inverts the backdrop color; painting with black produces no change.
Exclusion	Produces an effect similar to that of the Difference mode, but lower in contrast. Painting with white inverts the backdrop color; painting with black produces no change.

Table 7.3 lists the standard nonseparable blend modes. Their effects are described, but no mathematical formulas are given. These modes all entail conversion to and from an intermediate *HSL* (hue-saturation-luminance) representation. Since the nonseparable blend modes consider all color components in

combination, their computation depends on the blending color space in which the components are interpreted.

TABLE 7.3 Standard nonseparable blend modes

NAME	RESULT
Hue	Creates a color with the hue of the source color and the saturation and luminance of the backdrop color.
Saturation	Creates a color with the saturation of the source color and the hue and luminance of the backdrop color. Painting with this mode in an area of the backdrop that is a pure gray (no saturation) produces no change.
Color	Creates a color with the hue and saturation of the source color and the luminance of the backdrop color. This preserves the gray levels of the backdrop and is useful for coloring monochrome images or tinting color images.
Luminosity	Creates a color with the luminance of the source color and the hue and saturation of the backdrop color. This produces an inverse effect to that of the Color mode.

Note: An additional standard blend mode, **Compatible**, is a vestige of an earlier design and is no longer needed, but is still recognized for the sake of compatibility; its effect is equivalent to that of the **Normal** blend mode. See “Compatibility with Opaque Overprinting” on page 460 for further discussion.

7.2.5 Interpretation of Alpha

The color compositing formula

$$C_r = \left(1 - \frac{\alpha_s}{\alpha_r}\right) \times C_b + \frac{\alpha_s}{\alpha_r} \times [(1 - \alpha_b) \times C_s + \alpha_b \times B(C_b, C_s)]$$

produces a result color that is a weighted average of the backdrop color, the source color, and the blended $B(C_b, C_s)$ term, with the weighting determined by the backdrop and source alphas α_b and α_s . For the simplest blend mode, **Normal**, defined by

$$B(c_b, c_s) = c_s$$

the compositing formula collapses to a simple weighted average of the backdrop and source colors, controlled by the backdrop and source alpha values. For more interesting blend functions, the backdrop and source alphas control whether the effect of the blend mode is fully realized or is toned down by mixing the result with the backdrop and source colors.

The result alpha, α_r , is actually a computed result, described below in Section 7.2.6, “Shape and Opacity Computations.” The result color is normalized by the result alpha, ensuring that when this color and alpha are subsequently used together in another compositing operation, the color’s contribution will be correctly represented. Note that if α_r is zero, the result color is undefined.

The formula shown above is a simplification of the following one, which presents the relative contributions of backdrop, source, and blended colors in a more straightforward way:

$$\alpha_r \times C_r = [(1 - \alpha_s) \times \alpha_b \times C_b] + [(1 - \alpha_b) \times \alpha_s \times C_s] + [\alpha_b \times \alpha_s \times B(C_b, C_s)]$$

(The simplification requires a substitution based on the alpha compositing formula, which is presented in the next section.) Thus, mathematically, the backdrop and source alphas control the influence of the backdrop and source colors, respectively, while their product controls the influence of the blend function. An alpha value of $\alpha_s = 0.0$ or $\alpha_b = 0.0$ results in no blend mode effect; setting $\alpha_s = 1.0$ and $\alpha_b = 1.0$ results in maximum blend mode effect.

7.2.6 Shape and Opacity Computations

As stated earlier, the alpha values that control the compositing process are defined as the product of shape and opacity:

$$\begin{aligned}\alpha_b &= f_b \times q_b \\ \alpha_r &= f_r \times q_r \\ \alpha_s &= f_s \times q_s\end{aligned}$$

This section examines the various shape and opacity values individually. Once again, keep in mind that conceptually these values are computed for every point on the page.

Source Shape and Opacity

Shape and opacity values can come from several sources. The transparency model provides for three independent sources for each; however, the PDF representation imposes some limitations on the ability to specify all of these sources independently (see Section 7.5.3, “Specifying Shape and Opacity”).

- *Object shape.* Elementary objects such as strokes, fills, and text have an intrinsic shape, whose value is 1.0 for points inside the object and 0.0 outside. Similarly, an image with an explicit mask (see “Explicit Masking” on page 277) has a shape that is 1.0 in the unmasked portions and 0.0 in the masked portions. The shape of a group object is the union of the shapes of the objects it contains.

Note: Mathematically, elementary objects have “hard” edges, with a shape value of either 0.0 or 1.0 at every point. However, when such objects are rasterized to device pixels, the shape values along the boundaries may be anti-aliased, taking on fractional values representing fractional coverage of those pixels. When such anti-aliasing is performed, it is important to treat the fractional coverage as shape rather than opacity.

- *Mask shape.* Shape values for compositing an object can be taken from an additional source, or *soft mask*, independent of the object itself. (See Section 7.4, “Soft Masks,” for a discussion of how such a mask might be generated.) The use of a soft mask to modify the shape of an object or group, called *soft clipping*, can produce effects such as a gradual transition between an object and its backdrop, as in a vignette.
- *Constant shape.* The source shape can be modified at every point by a scalar *shape constant*. This is merely a convenience, since the same effect could be achieved with a shape mask whose value is the same everywhere.
- *Object opacity.* Elementary objects have an opacity of 1.0 everywhere. The opacity of a group object is the result of the opacity computations for all of the objects it contains.
- *Mask opacity.* Opacity values, like shape values, can be provided by a soft mask independent of the object being composited.

- *Constant opacity.* The source opacity can be modified at every point by a scalar *opacity constant*. It is useful to think of this value as the “current opacity,” analogous to the current color used when painting elementary objects.

All of these shape and opacity inputs range in value from 0.0 to 1.0, with a default value of 1.0. The intent is that any of the inputs will make the painting operation more transparent as it goes toward 0.0. If more than one input goes toward 0.0, the effect is compounded. This is achieved mathematically by simply multiplying the three inputs of each type, producing intermediate values called the *source shape* and the *source opacity*:

$$f_s = f_j \times f_m \times f_k$$

$$q_s = q_j \times q_m \times q_k$$

where the variables have the meanings shown in Table 7.4.

TABLE 7.4 Variables used in the source shape and opacity formulas

VARIABLE	MEANING
f_s	Source shape
f_j	Object shape
f_m	Mask shape
f_k	Constant shape
q_s	Source opacity
q_j	Object opacity
q_m	Mask opacity
q_k	Constant opacity

Note: When an object is painted with a tiling pattern, the object shape and object opacity for points in the object’s interior are determined by those of corresponding points in the pattern, rather than being 1.0 everywhere (see Section 7.5.6, “Patterns and Transparency”).

Result Shape and Opacity

In addition to a result color, the painting operation also computes an associated *result shape* and *result opacity*. These computations are based on the *union function*

$$\begin{aligned}\text{Union}(b, s) &= 1 - [(1 - b) \times (1 - s)] \\ &= b + s - (b \times s)\end{aligned}$$

where b and s are the backdrop and source values to be composited. This is a generalization of the conventional concept of union for opaque shapes, and can be thought of as an “inverted multiplication”—a multiplication with the inputs and outputs complemented. The result tends toward 1.0: if either input is 1.0, the result will be 1.0.

The result shape and opacity are given by

$$\begin{aligned}f_r &= \text{Union}(f_b, f_s) \\ q_r &= \frac{\text{Union}(f_b \times q_b, f_s \times q_s)}{f_r}\end{aligned}$$

where the variables have the meanings shown in Table 7.5.

TABLE 7.5 Variables used in the result shape and opacity formulas

VARIABLE	MEANING
f_r	Result shape
f_b	Backdrop shape
f_s	Source shape
q_r	Result opacity
q_b	Backdrop opacity
q_s	Source opacity

These formulas can be interpreted as follows:

- The result shape is simply the union of the backdrop and source shapes.
- The result opacity is the union of the backdrop and source opacities, weighted by their respective shapes. The result is then normalized by the result shape, ensuring that when this shape and opacity are subsequently used together in another compositing operation, the opacity's contribution will be correctly represented.

Since alpha is just the product of shape and opacity, it can easily be shown that

$$\alpha_r = \text{Union}(\alpha_b, \alpha_s)$$

This formula can be used whenever the independent shape and opacity results are not needed.

7.2.7 Summary of Basic Compositing Computations

Below is a summary of all the computations presented in this section. They are given in an order such that no variable is used before it is computed; also, some of the formulas have been rearranged to simplify them. See Tables 7.1, 7.4, and 7.5 above for the meanings of the variables used in these formulas.

$$\begin{aligned} \text{Union}(b, s) &= 1 - [(1 - b) \times (1 - s)] \\ &= b + s - (b \times s) \end{aligned}$$

$$f_s = f_j \times f_m \times f_k$$

$$q_s = q_j \times q_m \times q_k$$

$$f_r = \text{Union}(f_b, f_s)$$

$$\alpha_b = f_b \times q_b$$

$$\alpha_s = f_s \times q_s$$

$$\alpha_r = \text{Union}(\alpha_b, \alpha_s)$$

$$q_r = \frac{\alpha_r}{f_r}$$

$$C_r = \left(1 - \frac{\alpha_s}{\alpha_r}\right) \times C_b + \frac{\alpha_s}{\alpha_r} \times [(1 - \alpha_b) \times C_s + \alpha_b \times B(C_b, C_s)]$$

7.3 Transparency Groups

A *transparency group* is a sequence of consecutive objects in a transparency stack that are collected together and composited to produce a single color, shape, and opacity at each point. The result is then treated as if it were a single object for subsequent compositing operations. This facilitates creating independent pieces of artwork, each composed of multiple objects, and then combining them, possibly with additional transparency effects applied during the combination. Groups can be nested within other groups to form a tree-structured group hierarchy.

The objects contained within a group are treated as a separate transparency stack, called the *group stack*. The objects in the stack are composited against some initial backdrop (discussed later), producing a composite color, shape, and opacity for the group as a whole. The result is an object whose shape is the union of the shapes of its constituent objects and whose color and opacity are the result of the compositing operations. This object is then composited with the group's backdrop in the usual way.

In addition to its computed color, shape, and opacity, the group as a whole can have several further attributes:

- All of the input variables that affect the compositing computation for individual objects can also be applied when compositing the group with its backdrop. These include mask and constant shape, mask and constant opacity, and blend mode.
- The group can be *isolated* or *non-isolated*, determining the initial backdrop against which its stack is composited.
- The group can be *knockout* or *non-knockout*, determining whether the objects within its stack are composited with one another or only with the group's backdrop.
- An isolated group can specify its own blending color space, independent of that of the group's backdrop.
- Instead of being composited onto the current page, a group's results can be used as a source of shape or opacity values for creating a *soft mask* (see Section 7.4, "Soft Masks").

The next section introduces some new notation for dealing with group compositing. Subsequent sections describe the group compositing formulas for a non-isolated, non-knockout group and the special properties of isolated and knockout groups.

7.3.1 Notation for Group Compositing Computations

Since we are now dealing with multiple objects at a time, it is useful to have some notation for distinguishing among them. Accordingly, the variables introduced earlier are altered to include a second-level subscript denoting an object's position in the transparency stack. Thus, for example, C_{s_i} stands for "source color of the i th object in the stack." The subscript 0 represents the initial backdrop; subscripts 1 to n denote the bottommost to topmost objects in an n -element stack. In addition, the subscripts b and r are dropped from the variables $C_b, f_b, q_b, \alpha_b, C_r, f_r, q_r$, and α_r ; other variables retain their mnemonic subscripts.

These conventions permit the compositing formulas to be restated as recurrence relations among the elements of a stack. For instance, the result of the color compositing computation for object i is denoted by C_i (formerly C_r). This computation takes as one of its inputs the immediate backdrop color, which is the result of the color compositing computation for object $i - 1$; this is denoted by C_{i-1} (formerly C_b).

The revised formulas for a simple n -element stack (not including any groups) are, for $i = 1, \dots, n$:

$$f_{s_i} = f_{j_i} \times f_{m_i} \times f_{k_i}$$

$$q_{s_i} = q_{j_i} \times q_{m_i} \times q_{k_i}$$

$$\alpha_{s_i} = f_{s_i} \times q_{s_i}$$

$$\alpha_i = \text{Union}(\alpha_{i-1}, \alpha_{s_i})$$

$$f_i = \text{Union}(f_{i-1}, f_{s_i})$$

$$q_i = \frac{\alpha_i}{f_i}$$

$$C_i = \left(1 - \frac{\alpha_{s_i}}{\alpha_i}\right) \times C_{i-1} + \frac{\alpha_{s_i}}{\alpha_i} \times [(1 - \alpha_{i-1}) \times C_{s_i} + \alpha_{i-1} \times B_i(C_{i-1}, C_{s_i})]$$

where the variables have the meanings shown in Table 7.6. Compare these formulas with those shown in Section 7.2.7, “Summary of Basic Compositing Computations.”

TABLE 7.6 Revised variables for the basic compositing formulas

VARIABLE	MEANING
f_{s_i}	Source shape for object i
f_{j_i}	Object shape for object i
f_{m_i}	Mask shape for object i
f_{k_i}	Constant shape for object i
f_i	Result shape after compositing object i
q_{s_i}	Source opacity for object i
q_{j_i}	Object opacity for object i
q_{m_i}	Mask opacity for object i
q_{k_i}	Constant opacity for object i
q_i	Result opacity after compositing object i
α_{s_i}	Source alpha for object i
α_i	Result alpha after compositing object i
C_{s_i}	Source color for object i
C_i	Result color after compositing object i
$B_i(C_{i-1}, C_{s_i})$	Blend function for object i

7.3.2 Group Structure and Nomenclature

As stated earlier, the elements of a group are treated as a separate transparency stack, the group stack. These objects are composited against a selected initial backdrop (to be described) and the resulting color, shape, and opacity are then treated as if they belonged to a single object. The resulting object is in turn composited with the group’s backdrop in the usual way.

This computation entails interpreting the stack as a tree. For an n -element group that begins at position i in the stack, it treats the next n objects as an n -element substack, whose elements are given an independent numbering of 1 to n . These objects are then removed from the object numbering in the parent (containing) stack and replaced by the group object, numbered i , followed by the remaining objects to be painted on top of the group, renumbered starting at $i + 1$. This operation applies recursively to any nested subgroups. Henceforth, the term *element* (denoted E_i) refers to a member of some group; it can itself be either an individual object or a contained subgroup.

From the perspective of a particular element in a nested group, there are three different backdrops of interest:

- *The group backdrop* is the result of compositing all elements up to but not including the first element in the group. (This definition is altered if the parent group is a knockout group; see Section 7.3.5, “Knockout Groups.”)
- *The initial backdrop* is a backdrop that is selected for compositing the group’s first element. This is either the same as the group backdrop (for a non-isolated group) or a fully transparent backdrop (for an isolated group).
- *The immediate backdrop* is the result of compositing all elements in the group up to but not including the current element.

When all elements in a group have been composited, the result is treated as if the group were a single object, which is then composited with the group backdrop. (Note that this operation occurs whether the initial backdrop chosen for compositing the elements of the group was the group backdrop or a transparent backdrop. There is a special correction to ensure that the backdrop’s contribution to the overall result is applied only once.)

7.3.3 Group Compositing Computations

The color and opacity of a group are defined by the *group compositing function*:

$$\langle C, f, \alpha \rangle = \text{Composite}(C_0, \alpha_0, G)$$

where the variables have the meanings shown in Table 7.7.

TABLE 7.7 Arguments and results of the group compositing function

VARIABLE	MEANING
G	The transparency group: a compound object consisting of all elements E_1, \dots, E_n of the group—the n constituent objects' colors, shapes, opacities, and blend modes
C_0	Color of the group's backdrop
C	Computed color of the group, to be used as the source color when the group itself is treated as an object
f	Computed shape of the group, to be used as the object shape when the group itself is treated as an object
α_0	Alpha of the group's backdrop
α	Computed alpha of the group, to be used as the object alpha when the group itself is treated as an object

Note that the opacity is not given explicitly as an argument or result of this function. Almost all of the computations use the product of shape and opacity (alpha) rather than opacity by itself, so it is usually convenient to work directly with shape and alpha, rather than shape and opacity. When needed, the opacity can be computed by dividing the alpha by the associated shape.

The result of applying the group compositing function is then treated as if it were a single object, which in turn is composited with the group's backdrop according to the usual formulas. In those formulas, the color, shape, and alpha (C , f , and α) calculated by the group compositing function are used, respectively, as the source color C_s , the object shape f_j , and the object alpha α_j .

The group compositing formulas for a non-isolated, non-knockout group are defined as follows:

- Initialization:

$$f_{g_0} = \alpha_{g_0} = 0.0$$

- For each group element $E_i \in G$ ($i = 1, \dots, n$):

$$\langle C_{s_i}, f_{j_i}, \alpha_{j_i} \rangle = \begin{cases} \text{Composite}(C_{i-1}, \alpha_{i-1}, E_i) & \text{if } E_i \text{ is a group} \\ \text{intrinsic color, shape, and (shape} \times \text{opacity) of } E_i & \text{otherwise} \end{cases}$$

$$f_{s_i} = f_{j_i} \times f_{m_i} \times f_{k_i}$$

$$\alpha_{s_i} = \alpha_{j_i} \times (f_{m_i} \times q_{m_i}) \times (f_{k_i} \times q_{k_i})$$

$$f_{g_i} = \text{Union}(f_{g_{i-1}}, f_{s_i})$$

$$\alpha_{g_i} = \text{Union}(\alpha_{g_{i-1}}, \alpha_{s_i})$$

$$\alpha_i = \text{Union}(\alpha_0, \alpha_{g_i})$$

$$C_i = \left(1 - \frac{\alpha_{s_i}}{\alpha_i}\right) \times C_{i-1} + \frac{\alpha_{s_i}}{\alpha_i} \times ((1 - \alpha_{i-1}) \times C_{s_i} + \alpha_{i-1} \times B_i(C_{i-1}, C_{s_i}))$$

- Result:

$$C = C_n + (C_n - C_0) \times \left(\frac{\alpha_0}{\alpha_{g_n}} - \alpha_0\right)$$

$$f = f_{g_n}$$

$$\alpha = \alpha_{g_n}$$

where the variables have the meanings shown in Table 7.8 (in addition to those in Table 7.7 above).

For an element E_i that is an elementary object, the color, shape, and alpha values C_{s_i} , f_{j_i} , and α_{j_i} are intrinsic attributes of the object itself. For an element that is itself a group, the group compositing function is applied recursively to the subgroup and the resulting C , f , and α values are used for its C_{s_i} , f_{j_i} , and α_{j_i} in the calculations for the parent group.

TABLE 7.8 Variables used in the group compositing formulas

VARIABLE	MEANING
E_i	Element i of the group: a compound variable representing the element's color, shape, opacity, and blend mode
f_{s_i}	Source shape for element E_i
f_{j_i}	Object shape for element E_i
f_{m_i}	Mask shape for element E_i
f_{k_i}	Constant shape for element E_i
f_{g_i}	Group shape: the accumulated source shapes of group elements E_1 to E_i , excluding the initial backdrop
q_{m_i}	Mask opacity for element E_i
q_{k_i}	Constant opacity for element E_i
α_{s_i}	Source alpha for element E_i
α_{j_i}	Object alpha for element E_i : the product of its object shape and object opacity
α_{g_i}	Group alpha: the accumulated source alphas of group elements E_1 to E_i , excluding the initial backdrop
α_i	Accumulated alpha after compositing element E_i , including the initial backdrop
C_{s_i}	Source color for element E_i
C_i	Accumulated color after compositing element E_i , including the initial backdrop
$B_i(C_{i-1}, C_{s_i})$	Blend function for element E_i

Note that the elements of a group are composited onto a backdrop that includes the group's initial backdrop. This is done to achieve the correct effects of the blend modes, most of which are dependent on both the backdrop and source colors being blended. (This feature is what distinguishes non-isolated groups from isolated groups, discussed in the next section.)

Special attention should be directed to the formulas at the end that compute the final results, C , f , and α , of the group compositing function. Essentially, these formulas remove the contribution of the group backdrop from the computed results. This ensures that when the group itself is subsequently composited with that backdrop (possibly with additional shape or opacity inputs or a different blend mode), the backdrop's contribution is included only once.

For color, the backdrop removal is accomplished by an explicit calculation, whose effect is essentially the reverse of compositing with the **Normal** blend mode. The formula is a simplification of the following formulas, which present this operation more intuitively:

$$\phi_b = \frac{(1 - \alpha_{g_n}) \times \alpha_0}{\text{Union}(\alpha_0, \alpha_{g_n})}$$

$$C = \frac{C_n - \phi_b \times C_0}{1 - \phi_b}$$

where ϕ_b is the *backdrop fraction*, the relative contribution of the backdrop color to the overall color.

For shape and alpha, backdrop removal is accomplished by maintaining two sets of variables to hold the accumulated values. The group shape and alpha, f_{g_i} and α_{g_i} , accumulate only the shape and alpha of the group elements, excluding the group backdrop; their final values become the group results returned by the group compositing function. The complete alpha, α_i , includes the backdrop contribution as well; its value is used in the color compositing computations. (There is never any need to compute the corresponding complete shape, f_i , that includes the backdrop contribution.)

As a result of these corrections, the effect of compositing objects as a group is the same as that of compositing them separately (without grouping) if the following conditions hold:

- The group is non-isolated and has the same knockout attribute as its parent group (see Sections 7.3.4, “Isolated Groups,” and 7.3.5, “Knockout Groups”).
- When compositing the group's results with the group backdrop, the **Normal** blend mode is used and the shape and opacity inputs are always 1.0.

7.3.4 Isolated Groups

An *isolated group* is one whose elements are composited onto a fully transparent initial backdrop rather than onto the group's backdrop. The resulting source color, object shape, and object alpha for the group are therefore independent of the group backdrop. The only interaction with the group backdrop occurs when the group's computed color, shape, and alpha are then composited with it.

In particular, the special effects produced by the blend modes of objects within the group take into account only the intrinsic colors and opacities of those objects; they are not influenced by the group's backdrop. For example, applying the **Multiply** blend mode to an object in the group will produce a darkening effect on other objects lower in the group's stack, but not on the group's backdrop.

Plate 17 illustrates this effect for a group consisting of four overlapping circles in a light gray color ($C = M = Y = 0.0$; $K = 0.15$). The circles are painted within the group with opacity 1.0 in the **Multiply** blend mode; the group itself is painted against its backdrop in **Normal** blend mode. In the top row, the group is isolated and thus does not interact with the rainbow backdrop; in the bottom row, it is non-isolated and composites with the backdrop. The plate also illustrates the difference between knockout and non-knockout groups (see Section 7.3.5, "Knockout Groups").

The effect of an isolated group can be represented by a simple object that directly specifies a color, shape, and opacity at each point. This so-called "flattening" of an isolated group is sometimes useful for importing and exporting fully composited artwork in applications. Furthermore, a group that specifies an explicit blending color space must be an isolated group.

For an isolated group, the group compositing formulas are altered by simply adding one statement to the initialization:

$$\alpha_0 = 0.0 \quad \text{if the group is isolated}$$

That is, the initial backdrop on which the elements of the group are composited is transparent, rather than inherited from the group's backdrop. This substitution also makes C_0 undefined, but the normal compositing formulas take care of that. Also, the result computation for C automatically simplifies to $C = C_n$, since there is no backdrop contribution to be factored out.

7.3.5 Knockout Groups

In a knockout group, each individual element is composited with the group's initial backdrop, rather than with the stack of preceding elements in the group. When objects have binary shapes (1.0 for “inside,” 0.0 for “outside”), each object overwrites (“knocks out”) the effects of any earlier elements it overlaps within the same group. At any given point, only the topmost object enclosing the point contributes to the result color and opacity of the group as a whole.

Plate 17, already discussed above in Section 7.3.4, “Isolated Groups,” illustrates the difference between knockout and non-knockout groups. In the left column, the four overlapping circles are defined as a knockout group and therefore do not composite with each other within the group; in the right column, they form a non-knockout group and thus do composite with each other. In each column, the upper and lower figures depict an isolated and a non-isolated group, respectively.

This model is similar to the opaque imaging model, except that the “topmost object wins” rule applies to both the color and the opacity. Knockout groups are useful in composing a piece of artwork from a collection of overlapping objects, where the topmost object in any overlap completely obscures those beneath. At the same time, the topmost object interacts with the group's initial backdrop in the usual way, with its opacity and blend mode applied as appropriate.

The concept of “knockout” is generalized to accommodate fractional shape values. In that case, the immediate backdrop is only partially knocked out and replaced by only a fraction of the result of compositing the object with the initial backdrop.

The restated group compositing formulas deal with knockout groups by introducing a new variable, b , which is a subscript that specifies which previous result to use as the backdrop in the compositing computations: 0 in a knockout group or $i - 1$ in a non-knockout group. When $b = i - 1$, the formulas simplify to the ones given in Section 7.3.3, “Group Compositing Computations.”

In the general case, the computation proceeds in two stages:

1. Composite the object with the group's initial backdrop, but disregarding the object's shape and using a source shape value of 1.0 everywhere. This produces unnormalized temporary alpha and color results, α_t and C_t . (For color, this

computation is essentially the same as the unsimplified color compositing formula given in Section 7.2.5, “Interpretation of Alpha,” but using a source shape of 1.0.)

$$\alpha_t = \text{Union}(\alpha_{g_b}, q_{s_i})$$

$$C_t = (1 - q_{s_i}) \times \alpha_b \times C_b + q_{s_i} \times ((1 - \alpha_b) \times C_{s_i} + \alpha_b \times B_i(C_b, C_{s_i}))$$

2. Compute a weighted average of this result with the object’s immediate backdrop, using the source shape as the weighting factor. Then normalize the result color by the result alpha:

$$\alpha_{g_i} = (1 - f_{s_i}) \times \alpha_{g_{i-1}} + f_{s_i} \times \alpha_t$$

$$\alpha_i = \text{Union}(\alpha_0, \alpha_{g_i})$$

$$C_i = \frac{(1 - f_{s_i}) \times \alpha_{i-1} \times C_{i-1} + f_{s_i} \times \alpha_t}{\alpha_i}$$

This averaging computation is performed for both color and alpha. The formulas above show this averaging directly; those in Section 7.3.7, “Summary of Group Compositing Computations,” are slightly altered to use source shape and alpha rather than source shape and opacity, avoiding the need to compute a source opacity value explicitly. (Note that C_t there is slightly different from C_t above: it is premultiplied by f_{s_i} .)

The extreme values of the source shape produce the straightforward knockout effect. That is, a shape value of 1.0 (“inside”) yields the color and opacity that result from compositing the object with the initial backdrop. A shape value of 0.0 (“outside”) leaves the previous group results unchanged. The existence of the knockout feature is the main reason for maintaining a separate shape value, rather than only a single alpha that combines shape and opacity. The separate shape value must be computed in any group that is subsequently used as an element of a knockout group.

A knockout group can be isolated or non-isolated; that is, *isolated* and *knockout* are independent attributes. A non-isolated knockout group composites its top-most enclosing element with the group’s backdrop; an isolated knockout group composites the element with a transparent backdrop.

Note: When a non-isolated group is nested within a knockout group, the initial backdrop of the inner group is the same as that of the outer group; it is not the immediate backdrop of the inner group. This behavior, although perhaps unexpected, is a consequence of the group compositing formulas when $b = 0$.

7.3.6 Page Group

All of the elements painted directly onto a page—both top-level groups and top-level objects that are not part of any group—are treated as if they were contained in a transparency group P , which in turn is composited with a context-dependent backdrop. This group is called the *page group*.

The page group can be treated in two distinctly different ways:

- Ordinarily, the page is imposed directly on an output medium, such as paper or a display screen. The page group is treated as an isolated group, whose results are then composited with a backdrop color appropriate for the medium. The backdrop is nominally white, although varying according to the actual properties of the medium. However, some applications may choose to provide a different backdrop, such as a checkerboard or grid to aid in visualizing the effects of transparency in the artwork.
- A “page” of a PDF file can be treated as a graphics object to be used as an element of a page of some other document. This case arises, for example, when placing a PDF file containing a piece of artwork produced by Illustrator into a page layout produced by InDesign™. In this situation, the PDF “page” is not composited with the media color; instead, it is treated as an ordinary transparency group, which can be either isolated or non-isolated and is composited with its backdrop in the normal way.

The remainder of this section pertains only to the first use of the page group, where it is to be imposed directly on the medium.

The color C of the page at a given point is defined by a simplification of the general group compositing formula:

$$\langle C_g, f_g, \alpha_g \rangle = \text{Composite}(U, 0, P)$$

$$C = (1 - \alpha_g) \times W + \alpha_g \times C_g$$

where the variables have the meanings shown in Table 7.9. The first formula computes the color and alpha for the group given a transparent backdrop—in effect, treating P as an isolated group. The second formula composites the results with the context-dependent backdrop (using the equivalent of the **Normal** blend mode).

TABLE 7.9 Variables used in the page group compositing formulas

VARIABLE	MEANING
P	The page group, consisting of all elements E_1, \dots, E_n in the page's top-level stack
C_g	Computed color of the page group
f_g	Computed shape of the page group
α_g	Computed alpha of the page group
C	Computed color of the page
W	Initial color of the page (nominally white, but may vary depending on the properties of the medium or the needs of the application)
U	An undefined color (which is not used, since the α_0 argument of Composite is 0)

If not otherwise specified, the page group's color space is inherited from the native color space of the output device—that is, a device color space, such as **DeviceRGB** or **DeviceCMYK**. It is often preferable to specify an explicit color space, particularly a CIE-based space, to ensure more predictable results of the compositing computations within the page group. In this case, all page-level compositing is done in the specified color space, with the entire result then converted to the native color space of the output device before being composited with the context-dependent backdrop. This case also arises when the page is not actually being rendered but is converted to a “flattened” representation in an opaque imaging model, such as PostScript.

7.3.7 Summary of Group Compositing Computations

The following restatement of the group compositing formulas also takes isolated groups and knockout groups into account. See Tables 7.7 and 7.8 on pages 429 and 431 for the meanings of the variables.

$$\langle C, f, \alpha \rangle = \text{Composite}(C_0, \alpha_0, G)$$

- Initialization:

$$f_{g_0} = \alpha_{g_0} = 0$$

$$\alpha_0 = 0 \quad \text{if the group is isolated}$$

- For each group element $E_i \in G$ ($i = 1, \dots, n$):

$$b = \begin{cases} 0 & \text{if the group is knockout} \\ i - 1 & \text{otherwise} \end{cases}$$

$$\langle C_{s_i}, f_{j_i}, \alpha_{j_i} \rangle = \begin{cases} \text{Composite}(C_b, \alpha_b, E_i) & \text{if } E_i \text{ is a group} \\ \text{intrinsic color, shape, and (shape} \times \text{opacity) of } E_i & \text{otherwise} \end{cases}$$

$$f_{s_i} = f_{j_i} \times f_{m_i} \times f_{k_i}$$

$$\alpha_{s_i} = \alpha_{j_i} \times (f_{m_i} \times q_{m_i}) \times (f_{k_i} \times q_{k_i})$$

$$f_{g_i} = \text{Union}(f_{g_{i-1}}, f_{s_i})$$

$$\alpha_{g_i} = (1 - f_{s_i}) \times \alpha_{g_{i-1}} + (f_{s_i} - \alpha_{s_i}) \times \alpha_{g_b} + \alpha_{s_i}$$

$$\alpha_i = \text{Union}(\alpha_0, \alpha_{g_i})$$

$$C_t = (f_{s_i} - \alpha_{s_i}) \times \alpha_b \times C_b + \alpha_{s_i} \times ((1 - \alpha_b) \times C_{s_i} + \alpha_b \times B_i(C_b, C_{s_i}))$$

$$C_i = \frac{(1 - f_{s_i}) \times \alpha_{i-1} \times C_{i-1} + C_t}{\alpha_i}$$

- Result:

$$C = C_n + (C_n - C_0) \times \left(\frac{\alpha_0}{\alpha_{g_n}} - \alpha_0 \right)$$

$$f = f_{g_n}$$

$$\alpha = \alpha_{g_n}$$

Note: Once again, keep in mind that these formulas are in their most general form. They can be significantly simplified when some sources of shape and opacity are not present or when shape and opacity need not be maintained separately. Furthermore, in each specific type of group (isolated or not, knockout or not), some terms of these formulas cancel or drop out. An efficient implementation should use the simplified derived formulas.

7.4 Soft Masks

As stated in earlier sections, the shape and opacity values used in compositing an object can include components called the mask shape (f_m) and mask opacity (q_m), which originate from a source independent of the object itself. Such an independent source, called a *soft mask*, defines values that can vary across different points on the page. The word *soft* emphasizes that the mask value at a given point is not limited to just 0.0 or 1.0, but can take on intermediate fractional values as well. Such a mask is typically the only means of providing position-dependent opacity values, since elementary objects do not have intrinsic opacity of their own.

A mask used as a source of shape values is also called a *soft clip*, by analogy with the “hard” clipping path of the opaque imaging model (see Section 4.4.3, “Clipping Path Operators”). The soft clip is a generalization of the hard clip: a hard clip can be represented as a soft clip having shape values of 1.0 inside and 0.0 outside the clipping path. Everywhere inside a hard clipping path, the source object’s color replaces the backdrop; everywhere outside, the backdrop shows through unchanged. With a soft clip, by contrast, a gradual transition can be created between an object and its backdrop, as in a vignette.

A mask can be defined by creating a transparency group and painting objects into it, thereby defining color, shape, and opacity in the usual way. The resulting group can then be used to derive the mask in either of two ways, as described in the following sections.

7.4.1 Deriving a Soft Mask from Group Alpha

In the first method of defining a soft mask, the color, shape, and opacity of a transparency group G are first computed by the usual formula

$$\langle C, f, \alpha \rangle = \text{Composite}(C_0, \alpha_0, G)$$

where C_0 and α_0 represent an arbitrary backdrop whose value does not contribute to the eventual result. The C , f , and α results are the group's color, shape, and alpha, respectively, with the backdrop factored out.

The mask value at each point is then derived from the alpha of the group. Since the group's color is not used in this case, there is no need to compute it. The alpha value is passed through a separately specified transfer function, allowing the masking effect to be customized.

7.4.2 Deriving a Soft Mask from Group Luminosity

The second method of deriving a soft mask from a transparency group begins by compositing the group with a fully opaque backdrop of some selected color. The mask value at any given point is then defined to be the luminosity of the resulting color. This allows the mask to be derived from the shape and color of an arbitrary piece of artwork drawn with ordinary painting operators.

The color C used to create the mask from a group G is defined by

$$\langle C_g, f_g, \alpha_g \rangle = \text{Composite}(C_0, 1, G)$$

$$C = (1 - \alpha_g) \times C_0 + \alpha_g \times C_g$$

where C_0 is the selected backdrop color.

G can be any kind of group—isolated or not, knockout or not—producing various effects on the C result in each case. The color C is then converted to luminosity in one of the following ways, depending on the group's color space:

- For CIE-based spaces, convert to the CIE 1931 XYZ space and use the Y component as the luminosity. This produces a colorimetrically correct luminosity. In the case of a PDF **CalRGB** space, the formula is

$$Y = Y_A \times A^{G_R} + Y_B \times B^{G_G} + Y_C \times C^{G_B}$$

using components of the **Gamma** and **Matrix** entries of the color space dictionary (see Table 4.14 on page 185). An analogous computation applies to other CIE-based color spaces.

- For device color spaces, convert the color to **DeviceGray** by device-dependent means and use the resulting gray value as the luminosity, with no compensation for gamma or other color calibration. This method makes no pretense of colorimetric correctness; it merely provides a numerically simple means to produce continuous-tone mask values. Here are some recommended formulas for converting from **DeviceRGB** and **DeviceCMYK**, respectively:

$$Y = 0.30 \times R + 0.59 \times G + 0.11 \times B$$

$$\begin{aligned} Y &= 0.30 \times (1 - C) \times (1 - K) \\ &\quad + 0.59 \times (1 - M) \times (1 - K) \\ &\quad + 0.11 \times (1 - Y) \times (1 - K) \end{aligned}$$

Following this conversion, the result is passed through a separately specified transfer function, allowing the masking effect to be customized.

The backdrop color most likely to be useful is black, which causes any areas outside the group's shape to end up with zero luminosity values in the resulting mask. If the contents of the group are viewed as a positive mask, this produces the results that would be expected with respect to points outside the shape.

7.5 Specifying Transparency in PDF

The preceding sections have presented the transparent imaging model at an abstract level, with little mention of its representation in PDF. This section describes the facilities available for specifying transparency in PDF 1.4.

7.5.1 Specifying Source and Backdrop Colors

Single graphics objects, as defined in Section 4.1, “Graphics Objects,” are treated as elementary objects for transparency compositing purposes (subject to special treatment for text objects, as described in Section 5.2.7, “Text Knockout”). That is, all of a given object is considered to be one element of a transparency stack; portions of an object are not composited with one another, even if they are described in a way that would seem to cause overlaps (such as a self-intersecting path, combined fill and stroke of a path, or a shading pattern containing an overlap or fold-over). An object's source color C_s , used in the color compositing formula, is specified in the same way as in the opaque imaging model: via the current color in the graphics state or the source samples in an image. The backdrop color C_b is the result of previous painting operations.

7.5.2 Specifying Blending Color Space and Blend Mode

The blending color space is an attribute of the transparency group within which an object is painted; its specification is described below in Section 7.5.5, “Transparency Group XObjects.” The page as a whole is also treated as a group, the *page group* (see Section 7.3.6, “Page Group”), with a color space attribute of its own. If not otherwise specified, the page group’s color space is inherited from the native color space of the output device.

The blend mode $B(C_b, C_s)$ is determined by the *current blend mode* parameter in the graphics state (see Section 4.3, “Graphics State”), set via the **BM** entry in a graphics state parameter dictionary (Section 4.3.4, “Graphics State Parameter Dictionaries”). Its value is either a name object, designating one of the standard blend modes listed in Tables 7.2 and 7.3 on pages 417 and 419, or an array of such names. In the latter case, the viewer application should use the first blend mode in the array that it recognizes (or **Normal** if it recognizes none of them). This allows new blend modes to be introduced in the future while providing reasonable fallback behavior by viewer applications that do not recognize them. (See implementation note 52 in Appendix H.)

Note: The current blend mode always applies to process color components, but only sometimes to spot colorants; see “Blend Modes and Overprinting” on page 459 for details.

7.5.3 Specifying Shape and Opacity

As discussed under “Source Shape and Opacity” on page 421, the shape (f) and opacity (q) values used in the compositing computation can come from a variety of sources:

- The intrinsic shape (f_j) and opacity (q_j) of the object being composited
- A separate shape (f_m) or opacity (q_m) mask independent of the object itself
- A scalar shape (f_k) or opacity (q_k) constant to be added at every point

The following sections describe how each of these shape and opacity sources are specified in PDF.

Object Shape and Opacity

The shape value f_j of an object painted with PDF painting operators is defined as follows:

- For objects defined by a path or a glyph and painted in a uniform color with a path-painting or text-showing operator (Sections 4.4.2, “Path-Painting Operators,” and 5.3.2, “Text-Showing Operators”), the shape is always 1.0 inside and 0.0 outside the path.
- For images (Section 4.8, “Images”), the shape is nominally 1.0 inside the image rectangle and 0.0 outside it; this can be further modified by an explicit or color key mask (“Explicit Masking” on page 277 and “Color Key Masking” on page 277).
- For image masks (“Stencil Masking” on page 276), the shape is 1.0 for painted areas and 0.0 for masked areas.
- For objects painted with a tiling pattern (Section 4.6.2, “Tiling Patterns”) or a shading pattern (Section 4.6.3, “Shading Patterns”), the shape is further constrained by the objects that define the pattern (see Section 7.5.6, “Patterns and Transparency”).
- For objects painted with the **sh** operator (“Shading Operator” on page 232), the shape is 1.0 inside and 0.0 outside the bounds of the shading’s painting geometry, disregarding the **Background** entry in the shading dictionary (see “Shading Dictionaries” on page 233).

All elementary objects have an intrinsic opacity q_j of 1.0 everywhere. Any desired opacity less than 1.0 must be applied by means of an opacity mask or constant, as described in the following sections.

Mask Shape and Opacity

At most one mask input—called a *soft mask*, or *alpha mask*—can be provided to any PDF compositing operation. The mask can serve as a source of either shape (f_m) or opacity (q_m) values, depending on the setting of the *alpha source* parameter in the graphics state (see Section 4.3, “Graphics State”). This is a boolean flag, set with the **AIS** (“alpha is shape”) entry in a graphics state parameter dictionary (Section 4.3.4, “Graphics State Parameter Dictionaries”): **true** if the soft mask contains shape values, **false** for opacity.

The soft mask can be specified in either of two ways:

- The *current soft mask* parameter in the graphics state, set with the **SMask** entry in a graphics state parameter dictionary, contains a *soft-mask dictionary* (see “Soft-Mask Dictionaries” on page 445) defining the contents of the mask. The name **None** may be specified in place of a soft-mask dictionary, denoting the absence of a soft mask; in this case, the mask shape or opacity is implicitly 1.0 everywhere. (See implementation note 52 in Appendix H.)
- An image XObject can contain its own *soft-mask image* in the form of a subsidiary image XObject in the **SMask** entry of the image dictionary (see Section 4.8.4, “Image Dictionaries”). This mask, if present, overrides any explicit or color key mask specified by the image dictionary’s **Mask** entry; either form of mask in the image dictionary overrides the current soft mask in the graphics state. (See implementation note 53 in Appendix H.)

Note: The current soft mask in the graphics state is intended to be used to clip only a single object at a time (either an elementary object or a transparency group). If a soft mask is applied when painting two or more overlapping objects, the effect of the mask will multiply with itself in the area of overlap (except in a knockout group), producing a result shape or opacity that is probably not what is intended. To apply a soft mask to multiple objects, it is usually best to define the objects as a transparency group and apply the mask to the group as a whole. These considerations also apply to the current alpha constant (see the next section).

Constant Shape and Opacity

The *current alpha constant* parameter in the graphics state (see Section 4.3, “Graphics State”) specifies two scalar values—one for strokes and one for all other painting operations—to be used for the constant shape (f_k) or constant opacity (q_k) component in the color compositing formulas. This parameter can be thought of as analogous to the current color used when painting elementary objects. (Note, however, that the nonstroking alpha constant is also applied when painting a transparency group’s results onto its backdrop; see also implementation note 52 in Appendix H.)

The stroking and nonstroking alpha constants are set, respectively, by the **CA** and **ca** entries in a graphics state parameter dictionary (see Section 4.3.4, “Graphics State Parameter Dictionaries”). As described above for the soft mask, the alpha source flag in the graphics state determines whether the alpha constants are interpreted as shape values (**true**) or opacity values (**false**).

Note: The note at the end of “Mask Shape and Opacity,” above, applies to the current alpha constant parameter as well as the current soft mask.

7.5.4 Specifying Soft Masks

As noted under “Mask Shape and Opacity” on page 443, soft masks for use in compositing computations can be specified in either of two ways: as a soft-mask dictionary in the current soft mask parameter of the graphics state or as a soft-mask image associated with a sampled image. The following sections describe these two methods of specifying soft masks in more detail.

Soft-Mask Dictionaries

The most common way of defining a soft mask is with a *soft-mask dictionary* specified as the current soft mask in the graphics state (see Section 4.3, “Graphics State”). Table 7.10 shows the contents of this type of dictionary. (See implementation note 52 in Appendix H.)

The mask values are derived from those of a transparency group, using one of the two methods described in Sections 7.4.1, “Deriving a Soft Mask from Group Alpha,” and 7.4.2, “Deriving a Soft Mask from Group Luminosity.” The group is defined by a transparency group XObject (see Section 7.5.5, “Transparency Group XObjects”) designated by the **G** entry in the soft-mask dictionary. The **S** (subtype) entry specifies which of the two derivation methods to use:

- If the subtype is **Alpha**, the transparency group XObject **G** is evaluated to compute a group alpha only; the colors of the constituent objects are ignored and the color compositing computations are not performed. The transfer function **TR** is then applied to the computed group alpha to produce the mask values. Outside the bounding box of the transparency group, the mask value is the result of applying the transfer function to the input value 0.0.
- If the subtype is **Luminosity**, the transparency group XObject **G** is composited with a fully opaque backdrop whose color is everywhere defined by the soft-mask dictionary’s **BC** entry. The computed result color is then converted to a single-component luminosity value and the transfer function **TR** is applied to this luminosity to produce the mask values. Outside the transparency group’s bounding box, the mask value is derived by transforming the **BC** color to luminosity and applying the transfer function to the result.

TABLE 7.10 Entries in a soft-mask dictionary

KEY	TYPE	VALUE
Type	name	<i>(Optional)</i> The type of PDF object that this dictionary describes; if present, must be Mask for a soft-mask dictionary.
S	name	<i>(Required)</i> A subtype specifying the method to be used in deriving the mask values from the transparency group specified by the G entry: <ul style="list-style-type: none"> Alpha Use the group’s computed alpha, disregarding its color (see Section 7.4.1, “Deriving a Soft Mask from Group Alpha”). Luminosity Convert the group’s computed color to a single-component luminosity value (see Section 7.4.2, “Deriving a Soft Mask from Group Luminosity”).
G	stream	<i>(Required)</i> A transparency group XObject (see Section 7.5.5, “Transparency Group XObjects”) to be used as the source of alpha or color values for deriving the mask. If the subtype S is Luminosity , the group attributes dictionary must contain a CS entry defining the color space in which the compositing computation is to be performed.
BC	array	<i>(Optional)</i> An array of component values specifying the color to be used as the backdrop against which to composite the transparency group XObject G . This entry is consulted only if the subtype S is Luminosity . The array consists of n numbers, where n is the number of components in the color space specified by the CS entry in the group attributes dictionary (see Section 7.5.5, “Transparency Group XObjects”). Default value: the color space’s initial value, representing black.
TR	function or name	<i>(Optional)</i> A function object (see Section 3.9, “Functions”) specifying the transfer function to be used in deriving the mask values. The function accepts one input, the computed group alpha or luminosity (depending on the value of the subtype S), and returns one output, the resulting mask value. Both the input and output must be in the range 0.0 to 1.0; if the computed output falls outside this range, it is forced to the nearest valid value. The name Identity may be specified in place of a function object to designate the identity function. Default value: Identity .

The mask’s coordinate system is defined by concatenating the transformation matrix specified by the **Matrix** entry in the transparency group’s form dictionary (see Section 4.9.1, “Form Dictionaries”) with the current transformation matrix at the moment the soft mask is established in the graphics state with the **gs** operator.

Note: In a transparency group *XObject* that defines a soft mask, spot color components are never available, even if they are available in the group or page on which the soft mask is used. If the group *XObject*'s content stream specifies a **Separation** or **DeviceN** color space that uses spot color components, the alternate color space will be substituted (see “Separation Color Spaces” on page 201 and “DeviceN Color Spaces” on page 205).

Soft-Mask Images

The second way to define a soft mask is by associating a *soft-mask image* with an image *XObject*. This is a subsidiary image *XObject* specified in the **SMask** entry of the parent *XObject*'s image dictionary (see Section 4.8.4, “Image Dictionaries”; see also implementation note 53 in Appendix H). Entries in the subsidiary image dictionary for such a soft-mask image have the same format and meaning as in that of an ordinary image *XObject* (as described in Table 4.35 on page 267), subject to the restrictions listed in Table 7.11. This type of image dictionary can also optionally contain an additional entry, **Matte**, discussed below.

When an image is accompanied by a soft-mask image, it is sometimes advantageous for the image data to be *preblended* with some background color, called the *matte color*. Each image sample represents a weighted average of the original source color and the matte color, using the corresponding mask sample as the weighting factor. (This is a generalization of a technique commonly called “pre-multiplied alpha.”)

If the image data is preblended, the matte color must be specified by a **Matte** entry in the soft-mask image dictionary (see Table 7.12). The preblending computation, performed independently for each component, is as follows:

$$c' = m + \alpha \times (c - m)$$

where

c' is the value to be provided in the image source data

c is the original image component value

m is the matte color component value

α is the corresponding mask sample

*Note: This computation uses actual color component values, with the effects of the **Filter** and **Decode** transformations already performed. The computation is the same whether the color space is additive or subtractive.*

TABLE 7.11 Restrictions on the entries in a soft-mask image dictionary

KEY	RESTRICTION
Type	If present, must be XObject .
Subtype	Must be Image .
Width	If a Matte entry (see Table 7.12, below) is present, must be the same as the Width value of the parent image; otherwise independent of it. Both images are mapped to the unit square in user space (as are all images), whether or not the samples coincide individually.
Height	Same considerations as for Width .
ColorSpace	Required; must be DeviceGray .
BitsPerComponent	Required.
Intent	Ignored.
ImageMask	Must be false or absent.
Mask	Must be absent.
SMask	Must be absent.
Decode	Default value: [0 1].
Interpolate	Optional.
Alternates	Ignored.
Name	Ignored.
StructParent	Ignored.
ID	Ignored.
OPI	Ignored.

TABLE 7.12 Additional entry in a soft-mask image dictionary

KEY	TYPE	VALUE
Matte	array	(<i>Optional; PDF 1.4</i>) An array of component values specifying the matte color with which the image data in the parent image has been preblended. The array consists of n numbers, where n is the number of components in the color space specified by the ColorSpace entry in the parent image's image dictionary; the numbers must be valid color components in that color space. If this entry is absent, the image data is not preblended.

When preblended image data is used in transparency blending and compositing computations, the results are the same as if the original, unblended image data were used and no matte color were specified. In particular, the inputs to the blend function are the original color values. This may sometimes require the viewer application to invert the formula shown above in order to derive c from c' . If the resulting c value lies outside the range of color component values for the image color space, the results are unpredictable.

The preblending computation is done in the color space specified by the parent image's **ColorSpace** entry. This is independent of the group color space into which the image may be painted; if a color conversion is required, inversion of the preblending must precede the color conversion. If the image color space is an **Indexed** space (see "Indexed Color Spaces" on page 199), it is the color values in the color table (not the index values themselves) that are preblended.

7.5.5 Transparency Group XObjects

A transparency group is represented in PDF as a special type of group XObject (see Section 4.9.2, "Group XObjects") called a *transparency group XObject*. A group XObject is in turn a type of form XObject, distinguished by the presence of a **Group** entry in its form dictionary (see Section 4.9.1, "Form Dictionaries"). The value of this entry is a subsidiary *group attributes dictionary* defining the properties of the group. The format and meaning of the dictionary's contents are determined by its *group subtype*, specified by the dictionary's **S** entry; those for a transparency group (subtype **Transparency**) are shown in Table 7.13.

Note: A page object (see "Page Objects" on page 87) may also have a **Group** entry, whose value is a group attributes dictionary specifying the attributes of the page

group (see Section 7.3.6, “Page Group”). Some of the dictionary entries are interpreted slightly differently for a page group than for a transparency group XObject; see their descriptions in the table for details.

TABLE 7.13 Additional entries specific to a transparency group attributes dictionary

KEY	TYPE	VALUE
S	name	(Required) The <i>group subtype</i> , which identifies the type of group whose attributes this dictionary describes; must be Transparency for a transparency group.
CS	name or array	(Sometimes required, as discussed below) The group color space, which is used for the following purposes: <ul style="list-style-type: none"> • As the color space into which colors are converted when painted into the group • As the blending color space in which objects are composited within the group (see Section 7.2.3, “Blending Color Space”) • As the color space of the group as a whole when it in turn is painted as an object onto its backdrop

The group color space may be any device or CIE-based color space that treats its components as independent additive or subtractive values in the range 0.0 to 1.0, subject to the restrictions described in Section 7.2.3, “Blending Color Space.” These restrictions exclude **Lab** and lightness-chromaticity **ICCBased** color spaces, as well as the special color spaces **Pattern**, **Indexed**, **Separation**, and **DeviceN**. Device color spaces are subject to remapping according to the **Default-Gray**, **DefaultRGB**, and **DefaultCMYK** entries in the **ColorSpace** subdictionary of the current resource dictionary (see “Default Color Spaces” on page 194).

Ordinarily, the **CS** entry is allowed only for isolated transparency groups (those for which **I**, below, is **true**) and even then it is optional. However, this entry is required in the group attributes dictionary for any transparency group XObject that has no parent group or page from which to inherit—in particular, one that is the value of the **G** entry in a soft-mask dictionary of subtype **Luminosity** (see “Soft-Mask Dictionaries” on page 445).

In addition, it is always permissible to specify **CS** in the group attributes dictionary associated with a page object, even if **I** is **false** or absent. In the normal case in which the page is imposed directly on the output medium, the page group is effectively isolated regardless of the **I** value, and the specified **CS** value is therefore honored. But if the page is in turn used as an element of some other page and if the group is non-isolated, **CS** is ignored and the color space is inherited from the actual backdrop with which the page is composited (see Section 7.3.6, “Page Group”).

Default value: the color space of the parent group or page into which this transparency group is painted. (The parent’s color space in turn can be either explicitly specified or inherited.)

Note: For a transparency group XObject used as an annotation appearance (see Section 8.4.4, “Appearance Streams”), the default color space is inherited from the page on which the annotation appears.

I	boolean	<p>(Optional) A flag specifying whether the transparency group is isolated (see Section 7.3.4, “Isolated Groups”). If this flag is true, objects within the group are composited against a fully transparent initial backdrop; if false, they are composited against the group’s backdrop. Default value: false.</p> <p>In the group attributes dictionary for a page, the interpretation of this entry is slightly altered. In the normal case in which the page is imposed directly on the output medium, the page group is effectively isolated and the specified I value is ignored. But if the page is in turn used as an element of some other page, it is treated as if it were a transparency group XObject; the I value is interpreted in the normal way to determine whether the page group is isolated.</p>
K	boolean	<p>(Optional) A flag specifying whether the transparency group is a knockout group (see Section 7.3.5, “Knockout Groups”). If this flag is false, later objects within the group are composited with earlier ones with which they overlap; if true, they are composited with the group’s initial backdrop and overwrite (“knock out”) any earlier overlapping objects. Default value: false.</p>

The transparency group XObject’s content stream defines the graphics objects belonging to the group. Invoking the **Do** operator on the XObject executes its content stream and composites the resulting group color, shape, and opacity into the group’s parent group or page as if they had come from an elementary graphics object. When applied to a transparency group XObject, **Do** performs the following actions in addition to the normal ones for a form XObject (as described in Section 4.9, “Form XObjects”):

- If the transparency group is non-isolated (the value of the **I** entry in its group attributes dictionary is **false**), its initial backdrop, within the bounding box specified by the XObject’s **BBox** entry, is defined to be the accumulated color and alpha of the parent group or page—that is, the result of everything that has been painted in the parent up to that point. (However, if the parent is a knockout group, the initial backdrop is the same as that of the parent.) If the group is isolated (**I** is **true**), its initial backdrop is defined to be transparent.

- Before execution of the transparency group XObject's content stream, the current blend mode in the graphics state is initialized to **Normal**, the current stroking and nonstroking alpha constants to 1.0, and the current soft mask to **None**.

***Note:** The purpose of initializing these graphics state parameters at the beginning of execution is to ensure that they are not applied twice: once when member objects are painted into the group and again when the group itself is painted into the parent group or page.*

- Objects painted by operators in the transparency group XObject's content stream are composited into the group according to the rules described in Section 7.2.2, "Basic Compositing Formula." The knockout flag (**K**) in the group attributes dictionary and the transparency-related parameters of the graphics state contribute to this computation.
- If a group color space (**CS**) is specified in the group attributes dictionary, all painting operators convert source colors to that color space before compositing objects into the group, and the resulting color at each point is interpreted in that color space. If no group color space is specified, the prevailing color space is dynamically inherited from the parent group or page. (If not otherwise specified, the page group's color space is inherited from the native color space of the output device.)
- After execution of the transparency group XObject's content stream, the graphics state reverts to its former state before the invocation of the **Do** operator (as it does for any form XObject). The group's shape—the union of all objects painted into the group, clipped by the group XObject's bounding box—is then painted into the parent group or page, using the group's accumulated color and opacity at each point.

***Note:** If the **Do** operator is invoked more than once for a given transparency group XObject, each invocation is treated as a separate transparency group. That is, the result is as if the group were independently composited with the backdrop on each invocation. Viewer applications that perform caching of rendered form XObjects must take this requirement into account.*

The actions described above occur only for a transparency group XObject—a form XObject having a **Group** entry designating a group attributes subdictionary whose group subtype (**S**) is **Transparency**. An ordinary form XObject—one having no **Group** entry—is not subject to any grouping behavior for transparency purposes. That is, the graphics objects it contains are composited individually, just as if they were painted directly into the parent group or page.

7.5.6 Patterns and Transparency

In the transparent imaging model, the graphics objects making up the pattern cell of a tiling pattern (see Section 4.6.2, “Tiling Patterns”) can include transparent objects and transparency groups. Transparent compositing can occur both within the pattern cell and between it and the backdrop wherever the pattern is painted. Similarly, a shading pattern (Section 4.6.3, “Shading Patterns”) composites with its backdrop as if the shading dictionary were applied with the **sh** operator.

In both cases, the pattern definition is treated as if it were implicitly enclosed in a non-isolated transparency group: a non-knockout group for tiling patterns, a knockout group for shading patterns. The definition does *not* inherit the current values of the graphics state parameters at the time it is evaluated; these take effect only when the resulting pattern is later used to paint an object. Instead, the graphics state parameters are initialized as follows:

- As always for transparency groups, those parameters related to transparency (blend mode, soft mask, and alpha constant) are initialized to their standard default values.
- All other parameters are initialized to their values at the beginning of the content stream (such as a page or a form XObject) in which the pattern is defined as a resource; this is simply the normal behavior for all patterns, in both the opaque and transparent imaging models.
- In the case of a shading pattern, the parameter values may be augmented by the contents of the **ExtGState** entry in the pattern dictionary (see Section 4.6.3, “Shading Patterns”). Only those parameters that affect the **sh** operator, such as the current transformation matrix and rendering intent, are used; parameters that affect path-painting operators are not, since the execution of **sh** does not entail painting a path.
- If the shading dictionary has a **Background** entry, the pattern’s implicit transparency group is filled with the specified background color before the **sh** operator is invoked.

When the pattern is later used to paint a graphics object, the color, shape, and opacity values resulting from the evaluation of the pattern definition are used as the object’s source color (C_s), object shape (f_j), and object opacity (q_j) in the

transparency compositing formulas. This painting operation is subject to the values of the graphics state parameters in effect at the time, just as in painting an object with a constant color.

Unlike the opaque imaging model, in which the pattern cell of a tiling pattern can be evaluated once and then replicated indefinitely to fill the painted area, the effect in the general transparent case is as if the pattern definition were reexecuted independently for each tile, taking into account the color of the backdrop at each point. However, in the common case in which the pattern consists entirely of objects painted with the **Normal** blend mode, this behavior can be optimized by treating the pattern cell as if it were an isolated group. Since in this case the results depend only on the color, shape, and opacity of the pattern cell itself and not on those of the backdrop, the pattern cell can be evaluated once and then replicated, just as in opaque painting.

Note: In a raster-based implementation of tiling, it is important that all tiles together be treated as a single transparency group. This avoids artifacts due to multiple marking of pixels along the boundaries between adjacent tiles.

The foregoing discussion applies to both colored (**PaintType 1**) and uncolored (**PaintType 2**) tiling patterns. In the latter case, the restriction that an uncolored pattern's definition may not specify colors extends as well to any transparency group that the definition may include. There are no corresponding restrictions, however, on specifying transparency-related parameters in the graphics state.

7.6 Color Space and Rendering Issues

This section describes the interactions between transparency and other aspects of color specification and rendering in the Adobe imaging model.

7.6.1 Color Spaces for Transparency Groups

As discussed in Section 7.5.5, “Transparency Group XObjects,” a transparency group can either have an explicitly declared color space of its own or inherit that of its parent group. In either case, the colors of source objects within the group are converted to the group's color space, if necessary, and all blending and compositing computations are done in that space (see Section 7.2.3, “Blending Color Space”). The resulting colors are then interpreted in that color space when the group is subsequently composited with its backdrop.

Under this arrangement, it is envisioned that all or most of a given piece of artwork will be created in a single color space—most likely, the working color space of the application generating it. The use of multiple color spaces typically will arise only when assembling independently produced artwork onto a page. After all the artwork has been placed on the page, the conversion from the group’s color space to the page’s device color space will be done as the last step, without any further transparency compositing. The transparent imaging model does not require that this convention be followed, however; the reason for adopting it is to avoid the loss of color information and the introduction of errors resulting from unnecessary color space conversions.

Only an isolated group may have an explicitly declared color space of its own; non-isolated groups must inherit their color space from the parent group (subject to special treatment for the page group, as described in Section 7.3.6, “Page Group”). This is because the use of an explicit color space in a non-isolated group would require converting colors from the backdrop’s color space to that of the group in order to perform the compositing computations. Such conversion may not be possible (since some color conversions can be performed only in one direction), and even if possible, it would entail an excessive number of color conversions.

The choice of a group color space will have significant effects on the results that are produced. In particular:

- As noted in Section 7.2.3, “Blending Color Space,” the results of compositing in a device color space will be device-dependent; in order for the compositing computations to work in a device-independent way, the group’s color space must be CIE-based.
- A consequence of choosing a CIE-based group color space is that only CIE-based spaces can be used to specify the colors of objects within the group. This is because conversion from device to CIE-based colors is not possible in general; the defined conversions work only in the opposite direction. See below for further discussion.
- The compositing computations and blend functions generally compute linear combinations of color component values, on the assumption that the component values themselves are linear. For this reason, it is usually best to choose a group color space that has a linear gamma function. If a nonlinear color space is chosen, the results will still be well-defined, but the appearance may not

match the user's expectations. Note, in particular, that the CIE-based *sRGB* color space (see page 192) is nonlinear, and hence may be unsuitable for use as a group color space.

Note: Implementations of the transparent imaging model are advised to use as much precision as possible in representing colors during compositing computations and in the accumulated group results. To minimize the accumulation of roundoff errors and avoid additional errors arising from the use of linear group color spaces, more precision is needed for intermediate results than is typically used to represent either the original source data or the final rasterized results.

If a group's color space—whether specified explicitly or inherited from the parent group—is CIE-based, any use of device color spaces for painting objects is subject to special treatment. Device colors cannot be painted directly into such a group, since there is no generally defined method for converting them to the CIE-based color space. This problem arises in the following cases:

- **DeviceGray**, **DeviceRGB**, and **DeviceCMYK** color spaces, unless remapped to default CIE-based color spaces (see “Default Color Spaces” on page 194)
- Operators (such as **rg**) that specify a device color space implicitly, unless that space is remapped
- Special color spaces whose base or underlying space is a device color space, unless that space is remapped

It is recommended that the default color space remapping mechanism always be employed when defining a transparency group whose color space is CIE-based. If a device color is specified and is not remapped, it will be converted to the CIE-based color space in an implementation-dependent fashion, producing unpredictable results.

7.6.2 Spot Colors and Transparency

The foregoing discussion of color spaces has been concerned with *process colors*—those produced by combinations of an output device's process colorants. Process colors may be specified directly in the device's native color space (such as **DeviceCMYK**), or they may be produced by conversion from some other color space, such as a CIE-based (**CalRGB** or **ICCBased**) space. Whatever means is used to specify them, process colors are subject to conversion to and from the group's color space.

A *spot color* is an additional color component, independent of those used to produce process colors. It may represent either an additional separation to be produced or an additional colorant to be applied to the composite page (see “Separation Color Spaces” on page 201 and “DeviceN Color Spaces” on page 205). The color component value, or *tint*, for a spot color specifies the concentration of the corresponding spot colorant. Tints are conventionally represented as subtractive, rather than additive, values.

Spot colors are inherently device-dependent and are not always available. In the opaque imaging model, each use of a spot color component in a **Separation** or **DeviceN** color space is accompanied by an *alternate color space* and a *tint transformation function* for mapping tint values into that space. This enables the color to be approximated with process colorants when the corresponding spot colorant is not available on the device.

Spot colors can be accommodated straightforwardly in the transparent imaging model (except for issues relating to overprinting, discussed in Section 7.6.3, “Overprinting and Transparency”). When an object is painted transparently with a spot color component that is available in the output device, that color is composited with the corresponding spot color component of the backdrop, independently of the compositing that is performed for process colors. A spot color retains its own identity; it is not subject to conversion to or from the color space of the enclosing transparency group or page. If the object is an element of a transparency group, one of two things can happen:

- The group maintains a separate color value for each spot color component, independently of the group’s color space. In effect, the spot color passes directly through the group hierarchy to the device, with no color conversions performed; however, it is still subject to blending and compositing with other objects that use the same spot color.
- The spot color is converted to its alternate color space. The resulting color is then subject to the usual compositing rules for process colors. In particular, spot colors are never available in a transparency group XObject that is used to define a soft mask; the alternate color space will always be substituted in that case.

Only a single shape value and opacity value are maintained at each point in the computed group results; they apply to both process and spot color components. In effect, every object is considered to paint every existing color component, both process and spot. Where no value has been explicitly specified for a given com-

ponent in a given object, an additive value of 1.0 (or a subtractive tint value of 0.0) is assumed. For instance, when painting an object with a color specified in a **DeviceCMYK** or **ICCBased** color space, the process color components are painted as specified and the spot color components are painted with an additive value of 1.0. Likewise, when painting an object with a color specified in a **Separation** color space, the named spot color is painted as specified and all other components (both process colors and other spot colors) are painted with an additive value of 1.0. The consequences of this are discussed in Section 7.6.3, “Overprinting and Transparency.”

The opaque imaging model also allows process color components to be addressed individually, as if they were spot colors. For instance, it is possible to specify a **Separation** color space named **Cyan**, which paints just the cyan component on a *CMYK* output device. However, this capability is very difficult to extend to transparency groups. In general, the color components in a group are not the process colorants themselves, but are converted to process colorants only after the completion of all color compositing computations for the group (and perhaps some of its parent groups as well). For instance, if the group’s color space is **ICCBased**, the group has no **Cyan** component to be painted. Consequently, treating a process color component as if it were a spot color is permitted only within a group that inherits the native color space of the output device. Attempting to do so in a group that specifies its own color space will result in conversion of the requested spot color to its alternate color space.

7.6.3 Overprinting and Transparency

In the opaque imaging model, overprinting is controlled by two parameters of the graphics state: the *overprint parameter* and the *overprint mode* (see Section 4.5.6, “Overprint Control”). Painting an object causes some specific set of device colorants to be marked, as determined by the current color space and current color in the graphics state. The remaining colorants are either erased or left unchanged, depending on whether the overprint parameter is **false** or **true**. When the current color space is **DeviceCMYK**, the overprint mode parameter additionally enables this selective marking of colorants to be applied to individual color components according to whether the component value is zero or nonzero.

Because this model of overprinting deals directly with the painting of device colorants, independently of the color space in which source colors have been specified, it is highly device-dependent and primarily addresses production needs rather than design intent. Overprinting is usually reserved for opaque colorants

or for very dark colors, such as black. It is also invoked during late-stage production operations such as trapping (see Section 9.10.5, “Trapping Support”), when the actual set of device colorants has already been determined.

Consequently, it is best to think of transparency as taking place in appearance space, but overprinting of device colorants in device space. This means that colorant overprint decisions should be made at output time, based on the actual resultant colorants of any transparency compositing operation. On the other hand, effects similar to overprinting can be achieved in a device-independent manner by taking advantage of blend modes, as described in the next section.

Blend Modes and Overprinting

As stated in Section 7.6.2, “Spot Colors and Transparency,” each graphics object painted affects all existing color components: all of the process colorants in the transparency group’s color space as well as any available spot colorants. For color components whose value has not been specified, a source color value of 1.0 is assumed; when objects are fully opaque and the **Normal** blend mode is used, this has the effect of erasing those components. This treatment is consistent with the behavior of the opaque imaging model with the overprint parameter set to **false**.

The transparent imaging model defines some blend modes, such as **Darken**, that can be used to achieve effects similar to overprinting. The blend function for **Darken** is

$$B(c_b, c_s) = \min(c_b, c_s)$$

In this blend mode, the result of compositing will always be the same as the backdrop color when the source color is 1.0, as it is for all unspecified color components. When the backdrop is fully opaque, this leaves the result color unchanged from that of the backdrop. This is consistent with the behavior of the opaque imaging model with the overprint parameter set to **true**.

If the object or backdrop is not fully opaque, the actions described above are altered accordingly. That is, the erasing effect is reduced, and overprinting an object with a color value of 1.0 may affect the result color. While these results may or may not be useful, they lie outside the realm of the overprinting and erasing behavior defined in the opaque imaging model.

When process colors are overprinted or erased (because a spot color is being painted), the blending computations described above are done independently for each component in the group's color space. If that space is different from the native color space of the output device, its components are not the device's actual process colorants; the blending computations affect the process colorants only after the group's results are converted to the device color space. Thus the effect is different from that of overprinting or erasing the device's process colorants directly. On the other hand, this is a fully general operation that works uniformly, regardless of the type of object or of the computations that produced the source color.

The discussion so far has focused on those color components whose values are *not* specified and that are to be either erased or left unchanged. However, the **Normal** or **Darken** blend modes used for these purposes may not be suitable for use on those components whose color values *are* specified. In particular, using the **Darken** blend mode for such components would preclude overprinting a dark color with a lighter one. Moreover, some other blend mode may be specifically desired for those components.

The PDF graphics state specifies only one current blend mode parameter, which always applies to process colorants and sometimes to spot colorants as well. Specifically, only separable, white-preserving blend modes can be used for spot colors. A blend mode is *white-preserving* if its blend function B has the property that $B(1.0, 1.0) = 1.0$. (Of the standard separable blend modes listed in Table 7.2 on page 417, all except **Difference** and **Exclusion** are white-preserving.) If the specified blend mode is not separable and white-preserving, it applies only to process color components; the **Normal** blend mode is substituted for spot colors. This ensures that when objects accumulate in an isolated transparency group, the accumulated values for unspecified components remain 1.0 so long as only white-preserving blend modes are used. The group's results can then be overprinted using **Darken** (or other useful modes) while avoiding unwanted interactions with components whose values were never specified within the group.

Compatibility with Opaque Overprinting

Because the use of blend modes to achieve effects similar to overprinting does not make direct use of the overprint control parameters in the graphics state, such methods are usable only by transparency-aware applications. For compatibility with the methods of overprint control used in the opaque imaging model, a special blend mode, `CompatibleOverprint`, is provided that consults the overprint-

related graphics state parameters to compute its result. This mode applies only when painting elementary graphics objects (fills, strokes, text, images, and shadings). It is never invoked explicitly and is not identified by any PDF name object; rather, it is implicitly invoked whenever an elementary graphics object is painted while overprinting is enabled (that is, when the overprint parameter in the graphics state is **true**).

***Note:** Earlier designs of the transparent imaging model included an additional blend mode named **Compatible**, which explicitly invoked the **CompatibleOverprint** blend mode described here. Because **CompatibleOverprint** is now invoked implicitly whenever appropriate, it is never necessary to specify the **Compatible** blend mode for use in compositing. It is still recognized as a valid blend mode for the sake of compatibility, but is simply treated as equivalent to **Normal**.*

The value of the blend function $B(c_b, c_s)$ in the **CompatibleOverprint** mode is either c_b or c_s , depending on the setting of the overprint mode parameter, the current and group color spaces, and the source color value c_s :

- If the overprint mode is 1 (nonzero overprint mode) and the current color space and group color space are both **DeviceCMYK**, then only process color components with nonzero values replace the corresponding component values of the backdrop; all other component values leave the existing backdrop value unchanged. That is, the value of the blend function $B(c_b, c_s)$ is the source component c_s for any process (**DeviceCMYK**) color component whose (subtractive) color value is nonzero; otherwise it is the backdrop component c_b . For spot color components, the value is always c_b .
- In all other cases, the value of $B(c_b, c_s)$ is c_s for all color components specified in the current color space, otherwise c_b . For instance, if the current color space is **DeviceCMYK** or **CalRGB**, the value of the blend function is c_s for process color components and c_b for spot components. On the other hand, if the current color space is a **Separation** space representing a spot color component, the value is c_s for that spot component and c_b for all process components and all other spot components.

***Note:** In the descriptions above, the term current color space refers to the color space used for a painting operation. This may be specified by the current color space parameter in the graphics state (see Section 4.5.1, “Color Values”), implicitly by color operators such as **rg** (Section 4.5.7, “Color Operators”), or by the **ColorSpace** entry of an image **XObject** (Section 4.8.4, “Image Dictionaries”). In the case of an **Indexed** space, it refers to the base color space (see “Indexed Color Spaces” on page 199); like-*

wise for **Separation** and **DeviceN** spaces that revert to their alternate color space, as described under “Separation Color Spaces” on page 201 and “DeviceN Color Spaces” on page 205.

If the current blend mode when `CompatibleOverprint` is invoked is any mode other than **Normal**, the object being painted is implicitly treated as if it were defined in a non-isolated, non-knockout transparency group and painted using the `CompatibleOverprint` blend mode; the group’s results are then painted using the current blend mode in the graphics state.

*Note: It is not necessary to create such an implicit transparency group if the current blend mode is **Normal**; simply substituting the `CompatibleOverprint` blend mode while painting the object produces equivalent results. There are some additional cases in which the implicit transparency group can be optimized out.*

Plate 20 shows the effects of all four possible combinations of blending and overprinting, using the **Screen** blend mode in the **DeviceCMYK** color space. The label “overprint enabled” means that the overprint parameter in the graphics state is **true** and the overprint mode is 1. In the upper half of the figure, a light green oval is painted opaquely (opacity = 1.0) over a backdrop shading from pure yellow to pure magenta. In the lower half, the same object is painted with transparency (opacity = 0.5).

Special Path-Painting Considerations

The overprinting considerations discussed above also affect those path-painting operations that combine filling and stroking a path in a single operation. These include the **B**, **B***, **b**, and **b*** operators (see Section 4.4.2, “Path-Painting Operators”) and the painting of glyphs with text rendering mode 2 or 6 (Section 5.2.5, “Text Rendering Mode”). For transparency compositing purposes, the combined fill and stroke are treated as a single graphics object, as if they were enclosed in a transparency group. This implicit group is established and used as follows:

- If overprinting is enabled (the overprint parameter in the graphics state is **true**) and the current stroking and nonstroking alpha constants are equal, a non-isolated, non-knockout transparency group is established. Within the group, the fill and stroke are performed with an alpha value of 1.0 but with the `CompatibleOverprint` blend mode. The group results are then composited with the backdrop using the originally specified alpha and blend mode.

- In all other cases, a non-isolated knockout group is established. Within the group, the fill and stroke are performed with their respective prevailing alpha constants and the prevailing blend mode. The group results are then composited with the backdrop using an alpha value of 1.0 and the **Normal** blend mode.

Note that in the case of showing text with the combined filling and stroking text rendering modes, this behavior is independent of the text knockout parameter in the graphics state (see Section 5.2.7, “Text Knockout”).

The purpose of these rules is to avoid having a non-opaque stroke composite with the result of the fill in the region of overlap, which would produce a “double border” effect that is usually undesirable. The special case that applies when the overprint parameter is **true** is for backward compatibility with the overprinting behavior of the opaque imaging model. If a desired effect cannot be achieved with a combined filling and stroking operator or text rendering mode, it can be achieved by specifying the fill and stroke with separate path objects and an explicit transparency group.

***Note:** Overprinting of the stroke over the fill does not work in the second case described above (although either the fill or the stroke can still overprint the backdrop). Furthermore, if the overprint graphics state parameter is **true**, the results are discontinuous at the transition between equal and unequal values of the stroking and non-stroking alpha constants. For this reason, it is best not to use overprinting for combined filling and stroking operations if the stroking and nonstroking alpha constants are being varied independently.*

Summary of Overprinting Behavior

Tables 7.14 and 7.15 summarize the overprinting and erasing behavior in the opaque and transparent imaging models, respectively. Table 7.14 shows the overprinting rules used in the opaque model, as described in Section 4.5.6, “Overprint Control”; Table 7.15 shows the equivalent rules as implemented by the CompatibleOverprint blend mode in the transparent model. The names **OP** and **OPM** in the tables refer to the overprint and overprint mode parameters of the graphics state.

TABLE 7.14 Overprinting behavior in the opaque imaging model

SOURCE COLOR SPACE	AFFECTED COLOR COMPONENT	EFFECT ON COLOR COMPONENT		
		OP FALSE	OP TRUE, OPM 0	OP TRUE, OPM 1
DeviceCMYK, specified directly, not in a sampled image	$C, M, Y, \text{ or } K$	Paint source	Paint source	Paint source if $\neq 0.0$ Do not paint if $= 0.0$
	Process colorant other than $CMYK$	Paint source	Paint source	Paint source
	Spot colorant	Paint 0.0	Do not paint	Do not paint
Any process color space (including other cases of DeviceCMYK)	Process colorant	Paint source	Paint source	Paint source
	Spot colorant	Paint 0.0	Do not paint	Do not paint
Separation or DeviceN	Process colorant	Paint 0.0	Do not paint	Do not paint
	Spot colorant named in source space	Paint source	Paint source	Paint source
	Spot colorant not named in source space	Paint 0.0	Do not paint	Do not paint

TABLE 7.15 Overprinting behavior in the transparent imaging model

SOURCE COLOR SPACE	AFFECTED COLOR COMPONENT OF GROUP COLOR SPACE	VALUE OF BLEND FUNCTION $B(c_b, c_s)$ EXPRESSED AS TINT		
		OP FALSE	OP TRUE, OPM 0	OP TRUE, OPM 1
DeviceCMYK, specified directly, not in a sampled image	$C, M, Y, \text{ or } K$	c_s	c_s	c_s if $c_s \neq 0.0$ c_b if $c_s = 0.0$
	Process color component other than $CMYK$	c_s	c_s	c_s
	Spot colorant	$c_s (= 0.0)$	c_b	c_b

TABLE 7.15 Overprinting behavior in the transparent imaging model

SOURCE COLOR SPACE	AFFECTED COLOR COMPONENT OF GROUP COLOR SPACE	VALUE OF BLEND FUNCTION $B(c_b, c_s)$ EXPRESSED AS TINT		
		OP FALSE	OP TRUE, OPM 0	OP TRUE, OPM 1
Any process color space (including other cases of DeviceCMYK)	Process color component	c_s	c_s	c_s
	Spot colorant	$c_s (= 0.0)$	c_b	c_b
Separation or DeviceN	Process color component	$c_s (= 0.0)$	c_b	c_b
	Spot colorant named in source space	c_s	c_s	c_s
	Spot colorant not named in source space	$c_s (= 0.0)$	c_b	c_b
A group (not an elementary object)	All color components	c_s	c_s	c_s

Color component values are represented in these tables as subtractive tint values, because overprinting is typically applied to subtractive colorants such as inks, rather than to additive ones such as phosphors on a display screen. The CompatibleOverprint blend mode is therefore described as if it took subtractive arguments and returned subtractive results. In reality, however, CompatibleOverprint (like all blend modes) treats color components as additive values; subtractive components must be complemented before and after application of the blend function.

Note an important difference between the two tables. In Table 7.14, the process color components being discussed are the actual device colorants—the color components of the output device’s native color space (**DeviceGray**, **DeviceRGB**, or **DeviceCMYK**). In Table 7.15, the process color components are those of the group’s color space, which is not necessarily the same as that of the output device (and can even be something like **CalRGB** or **ICCBased**). For this reason, the process color components of the group color space cannot be treated as if they were spot colors in a **Separation** or **DeviceN** color space (see Section 7.6.2, “Spot

Colors and Transparency”). This difference between opaque and transparent overprinting and erasing rules arises only within a transparency group (including the page group, if its color space is different from the native color space of the output device). There is no difference in the treatment of spot color components.

Table 7.15 has one additional row at the bottom. It applies when painting an object that is itself a transparency group rather than an elementary object (fill, stroke, text, image, or shading). As stated in Section 7.6.2, “Spot Colors and Transparency,” a group is considered to paint all color components, both process and spot. Color components that were not explicitly painted by any object in the group have an additive color value of 1.0 (subtractive tint 0.0). Since no information is retained about which components were actually painted within the group, compatible overprinting is not possible in this case; the `CompatibleOverprint` blend mode reverts to **Normal**, with no consideration of the overprint and overprint mode parameters. (Note that a transparency-aware application can choose a more suitable blend mode, such as **Darken**, if it desires to produce an effect similar to overprinting.)

7.6.4 Rendering Parameters and Transparency

The opaque imaging model has several graphics state parameters dealing with the rendering of color: the current halftone (see Section 6.4.4, “Halftone Dictionaries”), transfer functions (Section 6.3, “Transfer Functions”), rendering intent (“Rendering Intents” on page 197), and black-generation and undercolor-removal functions (Section 6.2.3, “Conversion from DeviceRGB to DeviceCMYK”). All of these rendering parameters can be specified on a per-object basis; they control how a particular object will be rendered. When all objects are opaque, it is easy to define what this means. But when they are transparent, more than one object can contribute to the color at a given point; it is unclear which rendering parameters to apply in an area where transparent objects overlap. At the same time, the transparent imaging model should be consistent with the opaque model when only opaque objects are painted.

Furthermore, some of the rendering parameters—the halftone and transfer functions, in particular—can be applied only when the final color at a given point is known. In the presence of transparency, these parameters must be treated somewhat differently from those (rendering intent, black generation, and undercolor removal) that apply whenever colors must be converted from one color space to another. When objects are transparent, the rendering of an object does not occur when the object is specified, but at some later time; hence for rendering param-

eters in the former category, the implementation must keep track of the rendering parameters at each point from the time they are specified until the time the rendering actually occurs. This means that these rendering parameters must be associated with regions of the page rather than with individual objects.

Halftone and Transfer Function

The halftone and transfer function to be used at any given point on the page are those in effect at the time of painting the last (topmost) elementary graphics object enclosing that point, but only if the object is fully opaque. (Only elementary objects are relevant; the rendering parameters associated with a group object are ignored.) The *topmost object* at any point is defined to be the topmost elementary object in the entire page stack that has a nonzero object shape value (f_j) at that point (that is, for which the point is inside the object). An object is considered to be *fully opaque* if all of the following conditions hold at the time the object is painted:

- The current alpha constant in the graphics state (stroking or nonstroking, depending on the painting operation) is 1.0.
- The current blend mode in the graphics state is **Normal** (or **Compatible**, which is treated as equivalent to **Normal**).
- The current soft mask in the graphics state is **None**. If the object is an image XObject, there is no **SMask** entry in its image dictionary.
- The foregoing three conditions were also true at the time the **Do** operator was invoked for the group containing the object, as well as for any direct ancestor groups.
- If the current color is a tiling pattern, all objects in the definition of its pattern cell also satisfy the foregoing conditions.

Taken together, these conditions ensure that only the object itself contributes to the color at the given point, completely obscuring the backdrop. For portions of the page whose topmost object is not fully opaque or that are never painted at all, the default halftone and transfer function for the page are used.

Note: *If a graphics object is painted with overprinting enabled—that is, if the applicable (stroking or nonstroking) overprint parameter in the graphics state is **true**—the halftone and transfer function to use at a given point must be determined independently for each color component. Overprinting implicitly invokes the Compatible-*

Overprint blend mode (see “Compatibility with Opaque Overprinting” on page 460). An object is considered opaque for a given component only if Compatible-Overprint yields the source color (not the backdrop color) for that component.

Rendering Intent and Color Conversions

The rendering intent, black-generation, and undercolor-removal parameters need to be handled somewhat differently. The rendering intent influences the conversion from a CIE-based color space to a target color space, taking into account the target space’s color gamut (the range of colors it can reproduce). Whereas in the opaque imaging model the target space is always the native color space of the output device, in the transparent model it may instead be the group color space of a transparency group into which an object is being painted.

The rendering intent is needed at the moment such a conversion must be performed—that is, when painting an elementary or group object specified in a CIE-based color space into a parent group having a different color space. This differs from the current halftone and transfer function, whose values are used only when all color compositing has been completed and rasterization is being performed.

In all cases, the rendering intent to use for converting an object’s color (whether that of an elementary object or of a transparency group) is determined by the rendering intent parameter associated with the object. In particular:

- When painting an elementary object with a CIE-based color into a transparency group having a different color space, the rendering intent used is the current rendering intent in effect in the graphics state at the time of the painting operation.
- When painting a transparency group whose color space is CIE-based into a parent group having a different color space, the rendering intent used is the current rendering intent in effect at the time the **Do** operator is applied to the group.
- When the color space of the page group is CIE-based, the rendering intent used to convert colors to the native color space of the output device is the default rendering intent for the page.

***Note:** Since there may be one or more nested transparency groups having different CIE-based color spaces, the color of an elementary source object may be converted to the device color space in multiple stages, controlled by the rendering intent in effect*

at each stage. The proper choice of rendering intent at each stage depends on the relative gamuts of the source and target color spaces. It is specified explicitly by the document producer, not prescribed by the PDF specification, since no single policy for managing rendering intents is appropriate for all situations.

A similar approach works for the black-generation and undercolor-removal functions, which are applied only during conversion from **DeviceRGB** to **DeviceCMYK** color spaces:

- When painting an elementary object with a **DeviceRGB** color directly into a transparency group whose color space is **DeviceCMYK**, the functions used are the current black-generation and undercolor-removal functions in effect in the graphics state at the time of the painting operation.
- When painting a transparency group whose color space is **DeviceRGB** into a parent group whose color space is **DeviceCMYK**, the functions used are the ones in effect at the time the **Do** operator is applied to the group.
- When the color space of the page group is **DeviceRGB** and the native color space of the output device is **DeviceCMYK**, the functions used to convert colors to the device's color space are the default functions for the page.

7.6.5 PostScript Compatibility

Because the PostScript language does not support the transparent imaging model, PDF 1.4 viewer applications must have some means for converting the appearance of a document that uses transparency into a purely opaque description for printing on PostScript output devices. Similar techniques can also be used to convert such documents into a form that can be correctly viewed by PDF 1.3 and earlier viewers.

Converting the contents of a page from transparent to opaque form entails some combination of shape decomposition and prerendering to “flatten” the stack of transparent objects on the page, perform all the needed transparency computations, and describe the final appearance using opaque objects only. Whether the page contains transparent content needing to be flattened can be determined by straightforward analysis of the page's resources; it is not necessary to analyze the content stream itself. The conversion to opaque form is irreversible, since all information about how the transparency effects were produced is lost.

In order to perform the transparency computations properly, the viewer application needs to know the native color space of the output device. This is no problem when the viewer controls the output device directly. However, when generating PostScript output, the viewer has no way of knowing the native color space of the PostScript output device. An incorrect assumption will ruin the calibration of any CIE-based colors appearing on the page. This problem can be addressed in either of two ways:

- If the entire page consists of CIE-based colors, flatten the colors to a single CIE-based color space rather than to a device color space. The preferred color space for this purpose can easily be determined if the page has a group attributes dictionary (**Group** entry in the page object) specifying a CIE-based color space (see Section 7.5.5, “Transparency Group XObjects”).
- Otherwise, flatten the colors to some assumed device color space with pre-determined calibration. In the generated PostScript output, paint the flattened colors in a CIE-based color space having that calibration.

Because the choice between using spot colorants and converting them to an alternate color space affects the flattened results of process colors, a decision must also be made during PostScript conversion about the set of available spot colorants to assume. (This differs from strictly opaque painting, where the decision can be deferred until the generated PostScript code is executed.)

CHAPTER 8

Interactive Features

THIS CHAPTER DESCRIBES those features of PDF that allow a user to interact with a document on the screen, using the mouse and keyboard. These include:

- *Preference settings* to control the way the document is presented on the screen (Section 8.1, “Viewer Preferences”)
- *Navigation* facilities for moving through the document in a variety of ways (Sections 8.2, “Document-Level Navigation,” and 8.3, “Page-Level Navigation”)
- *Annotations* for adding text notes, sounds, movies, and other ancillary information to the document (Section 8.4, “Annotations”)
- *Actions* that can be triggered by specified events (Section 8.5, “Actions”)
- *Interactive forms* for gathering information from the user (Section 8.6, “Interactive Forms”)
- *Sounds* to be played through the computer’s speakers (Section 8.7, “Sounds”)
- *Movies* to be displayed on the screen (Section 8.8, “Movies”)

8.1 Viewer Preferences

The **ViewerPreferences** entry in a document’s catalog (see Section 3.6.1, “Document Catalog”) designates a *viewer preferences dictionary* (PDF 1.2) controlling the way the document is to be presented on the screen or in print. If no such dictionary is specified, viewer applications should behave in accordance with their own current user preference settings. Table 8.1 shows the contents of the viewer preferences dictionary. (See implementation note 54 in Appendix H.)

TABLE 8.1 Entries in a viewer preferences dictionary

KEY	TYPE	VALUE						
HideToolbar	boolean	<i>(Optional)</i> A flag specifying whether to hide the viewer application's tool bars when the document is active. Default value: false .						
HideMenubar	boolean	<i>(Optional)</i> A flag specifying whether to hide the viewer application's menu bar when the document is active. Default value: false .						
HideWindowUI	boolean	<i>(Optional)</i> A flag specifying whether to hide user interface elements in the document's window (such as scroll bars and navigation controls), leaving only the document's contents displayed. Default value: false .						
FitWindow	boolean	<i>(Optional)</i> A flag specifying whether to resize the document's window to fit the size of the first displayed page. Default value: false .						
CenterWindow	boolean	<i>(Optional)</i> A flag specifying whether to position the document's window in the center of the screen. Default value: false .						
DisplayDocTitle	boolean	<i>(Optional; PDF 1.4)</i> A flag specifying whether the window's title bar should display the document title taken from the Title entry of the document information dictionary (see Section 9.2.1, "Document Information Dictionary"). If false , the title bar should instead display the name of the PDF file containing the document. Default value: false .						
NonFullScreenPageMode	name	<p><i>(Optional)</i> The document's <i>page mode</i>, specifying how to display the document on exiting full-screen mode:</p> <table> <tr> <td>UseNone</td> <td>Neither document outline nor thumbnail images visible</td> </tr> <tr> <td>UseOutlines</td> <td>Document outline visible</td> </tr> <tr> <td>UseThumbs</td> <td>Thumbnail images visible</td> </tr> </table> <p>This entry is meaningful only if the value of the PageMode entry in the catalog dictionary (see Section 3.6.1, "Document Catalog") is FullScreen; it is ignored otherwise. Default value: UseNone.</p>	UseNone	Neither document outline nor thumbnail images visible	UseOutlines	Document outline visible	UseThumbs	Thumbnail images visible
UseNone	Neither document outline nor thumbnail images visible							
UseOutlines	Document outline visible							
UseThumbs	Thumbnail images visible							
Direction	name	<p><i>(Optional; PDF 1.3)</i> The predominant reading order for text:</p> <table> <tr> <td>L2R</td> <td>Left to right</td> </tr> <tr> <td>R2L</td> <td>Right to left (including vertical writing systems such as Chinese, Japanese, and Korean)</td> </tr> </table> <p>This entry has no direct effect on the document's contents or page numbering, but can be used to determine the relative positioning of pages when displayed side by side or printed <i>n</i>-up. Default value: L2R.</p>	L2R	Left to right	R2L	Right to left (including vertical writing systems such as Chinese, Japanese, and Korean)		
L2R	Left to right							
R2L	Right to left (including vertical writing systems such as Chinese, Japanese, and Korean)							

ViewArea	name	<p>(Optional; PDF 1.4) The name of the page boundary representing the area of a page to be displayed when viewing the document on the screen. The value is the key designating the relevant page boundary in the page object (see “Page Objects” on page 87 and Section 9.10.1, “Page Boundaries”). If the specified page boundary is not defined in the page object, its default value will be used, as specified in Table 3.18 on page 88. Default value: CropBox.</p> <p><i>Note:</i> This entry is intended primarily for use by prepress applications that interpret or manipulate the page boundaries as described in Section 9.10.1, “Page Boundaries.” Most PDF consumer applications will disregard it.</p>
ViewClip	name	<p>(Optional; PDF 1.4) The name of the page boundary to which the contents of a page are to be clipped when viewing the document on the screen. The value is the key designating the relevant page boundary in the page object (see “Page Objects” on page 87 and Section 9.10.1, “Page Boundaries”). If the specified page boundary is not defined in the page object, its default value will be used, as specified in Table 3.18 on page 88. Default value: CropBox.</p> <p><i>Note:</i> This entry is intended primarily for use by prepress applications that interpret or manipulate the page boundaries as described in Section 9.10.1, “Page Boundaries.” Most PDF consumer applications will disregard it.</p>
PrintArea	name	<p>(Optional; PDF 1.4) The name of the page boundary representing the area of a page to be rendered when printing the document. The value is the key designating the relevant page boundary in the page object (see “Page Objects” on page 87 and Section 9.10.1, “Page Boundaries”). If the specified page boundary is not defined in the page object, its default value will be used, as specified in Table 3.18 on page 88. Default value: CropBox.</p> <p><i>Note:</i> This entry is intended primarily for use by prepress applications that interpret or manipulate the page boundaries as described in Section 9.10.1, “Page Boundaries.” Most PDF consumer applications will disregard it.</p>
PrintClip	name	<p>(Optional; PDF 1.4) The name of the page boundary to which the contents of a page are to be clipped when printing the document. The value is the key designating the relevant page boundary in the page object (see “Page Objects” on page 87 and Section 9.10.1, “Page Boundaries”). If the specified page boundary is not defined in the page object, its default value will be used, as specified in Table 3.18 on page 88. Default value: CropBox.</p> <p><i>Note:</i> This entry is intended primarily for use by prepress applications that interpret or manipulate the page boundaries as described in Section 9.10.1, “Page Boundaries.” Most PDF consumer applications will disregard it.</p>

8.2 Document-Level Navigation

The features described in this section allow a PDF viewer application to present the user with an interactive, global overview of a document in either of two forms:

- As a hierarchical *outline* showing the document’s internal structure
- As a collection of *thumbnail images* representing the pages of the document in miniature form

Each item in the outline or each thumbnail image can then be associated with a corresponding *destination* in the document, allowing the user to jump directly to that destination by clicking with the mouse.

8.2.1 Destinations

A *destination* defines a particular view of a document, consisting of the following:

- The page of the document to be displayed
- The location of the document window on that page
- The magnification (zoom) factor to use when displaying the page

Destinations may be associated with outline items (see Section 8.2.2, “Document Outline”), annotations (“Link Annotations” on page 501), or actions (“Go-To Actions” on page 519 and “Remote Go-To Actions” on page 520). In each case, the destination specifies the view of the document to be presented when the outline item or annotation is opened or the action is performed. In addition, the optional **OpenAction** entry in a document’s catalog (Section 3.6.1, “Document Catalog”) may specify a destination to be displayed when the document is opened. A destination may be specified either explicitly, by an array of parameters defining its properties, or indirectly by name.

Explicit Destinations

Table 8.2 shows the allowed syntactic forms for specifying a destination explicitly in a PDF file. In each case, *page* is an indirect reference to a page object. All coordinate values (*left*, *right*, *top*, and *bottom*) are expressed in the default user space coordinate system. The page’s *bounding box* is the smallest rectangle enclosing all of its contents. (If any side of the bounding box lies outside the page’s crop box,

the corresponding side of the crop box is used instead; see Section 9.10.1, “Page Boundaries,” for further discussion of the crop box.)

TABLE 8.2 Destination syntax

SYNTAX	MEANING
[<i>page</i> /XYZ <i>left top zoom</i>]	Display the page designated by <i>page</i> , with the coordinates (<i>left</i> , <i>top</i>) positioned at the top-left corner of the window and the contents of the page magnified by the factor <i>zoom</i> . A null value for any of the parameters <i>left</i> , <i>top</i> , or <i>zoom</i> specifies that the current value of that parameter is to be retained unchanged. A <i>zoom</i> value of 0 has the same meaning as a null value.
[<i>page</i> /Fit]	Display the page designated by <i>page</i> , with its contents magnified just enough to fit the entire page within the window both horizontally and vertically. If the required horizontal and vertical magnification factors are different, use the smaller of the two, centering the page within the window in the other dimension.
[<i>page</i> /FitH <i>top</i>]	Display the page designated by <i>page</i> , with the vertical coordinate <i>top</i> positioned at the top edge of the window and the contents of the page magnified just enough to fit the entire width of the page within the window.
[<i>page</i> /FitV <i>left</i>]	Display the page designated by <i>page</i> , with the horizontal coordinate <i>left</i> positioned at the left edge of the window and the contents of the page magnified just enough to fit the entire height of the page within the window.
[<i>page</i> /FitR <i>left bottom right top</i>]	Display the page designated by <i>page</i> , with its contents magnified just enough to fit the rectangle specified by the coordinates <i>left</i> , <i>bottom</i> , <i>right</i> , and <i>top</i> entirely within the window both horizontally and vertically. If the required horizontal and vertical magnification factors are different, use the smaller of the two, centering the rectangle within the window in the other dimension.
[<i>page</i> /FitB]	(PDF 1.1) Display the page designated by <i>page</i> , with its contents magnified just enough to fit its bounding box entirely within the window both horizontally and vertically. If the required horizontal and vertical magnification factors are different, use the smaller of the two, centering the bounding box within the window in the other dimension.
[<i>page</i> /FitBH <i>top</i>]	(PDF 1.1) Display the page designated by <i>page</i> , with the vertical coordinate <i>top</i> positioned at the top edge of the window and the contents of the page magnified just enough to fit the entire width of its bounding box within the window.
[<i>page</i> /FitBV <i>left</i>]	(PDF 1.1) Display the page designated by <i>page</i> , with the horizontal coordinate <i>left</i> positioned at the left edge of the window and the contents of the page magnified just enough to fit the entire height of its bounding box within the window.

***Note:** No page object can be specified for a destination associated with a remote go-to action (see “Remote Go-To Actions” on page 520), because the destination page is in a different PDF document. In this case, the page parameter specifies a page number within the remote document instead of a page object in the current document.*

Named Destinations

Instead of being defined directly, using the explicit syntax shown in Table 8.2, a destination may be referred to indirectly, using a name object (*PDF 1.1*) or a string (*PDF 1.2*). This capability is especially useful when the destination is located in another PDF document. For example, a link to the beginning of Chapter 6 in another document might refer to the destination by a name such as `Chap6.begin` instead of giving an explicit page number in the other document. This would allow the location of the chapter opening to change within the other document without invalidating the link. If an annotation or outline item that refers to a named destination has an associated action, such as a remote go-to action (see “Remote Go-To Actions” on page 520) or a thread action (“Thread Actions” on page 522), the destination is in the file specified by the action’s **F** entry, if any; if there is no **F** entry, the destination is in the current file.

In PDF 1.1, the correspondence between name objects and destinations is defined by the **Dests** entry in the document catalog (see Section 3.6.1, “Document Catalog”). The value of this entry is a dictionary in which each key is a destination name and the corresponding value is either an array defining the destination, using the syntax shown in Table 8.2, or a dictionary with a **D** entry whose value is such an array. The latter form allows additional attributes to be associated with the destination, as well as enabling a go-to action (see “Go-To Actions” on page 519) to be used as the target of a named destination.

In PDF 1.2, the correspondence between strings and destinations is defined by the **Dests** entry in the document’s name dictionary (see Section 3.6.3, “Name Dictionary”). The value of this entry is a name tree (Section 3.8.4, “Name Trees”) mapping name strings to destinations. (The keys in the name tree may be treated as text strings for display purposes.) The destination value associated with a key in the name tree may be either an array or a dictionary, as described in the preceding paragraph.

Note: *The use of strings as destination names is a PDF 1.2 feature. If compatibility with earlier versions of PDF is required, only name objects may be used to refer to named destinations. Where such compatibility is not a consideration, however, applications that generate large numbers of named destinations should use the string form of representation instead, since there are essentially no implementation limits.*

8.2.2 Document Outline

A PDF document may optionally display a *document outline* on the screen, allowing the user to navigate interactively from one part of the document to another. The outline consists of a tree-structured hierarchy of *outline items* (sometimes called *bookmarks*), which serve as a “visual table of contents” to display the document’s structure to the user. The user can interactively open and close individual items by clicking them with the mouse. When an item is open, its immediate children in the hierarchy become visible on the screen; each child may in turn be open or closed, selectively revealing or hiding further parts of the hierarchy. When an item is closed, all of its descendants in the hierarchy are hidden. Clicking the text of any visible item with the mouse *activates* the item, causing the viewer application to jump to a destination or trigger an action associated with the item.

The root of a document’s outline hierarchy is an *outline dictionary* specified by the **Outlines** entry in the document catalog (see Section 3.6.1, “Document Catalog”). Table 8.3 shows the contents of this dictionary. Each individual outline item within the hierarchy is defined by an *outline item dictionary* (Table 8.4). The items at each level of the hierarchy form a linked list, chained together through their **Prev** and **Next** entries and accessed through the **First** and **Last** entries in the parent item (or in the outline dictionary in the case of top-level items). When displayed on the screen, the items at a given level appear in the order in which they occur in the linked list. (See also implementation note 55 in Appendix H.)

TABLE 8.3 Entries in the outline dictionary

KEY	TYPE	VALUE
Type	name	<i>(Optional)</i> The type of PDF object that this dictionary describes; if present, must be Outlines for an outline dictionary.
First	dictionary	<i>(Required; must be an indirect reference)</i> An outline item dictionary representing the first top-level item in the outline.
Last	dictionary	<i>(Required; must be an indirect reference)</i> An outline item dictionary representing the last top-level item in the outline.
Count	integer	<i>(Required if the document has any open outline entries)</i> The total number of open items at all levels of the outline. This entry should be omitted if there are no open outline items.

TABLE 8.4 Entries in an outline item dictionary

KEY	TYPE	VALUE
Title	text string	<i>(Required)</i> The text to be displayed on the screen for this item.
Parent	dictionary	<i>(Required; must be an indirect reference)</i> The parent of this item in the outline hierarchy. The parent of a top-level item is the outline dictionary itself.
Prev	dictionary	<i>(Required for all but the first item at each level; must be an indirect reference)</i> The previous item at this outline level.
Next	dictionary	<i>(Required for all but the last item at each level; must be an indirect reference)</i> The next item at this outline level.
First	dictionary	<i>(Required if the item has any descendants; must be an indirect reference)</i> The first of this item's immediate children in the outline hierarchy.
Last	dictionary	<i>(Required if the item has any descendants; must be an indirect reference)</i> The last of this item's immediate children in the outline hierarchy.
Count	integer	<i>(Required if the item has any descendants)</i> If the item is open, the total number of its open descendants at all lower levels of the outline hierarchy. If the item is closed, a negative integer whose absolute value specifies how many descendants would appear if the item were reopened.
Dest	name, string, or array	<i>(Optional; not permitted if an A entry is present)</i> The destination to be displayed when this item is activated (see Section 8.2.1, "Destinations"; see also implementation note 56 in Appendix H).
A	dictionary	<i>(Optional; PDF 1.1; not permitted if a Dest entry is present)</i> The action to be performed when this item is activated (see Section 8.5, "Actions").

SE	dictionary	<p>(Optional; PDF 1.3; must be an indirect reference) The structure element to which the item refers (see Section 9.6.1, “Structure Hierarchy”).</p> <p>Note: The ability to associate an outline item with a structure element (such as the beginning of a chapter) is a PDF 1.3 feature. For backward compatibility with earlier PDF versions, such an item should also specify a destination (Dest) corresponding to an area of a page where the contents of the designated structure element are displayed.</p>
C	array	<p>(Optional; PDF 1.4) An array of three numbers in the range 0.0 to 1.0, representing the components in the DeviceRGB color space of the color to be used for the outline entry’s text. Default value: [0.0 0.0 0.0].</p>
F	integer	<p>(Optional; PDF 1.4) A set of flags specifying style characteristics for displaying the outline item’s text (see Table 8.5). Default value: 0.</p>

The value of the outline item dictionary’s **F** entry (PDF 1.4) is an unsigned 32-bit integer containing flags specifying style characteristics for displaying the item. Bit positions within the flag word are numbered from 1 (low-order) to 32 (high-order). Table 8.5 shows the meanings of the flags; all undefined flag bits are reserved and must be set to 0.

TABLE 8.5 Outline item flags

BIT POSITION	NAME	MEANING
1	Italic	If set, display the item in italic.
2	Bold	If set, display the item in bold.

Example 8.1 shows a typical outline dictionary and outline item dictionary. See Appendix G for an example of a complete outline hierarchy.

Example 8.1

```

21 0 obj
  << /Count 6
    /First 22 0 R
    /Last 29 0 R
  >>
endobj

```

```
22 0 obj
  << /Title (Chapter 1)
    /Parent 21 0 R
    /Next 26 0 R
    /First 23 0 R
    /Last 25 0 R
    /Count 3
    /Dest [3 0 R /XYZ 0 792 0]
  >>
endobj
```

8.2.3 Thumbnail Images

A PDF document may define *thumbnail images* representing the contents of its pages in miniature form. A viewer application can then display these images on the screen, allowing the user to navigate to a page by clicking its thumbnail image with the mouse.

Note: Thumbnail images are not required, and may be included for some pages and not for others.

The thumbnail image for a page is an image XObject specified by the **Thumb** entry in the page object (see “Page Objects” on page 87). It has the usual structure for an image dictionary (Section 4.8.4, “Image Dictionaries”), but only the **Width**, **Height**, **ColorSpace**, **BitsPerComponent**, and **Decode** entries are significant; all of the other entries listed in Table 4.35 on page 267 are ignored if present. (If a **Subtype** entry is specified, its value must be **Image**.) The image’s color space must be either **DeviceGray** or **DeviceRGB**, or an **Indexed** space based on one of these. Example 8.2 shows a typical thumbnail image definition.

Example 8.2

```
12 0 obj
  << /Width 76
    /Height 99
    /ColorSpace /DeviceRGB
    /BitsPerComponent 8
    /Length 13 0 R
    /Filter [/ASCII85Decode /DCTDecode]
  >>
```



```

stream
s4IA>!"M;*Ddm8XA,IT0!!3,S!/(=R!<E3%!<N<(!WrK*!WrN,
...Omitted data...
endstream
endobj

13 0 obj                                % Length of stream
...
endobj

```

8.3 Page-Level Navigation

This section describes PDF facilities that allow the user to navigate from page to page within a document. These include:

- *Page labels* for numbering or otherwise identifying individual pages
- *Article threads*, which chain together items of content within the document that are logically connected but not physically sequential
- *Presentations* that display the document in the form of a “slide show,” advancing from one page to the next either automatically or under user control

For another important form of page-level navigation, see “Link Annotations” on page 501.

8.3.1 Page Labels

Each page in a PDF document is identified by an integer *page index* that expresses the page’s relative position within the document. In addition, a document may optionally define *page labels* (*PDF 1.3*) to identify each page visually on the screen or in print. Page labels and page indices need not coincide: the indices are fixed, running consecutively through the document starting from 0 for the first page, but the labels can be specified in any way that is appropriate for the particular document. For example, if the document begins with 12 pages of front matter numbered in roman numerals and the remainder of the document is numbered in arabic, then the first page would have a page index of 0 and a page label of i, the twelfth page would have index 11 and label xii, and the thirteenth page would have index 12 and label 1.

For purposes of page labeling, a document can be divided into *labeling ranges*, each of which is a series of consecutive pages using the same numbering system. Pages within a range are numbered sequentially in ascending order. A page's label consists of a numeric portion based on its position within its labeling range, optionally preceded by a *label prefix* denoting the range itself. For example, the pages in an appendix might be labeled with decimal numeric portions prefixed with the string A–; the resulting page labels would be A–1, A–2, and so on.

A document's labeling ranges are defined by the **PageLabels** entry in the document catalog (see Section 3.6.1, “Document Catalog”). The value of this entry is a number tree (Section 3.8.5, “Number Trees”), each of whose keys is the page index of the first page in a labeling range; the corresponding value is a *page label dictionary* defining the labeling characteristics for the pages in that range. The tree must include a value for page index 0. Table 8.6 shows the contents of a page label dictionary. (See implementation note 57 in Appendix H.)

Example 8.3 shows a document with pages labeled

i, ii, iii, iv, 1, 2, 3, A–8, A–9, ...

Example 8.3

```

1 0 obj
  << /Type /Catalog
    /PageLabels << /Nums [ 0 << /S /r >>      % A number tree containing
                        4 << /S /D >>        %   three page label dictionaries
                        7 << /S /D
                          /P (A–)
                          /St 8
                        >>
                      ]
                    >>
  ...
  >>
endobj

```

TABLE 8.6 Entries in a page label dictionary

KEY	TYPE	VALUE
Type	name	(<i>Optional</i>) The type of PDF object that this dictionary describes; if present, must be PageLabel for a page label dictionary.
S	name	(<i>Optional</i>) The numbering style to be used for the numeric portion of each page label: <ul style="list-style-type: none"> D Decimal arabic numerals R Uppercase roman numerals r Lowercase roman numerals A Uppercase letters (A to Z for the first 26 pages, AA to ZZ for the next 26, and so on) a Lowercase letters (a to z for the first 26 pages, aa to zz for the next 26, and so on) <p>There is no default numbering style; if no S entry is present, page labels will consist solely of a label prefix with no numeric portion. For example, if the P entry (below) specifies the label prefix Contents, each page will simply be labeled Contents with no page number. (If the P entry is also missing or empty, the page label will be an empty string.)</p>
P	text string	(<i>Optional</i>) The label prefix for page labels in this range.
St	integer	(<i>Optional</i>) The value of the numeric portion for the first page label in the range. Subsequent pages will be numbered sequentially from this value, which must be greater than or equal to 1. Default value: 1.

8.3.2 Articles

Some types of document may contain sequences of content items that are logically connected but not physically sequential. For example, a news story may begin on the first page of a newsletter and run over onto one or more nonconsecutive interior pages. To represent such sequences of physically discontinuous but logically related items, a PDF document may define one or more *articles* (PDF 1.1). The sequential flow of an article is defined by an *article thread*; the individual content items that make up the article are called *beads* on the thread. PDF viewer applications can provide navigation facilities to allow the user to follow a thread from one bead to the next.

The optional **Threads** entry in the document catalog (see Section 3.6.1, “Document Catalog”) holds an array of *thread dictionaries* (Table 8.7) defining the document’s articles. Each individual bead within a thread is represented by a *bead dictionary* (Table 8.8). The thread dictionary’s **F** entry points to the first bead in the thread; the beads are then chained together sequentially in a doubly linked list

through their **N** (next) and **V** (previous) entries. In addition, for each page on which article beads appear, the page object (see “Page Objects” on page 87) should contain a **B** entry whose value is an array of indirect references to the beads on the page, in drawing order.

TABLE 8.7 Entries in a thread dictionary

KEY	TYPE	VALUE
Type	name	<i>(Optional)</i> The type of PDF object that this dictionary describes; if present, must be Thread for a thread dictionary.
F	dictionary	<i>(Required; must be an indirect reference)</i> The first bead in the thread.
I	dictionary	<i>(Optional)</i> A thread information dictionary containing information about the thread, such as its title, author, and creation date. The contents of this dictionary are similar to those of the document information dictionary (see Section 9.2.1, “Document Information Dictionary”).

TABLE 8.8 Entries in a bead dictionary

KEY	TYPE	VALUE
Type	name	<i>(Optional)</i> The type of PDF object that this dictionary describes; if present, must be Bead for a bead dictionary.
T	dictionary	<i>(Required for the first bead of a thread; optional for all others; must be an indirect reference)</i> The thread to which this bead belongs. <i>Note: In PDF 1.1, this entry is permitted only for the first bead of a thread. In PDF 1.2 and higher, it is permitted for any bead but required only for the first.</i>
N	dictionary	<i>(Required; must be an indirect reference)</i> The next bead in the thread. In the last bead, this entry points to the first.
V	dictionary	<i>(Required; must be an indirect reference)</i> The previous bead in the thread. In the first bead, this entry points to the last.
P	dictionary	<i>(Required; must be an indirect reference)</i> The page object representing the page on which this bead appears.
R	rectangle	<i>(Required)</i> A rectangle specifying the location of this bead on the page.

Example 8.4 shows a thread with three beads.

Example 8.4

```
22 0 obj
  << /F 23 0 R
    /I << /Title (Man Bites Dog) >>
  >>
endobj

23 0 obj
  << /T 22 0 R
    /N 24 0 R
    /V 25 0 R
    /P 8 0 R
    /R [158 247 318 905]
  >>
endobj

24 0 obj
  << /T 22 0 R
    /N 25 0 R
    /V 23 0 R
    /P 8 0 R
    /R [322 246 486 904]
  >>
endobj

25 0 obj
  << /T 22 0 R
    /N 23 0 R
    /V 24 0 R
    /P 10 0 R
    /R [157 254 319 903]
  >>
endobj
```

8.3.3 Presentations

Some PDF viewer applications may allow a document to be displayed in the form of a *presentation* or “slide show,” advancing from one page to the next either automatically or under user control. A page object (see “Page Objects” on page 87) may contain two optional entries, **Dur** and **Trans** (*PDF 1.1*), to specify how to display that page in presentation mode.

The **Dur** entry in the page object specifies the page’s *display duration* (also called its *advance timing*): the maximum length of time, in seconds, that the page will be

displayed before the presentation automatically advances to the next page. (The user can advance the page manually before the specified time has expired.) If no **Dur** entry is specified in the page object, the page does not advance automatically.

The **Trans** entry in the page object contains a *transition dictionary* describing the style and duration of the visual transition to use when moving from another page to the given page during a presentation. Table 8.9 shows the contents of the transition dictionary. (Some of the entries shown are needed only for certain transition styles, as indicated in the table.)

TABLE 8.9 Entries in a transition dictionary

KEY	TYPE	VALUE
Type	name	(Optional) The type of PDF object that this dictionary describes; if present, must be Trans for a transition dictionary.
D	number	(Optional) The duration of the transition effect, in seconds. Default value: 1.
S	name	(Optional) The <i>transition style</i> to use when moving to this page from another during a presentation: <ul style="list-style-type: none"> Split Two lines sweep across the screen, revealing the new page. The lines may be either horizontal or vertical and may move inward from the edges of the page or outward from the center, as specified by the Dm and M entries, respectively. Blinds Multiple lines, evenly spaced across the screen, synchronously sweep in the same direction to reveal the new page. The lines may be either horizontal or vertical, as specified by the Dm entry. Horizontal lines move downward, vertical lines to the right. Box A rectangular box sweeps inward from the edges of the page or outward from the center, as specified by the M entry, revealing the new page. Wipe A single line sweeps across the screen from one edge to the other in the direction specified by the Di entry, revealing the new page. Dissolve The old page “dissolves” gradually to reveal the new one. Glitter Similar to Dissolve, except that the effect sweeps across the page in a wide band moving from one side of the screen to the other in the direction specified by the Di entry. R The new page simply replaces the old one with no special transition effect; the D entry is ignored.

Default value: R.

Dm	name	(<i>Optional; Split and Blinds transition styles only</i>) The dimension in which the specified transition effect occurs: H Horizontal V Vertical Default value: H.
M	name	(<i>Optional; Split and Box transition styles only</i>) The direction of motion for the specified transition effect: I Inward from the edges of the page O Outward from the center of the page Default value: I.
Di	number	(<i>Optional; Wipe and Glitter transition styles only</i>) The direction in which the specified transition effect moves, expressed in degrees counterclockwise starting from a left-to-right direction. (Note that this differs from the page object's Rotate entry, which is measured clockwise from the top.) Only the following values are valid: 0 Left to right 90 Bottom to top (Wipe only) 180 Right to left (Wipe only) 270 Top to bottom 315 Top-left to bottom-right (Glitter only) Default value: 0.

Figure 8.1 illustrates the relationship between transition duration (**D** in the transition dictionary) and display duration (**Dur** in the page object). Note that the transition duration specified for a page (page 2 in the figure) governs the transition *to* that page from another page; the transition *from* the page is governed by the next page's transition duration.




FIGURE 8.1 *Presentation timing*

Example 8.5 shows the presentation parameters for a page to be displayed for 5 seconds. Before the page is displayed, there is a 3.5-second transition in which two vertical lines sweep outward from the center to the edges of the page.

Example 8.5

```
10 0 obj
  << /Type /Page
    /Parent 4 0 R
    /Contents 16 0 R
    /Dur 5
    /Trans << /Type /Trans
      /D 3.5
      /S /Split
      /Dm /V
      /M /O
    >>
  >>
endobj
```

8.4 Annotations

An *annotation* associates an object such as a note, sound, or movie with a location on a page of a PDF document, or provides a means of interacting with the user via the mouse and keyboard. PDF includes a wide variety of standard annotation types, described in detail in Section 8.4.5, “Annotation Types.”

Many of the standard annotation types may be displayed in either the *open* or the *closed* state. When closed, they appear on the page in some distinctive form depending on the specific annotation type, such as an icon, a box, or a rubber stamp. When the user *activates* the annotation by clicking it with the mouse, it exhibits its associated object, such as by opening a pop-up window displaying a text note (Figure 8.2) or by playing a sound or a movie.

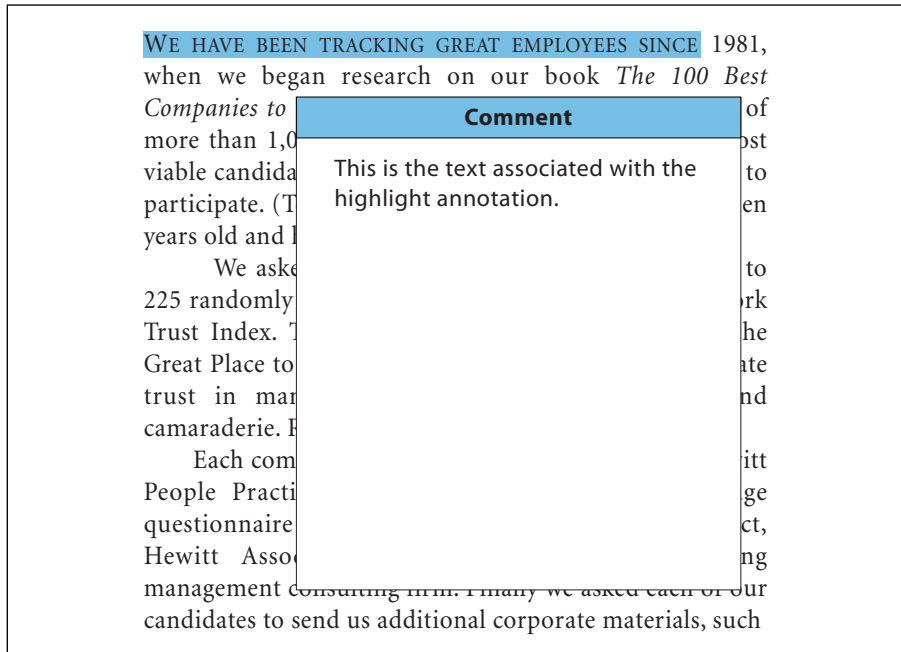


FIGURE 8.2 *Open annotation*

The behavior of each annotation type is implemented by a software module called an *annotation handler*. Handlers for the standard annotation types are built directly into the PDF viewer application; handlers for additional types can be supplied as plug-in extensions.

8.4.1 Annotation Dictionaries

The optional **Annots** entry in a page object (see “Page Objects” on page 87) holds an array of *annotation dictionaries*, each representing an annotation associated with the given page. Table 8.10 shows the required and optional entries that are common to all annotation dictionaries. The dictionary may contain additional entries specific to a particular annotation type; see the descriptions of individual annotation types in Section 8.4.5, “Annotation Types,” for details. (See also implementation note 58 in Appendix H.)

TABLE 8.10 Entries common to all annotation dictionaries

KEY	TYPE	VALUE
Type	name	<i>(Optional)</i> The type of PDF object that this dictionary describes; if present, must be Annot for an annotation dictionary.
Subtype	name	<i>(Required)</i> The type of annotation that this dictionary describes; see Table 8.14 on page 499 for specific values.
Contents	text string	<i>(Required or optional, depending on the annotation type)</i> Text to be displayed for the annotation or, if this type of annotation does not display text, an alternate description of the annotation’s contents in human-readable form. In either case, this text is useful when extracting the document’s contents in support of accessibility to disabled users or for other purposes (see Section 9.8.2, “Alternate Descriptions”).
P	dictionary	<i>(Optional; PDF 1.3; not used in FDF files)</i> An indirect reference to the page object with which this annotation is associated.
Rect	rectangle	<i>(Required)</i> The <i>annotation rectangle</i> , defining the location of the annotation on the page in default user space units.
NM	text	<i>(Optional; PDF 1.4)</i> The <i>annotation name</i> , a text string uniquely identifying it among all the annotations on its page.
M	date or string	<i>(Optional; PDF 1.1)</i> The date and time when the annotation was most recently modified. The preferred format is a date string as described in Section 3.8.2, “Dates,” but viewer applications should be prepared to accept and display a string in any format. (See implementation note 59 in Appendix H.)
F	integer	<i>(Optional; PDF 1.1)</i> A set of flags specifying various characteristics of the annotation (see Section 8.4.2, “Annotation Flags”). Default value: 0.
BS	dictionary	<i>(Optional; PDF 1.2)</i> A border style dictionary specifying the characteristics of the annotation’s border (see Section 8.4.3, “Border Styles”; see also implementation note 60 in Appendix H). <i>Note:</i> This entry also specifies the width and dash pattern for the lines drawn by line, square, circle, and ink annotations. See the note under Border (below) for additional information.
Border	array	<i>(Optional)</i> An array specifying the characteristics of the annotation’s border. The border is specified as a “rounded rectangle.” In PDF 1.0, the array consists of three numbers defining the horizontal corner radius, vertical corner radius, and border width, all in default user space units. If the corner radii are 0, the border has square (not rounded) corners; if the border width is 0, no border is drawn. (See implementation note 61 in Appendix H.)

In PDF 1.1, the array may have a fourth element, an optional *dash array* defining a pattern of dashes and gaps to be used in drawing the border. The dash array is specified in the same format as in the line dash pattern parameter of the graphics state (see “Line Dash Pattern” on page 155). For example, a **Border** value of [0 0 1 [3 2]] specifies a border 1 unit wide, with square corners, drawn with 3-unit dashes alternating with 2-unit gaps. Note that no dash phase is specified; the phase is assumed to be 0. (See implementation note 62 in Appendix H.)

Note: In PDF 1.2 or later, annotations may ignore this entry and use the **BS** entry (see above) to specify their border styles instead. In PDF 1.2 and 1.3, only widget annotations do so; in PDF 1.4, all of the standard annotation types except **Link** (see Table 8.14 on page 499) use **BS** rather than **Border** if both are present. For backward compatibility, however, **Border** is still supported for all annotation types.

Default value: [0 0 1].

AP	dictionary	(Optional; PDF 1.2) An <i>appearance dictionary</i> specifying how the annotation is presented visually on the page (see Section 8.4.4, “Appearance Streams”; see also implementation note 60 in Appendix H).
AS	name	(Required if the <i>appearance dictionary</i> AP contains one or more subdictionaries; PDF 1.2) The annotation’s <i>appearance state</i> , which selects the applicable appearance stream from an appearance subdictionary (see Section 8.4.4, “Appearance Streams”; see also implementation note 60 in Appendix H).
C	array	(Optional; PDF 1.1) An array of three numbers in the range 0.0 to 1.0, representing the components of a color in the DeviceRGB color space. This color will be used for the following purposes: <ul style="list-style-type: none"> • The background of the annotation’s icon when closed • The title bar of the annotation’s pop-up window • The border of a link annotation
CA	number	(Optional; PDF 1.4) The constant opacity value to be used in painting the annotation (see Sections 7.1, “Overview of Transparency,” and 7.2.6, “Shape and Opacity Computations”). This value applies to all visible elements of the annotation in its closed state (including its background and border), but not to the pop-up window that appears when the annotation is opened. The specified value is used as the initial alpha constant (both stroking and non-stroking) for interpreting the annotation’s appearance stream, if any (see Section 8.4.4, “Appearance Streams,” and “Constant Shape and Opacity” on page 444). The implicit blend mode (see Section 7.2.4, “Blend Mode”) is Normal . Default value: 1.0.

Note: If no explicit appearance stream is defined for the annotation, it will be painted by implementation-dependent means that do not necessarily conform to the Adobe imaging model; in this case, the effect of this entry is implementation-dependent as well.

Note: This entry is recognized by all of the standard annotation types listed in Table 8.14 on page 499 except **Link**, **Movie**, **Widget**, **PrinterMark**, and **TrapNet**.

T	text string	(Optional; PDF 1.1) The text label to be displayed in the title bar of the annotation's pop-up window when open and active.
Popup	dictionary	(Optional; PDF 1.3) An indirect reference to a pop-up annotation for entering or editing the text associated with this annotation.
A	dictionary	(Optional; PDF 1.1) An action to be performed when the annotation is activated (see Section 8.5, "Actions").
		Note: This entry is not permitted in link annotations if a Dest entry is present (see "Link Annotations" on page 501). Also note that the A entry in movie annotations has a different meaning (see "Movie Annotations" on page 510).
AA	dictionary	(Optional; PDF 1.2) An additional-actions dictionary defining the annotation's behavior in response to various trigger events (see Section 8.5.2, "Trigger Events"). At the time of publication, this entry is used only by widget annotations.
StructParent	integer	(Required if the annotation is a structural content item; PDF 1.3) The integer key of the annotation's entry in the structural parent tree (see "Finding Structure Elements from Content Items" on page 600).

8.4.2 Annotation Flags

The value of the annotation dictionary's **F** entry is an unsigned 32-bit integer containing flags specifying various characteristics of the annotation. Bit positions within the flag word are numbered from 1 (low-order) to 32 (high-order). Table 8.11 shows the meanings of the flags; all undefined flag bits are reserved and must be set to 0.

TABLE 8.11 Annotation flags

BIT POSITION	NAME	MEANING
1	Invisible	If set, do not display the annotation if it does not belong to one of the standard annotation types and no annotation handler is available. If clear, display such an unknown annotation using an appearance stream specified by its appearance dictionary, if any (see Section 8.4.4, “Appearance Streams”).
2	Hidden	(PDF 1.2) If set, do not display or print the annotation or allow it to interact with the user, regardless of its annotation type or whether an annotation handler is available. In cases where screen space is limited, the ability to hide and show annotations selectively can be used in combination with appearance streams (see Section 8.4.4, “Appearance Streams”) to display auxiliary pop-up information similar in function to online help systems. (See implementation note 63 in Appendix H.)
3	Print	(PDF 1.2) If set, print the annotation when the page is printed. If clear, never print the annotation, regardless of whether it is displayed on the screen. This can be useful, for example, for annotations representing interactive pushbuttons, which would serve no meaningful purpose on the printed page. (See implementation note 63 in Appendix H.)
4	NoZoom	(PDF 1.3) If set, do not scale the annotation’s appearance to match the magnification of the page. The location of the annotation on the page (defined by the upper-left corner of its annotation rectangle) remains fixed, regardless of the page magnification. See below for further discussion.
5	NoRotate	(PDF 1.3) If set, do not rotate the annotation’s appearance to match the rotation of the page. The upper-left corner of the annotation rectangle remains in a fixed location on the page, regardless of the page rotation. See below for further discussion.
6	NoView	(PDF 1.3) If set, do not display the annotation on the screen or allow it to interact with the user. The annotation may be printed (depending on the setting of the Print flag), but should be considered hidden for purposes of on-screen display and user interaction.
7	ReadOnly	(PDF 1.3) If set, do not allow the annotation to interact with the user. The annotation may be displayed or printed (depending on the settings of the NoView and Print flags), but should not respond to mouse clicks or change its appearance in response to mouse motions.

Note: This flag is ignored for widget annotations; its function is subsumed by the *ReadOnly* flag of the associated form field (see Table 8.50 on page 532).

If the `NoZoom` flag is set, the annotation always maintains the same fixed size on the screen and is unaffected by the magnification level at which the page itself is displayed. Similarly, if the `NoRotate` flag is set, the annotation retains its original orientation on the screen when the page is rotated (by changing the `Rotate` entry in the page object; see “Page Objects” on page 87).

In either case, the annotation’s position is determined by the coordinates of the upper-left corner of its annotation rectangle, as defined by the `Rect` entry in the annotation dictionary and interpreted in the default user space of the page. When the default user space is scaled or rotated, the positions of the other three corners of the annotation rectangle will be different in the altered user space than they were in the original user space. The viewer application performs this alteration automatically; however, it does not actually change the annotation’s `Rect` entry, which continues to describe the annotation’s relationship with the unscaled, unrotated user space.

For example, Figure 8.3 shows how an annotation whose `NoRotate` flag is set remains upright when the page it is on is rotated 90 degrees clockwise. The upper-left corner of the annotation remains at the same point in default user space; the annotation pivots around that point.

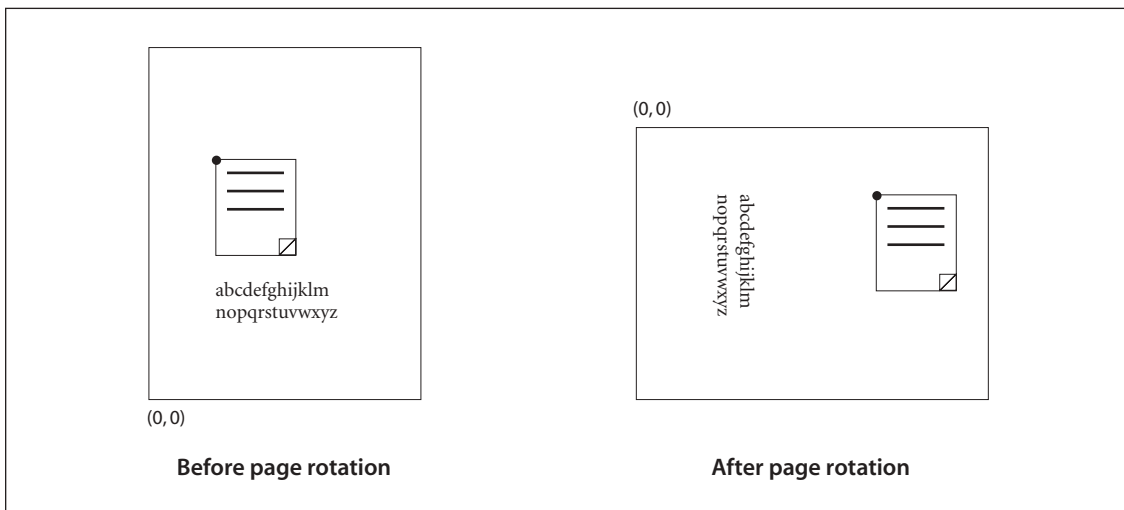


FIGURE 8.3 *Coordinate adjustment with the `NoRotate` flag*

8.4.3 Border Styles

An annotation may optionally be surrounded by a border when displayed or printed. If present, the border is drawn completely inside the annotation rectangle. In PDF 1.1, the characteristics of the border are specified by the **Border** entry in the annotation dictionary (see Table 8.10 on page 490). Beginning with PDF 1.2, some types of annotation may instead specify their border characteristics in a *border style dictionary* designated by the annotation's **BS** entry. Such dictionaries are also used to specify the width and dash pattern for the lines drawn by line, square, circle, and ink annotations. Table 8.12 summarizes the contents of the border style dictionary. If neither the **Border** nor the **BS** entry is present, the border is drawn as a solid line with a width of 1 point.

TABLE 8.12 Entries in a border style dictionary

KEY	TYPE	VALUE
Type	name	<i>(Optional)</i> The type of PDF object that this dictionary describes; if present, must be Border for a border style dictionary.
W	number	<i>(Optional)</i> The border width in points. If this value is 0, no border is drawn. Default value: 1.
S	name	<p><i>(Optional)</i> The border style:</p> <ul style="list-style-type: none"> S (Solid) A solid rectangle surrounding the annotation. D (Dashed) A dashed rectangle surrounding the annotation. The dash pattern is specified by the D entry (see below). B (Beveled) A simulated embossed rectangle that appears to be raised above the surface of the page. I (Inset) A simulated engraved rectangle that appears to be recessed below the surface of the page. U (Underline) A single line along the bottom of the annotation rectangle. <p>Other border styles may be defined in the future. (See implementation note 64 in Appendix H.) Default value: S.</p>
D	array	<i>(Optional)</i> A <i>dash array</i> defining a pattern of dashes and gaps to be used in drawing a dashed border (border style D above). The dash array is specified in the same format as in the line dash pattern parameter of the graphics state (see “Line Dash Pattern” on page 155). The dash phase is not specified and is assumed to be 0. For example, a D entry of [3 2] specifies a border drawn with 3-point dashes alternating with 2-point gaps. Default value: [3].

Note: In PDF 1.2 and 1.3, only widget annotations use border style dictionaries to specify their border styles; all of the other standard annotation types (see Table 8.14 on page 499) use the **Border** entry instead. In PDF 1.4, all except link annotations use border style dictionaries.

8.4.4 Appearance Streams

Beginning with PDF 1.2, an annotation can specify one or more *appearance streams* as an alternative to the simple border and color characteristics available in earlier versions. This allows the annotation to be presented visually on the page in different ways to reflect its interactions with the user. Each appearance stream is a form XObject (see Section 4.9, “Form XObjects”): a self-contained content stream to be rendered inside the annotation rectangle.

Before rendering an appearance stream, the viewer application establishes a coordinate system whose origin is at the lower-left corner of the annotation rectangle. Nominally, the coordinates in this space are in default user space units; however, the space may be scaled or rotated with respect to the actual page coordinate system, depending on the values of the annotation’s **NoZoom** and **NoRotate** flags (see Section 8.4.2, “Annotation Flags”). The form XObject’s **Matrix** entry should transform the form coordinate system into that of the annotation; the results are clipped to the annotation rectangle. All other elements of the graphics state are set to their default values.

In PDF 1.4, an annotation appearance can include transparency if the appearance stream defining it is a transparency group XObject (see Section 7.5.5, “Transparency Group XObjects”). This transparency group is composited with a backdrop consisting of the page content along with any previously painted annotations. The final compositing operation is performed using blend mode **Normal**, an alpha constant of 1.0 (or other value as specified by the **CA** entry in the annotation dictionary), and a soft mask of **None**.

Note: If a transparent annotation appearance is painted over an annotation that is drawn without using an appearance stream, the effect is implementation-dependent. This is because such annotations are sometimes drawn by means that do not conform to the Adobe imaging model. Also, the effect of highlighting a transparent annotation appearance is implementation-dependent.

An annotation can define as many as three separate appearances:

- The *normal appearance* is used when the annotation is not interacting with the user. This is also the appearance that is used for printing the annotation.
- The *rollover appearance* is used when the user moves the cursor into the annotation's active area without pressing the mouse button.
- The *down appearance* is used when the mouse button is pressed or held down within the annotation's active area.

Note: As used here, the term *mouse* denotes a generic pointing device that controls the location of a cursor on the screen and has at least one button that can be pressed, held down, and released. See Section 8.5.2, “Trigger Events,” for further discussion.

The normal, rollover, and down appearances are defined in an *appearance dictionary*, which in turn is the value of the **AP** entry in the annotation dictionary (see Table 8.10 on page 490). Table 8.13 shows the contents of the appearance dictionary.

TABLE 8.13 Entries in an appearance dictionary

KEY	TYPE	VALUE
N	stream or dictionary	<i>(Required)</i> The annotation's normal appearance.
R	stream or dictionary	<i>(Optional)</i> The annotation's rollover appearance. Default value: the value of the N entry.
D	stream or dictionary	<i>(Optional)</i> The annotation's down appearance. Default value: the value of the N entry.

Each entry in the appearance dictionary may contain either a single appearance stream or an *appearance subdictionary*. In the latter case, the subdictionary defines multiple appearance streams corresponding to different *appearance states* of the annotation.

For example, an annotation representing an interactive checkbox might have two appearance states named On and Off. Its appearance dictionary might be defined as follows:

```

/AP << /N << /On formXObject1
           /Off formXObject2
        >>
/D << /On formXObject3
       /Off formXObject4
     >>
>>

```

where *formXObject₁* and *formXObject₂* define the checkbox’s normal appearance in its checked and unchecked states, while *formXObject₃* and *formXObject₄* provide visual feedback, such as emboldening its outline, when the user clicks it with the mouse. (No **R** entry is defined because no special appearance is needed when the user moves the cursor over the checkbox without pressing the mouse button.) The choice between the checked and unchecked appearance states is determined by the **AS** entry in the annotation dictionary (see Table 8.10 on page 490).

***Note:** Some of the standard PDF annotation types, such as movie annotations—as well as all custom annotation types defined by third parties—are implemented through plug-in extensions. If the plug-in for a particular annotation type is not available, PDF viewer applications should display the annotation with its normal (**N**) appearance. Viewer applications should also attempt to provide reasonable behavior (such as displaying nothing at all) if an annotation’s **AS** entry designates an appearance state for which no appearance is defined in the appearance dictionary.*

For convenience in managing appearance streams that are used repeatedly, the **AP** entry in a PDF document’s name dictionary (see Section 3.6.3, “Name Dictionary”) can contain a name tree mapping name strings to appearance streams. The name strings have no standard meanings; no PDF objects refer to appearance streams by name.

8.4.5 Annotation Types

PDF supports the standard annotation types listed in Table 8.14. The following sections describe each of these types in detail. Plug-in extensions may add new annotation types, and further standard types may be added in the future. (See implementation note 65 in Appendix H.)

TABLE 8.14 Annotation types

ANNOTATION TYPE	DESCRIPTION
Text	Text annotation
Link	Link annotation
FreeText	(PDF 1.3) Free text annotation
Line	(PDF 1.3) Line annotation
Square	(PDF 1.3) Square annotation
Circle	(PDF 1.3) Circle annotation
Highlight	(PDF 1.3) Highlight annotation
Underline	(PDF 1.3) Underline annotation
Squiggly	(PDF 1.4) Squiggly-underline annotation
StrikeOut	(PDF 1.3) Strikeout annotation
Stamp	(PDF 1.3) Rubber stamp annotation
Ink	(PDF 1.3) Ink annotation
Popup	(PDF 1.3) Pop-up annotation
FileAttachment	(PDF 1.3) File attachment annotation
Sound	(PDF 1.2) Sound annotation
Movie	(PDF 1.2) Movie annotation
Widget	(PDF 1.2) Widget annotation
PrinterMark	(PDF 1.4) Printer's mark annotation
TrapNet	(PDF 1.3) Trap network annotation

Text Annotations

A *text annotation* represents a “sticky note” attached to a point in the PDF document. When closed, the annotation appears as an icon; when open, it displays a pop-up window containing the text of the note, in a font and size chosen by the viewer application. Text annotations do not scale and rotate with the page; they

behave as if the NoZoom and NoRotate annotation flags (see Table 8.11 on page 493) were always set. Table 8.15 shows the annotation dictionary entries specific to this type of annotation.

TABLE 8.15 Additional entries specific to a text annotation

KEY	TYPE	VALUE									
Subtype	name	<i>(Required)</i> The type of annotation that this dictionary describes; must be Text for a text annotation.									
Contents	text string	<i>(Required)</i> The text to be displayed in the pop-up window when the annotation is opened. Carriage returns may be used to separate the text into paragraphs.									
Open	boolean	<i>(Optional)</i> A flag specifying whether the annotation should initially be displayed open. Default value: false (closed).									
Name	name	<p><i>(Optional)</i> The name of an icon to be used in displaying the annotation. Viewer applications should provide predefined icon appearances for at least the following standard names:</p> <table border="0"> <tr> <td>Comment</td> <td>Key</td> <td>Note</td> </tr> <tr> <td>Help</td> <td>NewParagraph</td> <td>Paragraph</td> </tr> <tr> <td>Insert</td> <td></td> <td></td> </tr> </table> <p>Additional names may be supported as well. Default value: Note.</p> <p>Note: The annotation dictionary's AP entry, if present, takes precedence over the Name entry; see Table 8.10 on page 490 and Section 8.4.4, "Appearance Streams."</p>	Comment	Key	Note	Help	NewParagraph	Paragraph	Insert		
Comment	Key	Note									
Help	NewParagraph	Paragraph									
Insert											

Example 8.6 shows the definition of a text annotation.

Example 8.6

```

22 0 obj
  << /Type /Annot
    /Subtype /Text
    /Rect [266 116 430 204]
    /Contents (The quick brown fox ate the lazy mouse.)
  >>
endobj

```

Link Annotations

A *link annotation* represents either a hypertext link to a destination elsewhere in the document (see Section 8.2.1, “Destinations”) or an action to be performed (Section 8.5, “Actions”). Table 8.16 shows the annotation dictionary entries specific to this type of annotation.

TABLE 8.16 Additional entries specific to a link annotation

KEY	TYPE	VALUE
Subtype	name	<i>(Required)</i> The type of annotation that this dictionary describes; must be Link for a link annotation.
Contents	text string	<i>(Optional; PDF 1.4)</i> An alternate representation of the annotation’s contents in human-readable form, useful when extracting the document’s contents in support of accessibility to disabled users or for other purposes (see Section 9.8.2, “Alternate Descriptions”).
Dest	array, name, or string	<i>(Optional; not permitted if an A entry is present)</i> A destination to be displayed when the annotation is activated (see Section 8.2.1, “Destinations”; see also implementation note 66 in Appendix H).
H	name	<p><i>(Optional; PDF 1.2)</i> The annotation’s <i>highlighting mode</i>, the visual effect to be used when the mouse button is pressed or held down inside its active area:</p> <ul style="list-style-type: none"> N (None) No highlighting. I (Invert) Invert the contents of the annotation rectangle. O (Outline) Invert the annotation’s border. P (Push) Display the annotation’s down appearance, if any (see Section 8.4.4, “Appearance Streams”). If no down appearance is defined, offset the contents of the annotation rectangle to appear as if it were being “pushed” below the surface of the page. <p>A highlighting mode other than P overrides any down appearance defined for the annotation. Default value: I.</p> <p><i>Note: In PDF 1.1, highlighting is always done by inverting colors inside the annotation rectangle.</i></p>
PA	dictionary	<i>(Optional; PDF 1.3)</i> A URI action (see “URI Actions” on page 523) formerly associated with this annotation. When Web Capture (Section 9.9, “Web Capture”) changes an annotation from a URI to a go-to action (“Go-To Actions” on page 519), it uses this entry to save the data from the original URI action so that it can be changed back in case the target page for the go-to action is subsequently deleted.

Example 8.7 shows a link annotation that jumps to a destination elsewhere in the document.

Example 8.7

```
93 0 obj
  << /Type /Annot
    /Subtype /Link
    /Rect [71 717 190 734]
    /Border [16 16 1]
    /Dest [3 0 R /FitR -4 399 199 533]
  >>
endobj
```

Free Text Annotations

A *free text annotation* (PDF 1.3) displays text directly on the page. Unlike an ordinary text annotation (see “Text Annotations” on page 499), a free text annotation has no open or closed state; instead of being displayed in a pop-up window, the text is always visible. Table 8.17 shows the annotation dictionary entries specific to this type of annotation.

TABLE 8.17 Additional entries specific to a free text annotation

KEY	TYPE	VALUE
Subtype	name	<i>(Required)</i> The type of annotation that this dictionary describes; must be FreeText for a free text annotation.
Contents	text string	<i>(Required)</i> The text to be displayed.
DA	string	<i>(Required)</i> The default appearance string to be used in formatting the text (see “Variable Text” on page 533). <i>Note: The annotation dictionary’s AP entry, if present, takes precedence over the DA entry; see Table 8.10 on page 490 and Section 8.4.4, “Appearance Streams.”</i>
Q	integer	<i>(Optional; PDF 1.4)</i> A code specifying the form of <i>quadding</i> (justification) to be used in displaying the annotation’s text: <ul style="list-style-type: none"> 0 Left-justified 1 Centered 2 Right-justified Default value: 0 (left-justified).







Line Annotations

A *line annotation* (PDF 1.3) displays a single straight line on the page. When opened, it displays a pop-up window containing the text of the associated note. Table 8.18 shows the annotation dictionary entries specific to this type of annotation.

TABLE 8.18 Additional entries specific to a line annotation

KEY	TYPE	VALUE
Subtype	name	<i>(Required)</i> The type of annotation that this dictionary describes; must be Line for a line annotation.
Contents	text string	<i>(Required)</i> The text to be displayed in the pop-up window when the annotation is opened. Carriage returns may be used to separate the text into paragraphs.
L	array	<i>(Required)</i> An array of four numbers, $[x_1\ y_1\ x_2\ y_2]$, specifying the starting and ending coordinates of the line in default user space.
BS	dictionary	<i>(Optional)</i> A border style dictionary (see Table 8.12 on page 495) specifying the width and dash pattern to be used in drawing the line. <i>Note: The annotation dictionary's AP entry, if present, takes precedence over the L and BS entries; see Table 8.10 on page 490 and Section 8.4.4, "Appearance Streams."</i>
LE	array	<i>(Optional; PDF 1.4)</i> An array of two names specifying the line ending styles to be used in drawing the line. The first and second elements of the array specify the line ending styles for the endpoints defined, respectively, by the first and second pairs of coordinates, (x_1, y_1) and (x_2, y_2) , in the L array. Table 8.19 shows the possible values. Default value: [/None /None].
IC	array	<i>(Optional; PDF 1.4)</i> An array of three numbers in the range 0.0 to 1.0 specifying the components, in the DeviceRGB color space, of the <i>interior color</i> with which to fill the annotation's line endings (see Table 8.19). If this entry is absent, the interiors of the line endings are left transparent.

TABLE 8.19 Line ending styles

NAME	APPEARANCE	DESCRIPTION
Square		A square filled with the annotation's interior color, if any
Circle		A circle filled with the annotation's interior color, if any
Diamond		A diamond shape filled with the annotation's interior color, if any
OpenArrow		Two short lines meeting in an acute angle, forming an open arrowhead
ClosedArrow		Two short lines meeting in an acute angle as in the OpenArrow style (see above), connected by a third line to form a triangular closed arrowhead filled with the annotation's interior color, if any
None		No line ending

Square and Circle Annotations

Square and *circle* annotations (PDF 1.3) display, respectively, a rectangle or an ellipse on the page. When opened, they display a pop-up window containing the text of the associated note. The rectangle or ellipse is inscribed within the annotation rectangle defined by the annotation dictionary's **Rect** entry (see Table 8.10 on page 490); Figure 8.4 shows an example, using a border width of 18 points. Despite the names *square* and *circle*, the width and height of the annotation rectangle need not be equal. Table 8.20 shows the annotation dictionary entries specific to these types of annotation.




FIGURE 8.4 Square and circle annotations

TABLE 8.20 Additional entries specific to a square or circle annotation

KEY	TYPE	VALUE
Subtype	name	<i>(Required)</i> The type of annotation that this dictionary describes; must be Square or Circle for a square or circle annotation, respectively.
Contents	text string	<i>(Required)</i> The text to be displayed in the pop-up window when the annotation is opened. Carriage returns may be used to separate the text into paragraphs.
BS	dictionary	<i>(Optional)</i> A border style dictionary (see Table 8.12 on page 495) specifying the line width and dash pattern to be used in drawing the rectangle or ellipse. <i>Note:</i> The annotation dictionary's AP entry, if present, takes precedence over the Rect and BS entries; see Table 8.10 on page 490 and Section 8.4.4, "Appearance Streams."
IC	array	<i>(Optional; PDF 1.4)</i> An array of three numbers in the range 0.0 to 1.0 specifying the components, in the DeviceRGB color space, of the <i>interior color</i> with which to fill the annotation's rectangle or ellipse (see Table 8.19). If this entry is absent, the interior of the annotation is left transparent.

Markup Annotations

Markup annotations appear as highlights, underlines, strikeouts (*all PDF 1.3*), or jagged ("squiggly") underlines (*PDF 1.4*) in the text of a document. When opened, they display a pop-up window containing the text of the associated note.

Table 8.21 shows the annotation dictionary entries specific to these types of annotation.

TABLE 8.21 Additional entries specific to markup annotations

KEY	TYPE	VALUE
Subtype	name	<i>(Required)</i> The type of annotation that this dictionary describes; must be Highlight , Underline , Squiggly , or StrikeOut for a highlight, underline, squiggly-underline, or strikeout annotation, respectively.
Contents	text string	<i>(Required)</i> The text to be displayed in the pop-up window when the annotation is opened. Carriage returns may be used to separate the text into paragraphs.
QuadPoints	array	<p><i>(Required)</i> An array of $8 \times n$ numbers specifying the coordinates of n quadrilaterals in default user space. Each quadrilateral encompasses a word or group of contiguous words in the text underlying the annotation. The coordinates for each quadrilateral are given in the order</p> $x_1 \ y_1 \ x_2 \ y_2 \ x_3 \ y_3 \ x_4 \ y_4$ <p>specifying the quadrilateral's four vertices in counterclockwise order (see Figure 8.5). The text is oriented with respect to the edge connecting points (x_1, y_1) and (x_2, y_2). (See implementation note 67 in Appendix H.)</p> <p>Note: The annotation dictionary's AP entry, if present, takes precedence over the QuadPoints entry; see Table 8.10 on page 490 and Section 8.4.4, "Appearance Streams."</p>




FIGURE 8.5 *QuadPoints* specification

Rubber Stamp Annotations

A *rubber stamp annotation* (PDF 1.3) displays text or graphics intended to look as if they were stamped on the page with a rubber stamp. When opened, it displays a pop-up window containing the text of the associated note. Table 8.22 shows the annotation dictionary entries specific to this type of annotation.

TABLE 8.22 Additional entries specific to a rubber stamp annotation

KEY	TYPE	VALUE															
Subtype	name	<i>(Required)</i> The type of annotation that this dictionary describes; must be Stamp for a rubber stamp annotation.															
Contents	text string	<i>(Required)</i> The text to be displayed in the pop-up window when the annotation is opened. Carriage returns may be used to separate the text into paragraphs.															
Name	name	<p><i>(Optional)</i> The name of an icon to be used in displaying the annotation. Viewer applications should provide predefined icon appearances for at least the following standard names:</p> <table border="0"> <tr> <td>Approved</td> <td>Experimental</td> <td>NotApproved</td> </tr> <tr> <td>AsIs</td> <td>Expired</td> <td>NotForPublicRelease</td> </tr> <tr> <td>Confidential</td> <td>Final</td> <td>Sold</td> </tr> <tr> <td>Departmental</td> <td>ForComment</td> <td>TopSecret</td> </tr> <tr> <td>Draft</td> <td>ForPublicRelease</td> <td></td> </tr> </table>	Approved	Experimental	NotApproved	AsIs	Expired	NotForPublicRelease	Confidential	Final	Sold	Departmental	ForComment	TopSecret	Draft	ForPublicRelease	
Approved	Experimental	NotApproved															
AsIs	Expired	NotForPublicRelease															
Confidential	Final	Sold															
Departmental	ForComment	TopSecret															
Draft	ForPublicRelease																

Additional names may be supported as well. Default value: Draft.

Note: The annotation dictionary's **AP** entry, if present, takes precedence over the **Name** entry; see Table 8.10 on page 490 and Section 8.4.4, "Appearance Streams."

Ink Annotations

An *ink annotation* (PDF 1.3) represents a freehand "scribble" composed of one or more disjoint paths. When opened, it displays a pop-up window containing the text of the associated note. Table 8.23 shows the annotation dictionary entries specific to this type of annotation.

TABLE 8.23 Additional entries specific to an ink annotation

KEY	TYPE	VALUE
Subtype	name	<i>(Required)</i> The type of annotation that this dictionary describes; must be Ink for an ink annotation.
Contents	text string	<i>(Required)</i> The text to be displayed in the pop-up window when the annotation is opened. Carriage returns may be used to separate the text into paragraphs.
InkList	array	<i>(Required)</i> An array of n arrays, each representing a stroked path. Each array is a series of alternating horizontal and vertical coordinates in default user space, specifying points along the path. When drawn, the points are connected by straight lines or curves in an implementation-dependent way. (See implementation note 68 in Appendix H.)
BS	dictionary	<i>(Optional)</i> A border style dictionary (see Table 8.12 on page 495) specifying the line width and dash pattern to be used in drawing the paths. <i>Note:</i> The annotation dictionary's AP entry, if present, takes precedence over the InkList and BS entries; see Table 8.10 on page 490 and Section 8.4.4, "Appearance Streams."

Pop-up Annotations

A *pop-up annotation* (PDF 1.3) displays text in a pop-up window for entry and editing. It typically does not appear alone, but is associated with another annotation, its *parent annotation*, and is used for editing the parent's text. It has no appearance stream or associated actions of its own, and is identified by the **Popup** entry in the parent's annotation dictionary (see Table 8.10 on page 490). Table 8.24 shows the annotation dictionary entries specific to this type of annotation.

File Attachment Annotations

A *file attachment annotation* (PDF 1.3) contains a reference to a file, which typically will be embedded in the PDF file (see Section 3.10.3, "Embedded File Streams"). For example, a table of data might use a file attachment annotation to link to a spreadsheet file based on that data; activating the annotation will extract the embedded file and give the user an opportunity to view it or store it in the file system. Table 8.25 shows the annotation dictionary entries specific to this type of annotation.

TABLE 8.24 Additional entries specific to a pop-up annotation

KEY	TYPE	VALUE
Subtype	name	<i>(Required)</i> The type of annotation that this dictionary describes; must be Popup for a pop-up annotation.
Contents	text string	<i>(Optional; PDF 1.4)</i> An alternate representation of the annotation's contents in human-readable form, useful when extracting the document's contents in support of accessibility to disabled users or for other purposes (see Section 9.8.2, "Alternate Descriptions").
Parent	dictionary	<i>(Optional; must be an indirect reference)</i> The parent annotation with which this pop-up annotation is associated. <i>Note: If this entry is present, the parent annotation's Contents, M, C, and T entries (see Table 8.10 on page 490) override those of the pop-up annotation itself.</i>
Open	boolean	<i>(Optional)</i> A flag specifying whether the pop-up annotation should initially be displayed open. Default value: false (closed).

TABLE 8.25 Additional entries specific to a file attachment annotation

KEY	TYPE	VALUE				
Subtype	name	<i>(Required)</i> The type of annotation that this dictionary describes; must be FileAttachment for a file attachment annotation.				
FS	file specification	<i>(Required)</i> The file associated with this annotation.				
Contents	text string	<i>(Required)</i> The text to be displayed in the pop-up window when the annotation is opened. Carriage returns may be used to separate the text into paragraphs.				
Name	name	<i>(Optional)</i> The name of an icon to be used in displaying the annotation. Viewer applications should provide predefined icon appearances for at least the following standard names: <table style="margin-left: auto; margin-right: auto;"> <tr> <td>Graph</td> <td>PushPin</td> </tr> <tr> <td>Paperclip</td> <td>Tag</td> </tr> </table> Additional names may be supported as well. Default value: PushPin. <i>Note: The annotation dictionary's AP entry, if present, takes precedence over the Name entry; see Table 8.10 on page 490 and Section 8.4.4, "Appearance Streams."</i>	Graph	PushPin	Paperclip	Tag
Graph	PushPin					
Paperclip	Tag					

Sound Annotations

A *sound annotation* (PDF 1.2) is analogous to a text annotation, except that instead of a text note, it contains sound recorded from the computer’s microphone or imported from a file. When the annotation is activated, the sound is played. The annotation behaves like a text annotation in most ways, with a different icon (by default, a speaker) to indicate that it represents a sound. Table 8.26 shows the annotation dictionary entries specific to this type of annotation; sounds themselves are discussed in Section 8.7, “Sounds.”

TABLE 8.26 Additional entries specific to a sound annotation

KEY	TYPE	VALUE
Subtype	name	<i>(Required)</i> The type of annotation that this dictionary describes; must be Sound for a sound annotation.
Sound	stream	<i>(Required)</i> A sound object defining the sound to be played when the annotation is activated (see Section 8.7, “Sounds”).
Contents	text string	<i>(Optional)</i> Text to be displayed in a pop-up window for the annotation in place of the sound, useful when extracting the document’s contents in support of accessibility to disabled users or for other purposes (see Section 9.8.2, “Alternate Descriptions”).
Name	name	<i>(Optional)</i> The name of an icon to be used in displaying the annotation. Viewer applications should provide predefined icon appearances for at least the standard names Speaker and Microphone ; additional names may be supported as well. Default value: Speaker .

Note: The annotation dictionary’s **AP** entry, if present, takes precedence over the **Name** entry; see Table 8.10 on page 490 and Section 8.4.4, “Appearance Streams.”

Movie Annotations

A *movie annotation* (PDF 1.2) contains animated graphics and sound to be presented on the computer screen and through the speakers. When the annotation is activated, the movie is played. Table 8.27 shows the annotation dictionary entries specific to this type of annotation; movies themselves are discussed in Section 8.8, “Movies.” (See also implementation note 69 in Appendix H.)

TABLE 8.27 Additional entries specific to a movie annotation

KEY	TYPE	VALUE
Subtype	name	<i>(Required)</i> The type of annotation that this dictionary describes; must be Movie for a movie annotation.
Movie	dictionary	<i>(Required)</i> A movie dictionary describing the movie’s static characteristics (see Section 8.8, “Movies”).
Contents	text string	<i>(Optional; PDF 1.4)</i> An alternate representation of the annotation’s contents in human-readable form, useful when extracting the document’s contents in support of accessibility to disabled users or for other purposes (see Section 9.8.2, “Alternate Descriptions”).
A	boolean or dictionary	<i>(Optional)</i> A flag or dictionary specifying whether and how to play the movie when the annotation is activated. If this value is a dictionary, it is a movie activation dictionary (see Section 8.8, “Movies”) specifying how to play the movie; if it is the boolean value true , the movie should be played using default activation parameters; if it is false , the movie should not be played at all. Default value: true .

Widget Annotations

Interactive forms (see Section 8.6, “Interactive Forms”) use *widget annotations* (*PDF 1.2*) to represent the appearance of fields and to manage user interactions. As a convenience, when a field has only a single associated widget annotation, the contents of the field dictionary (Section 8.6.2, “Field Dictionaries”) and the annotation dictionary may be merged into a single dictionary containing entries that pertain to both a field and an annotation. (This presents no ambiguity, since the contents of the two kinds of dictionary do not conflict.) Table 8.28 shows the annotation dictionary entries specific to this type of annotation; interactive forms and fields are discussed at length in Section 8.6.

TABLE 8.28 Additional entries specific to a widget annotation

KEY	TYPE	VALUE
Subtype	name	<i>(Required)</i> The type of annotation that this dictionary describes; must be Widget for a widget annotation.
Contents	text string	<i>(Optional; PDF 1.4)</i> An alternate representation of the annotation’s contents in human-readable form, useful when extracting the document’s contents in support of accessibility to disabled users or for other purposes (see Section 9.8.2, “Alternate Descriptions”).
H	name	<p><i>(Optional)</i> The annotation’s <i>highlighting mode</i>, the visual effect to be used when the mouse button is pressed or held down inside its active area:</p> <ul style="list-style-type: none"> N (None) No highlighting. I (Invert) Invert the contents of the annotation rectangle. O (Outline) Invert the annotation’s border. P (Push) Display the annotation’s down appearance, if any (see Section 8.4.4, “Appearance Streams”). If no down appearance is defined, offset the contents of the annotation rectangle to appear as if it were being “pushed” below the surface of the page. T (Toggle) Same as P (which is preferred). <p>A highlighting mode other than P overrides any down appearance defined for the annotation. Default value: I.</p>
MK	dictionary	<p><i>(Optional)</i> An appearance characteristics dictionary to be used in constructing a dynamic appearance stream specifying the annotation’s visual presentation on the page; see “Variable Text” on page 533 for further discussion.</p> <p>Note: The name MK for this entry is of historical significance only and has no direct meaning.</p>

Printer’s Mark Annotations

A *printer’s mark annotation* (PDF 1.4) represents a graphic symbol, such as a registration target, color bar, or cut mark, added to a page to assist production personnel in identifying components of a multiple-plate job and maintaining consistent output during production. See Section 9.10.2, “Printer’s Marks,” for further discussion.

Trap Network Annotations

A *trap network annotation* (PDF 1.3) defines the trapping characteristics for a page of a PDF document. (*Trapping* is the process of adding marks to a page along color boundaries to avoid unwanted visual artifacts resulting from misregistration of colorants when the page is printed.) A page may have at most one trap network annotation, whose **Subtype** entry has the value **TrapNet** and which is always the last element in the page object's **Annots** array (see “Page Objects” on page 87). See Section 9.10.5, “Trapping Support,” for further discussion.

8.5 Actions

Instead of simply jumping to a destination in the document, an annotation or outline item can specify an *action* (PDF 1.1) for the viewer application to perform, such as launching an application, playing a sound, or changing an annotation's appearance state. The optional **A** entry in the annotation or outline item dictionary (see Tables 8.10 on page 490 and 8.4 on page 478) specifies an action to be performed when the annotation or outline item is activated; in PDF 1.2, a variety of other circumstances may trigger an action as well (see Section 8.5.2, “Trigger Events”). In addition, the optional **OpenAction** entry in a document's catalog (Section 3.6.1, “Document Catalog”) may specify an action to be performed when the document is opened. PDF includes a wide variety of standard action types, described in detail in Section 8.5.3, “Action Types.”

8.5.1 Action Dictionaries

An *action dictionary* defines the characteristics and behavior of an action. Table 8.29 shows the required and optional entries that are common to all action dictionaries. The dictionary may contain additional entries specific to a particular action type; see the descriptions of individual action types in Section 8.5.3, “Action Types,” for details.

TABLE 8.29 Entries common to all action dictionaries

KEY	TYPE	VALUE
Type	name	<i>(Optional)</i> The type of PDF object that this dictionary describes; if present, must be Action for an action dictionary.
S	name	<i>(Required)</i> The type of action that this dictionary describes; see Table 8.34 on page 518 for specific values.
Next	dictionary or array	<i>(Optional; PDF 1.2)</i> The next action, or sequence of actions, to be performed after this one. The value is either a single action dictionary or an array of action dictionaries to be performed in order; see below for further discussion.

The action dictionary's **Next** entry (*PDF 1.2*) allows sequences of actions to be chained together. For example, the effect of clicking a link annotation with the mouse might be to play a sound, jump to a new page, and start up a movie. Note that the **Next** entry is not restricted to a single action, but may contain an array of actions, each of which in turn may have a **Next** entry of its own. The actions may thus form a tree instead of a simple linked list. Actions within each **Next** array are executed in order, each followed in turn by any actions specified in *its* **Next** entry, and so on recursively. Viewer applications should attempt to provide reasonable behavior in anomalous situations; for example, self-referential actions should not be executed more than once, and actions that close the document or otherwise render the next action impossible should terminate the execution sequence. Applications should also provide some mechanism for the user to interrupt and manually terminate a sequence of actions.

Note: No action should modify its own action dictionary or any other in the action tree in which it resides. The effect of such modification on subsequent execution of actions in the tree is undefined.

8.5.2 Trigger Events

An annotation, page object, or (beginning with PDF 1.3) interactive form field may include an entry named **AA** that specifies an *additional-actions dictionary* (*PDF 1.2*), extending the set of events that can trigger the execution of an action. In PDF 1.4, the document catalog dictionary (see Section 3.6.1, “Document Cat-

alog”) may also contain an **AA** entry for trigger events affecting the document as a whole. Tables 8.30 to 8.33 show the contents of this type of dictionary. (See implementation notes 70 and 71 in Appendix H.)

TABLE 8.30 Entries in an annotation’s additional-actions dictionary

KEY	TYPE	VALUE
E	dictionary	<i>(Optional; PDF 1.2)</i> An action to be performed when the cursor enters the annotation’s active area.
X	dictionary	<i>(Optional; PDF 1.2)</i> An action to be performed when the cursor exits the annotation’s active area.
D	dictionary	<i>(Optional; PDF 1.2)</i> An action to be performed when the mouse button is pressed inside the annotation’s active area. (The name D stands for “down.”)
U	dictionary	<i>(Optional; PDF 1.2)</i> An action to be performed when the mouse button is released inside the annotation’s active area. (The name U stands for “up.”)
		<i>Note: For backward compatibility, the A entry in an annotation dictionary, if present, takes precedence over this entry (see Table 8.10 on page 490).</i>
Fo	dictionary	<i>(Optional; PDF 1.2; widget annotations only)</i> An action to be performed when the annotation receives the input focus.
Bl	dictionary	<i>(Optional; PDF 1.2; widget annotations only)</i> (Uppercase B, lowercase L) An action to be performed when the annotation loses the input focus. (The name Bl stands for “blurred.”)

TABLE 8.31 Entries in a page object’s additional-actions dictionary

O	dictionary	<i>(Optional; PDF 1.2)</i> An action to be performed when the page is opened (for example, when the user navigates to it from the next or previous page or via a link annotation or outline item). This action is independent of any that may be defined by the Open-Action entry in the document catalog (see Section 3.6.1, “Document Catalog”), and is executed after such an action. (See implementation note 72 in Appendix H.)
C	dictionary	<i>(Optional; PDF 1.2)</i> An action to be performed when the page is closed (for example, when the user navigates to the next or previous page or follows a link annotation or an outline item). This action applies to the page being closed, and is executed before any other page is opened. (See implementation note 72 in Appendix H.)

TABLE 8.32 Entries in a form field's additional-actions dictionary

K	dictionary	<i>(Optional; PDF 1.3)</i> A JavaScript action to be performed when the user types a keystroke into a text field or combo box or modifies the selection in a scrollable list box. This allows the keystroke to be checked for validity and rejected or modified.
F	dictionary	<i>(Optional; PDF 1.3)</i> A JavaScript action to be performed before the field is formatted to display its current value. This allows the field's value to be modified before formatting.
V	dictionary	<i>(Optional; PDF 1.3)</i> A JavaScript action to be performed when the field's value is changed. This allows the new value to be checked for validity. (The name V stands for "validate.")
C	dictionary	<i>(Optional; PDF 1.3)</i> A JavaScript action to be performed in order to recalculate the value of this field when that of another field changes. (The name C stands for "calculate.") The order in which the document's fields are recalculated is defined by the CO entry in the interactive form dictionary (see Section 8.6.1, "Interactive Form Dictionary").

TABLE 8.33 Entries in the document catalog's additional-actions dictionary

DC	dictionary	<i>(Optional; PDF 1.4)</i> A JavaScript action to be performed before closing a document. (The name DC stands for "document close.")
WS	dictionary	<i>(Optional; PDF 1.4)</i> A JavaScript action to be performed before saving a document. (The name WS stands for "will save.")
DS	dictionary	<i>(Optional; PDF 1.4)</i> A JavaScript action to be performed after saving a document. (The name DS stands for "did save.")
WP	dictionary	<i>(Optional; PDF 1.4)</i> A JavaScript action to be performed before printing a document. (The name WP stands for "will print.")
DP	dictionary	<i>(Optional; PDF 1.4)</i> A JavaScript action to be performed after printing a document. (The name DP stands for "did print.")

For purposes of the trigger events **E** (enter), **X** (exit), **D** (down), and **U** (up), the term *mouse* denotes a generic pointing device with the following characteristics:

- A selection button that can be *pressed*, *held down*, and *released*. If there is more than one mouse button, this is typically the left button.
- A notion of *location*—that is, an indication of where on the screen the device is pointing. This is typically denoted by a screen cursor.
- A notion of *focus*—that is, which element in the document is currently interacting with the user. In many systems, this is denoted by a blinking caret, a focus rectangle, or a color change.

PDF viewer applications must ensure the presence of such a device in order for the corresponding actions to be executed correctly. Mouse-related trigger events are subject to the following constraints:

- An **E** (enter) event can occur only when the mouse button is up.
- An **X** (exit) event cannot occur without a preceding **E** event.
- A **U** (up) event cannot occur without a preceding **E** and **D** event.
- In the case of overlapping or nested annotations, entering a second annotation's active area causes an **X** event to occur for the first annotation.

Note: The field-related trigger events **K** (keystroke), **F** (format), **V** (validate), and **C** (calculate) are not defined for button fields (see “Button Fields” on page 538). Note that the effects of an action triggered by one of these events are limited only by the action itself and can occur outside the described scope of the event. For example, even though the **F** event is used to trigger actions that format field values prior to display, it is possible for an action triggered by this event to perform a calculation or make any other modification to the document.

Note also that these field-related trigger events can occur either through user interaction or programmatically, such as in response to the **NeedAppearances** entry in the interactive form dictionary (see Section 8.6.1, “Interactive Form Dictionary”), importation of FDF data (Section 8.6.6, “Forms Data Format”), or JavaScript actions (“JavaScript Actions” on page 556). For example, the user's modifying a field value can trigger a cascade of calculations and further formatting and validation for other fields in the document.

8.5.3 Action Types

PDF supports the standard action types listed in Table 8.34. The following sections describe each of these types in detail. Plug-in extensions may add new action types.

TABLE 8.34 Action types

ACTION TYPE	DESCRIPTION
GoTo	Go to a destination in the current document.
GoToR	(“Go-to remote”) Go to a destination in another document.
Launch	Launch an application, usually to open a file.
Thread	Begin reading an article thread.
URI	Resolve a uniform resource identifier.
Sound	(PDF 1.2) Play a sound.
Movie	(PDF 1.2) Play a movie.
Hide	(PDF 1.2) Set an annotation’s Hidden flag.
Named	(PDF 1.2) Execute an action predefined by the viewer application.
SubmitForm	(PDF 1.2) Send data to a uniform resource locator.
ResetForm	(PDF 1.2) Set fields to their default values.
ImportData	(PDF 1.2) Import field values from a file.
JavaScript	(PDF 1.3) Execute a JavaScript script.

***Note:** Previous versions of the PDF specification described an action type known as the set-state action; this type of action is now considered obsolete and its use is no longer recommended. An additional action type, the no-op action, was defined in PDF 1.2 but never implemented; it is no longer defined and should be ignored.*

Go-To Actions

A *go-to action* changes the view to a specified destination (page, location, and magnification factor). Table 8.35 shows the action dictionary entries specific to this type of action.

TABLE 8.35 Additional entries specific to a go-to action

KEY	TYPE	VALUE
S	name	<i>(Required)</i> The type of action that this dictionary describes; must be GoTo for a go-to action.
D	name, string, or array	<i>(Required)</i> The destination to jump to (see Section 8.2.1, “Destinations”).

Specifying a go-to action in the **A** entry of a link annotation or outline item (see Tables 8.16 on page 501 and 8.4 on page 478) has the same effect as specifying the destination directly via the **Dest** entry. For example, the link annotation shown in Example 8.8, which uses a go-to action, has the same effect as the one in Example 8.7 on page 502, which specifies the destination directly. However, the go-to action is less compact and is not compatible with PDF 1.0, so using a direct destination is preferable.

Example 8.8

```

93 0 obj
  << /Type /Annot
    /Subtype /Link
    /Rect [71 717 190 734]
    /Border [16 16 1]
    /A << /Type /Action
      /S /GoTo
      /D [3 0 R /FitR -4 399 199 533]
    >>
  >>
endobj

```

Remote Go-To Actions

A *remote go-to action* is similar to an ordinary go-to action, but jumps to a destination in another PDF file instead of the current file. Table 8.36 shows the action dictionary entries specific to this type of action.

TABLE 8.36 Additional entries specific to a remote go-to action

KEY	TYPE	VALUE
S	name	<i>(Required)</i> The type of action that this dictionary describes; must be GoToR for a remote go-to action.
F	file specification	<i>(Required)</i> The file in which the destination is located.
D	name, string, or array	<i>(Required)</i> The destination to jump to (see Section 8.2.1, “Destinations”). If the value is an array defining an explicit destination (as described under “Explicit Destinations” on page 474), its first element must be a page number within the remote document rather than an indirect reference to a page object in the current document. The first page is numbered 0.
NewWindow	boolean	<i>(Optional; PDF 1.2)</i> A flag specifying whether to open the destination document in a new window. If this flag is false , the destination document will replace the current document in the same window. If this entry is absent, the viewer application should behave in accordance with the current user preference.

Launch Actions

A *launch action* launches an application or opens or prints a document. Table 8.37 shows the action dictionary entries specific to this type of action.

The optional **Win**, **Mac**, and **Unix** entries allow the action dictionary to include platform-specific parameters for launching the designated application. If no such entry is present for the given platform, the **F** entry is used instead. Table 8.38 shows the platform-specific launch parameters for the Windows platform; those for the Mac OS and UNIX platforms are not yet defined at the time of publication.

TABLE 8.37 Additional entries specific to a launch action

KEY	TYPE	VALUE
S	name	<i>(Required)</i> The type of action that this dictionary describes; must be Launch for a launch action.
F	file specification	<i>(Required if none of the entries Win, Mac, or Unix is present)</i> The application to be launched or the document to be opened or printed. If this entry is absent and the viewer application does not understand any of the alternative entries, it should do nothing.
Win	dictionary	<i>(Optional)</i> A dictionary containing Windows-specific launch parameters (see Table 8.38; see also implementation note 73 in Appendix H).
Mac	(undefined)	<i>(Optional)</i> Mac OS-specific launch parameters; not yet defined.
Unix	(undefined)	<i>(Optional)</i> UNIX-specific launch parameters; not yet defined.
NewWindow	boolean	<i>(Optional; PDF 1.2)</i> A flag specifying whether to open the destination document in a new window. If this flag is false , the destination document will replace the current document in the same window. If this entry is absent, the viewer application should behave in accordance with the current user preference. This entry is ignored if the file designated by the F entry is not a PDF document.

TABLE 8.38 Entries in a Windows launch parameter dictionary

KEY	TYPE	VALUE				
F	string	<i>(Required)</i> The file name of the application to be launched or the document to be opened or printed, in standard Windows pathname format. If the name string includes a backslash character (\), the backslash must itself be preceded by a backslash. <i>Note: This value must be a simple string; it is not a file specification.</i>				
D	string	<i>(Optional)</i> A string specifying the default directory in standard DOS syntax.				
O	string	<i>(Optional)</i> A string specifying the operation to perform: <table border="0"> <tr> <td style="padding-right: 10px;">open</td> <td>Open a document.</td> </tr> <tr> <td>print</td> <td>Print a document.</td> </tr> </table> <p>If the F entry designates an application instead of a document, this entry is ignored and the application is launched. Default value: open.</p>	open	Open a document.	print	Print a document.
open	Open a document.					
print	Print a document.					
P	string	<i>(Optional)</i> A parameter string to be passed to the application designated by the F entry. This entry should be omitted if F designates a document.				

Thread Actions

A *thread action* jumps to a specified bead on an article thread (see Section 8.3.2, “Articles”), in either the current document or a different one. Table 8.39 shows the action dictionary entries specific to this type of action.

TABLE 8.39 Additional entries specific to a thread action

KEY	TYPE	VALUE
S	name	<i>(Required)</i> The type of action that this dictionary describes; must be Thread for a thread action.
F	file specification	<i>(Optional)</i> The file containing the desired thread. If this entry is absent, the thread is in the current file.
D	dictionary, integer, or text string	<p><i>(Required)</i> The desired destination thread, specified in one of the following forms:</p> <ul style="list-style-type: none"> • An indirect reference to a thread dictionary (see Section 8.3.2, “Articles”). In this case, the thread must be in the current file. • The index of the thread within the Threads array of its document’s catalog (see Section 3.6.1, “Document Catalog”). The first thread in the array has index 0. • The title of the thread, as specified in its thread information dictionary (see Table 8.7 on page 484). If two or more threads have the same title, the one appearing first in the document catalog’s Threads array will be used.
B	dictionary or integer	<p><i>(Optional)</i> The desired bead in the destination thread, specified in one of the following forms:</p> <ul style="list-style-type: none"> • An indirect reference to a bead dictionary (see Section 8.3.2, “Articles”). In this case, the thread must be in the current file. • The index of the bead within its thread. The first bead in a thread has index 0.

URI Actions

A *uniform resource identifier* (URI) is a string that identifies (*resolves to*) a resource on the Internet—typically a file that is the destination of a hypertext link, although it can also resolve to a query or other entity. A *URI action* causes a URI to be resolved. Table 8.40 shows the action dictionary entries specific to this type of action. (See implementation notes 74 and 75 in Appendix H.)

TABLE 8.40 Additional entries specific to a URI action

KEY	TYPE	VALUE
S	name	(<i>Required</i>) The type of action that this dictionary describes; must be URI for a URI action.
URI	string	(<i>Required</i>) The uniform resource identifier to resolve, encoded in 7-bit ASCII.
IsMap	boolean	(<i>Optional</i>) A flag specifying whether to track the mouse position when the URI is resolved (see below). Default value: false . This entry applies only to actions triggered by the user's clicking an annotation; it is ignored for actions associated with outline items or with a document's OpenAction entry.

If the **IsMap** flag is **true** and the user has triggered the URI action by clicking an annotation with the mouse, the coordinates of the mouse position at the time the action is performed should be transformed from device space to user space and then offset relative to the upper-left corner of the annotation rectangle (that is, the value of the **Rect** entry in the annotation with which the URI action is associated). For example, if the mouse coordinates in user space are (x_m, y_m) and the annotation rectangle extends from (ll_x, ll_y) at the lower-left to (ur_x, ur_y) at the upper-right, the final coordinates (x_f, y_f) are as follows:

$$x_f = x_m - ll_x$$

$$y_f = ur_y - y_m$$

If the resulting coordinates (x_f, y_f) are fractional, they should be rounded to the nearest integer values. They are then appended to the URI to be resolved, separated by commas and preceded by a question mark. For example:

`http://www.adobe.com/intro?100,200`

To support URI actions, a PDF document's catalog (see Section 3.6.1, "Document Catalog") may include a **URI** entry whose value is a *URI dictionary*. At the time of publication, only one entry is defined for such a dictionary (see Table 8.41).

TABLE 8.41 Entry in a URI dictionary

KEY	TYPE	VALUE
Base	string	<i>(Optional)</i> The <i>base URI</i> to be used in resolving relative URI references. URI actions within the document may specify URIs in partial form, to be interpreted relative to this base address. If no base URI is specified, such partial URIs will be interpreted relative to the location of the document itself. The use of this entry is parallel to that of the body element <BASE>, as described in section 2.7.2 of Internet RFC 1866, <i>Hyper-text Markup Language 2.0 Proposed Standard</i> (see the Bibliography).

The **Base** entry allows the URI of the document itself to be recorded in situations in which the document may be accessed out of context. For example, if a document has been moved to a new location but contains relative links to other documents that have not, the **Base** entry could be used to refer such links to the true location of the other documents, rather than that of the moved document.

Sound Actions

A *sound action* (PDF 1.2) plays a sound through the computer's speakers. Table 8.42 shows the action dictionary entries specific to this type of action; sounds themselves are discussed in Section 8.7, "Sounds."

TABLE 8.42 Additional entries specific to a sound action

KEY	TYPE	VALUE
S	name	<i>(Required)</i> The type of action that this dictionary describes; must be Sound for a sound action.
Sound	stream	<i>(Required)</i> A sound object defining the sound to be played (see Section 8.7, “Sounds”; see also implementation note 76 in Appendix H).
Volume	number	<i>(Optional)</i> The volume at which to play the sound, in the range -1.0 to 1.0 . Higher values denote greater volume; negative values mute the sound. Default value: 1.0 .
Synchronous	boolean	<i>(Optional)</i> A flag specifying whether to play the sound synchronously or asynchronously. If this flag is true , the viewer application will retain control, allowing no further user interaction other than canceling the sound, until the sound has been completely played. Default value: false .
Repeat	boolean	<i>(Optional)</i> A flag specifying whether to repeat the sound indefinitely. If this entry is present, the Synchronous entry is ignored. Default value: false .
Mix	boolean	<i>(Optional)</i> A flag specifying whether to mix this sound with any other sound already playing. If this flag is false , any previously playing sound will be stopped before starting this sound; this can be used to stop a repeating sound (see Repeat , above). Default value: false .

Movie Actions

A *movie action* (PDF 1.2) can be used to play a movie in a floating window or within the annotation rectangle of a movie annotation (see “Movie Annotations” on page 510 and Section 8.8, “Movies”). The movie annotation must be associated with the page that is the destination of the link annotation or outline item containing the movie action, or with the page object with which the action is associated. (See implementation note 77 in Appendix H.)

The contents of a movie action dictionary are identical to those of a movie activation dictionary (see Table 8.80 on page 571), with the additional entries shown in Table 8.43. The contents of the activation dictionary associated with the movie annotation are ignored; the information specified in the movie action dictionary is used instead.

TABLE 8.43 Additional entries specific to a movie action

KEY	TYPE	VALUE								
S	name	<i>(Required)</i> The type of action that this dictionary describes; must be Movie for a movie action.								
Annot	dictionary	<i>(Optional)</i> An indirect reference to a movie annotation identifying the movie to be played.								
T	text string	<i>(Optional)</i> The title of a movie annotation identifying the movie to be played. <i>Note:</i> The dictionary must include either an Annot or a T entry, but not both.								
Operation	name	<i>(Optional)</i> The operation to be performed on the movie: <table border="0" style="margin-left: 2em;"> <tr> <td>Play</td> <td>Start playing the movie, using the play mode specified by the dictionary's Mode entry (see Table 8.79 on page 571). If the movie is currently paused, it is repositioned to the beginning before playing (or to the starting point specified by the dictionary's Start entry, if present).</td> </tr> <tr> <td>Stop</td> <td>Stop playing the movie.</td> </tr> <tr> <td>Pause</td> <td>Pause a playing movie.</td> </tr> <tr> <td>Resume</td> <td>Resume a paused movie.</td> </tr> </table> <p>Default value: Play.</p>	Play	Start playing the movie, using the play mode specified by the dictionary's Mode entry (see Table 8.79 on page 571). If the movie is currently paused, it is repositioned to the beginning before playing (or to the starting point specified by the dictionary's Start entry, if present).	Stop	Stop playing the movie.	Pause	Pause a playing movie.	Resume	Resume a paused movie.
Play	Start playing the movie, using the play mode specified by the dictionary's Mode entry (see Table 8.79 on page 571). If the movie is currently paused, it is repositioned to the beginning before playing (or to the starting point specified by the dictionary's Start entry, if present).									
Stop	Stop playing the movie.									
Pause	Pause a playing movie.									
Resume	Resume a paused movie.									

Hide Actions

A *hide action* (PDF 1.2) hides or shows one or more annotations on the screen by setting or clearing their Hidden flags (see Section 8.4.2, “Annotation Flags”). This type of action can be used in combination with appearance streams and trigger events (Sections 8.4.4, “Appearance Streams,” and 8.5.2, “Trigger Events”) to display pop-up help information on the screen. For example, the **E** (enter) and **X** (exit) trigger events in an annotation's additional-actions dictionary can be used to show and hide the annotation when the user rolls the cursor in and out of its active area on the page; this can be used to pop up a help label, or “tool tip,” describing the effect of clicking the mouse at that location on the page. Table 8.44 shows the action dictionary entries specific to this type of action. (See implementation notes 78 and 79 in Appendix H.)

TABLE 8.44 Additional entries specific to a hide action

KEY	TYPE	VALUE
S	name	<i>(Required)</i> The type of action that this dictionary describes; must be Hide for a hide action.
T	dictionary, string, or array	<i>(Required)</i> The annotation or annotations to be hidden or shown, specified in any of the following forms: <ul style="list-style-type: none"> • An indirect reference to an annotation dictionary • A string giving the fully qualified field name of an interactive form field whose associated widget annotation or annotations are to be affected (see “Field Names” on page 532) • An array of such dictionaries or strings
H	boolean	<i>(Optional)</i> A flag indicating whether to hide the annotation (true) or show it (false). Default value: true .

Named Actions

Table 8.45 lists several *named actions* (PDF 1.2) that PDF viewer applications are expected to support; further names may be added in the future. (See implementation notes 80 and 81 in Appendix H.)

TABLE 8.45 Named actions

NAME	ACTION
NextPage	Go to the next page of the document.
PrevPage	Go to the previous page of the document.
FirstPage	Go to the first page of the document.
LastPage	Go to the last page of the document.

Note: Viewer applications may support additional, nonstandard named actions, but any document using them will not be portable. If the viewer encounters a named action that is inappropriate for a viewing platform, or if the viewer does not recognize the name, it should take no action.

Table 8.46 shows the action dictionary entries specific to named actions.

TABLE 8.46 Additional entries specific to named actions

KEY	TYPE	VALUE
S	name	(Required) The type of action that this dictionary describes; must be Named for a named action.
N	name	(Required) The name of the action to be performed (see Table 8.45).

Form Actions

Four additional action types (submit-form, reset-form, import-data, and JavaScript) are defined for use with interactive forms. See Section 8.6.4, “Form Actions,” for details.

8.6 Interactive Forms

An *interactive form* (PDF 1.2)—sometimes referred to as an *AcroForm*—is a collection of *fields* for gathering information interactively from the user. A PDF document may contain any number of fields appearing on any combination of pages, all of which make up a single, global interactive form spanning the entire document. Arbitrary subsets of these fields can be imported or exported from the document; see Section 8.6.4, “Form Actions.”

Note: *Interactive forms should not be confused with form XObjects (see Section 4.9, “Form XObjects”). Despite the similarity of names, the two are different, unrelated types of object.*

Each field in a document’s interactive form is defined by a *field dictionary* (see Section 8.6.2, “Field Dictionaries”). For purposes of definition and naming, the fields can be organized hierarchically and can inherit attributes from their ancestors in the field hierarchy. A field’s children in the hierarchy may also include widget annotations (see “Widget Annotations” on page 511) that define its appearance on the page; such a field is called a *terminal field*.

As a convenience, when a field has only a single associated widget annotation, the contents of the field dictionary and the annotation dictionary (Section 8.4.1, “Annotation Dictionaries”) may be merged into a single dictionary containing

entries that pertain to both a field and an annotation. (This presents no ambiguity, since the contents of the two kinds of dictionary do not conflict.) If such an object defines an appearance stream, the appearance must be consistent with the object's current value as a field.

Note: *Fields containing text whose contents are not known in advance may need to construct their appearance streams dynamically instead of defining them statically in an appearance dictionary; see “Variable Text” on page 533.*

8.6.1 Interactive Form Dictionary

The contents and properties of a document's interactive form are defined by an *interactive form dictionary* that is referenced from the **AcroForm** entry in the document catalog (see Section 3.6.1, “Document Catalog”). Table 8.47 shows the contents of this dictionary.

TABLE 8.47 Entries in the interactive form dictionary

KEY	TYPE	VALUE
Fields	array	<i>(Required)</i> An array of references to the document's <i>root fields</i> (those with no ancestors in the field hierarchy).
NeedAppearances	boolean	<i>(Optional)</i> A flag specifying whether to construct appearance streams and appearance dictionaries for all widget annotations in the document (see “Variable Text” on page 533). Default value: false .
SigFlags	integer	<i>(Optional; PDF 1.3)</i> A set of flags specifying various document-level characteristics related to signature fields (see Table 8.48, below, and “Signature Fields” on page 547). Default value: 0.
CO	array	<i>(Required if any fields in the document have additional-actions dictionaries containing a C entry; PDF 1.3)</i> An array of indirect references to field dictionaries with calculation actions, defining the <i>calculation order</i> in which their values will be recalculated when the value of any field changes (see Section 8.5.2, “Trigger Events”).
DR	dictionary	<i>(Optional)</i> A document-wide default value for the DR attribute of variable text fields (see “Variable Text” on page 533).
DA	string	<i>(Optional)</i> A document-wide default value for the DA attribute of variable text fields (see “Variable Text” on page 533).
Q	integer	<i>(Optional)</i> A document-wide default value for the Q attribute of variable text fields (see “Variable Text” on page 533).

The value of the interactive form dictionary’s **SigFlags** entry is an unsigned 32-bit integer containing flags specifying various document-level characteristics related to signature fields (see “Signature Fields” on page 547). Bit positions within the flag word are numbered from 1 (low-order) to 32 (high-order). Table 8.48 shows the meanings of the flags; all undefined flag bits are reserved and must be set to 0.

TABLE 8.48 Signature flags

BIT POSITION	NAME	MEANING
1	SignaturesExist	If set, the document contains at least one signature field. This flag allows a viewer application to enable user interface items (such as menu items or pushbuttons) related to signature processing without having to scan the entire document for the presence of signature fields.
2	AppendOnly	If set, the document contains signatures that may be invalidated if the file is saved (written) in a way that alters its previous contents, such as with the “optimize” option. Merely updating the file by appending new information to the end of the previous version is safe (see Section G.6, “Updating Example”). Viewer applications can use this flag to present a user requesting an optimized save with an additional alert box warning that signatures will be invalidated and requiring explicit confirmation before continuing with the operation.

8.6.2 Field Dictionaries

Each field in a document’s interactive form is defined by a *field dictionary*, which must be an indirect object. The field dictionaries may be organized hierarchically into one or more tree structures. Many field attributes are *inheritable*, meaning that if they are not explicitly specified for a given field, their values are taken from those of its parent in the field hierarchy. Such inheritable attributes are designated as such in the tables below; the designation (*Required; inheritable*) means that an attribute must be defined for every field, whether explicitly in its own field dictionary or by inheritance from an ancestor in the hierarchy. Table 8.49 shows those entries that are common to all field dictionaries, regardless of type; those that pertain only to a particular type of field are described in the relevant sections below.

TABLE 8.49 Entries common to all field dictionaries

KEY	TYPE	VALUE
FT	name	<p><i>(Required for terminal fields; inheritable)</i> The type of field that this dictionary describes:</p> <p>Btn Button (see “Button Fields” on page 538) Tx Text (see “Text Fields” on page 543) Ch Choice (see “Choice Fields” on page 545) Sig (PDF 1.3) Signature (see “Signature Fields” on page 547)</p> <p><i>Note: This entry may be present in a nonterminal field (one whose descendants are themselves fields) in order to provide an inheritable FT value. However, a nonterminal field does not logically have a type of its own; it is merely a container for inheritable attributes that are intended for descendant terminal fields of any type.</i></p>
Parent	dictionary	<p><i>(Required if this field is the child of another in the field hierarchy; absent otherwise)</i> The field that is the immediate parent of this one (the field, if any, whose Kids array includes this field). A field can have at most one parent; that is, it can be included in the Kids array of at most one other field.</p>
Kids	array	<p><i>(Optional)</i> An array of indirect references to the immediate children of this field.</p>
T	text string	<p><i>(Optional)</i> The partial field name (see “Field Names,” below; see also implementation notes 82 and 83 in Appendix H).</p>
TU	text string	<p><i>(Optional; PDF 1.3)</i> An alternate field name, to be used in place of the actual field name wherever the field must be identified in the user interface (such as in error or status messages referring to the field). This text is also useful when extracting the document’s contents in support of accessibility to disabled users or for other purposes (see Section 9.8.2, “Alternate Descriptions”).</p>
TM	text string	<p><i>(Optional; PDF 1.3)</i> The <i>mapping name</i> to be used when exporting interactive form field data from the document.</p>
Ff	integer	<p><i>(Optional; inheritable)</i> A set of flags specifying various characteristics of the field (see Table 8.50). Default value: 0.</p>
V	(various)	<p><i>(Optional; inheritable)</i> The field’s value, whose format varies depending on the field type; see the descriptions of individual field types for further information.</p>

DV	(various)	(<i>Optional; inheritable</i>) The default value to which the field reverts when a reset-form action is executed (see “Reset-Form Actions” on page 554). The format of this value is the same as that of V .
AA	dictionary	(<i>Optional; PDF 1.2</i>) An additional-actions dictionary defining the field’s behavior in response to various trigger events (see Section 8.5.2, “Trigger Events”). This entry has exactly the same meaning as the AA entry in an annotation dictionary (see Section 8.4.1, “Annotation Dictionaries”).

The value of the field dictionary’s **Ff** entry is an unsigned 32-bit integer containing flags specifying various characteristics of the field. Bit positions within the flag word are numbered from 1 (low-order) to 32 (high-order). The flags shown in Table 8.50 are common to all types of field; flags that apply only to specific field types are discussed in the sections describing those types. All undefined flag bits are reserved and must be set to 0.

TABLE 8.50 Field flags common to all field types

BIT POSITION	NAME	MEANING
1	ReadOnly	If set, the user may not change the value of the field. Any associated widget annotations will not interact with the user; that is, they will not respond to mouse clicks or change their appearance in response to mouse motions. This flag is useful for fields whose values are computed or imported from a database.
2	Required	If set, the field must have a value at the time it is exported by a submit-form action (see “Submit-Form Actions” on page 550).
3	NoExport	If set, the field must not be exported by a submit-form action (see “Submit-Form Actions” on page 550).

Field Names

The **T** entry in the field dictionary (see Table 8.49 on page 531) holds a text string defining the field’s *partial field name*. The *fully qualified field name* is not explicitly defined, but is constructed from the partial field names of the field and all of its ancestors. For a field with no parent, the partial and fully qualified names are the same; for a field that is the child of another field, the fully qualified name is

formed by appending the child field's partial name to the parent's fully qualified name, separated by a period (.):

parent's_full_name.child's_partial_name

For example, if a field with the partial field name `PersonalData` has a child whose partial name is `Address`, which in turn has a child with the partial name `ZipCode`, then the fully qualified name of this last field would be

`PersonalData.Address.ZipCode`

Thus all fields descended from a common ancestor will share the ancestor's fully qualified field name as a common prefix in their own fully qualified names.

It is possible for different field dictionaries to have the same fully qualified field name if they are descendants of a common ancestor with that name and have no partial field names (**T** entries) of their own. Such field dictionaries are different representations of the same underlying field; they should differ only in properties that specify their visual appearance. In particular, field dictionaries with the same fully qualified field name must have the same field type (**FT**), value (**V**), and default value (**DV**).

Variable Text

When the contents and properties of a field are known in advance, its visual appearance can be specified by an appearance stream defined in the PDF file itself (see Section 8.4.4, “Appearance Streams,” and “Widget Annotations” on page 511). In some cases, however, the field may contain text whose value is not known until viewing time. Examples include text fields to be filled in with text typed by the user from the keyboard and scrollable list boxes whose contents are determined interactively at the time the document is displayed.

In such cases, the PDF document cannot provide a statically defined appearance stream for displaying the field; rather, the viewer application must construct an appearance stream dynamically at viewing time. The dictionary entries shown in Table 8.51 provide general information about the field's appearance that can be combined with the specific text it contains to construct an appearance stream.

TABLE 8.51 Additional entries common to all fields containing variable text

KEY	TYPE	VALUE
DR	dictionary	<i>(Required; inheritable)</i> A resource dictionary (see Section 3.7.2, “Resource Dictionaries”) containing default resources (such as fonts, patterns, or color spaces) to be used by the appearance stream. At a minimum, this dictionary must contain a Font entry specifying the resource name and font dictionary of the default font for displaying the field’s text. (See implementation note 84 in Appendix H.)
DA	string	<i>(Required; inheritable)</i> The <i>default appearance string</i> , containing a sequence of valid page-content graphics or text state operators defining such properties as the field’s text size and color.
Q	integer	<i>(Optional; inheritable)</i> A code specifying the form of <i>quadding</i> (justification) to be used in displaying the text: <ul style="list-style-type: none"> 0 Left-justified 1 Centered 2 Right-justified Default value: 0 (left-justified).

The new appearance stream becomes the normal appearance (**N**) in the appearance dictionary associated with the field’s widget annotation (see Table 8.13 on page 497). (If the widget annotation has no appearance dictionary, the viewer application must create one and store it in the annotation dictionary’s **AP** entry.) The appearance stream—which, like all appearance streams, is a form XObject—has the contents of its form dictionary initialized as follows:

- The resource dictionary (**Resources**) is set to the value of the field dictionary’s **DR** entry.
- The lower-left corner of the bounding box (**BBox**) is set to coordinates (0, 0) in the form coordinate system. The box’s top and right coordinates are taken from the dimensions of the annotation rectangle (the **Rect** entry in the widget annotation dictionary).
- All other entries in the appearance stream’s form dictionary are set to their default values (see Section 4.9, “Form XObjects”).

The contents of the appearance stream are as follows:

/Tx	BMC	% Begin marked content with tag Tx
	q	% Save graphics state
		... <i>Any required graphics state changes, such as clipping...</i>
	BT	% Begin text object
		... <i>Default appearance string (DA)...</i>
		... <i>Text-positioning and text-showing operators to show the variable text...</i>
	ET	% End text object
	Q	% Restore graphics state
	EMC	% End marked content

The **BMC** (begin marked content) and **EMC** (end marked content) operators are discussed in Section 9.5, “Marked Content”; **q** (save graphics state) and **Q** (restore graphics state) in Section 4.3.3, “Graphics State Operators”; and **BT** (begin text object) and **ET** (end text object) in Section 5.3, “Text Objects.”

The default appearance string (**DA**) contains any graphics state or text state operators needed to establish the graphics state parameters, such as text size and color, for displaying the field’s variable text. Only operators that are allowed within text objects may occur in this string (see Figure 4.1 on page 135). At a minimum, the string must include a **Tf** (text font) operator along with its two operands, *font* and *size*. The specified *font* value must match a resource name in the **Font** entry of the default resource dictionary (**DR**). A zero value for *size* means that the font is to be *auto-sized*: its size is computed as a function of the height of the annotation rectangle.

The default appearance string should contain at most one **Tm** (text matrix) operator. If this operator is present, the viewer application should replace the horizontal and vertical translation components with positioning values it determines to be appropriate, based on the field value, the quadding (**Q**) attribute, and any layout rules it employs. If the default appearance string contains no **Tm** operator, the viewer should insert one in the appearance stream, with appropriate horizontal and vertical translation components, after the default appearance string and before the text-positioning and text-showing operators for the variable text.

To update an existing appearance stream to reflect a new field value, the viewer application should first copy any needed resources from the field’s **DR** dictionary into the stream’s **Resources** dictionary. (If the **DR** and **Resources** dictionaries contain resources with the same name, the one already in the **Resources** dictionary should be left intact, *not* replaced with the corresponding value from the **DR** dic-

tionary.) The viewer application should then replace the existing contents of the appearance stream from

```
/Tx BMC
```

to the matching

```
EMC
```

with the corresponding new contents as shown above. (If the existing appearance stream contains no marked content with tag Tx, the new contents should be appended to the end of the original stream.)

The optional **MK** entry in the field’s widget annotation dictionary (see Table 8.28 on page 512) can be used to provide an *appearance characteristics dictionary* (PDF 1.2) containing additional information for constructing the annotation’s appearance stream. Table 8.52 shows the contents of this dictionary.

TABLE 8.52 Entries in an appearance characteristics dictionary

KEY	TYPE	VALUE
R	integer	(Optional) The number of degrees by which the widget annotation is rotated counterclockwise relative to the page. The value must be a multiple of 90. Default value: 0.
BC	array	(Optional) An array of numbers in the range 0.0 to 1.0 specifying the color of the widget annotation’s border. The number of array elements determines the color space in which the color is defined: <ul style="list-style-type: none"> 0 No color; transparent 1 DeviceGray 3 DeviceRGB 4 DeviceCMYK
BG	array	(Optional) An array of numbers in the range 0.0 to 1.0 specifying the color of the widget annotation’s background. The number of array elements determines the color space, as described above for BC .
CA	text string	(Optional; <i>button fields only</i>) The widget annotation’s <i>normal caption</i> , displayed when it is not interacting with the user.

Note: Unlike the remaining entries listed below, which apply only to widget annotations associated with pushbutton fields (see “Pushbuttons” on page 539), the **CA** entry can be used with any type of button field, including checkboxes (“Checkboxes” on page 539) and radio buttons (“Radio Buttons” on page 540).

RC	text string	<i>(Optional; pushbutton fields only)</i> The widget annotation's <i>rollover caption</i> , displayed when the user rolls the cursor into its active area without pressing the mouse button.
AC	text string	<i>(Optional; pushbutton fields only)</i> The widget annotation's <i>alternate (down) caption</i> , displayed when the mouse button is pressed within its active area.
I	stream	<i>(Optional; pushbutton fields only; must be an indirect reference)</i> A form XObject defining the widget annotation's <i>normal icon</i> , displayed when it is not interacting with the user.
RI	stream	<i>(Optional; pushbutton fields only; must be an indirect reference)</i> A form XObject defining the widget annotation's <i>rollover icon</i> , displayed when the user rolls the cursor into its active area without pressing the mouse button.
IX	stream	<i>(Optional; pushbutton fields only; must be an indirect reference)</i> A form XObject defining the widget annotation's <i>alternate (down) icon</i> , displayed when the mouse button is pressed within its active area.
IF	dictionary	<i>(Optional; pushbutton fields only)</i> An icon fit dictionary (see Table 8.73 on page 566) specifying how to display the widget annotation's icon within its annotation rectangle. If present, the icon fit dictionary applies to all of the annotation's icons (normal, rollover, and alternate).
TP	integer	<p><i>(Optional; pushbutton fields only)</i> A code indicating where to position the text of the widget annotation's caption relative to its icon:</p> <ul style="list-style-type: none"> 0 No icon; caption only 1 No caption; icon only 2 Caption below the icon 3 Caption above the icon 4 Caption to the right of the icon 5 Caption to the left of the icon 6 Caption overlaid directly on the icon

Default value: 0.

8.6.3 Field Types

Interactive forms support the following field types:

- *Button fields* represent interactive controls on the screen that the user can manipulate with the mouse. They include *pushbuttons*, *checkboxes*, and *radio buttons*.

- *Text fields* are boxes or spaces in which the user can enter text from the keyboard.
- *Choice fields* contain several text items, at most one of which may be selected as the field value. They include scrollable *list boxes* and *combo boxes*.
- *Signature fields* represent electronic “signatures” for authenticating the identity of a user. They can include purely mathematical authentication methods such as public/private-key encrypted document digests, or biometric forms of identification such as handwritten signatures, fingerprints, or retinal scans.

The following sections describe each of these field types in detail. Further types may be added in the future.

Button Fields

A *button field* (field type **Btn**) represents an interactive control on the screen that the user can manipulate with the mouse. It may be any of the following:

- A *pushbutton* is a purely interactive control that responds immediately to user input without retaining a permanent value.
- A *checkbox* toggles between two states, on and off.
- *Radio buttons* are a set of related toggles, at most one of which may be on at any given time; selecting any one of the buttons automatically deselects all the others.

The various types of button field are distinguished by flags in the **Ff** entry, as shown in Table 8.53.

TABLE 8.53 Field flags specific to button fields

BIT POSITION	NAME	MEANING
17	Pushbutton	If set, the field is a pushbutton that does not retain a permanent value.
16	Radio	If set, the field is a set of radio buttons; if clear, the field is a checkbox. This flag is meaningful only if the Pushbutton flag is clear.
15	NoToggleToOff	<i>(Radio buttons only)</i> If set, exactly one radio button must be selected at all times; clicking the currently selected button has no effect. If clear, clicking the selected button deselects it, leaving no button selected.

Pushbuttons

The simplest type of field is a *pushbutton field*, which has a field type of **Btn** and the Pushbutton flag (see Table 8.53) set. Because this type of button retains no permanent value, it does not use the **V** and **DV** entries in the field dictionary (see Table 8.49 on page 531).

Checkboxes

A *checkbox field* toggles between two states, on and off, when manipulated by the user with the mouse or keyboard. Its field type is **Btn** and its Pushbutton and Radio flags (see Table 8.53) are both clear. Each state can have a separate appearance, defined by an appearance stream in the appearance dictionary of the field's widget annotation (see Section 8.4.4, "Appearance Streams"). The **V** entry in the field dictionary (see Table 8.49 on page 531) holds a name object representing the checkbox's appearance state, which is used to select the appropriate appearance from the appearance dictionary. The appearance for the off state is optional, but if present must be stored in the appearance dictionary under the name **Off**. The recommended name for the on state is **Yes**, but this is not required. Example 8.9 shows a typical checkbox definition.

Example 8.9

```

1 0 obj
  << /FT /Btn
    /T (Urgent)
    /V /Yes
    /AS /Yes
    /AP << /N << /Yes 2 0 R >>
      >>
  >>
endobj

2 0 obj
  << /Resources 2 0 R
    /Length 104
  >>

```

```

stream
  q
  0 0 1 rg
  BT
  /ZaDb 12 Tf
  0 0 Td
  (8) Tj
  ET
  Q
endstream
endobj

```

Because they are limited to the Latin character set, name objects cannot be used to represent the values of checkbox fields in non-Latin writing systems. To address this problem, PDF 1.4 provides an additional, optional entry, **Opt**, in the field dictionary for checkbox fields (see Table 8.54). When present, this entry holds a human-readable text string to be used in place of the **V** entry to denote the value of the field. This allows the use of Unicode encoding to specify field values using non-Latin characters.

TABLE 8.54 Additional entry specific to a checkbox field

KEY	TYPE	VALUE
Opt	text string	<i>(Optional; inheritable; PDF 1.4)</i> A text string to be used in place of the V entry for the value of the field.

Radio Buttons

A *radio button field* is a set of related toggle controls, at most one of which may be on at any given time; selecting any one of the buttons automatically deselects all the others. The field type is **Btn**, the Pushbutton flag (see Table 8.53 on page 538) is clear, and the Radio flag is set. This type of button field has an additional flag, **NoToggleToOff**, which specifies, if set, that exactly one of the radio buttons must be selected at all times. In this case, clicking the currently selected button has no effect; if the **NoToggleToOff** flag is clear, clicking the selected button deselects it, leaving no button selected.

The **Kids** entry in the radio button field's field dictionary (see Table 8.49 on page 531) holds an array of widget annotations representing the individual buttons in the set, each of which is a separate checkbox field in its own right. The parent field's **V** entry holds a name object corresponding to the appearance state of whichever child field is currently in the on state; the default value for this entry is Off. Example 8.10 shows the object definitions for a set of radio buttons.

Example 8.10

```

10 0 obj                                % Radio button field
  << /FT /Btn
    /Ff ...                               % ...Radio flag = 1, Pushbutton = 0...
    /T (Credit card)
    /V /MasterCard
    /Kids [ 11 0 R
            12 0 R
          ]
  >>
endobj

11 0 obj                                % First checkbox
  << /Parent 10 0 R
    /AS /MasterCard
    /AP << /N << /MasterCard 8 0 R
      /Off 9 0 R
    >>
  >>
endobj

12 0 obj                                % Second checkbox
  << /Parent 10 0 R
    /AS /Off
    /AP << /N << /Visa 8 0 R
      /Off 9 0 R
    >>
  >>
endobj

8 0 obj                                  % Appearance stream for "on" state
  << /Resources 20 0 R
    /Length 104
  >>

```

```

stream
  q
    0 0 1 rg
  BT
    /ZaDb 12 Tf
    0 0 Td
    (8) Tj
  ET
Q
endstream
endobj

9 0 obj                                % Appearance stream for "off" state
  << /Resources 20 0 R
    /Length 104
  >>
stream
  q
    0 0 1 rg
  BT
    /ZaDb 12 Tf
    0 0 Td
    (4) Tj
  ET
Q
endstream
endobj

```

Like a checkbox field, a radio button field can use the optional **Opt** entry in the field dictionary (*PDF 1.4*) to define export values for its constituent radio buttons using Unicode encoding for non-Latin characters (see Table 8.55). For a radio button field, **Opt** holds an array of text strings corresponding to the widget annotations representing the individual radio buttons in the field's **Kids** array. The name object in the field's **V** entry consists of a sequence of digits giving the numerical position of the currently selected radio button within the **Kids** array, with /0 denoting the first button in the array. This feature can also be used to define the same export value for multiple radio buttons, using the different computer-generated values in the **V** entry to distinguish among them.

TABLE 8.55 Additional entry specific to a radio button field

KEY	TYPE	VALUE
Opt	array	<i>(Optional; inheritable; PDF 1.4)</i> An array of text strings to be used in place of the V entries for the values of the widget annotations representing the individual radio buttons. Each element in the array represents the export value of the corresponding widget annotation in the Kids array of the radio button field.

Text Fields

A *text field* (field type **Tx**) is a box or space in which the user can enter text from the keyboard. The text may be restricted to a single line or may be allowed to span multiple lines, depending on the setting of the Multiline flag in the field dictionary's **Ff** entry. Table 8.56 shows the flags pertaining to this type of field.

TABLE 8.56 Field flags specific to text fields

BIT POSITION	NAME	MEANING
13	Multiline	If set, the field may contain multiple lines of text; if clear, the field's text is restricted to a single line.
14	Password	If set, the field is intended for entering a secure password that should not be echoed visibly to the screen. Characters typed from the keyboard should instead be echoed in some unreadable form, such as asterisks or bullet characters. To protect password confidentiality, viewer applications should never store the value of the text field in the PDF file if this flag is set.
21	FileSelect	<i>(PDF 1.4)</i> If set, the text entered in the field represents the pathname of a file whose contents are to be submitted as the value of the field.
23	DoNotSpellCheck	<i>(PDF 1.4)</i> If set, the text entered in the field will not be spell-checked.
24	DoNotScroll	<i>(PDF 1.4)</i> If set, the field will not scroll (horizontally for single-line fields, vertically for multiple-line fields) to accommodate more text than will fit within its annotation rectangle. Once the field is full, no further text will be accepted.

The field's text is held in a text string in the **V** (value) entry of the field dictionary. The contents of this text string are used to construct an appearance stream for displaying the field, as described under "Variable Text" on page 533; the text is presented in a single style (font, size, color, and so forth), as specified by the **DA** (default appearance) string.

If the FileSelect flag (*PDF 1.4*) is set, the field functions as a *file-select control*. In this case, the field's text represents the pathname of a file whose contents are to be submitted as the field's value:

- For fields submitted in HTML Form format, the submission uses the MIME content type `multipart/form-data`, as described in Internet RFC 2045, *Multipurpose Internet Mail Extensions (MIME), Part One: Format of Internet Message Bodies* (see the Bibliography).
- For Forms Data Format (FDF) submission, the value of the **V** entry in the FDF field dictionary (see "FDF Fields" on page 564) is a file specification (Section 3.10, "File Specifications") identifying the selected file.
- XML format is not supported for file-select controls, so no value is submitted in this case.

Besides the usual entries common to all types of field (see Table 8.49 on page 531), the field dictionary for a text field can contain the additional entry shown in Table 8.57.

TABLE 8.57 Additional entry specific to a text field

KEY	TYPE	VALUE
MaxLen	integer	<i>(Optional; inheritable)</i> The maximum length of the field's text, in characters.

Example 8.11 shows the object definitions for a typical text field.

Example 8.11

```

6 0 obj
  << /FT /Tx
    /Ff ...                               % Set Multiline flag
    /T (Silly prose)
    /DR 21 0 R
    /DA (0 0 1 rg /Ti 12 Tf)
    /V (The quick brown fox ate the lazy mouse)
    /AP << /N 5 0 R >>
  >>
endobj

5 0 obj
  << /Resources 21 0 R
    /Length 172
  >>
stream
  /Tx BMC
  BT
    0 0 1 rg
    /Ti 12 Tf
    1 0 0 1 100 100 Tm
    0 0 Td
    (The quick brown fox ) Tj
    0 -13 Td
    (ate the lazy mouse.) Tj
  ET
EMC
endstream
endobj

```

Choice Fields

A *choice field* (field type **Ch**) contains several text items, one or more of which may be selected as the field value. The items may be presented to the user in either of two forms:

- A scrollable *list box*
- A *combo box* consisting of a drop list optionally accompanied by an editable text box in which the user can type a value other than the predefined choices

TABLE 8.58 Field flags specific to choice fields

BIT POSITION	NAME	MEANING
18	Combo	If set, the field is a combo box; if clear, the field is a list box.
19	Edit	If set, the combo box includes an editable text box as well as a drop list; if clear, it includes only a drop list. This flag is meaningful only if the Combo flag is set.
20	Sort	If set, the field's option items should be sorted alphabetically. This flag is intended for use by form authoring tools, not by PDF viewer applications; viewers should simply display the options in the order in which they occur in the Opt array (see Table 8.59).
22	MultiSelect	(PDF 1.4) If set, more than one of the field's option items may be selected simultaneously; if clear, no more than one item at a time may be selected.
23	DoNotSpellCheck	(PDF 1.4) If set, the text entered in the field will not be spell-checked. This flag is meaningful only if the Combo and Edit flags are both set.

The various types of choice field are distinguished by flags in the **Ff** entry, as shown in Table 8.58. Table 8.59 shows the field dictionary entries specific to choice fields.

TABLE 8.59 Additional entries specific to a choice field

KEY	TYPE	VALUE
Opt	array	(Required; inheritable) An array of options to be presented to the user. Each element of the array is either a text string representing one of the available options or a two-element array consisting of a text string together with a default appearance string for constructing the item's appearance dynamically at viewing time (see "Variable Text" on page 533; see also implementation note 85 in Appendix H).
Tl	integer	(Optional; inheritable) For scrollable list boxes, the <i>top index</i> (the index in the Opt array of the first option visible in the list).
I	array	(Sometimes required, otherwise optional; inheritable; PDF 1.4) For choice fields that allow multiple selection (MultiSelect flag set), an array of integers, sorted in ascending order, representing the zero-based indices in the Opt array of the currently selected option items. This entry is required when two or more elements in the Opt array have different names but the same export value, or when the value of the choice field is an array; in other cases, it is permitted but not required. If the items identified by this entry differ from those in the V entry of the field dictionary (see below), the V entry takes precedence.

The **Opt** array specifies the list of options to be presented to the user. Each option is represented by a text string to be displayed on the screen as the name of the option. The corresponding element of the **Opt** array may contain either this text string by itself or a two-element array with the text string as its first element. In the latter case, the second element of the array is a default appearance string to be used in constructing a dynamic appearance stream, as described under “Variable Text” on page 533.

The field dictionary’s **V** (value) entry (see Table 8.49 on page 531) identifies the item or items currently selected in the choice field. If the field does not allow multiple selection—that is, if the `MultiSelect` flag (*PDF 1.4*) is not set—or if multiple selection is supported but only one item is currently selected, **V** is a text string representing the name of the selected item, as given in the field dictionary’s **Opt** array; if multiple items are selected, it is an array of such strings. (For items represented in the **Opt** array by a two-element array, the name string is the first of the two array elements.) The default value of **V** is **null**, indicating that no item is currently selected.

Example 8.12 shows a typical choice field definition.

Example 8.12

```
<< /FT /Ch
  /Ff ...
  /T (Body Color)
  /V (Blue)
  /Opt [ (Red)
        (My favorite color)
        (Blue)
      ]
>>
```

Signature Fields

A *signature field* (*PDF 1.3*) represents an electronic “signature” for authenticating the identity of a user. The signature may be purely mathematical, such as a public/private-key encrypted document digest, or it may be a biometric form of identification such as a handwritten signature, fingerprint, or retinal scan. The specific form of authentication used is implemented by a plug-in *signature handler*. Third-party handler writers are encouraged to register their handler names with Adobe; see Appendix E.

Note: The specification for public-key digital signature authentication is available in the Adobe document PDF Public-Key Digital Signature and Encryption Specification (see the Bibliography).

The field dictionary representing a signature field contains the standard entries described in Table 8.49 on page 531. The field type (**FT**) is **Sig**, the field value (**V**) is a *signature dictionary* specifying various attributes of the signature field, and the default value (**DV**) is a signature dictionary containing default values for reinitializing the contents of the **V** dictionary in response to a reset-form action (see “Reset-Form Actions” on page 554). Filling in (“signing”) the signature field entails updating at least the **V** entry, and usually also the **AP** entry of the associated widget annotation. Exporting a signature field typically exports the **T**, **V**, and **AP** entries.

Table 8.60 shows the contents of the signature dictionary. Signature handlers are free to use or omit those entries that are marked optional in the table, but are encouraged to use them in a standard way if they are used at all. In addition, specific signature handlers may add private entries of their own. To avoid name duplication, it is suggested that the keys for all such private entries be prefixed with the registered handler name followed by a period (.).

The **DV** (default value) entry in the field dictionary may specify a signature dictionary for an unsigned field that is preloaded with default values for some entries. Signature handlers are free to use or ignore these values in initializing the **V** dictionary for a signed field. The default dictionary may include default values for private entries belonging to multiple handlers; a given handler should use only those entries that are pertinent to itself and strip out the others.

Like any other field, a signature field may actually be described by a widget annotation dictionary containing entries pertaining to an annotation as well as a field (see “Widget Annotations” on page 511). The annotation rectangle (**Rect**) in such a dictionary gives the position of the field on its page. For signature fields (such as PPK signatures) that are not intended to be visible, the annotation rectangle may have zero height and width.

TABLE 8.60 Entries in a signature dictionary

KEY	TYPE	VALUE
Type	name	<i>(Optional)</i> The type of PDF object that this dictionary describes; if present, must be Sig for a signature dictionary.
Filter	name	<i>(Required; inheritable)</i> The name of the signature handler to be used for authenticating the field's contents, such as Adobe.PPKLite , Entrust.PPKEF , CICI.SignIt , or VeriSign.PPKVS .
SubFilter	name	<i>(Optional)</i> The name of a specific submethod of the specified handler.
ByteRange	array	<i>(Required)</i> An array of pairs of integers (starting byte offset, length in bytes) describing the exact byte range for the digest calculation. Multiple discontinuous byte ranges may be used to describe a digest that does not include the signature token itself.
Contents	string	<i>(Required)</i> The encrypted signature token.
Name	text string	<i>(Optional)</i> The name of the person or authority signing the document.
M	date	<i>(Optional)</i> The time of signing. Depending on the signature handler, this may be a normal unverified computer time or a time generated in a verifiable way from a secure time server.
Location	text string	<i>(Optional)</i> The CPU host name or physical location of the signing.
Reason	text string	<i>(Optional)</i> The reason for the signing, such as (I agree...).

The appearance dictionary (**AP**) of a signature field's widget annotation defines the field's visual appearance on the page (see Section 8.4.4, "Appearance Streams"). For a signature field that has not been filled in, the normal appearance is usually blank. If no plug-in signature handler is available for a given signing method, the normal appearance for a signature that has been filled in should simply show an appropriate representation of the signature: text, ink strokes representing a handwritten signature, a bitmap of a fingerprint, or any other representation appropriate for the particular signing method. If an appropriate signature handler is available, three appearances are possible:

- The *unvalidated* appearance represents a signature that has not yet been validated by the signature handler.
- The *valid* appearance represents a signature that has been validated and accepted by the signature handler.

- The *invalid* appearance represents a signature that has been validated and rejected by the signature handler.

When the signature handler is invoked to validate the signature, it should alter the signature's appearance from unvalidated to either valid or invalid. It is suggested that the unvalidated appearance contain an overprinted yellow question mark, the invalid appearance an overprinted red X, and the valid appearance the logo of the signature handler underprinted as a watermark.

8.6.4 Form Actions

Interactive forms support four special types of action in addition to those described in Section 8.5.3, "Action Types":

- *Submit-form actions* transmit the names and values of selected interactive form fields to a specified uniform resource locator (URL), presumably the address of a World Wide Web server that will process them and send back a response.
- *Reset-form actions* reset selected interactive form fields to their default values.
- *Import-data actions* import Forms Data Format (FDF) data into the document's interactive form from a specified file.
- *JavaScript actions* (PDF 1.3) cause a script to be compiled and executed by the JavaScript interpreter.

Submit-Form Actions

A *submit-form action* transmits the names and values of selected interactive form fields to a specified uniform resource locator (URL), presumably the address of a World Wide Web server that will process them and send back a response. Table 8.61 shows the action dictionary entries specific to this type of action.

The value of the action dictionary's **Flags** entry is an unsigned 32-bit integer containing flags specifying various characteristics of the action. Bit positions within the flag word are numbered from 1 (low-order) to 32 (high-order). Table 8.62 shows the meanings of the flags; all undefined flag bits are reserved and must be set to 0.

TABLE 8.61 Additional entries specific to a submit-form action

KEY	TYPE	VALUE
S	name	<i>(Required)</i> The type of action that this dictionary describes; must be SubmitForm for a submit-form action.
F	file specification	<i>(Required)</i> A URL file specification (see Section 3.10.4, “URL Specifications”) giving the uniform resource locator (URL) of the script at the Web server that will process the submission.
Fields	array	<p><i>(Optional)</i> An array identifying which fields to include in the submission or which to exclude, depending on the setting of the Include/Exclude flag in the Flags entry (see Table 8.62). Each element of the array is either an indirect reference to a field dictionary or (<i>PDF 1.3</i>) a string representing the fully qualified name of a field. Elements of both kinds may be mixed in the same array.</p> <p>If this entry is omitted, the Include/Exclude flag is ignored; all fields in the document’s interactive form are submitted except those whose NoExport flag (see Table 8.50 on page 532) is set. (Fields with no values may also be excluded, depending on the setting of the IncludeNoValueFields flag; see Table 8.62.) See the text following Table 8.62 for further discussion.</p>
Flags	integer	<i>(Optional; inheritable)</i> A set of flags specifying various characteristics of the action (see Table 8.62). Default value: 0.

TABLE 8.62 Flags for submit-form actions

BIT POSITION	NAME	MEANING
1	Include/Exclude	If clear, the Fields array (see Table 8.61) specifies which fields to include in the submission. (All descendants of the specified fields in the field hierarchy are submitted as well.) If set, the Fields array tells which fields to exclude; all fields in the document’s interactive form are submitted <i>except</i> those listed in the Fields array and those whose NoExport flag (see Table 8.50 on page 532) is set.
2	IncludeNoValueFields	If set, all fields designated by the Fields array and the Include/Exclude flag are submitted, regardless of whether they have a value (V entry in the field dictionary); for fields without a value, only the field name is transmitted. If clear, fields without a value are not submitted.

- | | | |
|---|--------------------|---|
| 3 | ExportFormat | <p>If set, field names and values are submitted in HTML Form format. If clear, they are submitted in Forms Data Format (FDF); see Section 8.6.6, “Forms Data Format.” This flag is meaningful only when the XML flag is clear; if XML is set, this flag is ignored.</p> |
| 4 | GetMethod | <p>If set, field names and values are submitted using an HTTP GET request; if clear, they are submitted using a POST request. This flag is meaningful only when the ExportFormat flag is set; if ExportFormat is clear, this flag must also be clear.</p> |
| 5 | SubmitCoordinates | <p>If set, the coordinates of the mouse click that caused the submit-form action are transmitted as part of the form data. The coordinate values are relative to the upper-left corner of the field’s widget annotation rectangle. They are represented in the data in the format</p> <p style="text-align: center;"><i>name.x=xval&name.y=yval</i></p> <p>where <i>name</i> is the field’s mapping name (TM in the field dictionary) if present, otherwise the field name. If the value of the TM entry is a single space character, both the name and the dot following it are suppressed, resulting in the format</p> <p style="text-align: center;"><i>x=xval&y=yval</i></p> <p>This flag is meaningful only when the ExportFormat flag is set; if ExportFormat is clear, this flag must also be clear.</p> |
| 6 | XML | <p>(PDF 1.4) If set, field names and values are submitted in XML format; if clear, they are submitted in HTML Form format or Forms Data Format (FDF), according to the value of the ExportFormat flag.</p> |
| 7 | IncludeAppendSaves | <p>(PDF 1.4) Meaningful only when the form is being submitted in Forms Data Format (that is, when both the XML and ExportFormat flags are clear). If set, the submitted FDF file includes the contents of all incremental updates to the underlying PDF document, as contained in the Differences entry in the FDF dictionary (see Table 8.69 on page 561); if clear, the incremental updates are not included.</p> |
| 8 | IncludeAnnotations | <p>(PDF 1.4) Meaningful only when the form is being submitted in Forms Data Format (that is, when both the XML and ExportFormat flags are clear). If set, the submitted FDF file includes all annotations in the underlying PDF document; if clear, the annotations are not included.</p> |

9	SubmitPDF	<i>(PDF 1.4)</i> If set, the document is submitted in PDF format, using the MIME content type <code>application/pdf</code> (described in Internet RFC 2045, <i>Multipurpose Internet Mail Extensions (MIME), Part One: Format of Internet Message Bodies</i> ; see the Bibliography). If this flag is set, all other flags are ignored except <code>GetMethod</code> .
10	CanonicalFormat	<i>(PDF 1.4)</i> If set, any submitted field values representing dates are converted to the standard format described in Section 3.8.2, “Dates.” (The interpretation of a form field as a date is not specified explicitly in the field itself, but only in the JavaScript code that processes it.)
11	ExclNonUserAnnots	<i>(PDF 1.4)</i> Meaningful only when the form is being submitted in Forms Data Format (that is, when both the XML and Export-Format flags are clear) and the <code>IncludeAnnotations</code> flag is set. If set, will include only those annotations whose T entry matches the name of the current user, as determined by the remote server to which the form is being submitted. (The T entry, which specifies the text label to be displayed in the title bar of the annotation’s pop-up window, is assumed to represent the name of the user authoring the annotation.) This allows multiple users to collaborate in annotating a single remote PDF document without affecting one another’s annotations.
12	ExclFKey	<i>(PDF 1.4)</i> Meaningful only when the form is being submitted in Forms Data Format (that is, when both the XML and Export-Format flags are clear). If set, the submitted FDF will exclude the F entry.

The set of fields whose names and values are to be submitted is defined by the **Fields** array in the action dictionary (Table 8.61) together with the `Include/Exclude` and `IncludeNoValueFields` flags in the **Flags** entry (Table 8.62). Each element of the **Fields** array identifies an interactive form field, either by an indirect reference to its field dictionary or *(PDF 1.3)* by its fully qualified field name (see “Field Names” on page 532). If the `Include/Exclude` flag is clear, the submission consists of all fields listed in the **Fields** array, along with any descendants of those fields in the field hierarchy. If the `Include/Exclude` flag is set, the submission consists of all fields in the document’s interactive form *except* those listed in the **Fields** array.

Note: The `NoExport` flag in the field dictionary's **Ff** entry (see Tables 8.49 on page 531 and 8.50 on page 532) takes precedence over the action's **Fields** array and `Include/Exclude` flag. Fields whose `NoExport` flag is set are never included in a `submit-form` action.

Field names and values may be submitted in any of three formats—HTML Form format, XML, or Forms Data Format (FDF)—depending on the settings of the action's `ExportFormat` and `XML` flags. (The first two formats are described, respectively, in Internet RFC 1866, *Hypertext Markup Language 2.0 Proposed Standard*, and in the World Wide Web Consortium document *Extensible Markup Language (XML) 1.0*; see the Bibliography. FDF is described in Section 8.6.6, “Forms Data Format”; see also implementation note 86 in Appendix H.) The name submitted for each field is its fully qualified name (see “Field Names” on page 532), and the value is that specified by the **V** entry in its field dictionary.

Note: For `pushbutton` fields submitted in FDF, the value submitted is that of the **AP** entry in the field's widget annotation dictionary. If the `submit-form` action dictionary contains no **Fields** entry, such `pushbutton` fields are not submitted at all.

Fields with no value (that is, whose field dictionary does not contain a **V** entry) are ordinarily not included in the submission. The `submit-form` action's `IncludeNoValueFields` flag overrides this behavior; if this flag is set, such valueless fields are included in the submission by name only, with no associated value.

Reset-Form Actions

A `reset-form` action resets selected interactive form fields to their default values; that is, it sets the value of the **V** entry in the field dictionary to that of the **DV** entry (see Table 8.49 on page 531). If no default value is defined for a field, its **V** entry is removed. For fields that can have no value (such as `pushbuttons`), the action has no effect. Table 8.63 shows the action dictionary entries specific to this type of action.

The value of the action dictionary's **Flags** entry is an unsigned 32-bit integer containing flags specifying various characteristics of the action. Bit positions within the flag word are numbered from 1 (low-order) to 32 (high-order). At the time of publication, only one flag is defined for this type of action; Table 8.64 shows its meaning. All undefined flag bits are reserved and must be set to 0.

TABLE 8.63 Additional entries specific to a reset-form action

KEY	TYPE	VALUE
S	name	<i>(Required)</i> The type of action that this dictionary describes; must be ResetForm for a reset-form action.
Fields	array	<i>(Optional)</i> An array identifying which fields to reset or which to exclude from resetting, depending on the setting of the Include/Exclude flag in the Flags entry (see Table 8.64). Each element of the array is either an indirect reference to a field dictionary or (<i>PDF 1.3</i>) a string representing the fully qualified name of a field. Elements of both kinds may be mixed in the same array. If this entry is omitted, the Include/Exclude flag is ignored; all fields in the document's interactive form are reset.
Flags	integer	<i>(Optional; inheritable)</i> A set of flags specifying various characteristics of the action (see Table 8.64). Default value: 0.

TABLE 8.64 Flag for reset-form actions

BIT POSITION	NAME	MEANING
1	Include/Exclude	If clear, the Fields array (see Table 8.63) specifies which fields to reset. (All descendants of the specified fields in the field hierarchy are reset as well.) If set, the Fields array tells which fields to exclude from resetting; all fields in the document's interactive form are reset <i>except</i> those listed in the Fields array.

Import-Data Actions

An *import-data action* imports Forms Data Format (FDF) data into the document's interactive form from a specified file (see Section 8.6.6, "Forms Data Format"). Table 8.65 shows the action dictionary entries specific to this type of action.

TABLE 8.65 Additional entries specific to an import-data action

KEY	TYPE	VALUE
S	name	<i>(Required)</i> The type of action that this dictionary describes; must be ImportData for an import-data action.
F	file specification	<i>(Required)</i> The FDF file from which to import the data. (See implementation notes 87 and 88 in Appendix H.)

JavaScript Actions

A *JavaScript action* (PDF 1.3) causes a script to be compiled and executed by the JavaScript interpreter. Depending on the nature of the script, this can cause various interactive form fields in the document to update their values or change their visual appearances. Netscape Communications Corporation's *Client-Side JavaScript Reference* and Adobe Technical Note #5186, *Acrobat Forms JavaScript Object Specification* (see the Bibliography) give details on the contents and effects of JavaScript scripts. Table 8.66 shows the action dictionary entries specific to this type of action.

TABLE 8.66 Additional entries specific to a JavaScript action

KEY	TYPE	VALUE
S	name	<i>(Required)</i> The type of action that this dictionary describes; must be JavaScript for a JavaScript action.
JS	string or stream	<i>(Required)</i> A string or stream containing the JavaScript script to be executed. <i>Note: PDFDocEncoding or Unicode encoding (the latter identified by the Unicode prefix U+FEFF) is used to encode the contents of the string or stream. (See implementation note 89 in Appendix H.)</i>

To support the use of parameterized function calls in JavaScript scripts, the **JavaScript** entry in a PDF document's name dictionary (see Section 3.6.3, "Name Dictionary") can contain a name tree mapping name strings to document-level JavaScript actions. When the document is opened, all of the actions in this name tree are executed, defining JavaScript functions for use by other scripts in the document.

Note: The name strings associated with individual JavaScript actions in the name dictionary serve merely as a convenient means for organizing and packaging scripts. The names are arbitrary and need not bear any relation to the JavaScript name space itself.

8.6.5 Named Pages

The optional **Pages** entry (*PDF 1.3*) in a document's name dictionary (see Section 3.6.3, "Name Dictionary") contains a name tree that maps name strings to individual pages within the document. Naming a page allows it to be referenced in two different ways:

- An import-data action can add the named page to the document into which FDF is being imported, either as a page or as a button appearance.
- A script executed by a JavaScript action can add the named page to the current document as a regular page.

A named page that is to be visible to the user should be left in the page tree (see Section 3.6.2, "Page Tree"), with a reference to it in the appropriate leaf node of the name dictionary's **Pages** tree. If the page is not to be displayed by the viewer application, it should be referenced from the name dictionary's **Templates** tree instead. Such invisible pages should have an object type of **Template** rather than **Page**, and should have no **Parent** or **B** entry (see Table 3.18 on page 88). Regardless of whether the page is named in the **Pages** or **Templates** tree or whether it is added to a document by an import-data or JavaScript action, the new copy is not itself named.

8.6.6 Forms Data Format

This section describes Forms Data Format (FDF), the file format used for interactive form data (*PDF 1.2*). FDF is used when submitting form data to a server, receiving the response, and incorporating it into the interactive form. It can also be used to export form data to stand-alone files that can be stored, transmitted electronically, and imported back into the corresponding PDF interactive form. In addition, beginning in PDF 1.3, it can be used to define a container for annotations that are separate from the PDF document to which they apply.

FDF is based on PDF; it uses the same syntax (see Section 3.1, “Lexical Conventions”) and basic object types (Section 3.2, “Objects”), and has essentially the same file structure (Section 3.4, “File Structure”). However, it differs from PDF in the following ways:

- The cross-reference table (Section 3.4.3, “Cross-Reference Table”) is optional.
- FDF files cannot be updated (see Section 3.4.5, “Incremental Updates”); objects can only be of generation 0, and no two objects can have the same object number.
- The document structure is much simpler than PDF, since the body of an FDF document consists of only one required object.
- The length of a stream may not be specified by an indirect object.

FDF uses the MIME content type `application/vnd.fdf`. On the Windows and UNIX platforms, FDF files have the extension `.fdf`; on Mac OS, they have file type 'FDF '.

FDF File Structure

An FDF file is structured in essentially the same way as a PDF file, but need contain only those elements required for the export and import of interactive form and annotation data. It consists of three required elements and one optional one (see Figure 8.6):

- A one-line *header* identifying the version number of the PDF specification to which the file conforms
- A *body* containing the objects that make up the content of the file
- An optional *cross-reference table* containing information about the objects in the file
- A *trailer* giving the location of various objects within the body of the file




FIGURE 8.6 *FDF file structure*

FDF Header

The first line of an FDF file is a *header*, originally intended to identify the version of the PDF specification to which the file conforms. However, for historical reasons, this version number is now frozen and must read

```
%PDF-1.2
```

The true version number is now given by the **Version** entry in the FDF catalog dictionary (see “FDF Catalog,” below; see also implementation note 90 in Appendix H).

FDF Body

The *body* of an FDF file consists of a sequence of indirect objects representing the file’s catalog (see “FDF Catalog” on page 560), together with any additional objects that the catalog may reference. The objects are of the same basic types

described in Section 3.2, “Objects.” Just as in PDF, objects in FDF can be direct or indirect.

FDF Trailer

The *trailer* of an FDF file enables an application reading the file to find significant objects quickly within the body of the file. The last line of the file contains only the end-of-file marker, %%EOF. This is preceded by the *FDF trailer dictionary*, consisting of the keyword **trailer** followed by a series of one or more key-value pairs enclosed in double angle brackets (<<...>>). The only required key is **Root**, whose value is an indirect reference to the file’s catalog dictionary (see Table 8.67). The trailer may optionally contain additional entries for objects that are referenced from within the catalog.

TABLE 8.67 Entry in the FDF trailer dictionary

KEY	TYPE	VALUE
Root	dictionary	(Required; must be an indirect reference) The catalog object for this FDF file (see “FDF Catalog,” below).

Thus the trailer has the following overall structure:

```
trailer
  << Root c 0 R
      key2 value2
      ...
      keyn valuen
  >>
%%EOF
```

where *c* is the object number of the file’s catalog dictionary.

FDF Catalog

The root node of an FDF file’s object hierarchy is the *catalog* dictionary, located via the **Root** entry in the file’s trailer dictionary (see “FDF Trailer,” above). As shown in Table 8.68, the only required entry in the catalog is **FDF**; its value is an *FDF dictionary* (Table 8.69), which in turn contains references to other objects

describing the file's contents. The catalog may also contain an optional **Version** entry identifying the version of the PDF specification to which this FDF file conforms.

TABLE 8.68 Entries in the FDF catalog dictionary

KEY	TYPE	VALUE
Version	name	<p>(Optional; PDF 1.4) The version of the PDF specification to which this FDF file conforms (for example, 1.4), if later than the version specified in the file's header (see "FDF Header" on page 559). If the header specifies a later version, or if this entry is absent, the document conforms to the version specified in the header.</p> <p>Note: The value of this entry is a name object, not a number, and so must be preceded by a slash character (/) when written in the FDF file (for example, /1.4).</p>
FDF	dictionary	(Required) The FDF dictionary for this file (see Table 8.69).

TABLE 8.69 Entries in the FDF dictionary

KEY	TYPE	VALUE
F	file specification	(Optional) The <i>source file</i> or <i>target file</i> : the PDF document file that this FDF file was exported from or is intended to be imported into.
ID	array	(Optional) An array of two strings constituting a file identifier (see Section 9.3, "File Identifiers") for the source or target file designated by F , taken from the ID entry in the file's trailer dictionary (see Section 3.4.4, "File Trailer").
Fields	array	(Optional) An array of FDF field dictionaries (see "FDF Fields" on page 564) describing the root fields (those with no ancestors in the field hierarchy) to be exported or imported. This entry and the Pages entry may not both be present.
Status	string	(Optional) A status string to be displayed indicating the result of an action, typically a submit-form action (see "Submit-Form Actions" on page 550). The string is encoded with PDFDocEncoding . (See implementation note 91 in Appendix H.) This entry and the Pages entry may not both be present.
Pages	array	(Optional; PDF 1.3) An array of FDF page dictionaries (see "FDF Pages" on page 566) describing new pages to be added to a PDF target document. The Fields and Status entries may not be present together with this entry.

Encoding	name	<i>(Optional; PDF 1.3)</i> The encoding to be used for any FDF field value or option (V or Opt in the field dictionary; see Table 8.72 on page 564) that is a string and does not begin with the Unicode prefix U+FEFF. (See implementation note 92 in Appendix H.) Default value: PDFDocEncoding .
Annots	array	<i>(Optional; PDF 1.3)</i> An array of FDF annotation dictionaries (see “FDF Annotation Dictionaries” on page 568). The array can include annotations of any of the standard types listed in Table 8.14 on page 499 except Link , Movie , Widget , PrinterMark , and TrapNet .
Differences	stream	<i>(Optional; PDF 1.4)</i> A stream containing all the bytes in all incremental updates made to the underlying PDF document since it was opened (see Section 3.4.5, “Incremental Updates”). If a submit-form action submitting the document to a remote server in FDF format has its <code>IncludeAppendSaves</code> flag set (see “Submit-Form Actions” on page 550), the contents of this stream are included in the submission. This allows any digital signatures (see “Signature Fields” on page 547) to be transmitted to the server. An incremental update is automatically performed just before the submission takes place, in order to capture all changes made to the document. Note that the submission always includes the full set of incremental updates back to the time the document was first opened, even if some of them may already have been included in intervening submissions. <i>Note: Although a Fields or Annots entry (or both) may be present along with Differences, there is no guarantee that their contents will be consistent with it. In particular, if Differences contains a digital signature, only the values of the form fields given in the Differences stream can be considered trustworthy under that signature.</i>
Target	string	<i>(Optional; PDF 1.4)</i> The name of a browser frame in which the underlying PDF document is to be opened. This mimics the behavior of the <code>target</code> attribute in HTML <code><href></code> tags.
EmbeddedFDFs	array	<i>(Optional; PDF 1.4)</i> An array of file specifications (see Section 3.10, “File Specifications”) representing other FDF files embedded within this one (Section 3.10.3, “Embedded File Streams”).
JavaScript	dictionary	<i>(Optional; PDF 1.4)</i> A <i>JavaScript dictionary</i> (see Table 8.71) defining document-level JavaScript scripts.

Embedded FDF files specified in the FDF dictionary’s **EmbeddedFDFs** entry may optionally be encrypted. Besides the usual entries for an embedded file stream, the stream dictionary representing such an encrypted FDF file must contain the

additional entry shown in Table 8.70 to identify the revision number of the FDF encryption algorithm used to encrypt the file. Although the FDF encryption mechanism is separate from the one for PDF file encryption described in Section 3.5, “Encryption,” revision 1 (the only one defined at the time of publication) uses a similar RC4 encryption algorithm based on a 40-bit encryption key. The key is computed via an MD5 hash, using a padded user-supplied password as input. The computation is identical to steps 1 and 2 of Algorithm 3.2 on page 78; the first 5 bytes of the result are the encryption key for the embedded FDF file.

TABLE 8.70 Additional entry in an embedded file stream dictionary for an encrypted FDF file

KEY	TYPE	VALUE
EncryptionRevision	integer	<i>(Required if the FDF file is encrypted; PDF 1.4)</i> The revision number of the FDF encryption algorithm used to encrypt the file. The only valid value defined at the time of publication is 1.

The **JavaScript** entry in the FDF dictionary holds a *JavaScript dictionary* containing JavaScript scripts that are defined globally at the document level, rather than associated with individual fields. The dictionary can contain scripts defining JavaScript functions for use by other scripts in the document, as well as scripts to be executed immediately before and after the FDF file is imported. Table 8.71 shows the contents of this dictionary.

TABLE 8.71 Entries in the JavaScript dictionary

KEY	TYPE	VALUE
Before	string or stream	<i>(Optional)</i> A string or stream containing a JavaScript script to be executed just before the FDF file is imported.
After	string or stream	<i>(Optional)</i> A string or stream containing a JavaScript script to be executed just after the FDF file is imported.
Doc	array	<i>(Optional)</i> An array defining additional JavaScript scripts to be added to those defined in the JavaScript entry of the document’s name dictionary (see Section 3.6.3, “Name Dictionary”). The array contains an even number of elements, organized in pairs. The first element of each pair is a name and the second is a string or stream defining the script corresponding to that name. Each of the defined scripts will be added to those already defined in the name dictionary and then executed before the script defined in the Before entry is executed. As described in “JavaScript Actions” on page 556, these scripts are used to define JavaScript functions for use by other scripts in the document.

FDF Fields

Each field in an FDF file is described by an *FDF field dictionary*. Table 8.72 shows the contents of this type of dictionary. Most of the entries have the same form and meaning as the corresponding entries in a field dictionary (Tables 8.49 on page 531, 8.51 on page 534, 8.57 on page 544, and 8.59 on page 546) or a widget annotation dictionary (Tables 8.10 on page 490 and 8.28 on page 512). Unless otherwise indicated in the table, importing a field causes the values of the entries in the FDF field dictionary to replace those of the corresponding entries in the field with the same fully qualified name in the target document. (See implementation notes 93–98 in Appendix H.)

TABLE 8.72 Entries in an FDF field dictionary

KEY	TYPE	VALUE
Kids	array	<i>(Optional)</i> An array containing the immediate children of this field. <i>Note:</i> Unlike the children of fields in a PDF file, which must be specified as indirect object references, those of an FDF field may be either direct or indirect objects.
T	text string	<i>(Required)</i> The partial field name (see “Field Names” on page 532).
V	(various)	<i>(Optional)</i> The field’s value, whose format varies depending on the field type; see the descriptions of individual field types in Section 8.6.3 for further information.
Ff	integer	<i>(Optional)</i> A set of flags specifying various characteristics of the field (see Tables 8.50 on page 532, 8.53 on page 538, 8.56 on page 543, and 8.58 on page 546). When imported into an interactive form, the value of this entry replaces that of the Ff entry in the form’s corresponding field dictionary. If this field is present, the SetFf and ClrFf entries, if any, are ignored.
SetFf	integer	<i>(Optional)</i> A set of flags to be set (turned on) in the Ff entry of the form’s corresponding field dictionary. Bits equal to 1 in SetFf cause the corresponding bits in Ff to be set to 1. This entry is ignored if an Ff entry is present in the FDF field dictionary.
ClrFf	integer	<i>(Optional)</i> A set of flags to be cleared (turned off) in the Ff entry of the form’s corresponding field dictionary. Bits equal to 1 in ClrFf cause the corresponding bits in Ff to be set to 0. If a SetFf entry is also present in the FDF field dictionary, it is applied before this entry. This entry is ignored if an Ff entry is present in the FDF field dictionary.

F	integer	<i>(Optional)</i> A set of flags specifying various characteristics of the field’s widget annotation (see Section 8.4.2, “Annotation Flags”). When imported into an interactive form, the value of this entry replaces that of the F entry in the form’s corresponding annotation dictionary. If this field is present, the SetF and ClrF entries, if any, are ignored.
SetF	integer	<i>(Optional)</i> A set of flags to be set (turned on) in the F entry of the form’s corresponding widget annotation dictionary. Bits equal to 1 in SetF cause the corresponding bits in F to be set to 1. This entry is ignored if an F entry is present in the FDF field dictionary.
ClrF	integer	<i>(Optional)</i> A set of flags to be cleared (turned off) in the F entry of the form’s corresponding widget annotation dictionary. Bits equal to 1 in ClrF cause the corresponding bits in F to be set to 0. If a SetF entry is also present in the FDF field dictionary, it is applied before this entry. This entry is ignored if an F entry is present in the FDF field dictionary.
AP	dictionary	<i>(Optional)</i> An appearance dictionary specifying the appearance of a pushbutton field (see “Pushbuttons” on page 539). The appearance dictionary’s contents are as shown in Table 8.13 on page 497, except that the values of the N , R , and D entries must all be streams.
APRef	dictionary	<i>(Optional; PDF 1.3)</i> A dictionary holding references to external PDF files containing the pages to use for the appearances of a pushbutton field. This dictionary is similar to an appearance dictionary (see Table 8.13 on page 497), except that the values of the N , R , and D entries must all be named page reference dictionaries (Table 8.76 on page 568). This entry is ignored if an AP entry is present.
IF	dictionary	<i>(Optional; PDF 1.3; button fields only)</i> An icon fit dictionary (see Table 8.73) specifying how to display a button field’s icon within the annotation rectangle of its widget annotation.
Opt	array	<p><i>(Required; choice fields only)</i> An array of options to be presented to the user. Each element of the array can take either of two forms:</p> <ul style="list-style-type: none"> • A text string representing one of the available options • A two-element array consisting of a text string representing one of the available options and a default appearance string for constructing the item’s appearance dynamically at viewing time (see “Variable Text” on page 533)
A	dictionary	<i>(Optional)</i> An action to be performed when this field’s widget annotation is activated (see Section 8.5, “Actions”).
AA	dictionary	<i>(Optional)</i> An additional-actions dictionary defining the field’s behavior in response to various trigger events (see Section 8.5.2, “Trigger Events”).

In an FDF field dictionary representing a button field, the optional **IF** entry holds an *icon fit dictionary* (PDF 1.3) specifying how to display the button's icon within the annotation rectangle of its widget annotation. Table 8.73 shows the contents of this type of dictionary.

TABLE 8.73 Entries in an icon fit dictionary

KEY	TYPE	VALUE
SW	name	<p>(Required) The circumstances under which the icon should be scaled inside the annotation rectangle:</p> <ul style="list-style-type: none"> A Always scale. B Scale only when the icon is bigger than the annotation rectangle. S Scale only when the icon is smaller than the annotation rectangle. N Never scale. <p>Default value: A.</p>
S	name	<p>(Required) The type of scaling to use:</p> <ul style="list-style-type: none"> A <i>Anamorphic scaling</i>: scale the icon to fill the annotation rectangle exactly, without regard to its original aspect ratio (ratio of width to height). P <i>Proportional scaling</i>: scale the icon to fit the width or height of the annotation rectangle while maintaining the icon's original aspect ratio. If the required horizontal and vertical scaling factors are different, use the smaller of the two, centering the icon within the annotation rectangle in the other dimension. <p>Default value: P.</p>
A	array	<p>(Required) An array of two numbers between 0.0 and 1.0 indicating the fraction of left-over space to allocate at the left and bottom of the icon. A value of [0.0 0.0] positions the icon at the bottom-left corner of the annotation rectangle; a value of [0.5 0.5] centers it within the rectangle. This entry is used only if the icon is scaled proportionally. Default value: [0.5 0.5].</p>

FDF Pages

The optional **Pages** field in an FDF dictionary (see Table 8.69 on page 561) contains an array of *FDF page dictionaries* (PDF 1.3) describing new pages to be added to the target document. Table 8.74 shows the contents of this type of dictionary.

TABLE 8.74 Entries in an FDF page dictionary

KEY	TYPE	VALUE
Templates	array	<i>(Required)</i> An array of <i>FDF template dictionaries</i> (see Table 8.75) describing the named pages that serve as templates on the page.
Info	dictionary	<i>(Optional)</i> An <i>FDF page information dictionary</i> containing additional information about the page. At the time of publication, no entries have been defined for this dictionary.

An *FDF template dictionary* contains information describing a named page that serves as a template. Table 8.75 shows the contents of this type of dictionary.

TABLE 8.75 Entries in an FDF template dictionary

KEY	TYPE	VALUE
TRef	dictionary	<i>(Required)</i> A named page reference dictionary (see Table 8.76) specifying the location of the template.
Fields	array	<i>(Optional)</i> An array of references to FDF field dictionaries (see Table 8.72 on page 564) describing the root fields to be imported (those with no ancestors in the field hierarchy).
Rename	boolean	<i>(Optional)</i> A flag specifying whether fields imported from the template may be renamed in the event of name conflicts with existing fields; see below for further discussion. Default value: true .

The names of fields imported from a template may sometimes conflict with those of existing fields in the target document. This can occur, for example, if the same template page is imported more than once or if two different templates have fields with the same names. If the **Rename** flag in the FDF template dictionary is **true**, fields with such conflicting names are renamed to guarantee their uniqueness. If **Rename** is **false**, the fields are not renamed; this results in multiple fields with the same name in the target document. Each time the FDF file provides attributes for a given field name, all fields with that name will be updated. (See implementation notes 99 and 100 in Appendix H.)

The **TRef** entry in an FDF template dictionary holds a *named page reference dictionary* describing the location of external templates or page elements. Table 8.76 shows the contents of this type of dictionary.

TABLE 8.76 Entries in an FDF named page reference dictionary

KEY	TYPE	VALUE
Name	string	<i>(Required)</i> The name of the referenced page.
F	file specification	<i>(Optional)</i> The file containing the named page. If this key is absent, it is assumed that the page resides in the associated PDF file.

FDF Annotation Dictionaries

Each annotation dictionary in an FDF file must have a **Page** entry (see Table 8.77) indicating the page of the source document to which the annotation is attached.

TABLE 8.77 Additional entry for annotation dictionaries in an FDF file

KEY	TYPE	VALUE
Page	integer	<i>(Required for annotations in FDF files)</i> The ordinal page number on which this annotation should appear, where page 0 is the first page.

8.7 Sounds

A *sound object* (PDF 1.2) is a stream containing sample values that define a sound to be played through the computer's speakers. The **Sound** entry in a sound annotation or sound action dictionary (see Tables 8.26 on page 510 and 8.42 on page 525) identifies a sound object representing the sound to be played when the annotation is activated.

Since a sound object is a stream, it can contain any of the standard entries common to all streams, as described in Table 3.4 on page 38. In particular, if it contains an **F** (file specification) entry, then the sound is defined in an external file.

This must be a self-describing sound file, containing all information needed to render the sound; no additional information need be present in the PDF file.

Note: The *AIFF*, *AIFF-C* (Mac OS), *RIFF* (.wav), and *snd* (.au) file formats are all self-describing.

If no **F** entry is present, the sound object itself contains the sample data and all other information needed to define the sound. Table 8.78 shows the additional dictionary entries specific to a sound object.

TABLE 8.78 Additional entries specific to a sound object

KEY	TYPE	VALUE								
Type	name	(Optional) The type of PDF object that this dictionary describes; if present, must be Sound for a sound object.								
R	number	(Required) The sampling rate, in samples per second.								
C	integer	(Optional) The number of sound channels. Default value: 1. (See implementation note 101 in Appendix H.)								
B	integer	(Optional) The number of bits per sample value per channel. Default value: 8.								
E	name	(Optional) The encoding format for the sample data: <table data-bbox="442 926 1135 1067" style="margin-left: 20px;"> <tr> <td>Raw</td> <td>Unspecified or unsigned values in the range 0 to $2^B - 1$</td> </tr> <tr> <td>Signed</td> <td>Twos-complement values</td> </tr> <tr> <td>muLaw</td> <td>μ-law-encoded samples</td> </tr> <tr> <td>ALaw</td> <td>A-law-encoded samples</td> </tr> </table> <p>Default value: Raw.</p>	Raw	Unspecified or unsigned values in the range 0 to $2^B - 1$	Signed	Twos-complement values	muLaw	μ -law-encoded samples	ALaw	A-law-encoded samples
Raw	Unspecified or unsigned values in the range 0 to $2^B - 1$									
Signed	Twos-complement values									
muLaw	μ -law-encoded samples									
ALaw	A-law-encoded samples									
CO	name	(Optional) The sound compression format used on the sample data. (Note that this is separate from any stream compression specified by the sound object's Filter entry; see Table 3.4 on page 38 and Section 3.3, "Filters.") If this entry is absent, then no sound compression has been used; the data contains sampled waveforms to be played at R samples per second per channel.								
CP	(various)	(Optional) Optional parameters specific to the sound compression format used. <p>Note: At the time of publication, no standard values have been defined for the CO and CP entries.</p>								

Sample values are stored in the stream with the most significant bits first (“big-endian” order for samples larger than 8 bits). Samples that are not a multiple of 8 bits are packed into consecutive bytes, starting at the most significant end. If a sample extends across a byte boundary, the most significant bits are placed in the first byte, followed by less significant bits in subsequent bytes. For dual-channel stereophonic sounds, the samples are stored in an interleaved format, with each sample value for the left channel (channel 1) preceding the corresponding sample for the right (channel 2).

To maximize the portability of PDF documents containing embedded sounds, Adobe recommends that PDF viewer applications and plug-in extensions support at least the following formats (assuming the platform has sufficient hardware and OS support to play sounds at all):

R	8000, 11,025, or 22,050 samples per second
C	1 or 2 channels
B	8 or 16 bits per channel
E	Raw, Signed, or muLaw encoding

If the encoding (**E**) is Raw or Signed, then **R** must be 11,025 or 22,050 samples per channel. If the encoding is muLaw, then **R** must be 8000 samples per channel, **C** must be 1 channel, and **B** must be 8 bits per channel. Sound players should be prepared to convert between formats, downsample rates, and combine channels as necessary to render sound on the target platform.

8.8 Movies

PDF includes the ability to embed *movies* within a document by means of movie annotations (see “Movie Annotations” on page 510). Despite the name, a movie may consist entirely of sound with no visible images to be displayed on the screen. The **Movie** and **A** (activation) entries in the movie annotation dictionary refer, respectively, to a *movie dictionary* (Table 8.79) describing the static characteristics of the movie and a *movie activation dictionary* (Table 8.80) specifying how it should be presented.

TABLE 8.79 Entries in a movie dictionary

KEY	TYPE	VALUE
F	file specification	<i>(Required)</i> A file specification identifying a self-describing movie file. <i>Note:</i> The format of a “self-describing movie file” is left unspecified, and there is no guarantee of portability.
Aspect	array	<i>(Optional)</i> The width and height of the movie’s bounding box, in pixels, specified as [<i>width height</i>]. This entry should be omitted for a movie consisting entirely of sound with no visible images.
Rotate	integer	<i>(Optional)</i> The number of degrees by which the movie is rotated clockwise relative to the page. The value must be a multiple of 90. Default value: 0.
Poster	boolean or stream	<i>(Optional)</i> A flag or stream specifying whether and how to display a <i>poster image</i> representing the movie. If this value is a stream, it contains an image XObject (see Section 4.8, “Images”) to be displayed as the poster; if it is the boolean value true , the poster image should be retrieved from the movie file itself; if it is false , no poster should be displayed. Default value: false .

TABLE 8.80 Entries in a movie activation dictionary

KEY	TYPE	VALUE
Start	(various)	<i>(Optional)</i> The starting time of the movie segment to be played. Movie time values are expressed in units of time based on a <i>time scale</i> , which defines the number of units per second; the default time scale is defined in the movie data itself. The starting time is nominally a 64-bit integer, specified as follows: <ul style="list-style-type: none"> • If it is representable as an integer (subject to the implementation limit for integers, as described in Appendix C), it should be specified as such. • If it is not representable as an integer, it should be specified as an 8-byte string representing a 64-bit twos-complement integer, most significant byte first. • If it is expressed in a time scale different from that of the movie itself, it is represented as an array of two values: an integer or string denoting the starting time, as above, followed by an integer specifying the time scale in units per second. <p>If this entry is omitted, the movie is played from the beginning.</p>
Duration	(various)	<i>(Optional)</i> The duration of the movie segment to be played, specified in the same form as Start . Negative values specify that the movie is to be played backward. If this entry is omitted, the movie is played to the end.

Rate	number	<p>(<i>Optional</i>) The initial speed at which to play the movie. If the value of this entry is negative, the movie is played backward with respect to Start and Duration. Default value: 1.0.</p>								
Volume	number	<p>(<i>Optional</i>) The initial sound volume at which to play the movie, in the range -1.0 to 1.0. Higher values denote greater volume; negative values mute the sound. Default value: 1.0.</p>								
ShowControls	boolean	<p>(<i>Optional</i>) A flag specifying whether to display a movie controller bar while playing the movie. Default value: false.</p>								
Mode	name	<p>(<i>Optional</i>) The <i>play mode</i> for playing the movie:</p> <table border="0" style="margin-left: 2em;"> <tr> <td>Once</td> <td>Play once and stop.</td> </tr> <tr> <td>Open</td> <td>Play and leave the movie controller bar open.</td> </tr> <tr> <td>Repeat</td> <td>Play repeatedly from beginning to end until stopped.</td> </tr> <tr> <td>Palindrome</td> <td>Play continuously forward and backward until stopped.</td> </tr> </table> <p>Default value: Once.</p>	Once	Play once and stop.	Open	Play and leave the movie controller bar open.	Repeat	Play repeatedly from beginning to end until stopped.	Palindrome	Play continuously forward and backward until stopped.
Once	Play once and stop.									
Open	Play and leave the movie controller bar open.									
Repeat	Play repeatedly from beginning to end until stopped.									
Palindrome	Play continuously forward and backward until stopped.									
Synchronous	boolean	<p>(<i>Optional</i>) A flag specifying whether to play the movie synchronously or asynchronously. If this value is true, the movie player will retain control until the movie is completed or dismissed by the user; if false, it will return control to the viewer application immediately after starting the movie. Default value: false.</p>								
FWScale	array	<p>(<i>Optional</i>) The magnification (zoom) factor at which to play the movie. The presence of this entry implies that the movie is to be played in a floating window; if the entry is absent, it will be played in the annotation rectangle.</p> <p>The value of the entry is an array of two integers, [<i>numerator denominator</i>], denoting a rational magnification factor for the movie. The final window size, in pixels, is</p> $(\textit{numerator} \div \textit{denominator}) \times \mathbf{Aspect}$ <p>where the value of Aspect is taken from the movie dictionary (see Table 8.79).</p>								
FWPosition	array	<p>(<i>Optional</i>) For floating play windows, the relative position of the window on the screen. The value is an array of two numbers</p> $[\textit{horiz} \textit{vert}]$ <p>each in the range 0.0 to 1.0, denoting the relative horizontal and vertical position of the movie window with respect to the screen. For example, the value [0.5 0.5] centers the window on the screen. Default value: [0.5 0.5].</p>								

CHAPTER 9

Document Interchange

THE FEATURES DESCRIBED in this chapter do not affect the final appearance of a document. Rather, they allow it to include higher-level information that is useful for the interchange of documents among applications. They include:

- *Procedure sets* (Section 9.1) that define the implementation of PDF operators
- *Metadata* (Section 9.2) consisting of general information about a document or a component of a document, such as its title, author, and creation and modification dates
- *File identifiers* (Section 9.3) for reliable reference from one PDF file to another
- *Page-piece dictionaries* (Section 9.4) allowing an application to embed private data in a PDF document for its own use
- *Marked-content* operators (Section 9.5) for identifying portions of a content stream and associating them with additional properties or externally specified objects
- *Logical structure* facilities (Section 9.6) for imposing a hierarchical organization on the content of a document
- *Tagged PDF* (Section 9.7), a set of conventions for using the marked content and logical structure facilities to facilitate the extraction and reuse of a document's content for other purposes
- Various ways of increasing the *accessibility* of a document to disabled users (Section 9.8), including the identification of the natural language in which it is written (such as English or Spanish) for the benefit of a text-to-speech engine
- The *Web Capture* plug-in extension (Section 9.9), which creates PDF files from Internet-based or locally resident HTML, PDF, GIF, JPEG, and ASCII text files

- Facilities supporting prepress production workflows (Section 9.10), such as the specification of *page boundaries* and the generation of *printer's marks*, *color separations*, *output intents*, *traps*, and low-resolution *proxies* for high-resolution images

9.1 Procedure Sets

The PDF operators used in content streams are grouped into categories of related operators called *procedure sets* (see Table 9.1). Each procedure set corresponds to a named resource containing the implementations of the operators in that procedure set. The **ProcSet** entry in a content stream's resource dictionary (see Section 3.7.2, "Resource Dictionaries") holds an array consisting of the names of the procedure sets used in that content stream. These procedure sets are used only when the content stream is printed to a PostScript output device; the names identify PostScript procedure sets that must be sent to the device to interpret the PDF operators in the content stream. Each element of this array must be one of the predefined names shown in Table 9.1. (See implementation note 102 in Appendix H.)

TABLE 9.1 Predefined procedure sets

NAME	CATEGORY OF OPERATORS
PDF	Painting and graphics state
Text	Text
ImageB	Grayscale images or image masks
ImageC	Color images
ImageI	Indexed (color-table) images

Note: Beginning with PDF 1.4, this feature is considered obsolete. For compatibility with existing viewer applications, PDF producer applications should continue to specify procedure sets (preferably all of those listed in Table 9.1, unless it is known that fewer are needed). However, viewer applications should not depend on the correctness of this information.

9.2 Metadata

A PDF document may include general information such as the document's title, author, and creation and modification dates. Such global information about the document itself (as opposed to its content or structure) is called *metadata*, and is intended to assist in cataloguing and searching for documents in external databases. A document's metadata may also be added or changed by users or plug-in extensions (see implementation note 103 in Appendix H). Beginning with PDF 1.4, metadata can also be specified for individual components of a document.

Metadata can be stored in a PDF document in either of the following ways:

- In a *document information dictionary* associated with the document (Section 9.2.1)
- In a *metadata stream* (PDF 1.4) associated with the document or a component of the document (Section 9.2.2)

9.2.1 Document Information Dictionary

The optional **Info** entry in the trailer of a PDF file (see Section 3.4.4, “File Trailer”) can hold a *document information dictionary* containing metadata for the document; Table 9.2 shows its contents. Any entry whose value is not known should be omitted from the dictionary, rather than included with an empty string as its value.

Some plug-in extensions may choose to permit searches on the contents of the document information dictionary. To facilitate browsing and editing, all keys in the dictionary are fully spelled out, not abbreviated. New keys should be chosen with care so that they make sense to users.

Note: *Although viewer applications can store custom metadata in the document information dictionary, it is inappropriate to store private content or structural information there; such information should be stored in the document catalog instead (see Section 3.6.1, “Document Catalog”).*

TABLE 9.2 Entries in the document information dictionary

KEY	TYPE	VALUE
Title	text string	<i>(Optional; PDF 1.1)</i> The document's title.
Author	text string	<i>(Optional)</i> The name of the person who created the document.
Subject	text string	<i>(Optional; PDF 1.1)</i> The subject of the document.
Keywords	text string	<i>(Optional; PDF 1.1)</i> Keywords associated with the document.
Creator	text string	<i>(Optional)</i> If the document was converted to PDF from another format, the name of the application (for example, Adobe FrameMaker [®]) that created the original document from which it was converted.
Producer	text string	<i>(Optional)</i> If the document was converted to PDF from another format, the name of the application (for example, Acrobat Distiller) that converted it to PDF.
CreationDate	date	<i>(Optional)</i> The date and time the document was created, in human-readable form (see Section 3.8.2, "Dates").
ModDate	date	<i>(Optional; PDF 1.1)</i> The date and time the document was most recently modified, in human-readable form (see Section 3.8.2, "Dates").
Trapped	name	<p><i>(Optional; PDF 1.3)</i> A name object indicating whether the document has been modified to include trapping information (see Section 9.10.5, "Trapping Support"):</p> <p>True The document has been fully trapped; no further trapping is needed. (Note that this is the name True, not the boolean value true.)</p> <p>False The document has not yet been trapped; any desired trapping must still be done. (Note that this is the name False, not the boolean value false.)</p> <p>Unknown Either it is unknown whether the document has been trapped or it has been partly but not yet fully trapped; some additional trapping may still be needed.</p> <p>Default value: Unknown.</p> <p>The value of this entry may be set automatically by the software creating the document's trapping information or may be known only to a human operator and entered manually.</p>

Example 9.1 shows a typical document information dictionary.

Example 9.1

```
1 0 obj
  << /Title (PostScript Language Reference, Third Edition)
    /Author (Adobe Systems Incorporated)
    /Creator (Adobe® FrameMaker® 5.5.3 for Power Macintosh)
    /Producer (Acrobat® Distiller™ 3.01 for Power Macintosh)
    /CreationDate (D:19970915110347-08'00')
    /ModDate (D:19990209153925-08'00')
  >>
endobj
```

9.2.2 Metadata Streams

Metadata, both for an entire document and for components within a document, can be stored in PDF streams called *metadata streams* (PDF 1.4). The advantages of metadata streams over the document information dictionary include the following:

- PDF-based workflows often embed metadata-bearing artwork as components within larger documents. Metadata streams provide a standard way of preserving the metadata of these components for examination downstream. PDF-aware applications should be able to derive a list of all metadata-bearing document components from the PDF document itself.
- PDF documents are often made available on the World Wide Web or in other environments, where many tools routinely examine, catalog, and classify documents. These tools should be able to understand the self-contained description of the document even if they do not understand PDF.

Besides the usual entries common to all stream dictionaries (see Table 3.4 on page 38), the metadata stream dictionary contains the additional entries listed in Table 9.3.

The contents of a metadata stream is the metadata represented in Extensible Markup Language (XML). This information will be visible as plain text to tools that are not PDF-aware only if the metadata stream is both unfiltered and unencrypted.

TABLE 9.3 Additional entries in a metadata stream dictionary

KEY	TYPE	VALUE
Type	name	<i>(Required)</i> The type of PDF object that this dictionary describes; must be Metadata for a metadata stream.
Subtype	name	<i>(Required)</i> The type of metadata stream that this dictionary describes; must be XML .

The format of the XML representing the metadata is defined as part of a framework called the Extensible Metadata Platform (XMP™) and described in the Adobe document *XMP: Extensible Metadata Platform* (see the Bibliography). This framework provides a way to use XML to represent metadata describing documents and their components, and is intended to be adopted by a wider class of applications than just those that process PDF. It includes a method to embed XML data within non-XML data files in a platform-independent format that can be easily located and accessed by simple scanning rather than requiring the document file to be parsed.

A metadata stream can be attached to a document through the **Metadata** entry in the document catalog (see Chapter 3.6.1, “Document Catalog,” and also see implementation note 104 in Appendix H). In addition, most PDF document components represented as a stream or dictionary can have a **Metadata** entry (see Table 9.4).

TABLE 9.4 Additional entry for components having metadata

KEY	TYPE	VALUE
Metadata	stream	<i>(Optional; PDF 1.4)</i> A metadata stream containing metadata for the component.

In general, a PDF stream or dictionary can have metadata attached to it as long as the stream or dictionary represents an actual information resource, as opposed to serving as an implementation artifact. Table 9.5 lists the PDF constructs that are considered implementational (and hence cannot have associated metadata), and indicates where further information about each one can be found.

TABLE 9.5 PDF constructs that do not take metadata

CONSTRUCT	SECTION
File trailer dictionary	3.4.4
Encryption dictionary	3.5
Document catalog	3.6.1
Page tree node	3.6.2
Name dictionary	3.6.3
Resource dictionary	3.7.2
Name tree node	3.8.4
Number tree node	3.8.5
Function dictionary	3.9
File specification dictionary	3.10.2
Graphics state parameter dictionary	4.3.4
Color space dictionary and DeviceN color space attributes dictionary	4.5
Group attributes dictionary	4.9.2
Font dictionary and CIDFont dictionary	5.4
Encoding dictionary	5.5.5
Font descriptor	5.7
Destination	8.2.1
Action dictionary	8.5.1
Document information dictionary	9.2.1
Structure tree root	9.6.1
Role map	9.6.2
Marked-content reference and object reference dictionary	9.6.3
Class map	9.6.4
Objects belonging to the various Web Capture databases	9.9
Separation dictionary	9.10.3
OPI dictionary	9.10.6

For the remaining PDF constructs, there is sometimes ambiguity about exactly which stream or dictionary should bear the **Metadata** entry. Such cases are to be resolved so that the metadata is attached as close as possible to the object that actually stores the data resource described. For example, metadata describing a tiling pattern should be attached to the pattern stream's dictionary, but a shading should have metadata attached to the shading dictionary itself, rather than to the shading pattern dictionary that refers to it. Similarly, metadata describing an **ICCBased** color space should be attached to the ICC profile stream describing it, and metadata for fonts should be attached to font file streams rather than to font dictionaries.

In tables describing document components in this book, the **Metadata** entry is listed only for those in which it is most likely to be used; keep in mind, however, that this entry may appear in any component represented as a stream or dictionary except those listed in Table 9.5.

In addition, metadata can also be associated with marked content within a content stream. This association is created by including an entry in the property list dictionary whose key is **Metadata** and whose value is the metadata stream dictionary. Because this construct refers to an object outside the content stream, the property list must be referred to indirectly as a named resource (see Section 9.5.1, "Property Lists").

9.3 File Identifiers

PDF files may contain references to other PDF files (see Section 3.10, "File Specifications"). Simply storing a file name, however, even in a platform-independent format, does not guarantee that the file can be found. Even if the file still exists and its name has not been changed, different server software applications may identify it in different ways. For example, servers running on DOS platforms must convert all file names to 8 characters and a 3-character extension; different servers may use different strategies for converting longer file names to this format.

External file references can be made more reliable by including a *file identifier* (*PDF 1.1*) in the file itself and using it in addition to the normal platform-based file designation. Matching the identifier in the file reference with the one in the file itself confirms whether the desired file was found.

File identifiers are defined by the optional **ID** entry in a PDF file’s trailer dictionary (see Section 3.4.4, “File Trailer”; see also implementation note 105 in Appendix H). The value of this entry is an array of two strings. The first string is a permanent identifier based on the contents of the file at the time it was originally created, and does not change when the file is incrementally updated. The second string is a changing identifier based on the file’s contents at the time it was last updated. When a file is first written, both identifiers are set to the same value. If both identifiers match when a file reference is resolved, it is very likely that the correct file has been found; if only the first identifier matches, then a different version of the correct file has been found.

To help ensure the uniqueness of file identifiers, it is recommended that they be computed using a message digest algorithm such as MD5 (described in Internet RFC 1321, *The MD5 Message-Digest Algorithm*; see the Bibliography), using the following information (see implementation note 106 in Appendix H):

- The current time
- A string representation of the file’s location, usually a pathname
- The size of the file in bytes
- The values of all entries in the file’s document information dictionary (see Section 9.2.1, “Document Information Dictionary”)

9.4 Page-Piece Dictionaries

The optional **PieceInfo** entry in a page object or a form dictionary (see Tables 3.18 on page 88 and 4.41 on page 284) can hold a *page-piece dictionary* (PDF 1.3) containing private application data associated with a PDF page or a form XObject. Applications can use this dictionary as a place to store any private data they wish in connection with that page or form. Such private data can convey information meaningful to the application that produces it (such as information on object grouping for a graphics editor or the layer information used by Adobe Photoshop®), but is typically ignored by general-purpose PDF viewer applications.

As Table 9.6 shows, a page-piece dictionary may contain any number of entries, each keyed by the name of a distinct application or of a “well-known” data type recognized by a family of applications. The value associated with each key is an *application data dictionary* containing the private data to be used by the applica-

tion. The **Private** entry may have a value of any data type, but typically it will be a dictionary containing all of the private data needed by the application other than the actual content of the page or form.

TABLE 9.6 Entries in a page-piece dictionary

KEY	TYPE	VALUE
<i>any application name or well-known data type</i>	dictionary	An application data dictionary (see Table 9.7).

TABLE 9.7 Entries in an application data dictionary

KEY	TYPE	VALUE
LastModified	date	<i>(Required)</i> The date and time when the contents of the page or form were most recently modified by this application.
Private	(any)	<i>(Optional)</i> Any private data appropriate to the application, typically in the form of a dictionary.

The **LastModified** entry indicates when this application last altered the content of the page or form. If the page-piece dictionary contains several application data dictionaries, their modification dates can be compared with those in the corresponding entry of the page object or form dictionary (see Tables 3.18 on page 88 and 4.41 on page 284) to ascertain which application data dictionary corresponds to the current content of the page or form. Because some platforms may use only an approximate value for the date and time or may not deal correctly with differing time zones, modification dates are compared only for equality and not for sequential ordering.

*Note: It is possible for two or more application data dictionaries to have the same modification date. Applications can use this capability to define multiple or extended versions of the same data format. For example, suppose that earlier versions of an application use an application data dictionary named **PictureEdit**, while later versions of the same application extend the data to include additional items not previously used. The original data could continue to be kept in the **PictureEdit** dictionary, with the additional items placed in a new dictionary named **PictureEditExtended**. This allows the earlier versions of the application to continue to work as before, while later versions are able to locate and use the extended data items.*

9.5 Marked Content

Marked-content operators (PDF 1.2) identify a portion of a PDF content stream as a *marked-content element* of interest to a particular application or PDF plug-in extension. Marked-content elements and the operators that mark them fall into two categories:

- The **MP** and **DP** operators designate a single *marked-content point* in the content stream.
- The **BMC**, **BDC**, and **EMC** operators bracket a *marked-content sequence* of objects within the content stream. Note that this is not simply a sequence of bytes in the content stream, but of complete graphics objects. Each object is fully qualified by the parameters of the graphics state in which it is rendered.

A graphics application, for example, might use marked content to identify a set of related objects as a “group” to be processed as a single unit. A text-processing application might use it to maintain a connection between a footnote marker in the body of a document and the corresponding footnote text at the bottom of the page. Table 9.8 summarizes the marked-content operators.

All marked-content operators except **EMC** take a *tag* operand indicating the role or significance of the marked-content element to the processing application. All such tags must be registered with Adobe Systems (see Appendix E) to avoid conflicts between different applications marking the same content stream. In addition to the tag operand, the **DP** and **BDC** operators specify a *property list* containing further information associated with the marked content. Property lists are discussed further in Section 9.5.1, “Property Lists.”

Marked-content operators may appear only *between* graphics objects in the content stream; they may not occur within a graphics object or between a graphics state operator and its operands. Marked-content sequences may be nested one within another, but each sequence must be entirely contained within a single content stream; it may not cross page boundaries, for example.

Note: The **Contents** entry of a page object (see “Page Objects” on page 87), which may be either a single stream or an array of streams, is considered a single stream with respect to marked-content sequences.

TABLE 9.8 Marked-content operators

OPERANDS	OPERATOR	DESCRIPTION
<i>tag</i>	MP	Designate a marked-content point. <i>tag</i> is a name object indicating the role or significance of the point.
<i>tag properties</i>	DP	Designate a marked-content point with an associated property list. <i>tag</i> is a name object indicating the role or significance of the point; <i>properties</i> is either an inline dictionary containing the property list or a name object associated with it in the Properties subdictionary of the current resource dictionary (see Section 9.5.1, “Property Lists”).
<i>tag</i>	BMC	Begin a marked-content sequence terminated by a balancing EMC operator. <i>tag</i> is a name object indicating the role or significance of the sequence.
<i>tag properties</i>	BDC	Begin a marked-content sequence with an associated property list, terminated by a balancing EMC operator. <i>tag</i> is a name object indicating the role or significance of the sequence; <i>properties</i> is either an inline dictionary containing the property list or a name object associated with it in the Properties subdictionary of the current resource dictionary (see Section 9.5.1, “Property Lists”).
—	EMC	End a marked-content sequence begun by a BMC or BDC operator.

When the marked-content operators **BMC**, **BDC**, and **EMC** are combined with the text object operators **BT** and **ET** (see Section 5.3, “Text Objects”), each pair of matching operators (**BMC...EMC**, **BDC...EMC**, or **BT...ET**) must be properly (separately) nested. That is, the sequences

BMC		BT
BT		BMC
...	and	...
ET		EMC
EMC		ET

are valid, but

BMC		BT
BT		BMC
...	and	...
EMC		ET
BT		EMC

are not.

9.5.1 Property Lists

The marked-content operators **DP** and **BDC** associate a *property list* with a marked-content element within a content stream. This is a dictionary containing private information meaningful to the program (application or plug-in extension) creating the marked content. Although the dictionary may contain any entries the program wishes to place there, it is suggested that any particular program use the entries in a consistent way; for example, the values associated with a given key should always be of the same type (or small set of types).

If all of the values in a property list dictionary are direct objects, the dictionary may be written inline in the content stream as a direct object. If any of the values are indirect references to objects outside the content stream, the property list dictionary must instead be defined as a named resource in the **Properties** sub-dictionary of the current resource dictionary (see Section 3.7.2, “Resource Dictionaries”) and then referenced by name as the *properties* operand of the **DP** or **BDC** operator.

9.5.2 Marked Content and Clipping

Some PDF path and text objects are defined purely for their effect on the current clipping path, without themselves actually being painted on the page. This occurs when a path object is defined using the operator sequence **W n** or **W* n** (see Section 4.4.3, “Clipping Path Operators”) or when a text object is painted in text rendering mode 7 (see Section 5.2.5, “Text Rendering Mode”). Such clipped, unpainted path or text objects are called *clipping objects*. When a clipping object falls within a marked-content sequence, it is not considered part of the sequence unless the entire sequence consists only of clipping objects. In Example 9.2, for instance, the marked-content sequence tagged Clip includes the text string (Clip me), but not the rectangular path that defines the clipping boundary.

Example 9.2

```
/Clip BMC
  100 100 10 10 re W n          % Clipping path
  (Clip me) Tj                 % Object to be clipped
EMC
```

Only when a marked-content sequence consists entirely of clipping objects are the clipping objects considered part of the sequence. In this case, the sequence is

known as a *marked clipping sequence*. Such sequences may be nested. In Example 9.3, for instance, multiple lines of text are used to clip a subsequent graphics object (in this case, a filled path). Each line of text is bracketed within a separate marked clipping sequence, tagged `Pgf`; the entire series is bracketed in turn by an outer marked clipping sequence, tagged `Clip`. Note, however, that the marked-content sequence tagged `ClippedText` is *not* a marked clipping sequence, since it contains a filled rectangular path that is not a clipping object. The clipping objects belonging to the `Clip` and `Pgf` sequences are therefore not considered part of the `ClippedText` sequence.

Example 9.3

```

/ClippedText BMC
  /Clip <<...>>
    BDC
      BT
        7 Tr                % Begin text clip mode
        /Pgf BMC
          (Line 1) Tj
        EMC
        /Pgf BMC
          (Line) '
          ( 2) Tj
        EMC
      ET                    % Set current text clip
    EMC
    100 100 10 10 re f     % Filled path
  EMC

```

The precise rules governing marked clipping sequences are as follows:

- A *clipping object* is a path object ended by the operator sequence `W n` or `W* n` or a text object painted in text rendering mode 7.
- An *invisible graphics object* is a path object ended by the operator `n` only (with no preceding `W` or `W*`) or a text object painted in text rendering mode 3.
- A *visible graphics object* is a path object ended by any operator other than `n`, a text object painted in any text rendering mode other than 3 or 7, or any `XObject` invoked by the `Do` operator.
- An *empty marked-content element* is a marked-content point or a marked-content sequence that encloses no graphics objects.

- A *marked clipping sequence* is a marked-content sequence that contains at least one clipping object and no visible graphics objects.
- Clipping objects and marked clipping sequences are considered part of an enclosing marked-content sequence only if it is a marked clipping sequence.
- Invisible graphics objects and empty marked-content elements are always considered part of an enclosing marked-content sequence, regardless of whether it is a marked clipping sequence.
- The **q** (save) and **Q** (restore) operators may not occur within a marked clipping sequence.

Example 9.4 illustrates the application of these rules. Marked-content sequence S4 is a marked clipping sequence, because it contains a clipping object (clipping path 2) and no visible graphics objects. Clipping path 2 is therefore considered part of sequence S4. Marked-content sequences S1, S2, and S3 are *not* marked clipping sequences, since they each include at least one visible graphics object. Thus clipping paths 1 and 2 are not part of any of these three sequences.

Example 9.4

```

/S1 BMC
  /S2 BMC
    /S3 BMC
      0 0 m
      100 100 l
      0 100 l W n          % Clipping path 1
      0 0 m
      200 200 l
      0 100 l f          % Filled path
    EMC
  /S4 BMC
    0 0 m
    300 300 l
    0 100 l W n          % Clipping path 2
  EMC
EMC
100 100 10 10 re f      % Filled path
EMC

```

In Example 9.5, marked-content sequence S1 is a marked clipping sequence, because the only graphics object it contains is a clipping path. Thus the empty marked-content sequence S3 and the marked-content point P1 are both part of sequence S2, and S2, S3, and P1 are all part of sequence S1.

Example 9.5

```

/S1 BMC
  ...Clipping path...
/S2 BMC
  /S3 BMC
    EMC
  /P1 DP
    EMC
  EMC
EMC

```

In Example 9.6, marked-content sequences S1 and S4 are marked clipping sequences, because the only object they contain is a clipping path. Hence the clipping path is part of sequences S1 and S4; S3 is part of S2; and S2, S3, and S4 are all part of S1.

Example 9.6

```

/S1 BMC
  /S2 BMC
    /S3 BMC
      EMC
    EMC
  /S4 BMC
    ...Clipping path...
  EMC
EMC

```

9.6 Logical Structure

PDF's *logical structure* facilities (*PDF 1.3*) provide a mechanism for incorporating structural information about a document's content into a PDF file. Such information might include, for example, the organization of the document into chapters and sections or the identification of special elements such as figures,

tables, and footnotes. The logical structure facilities are extensible, allowing applications that produce PDF files to choose what structural information to include and how to represent it, while enabling PDF consumers to navigate a file without knowing the producer's structural conventions.

PDF logical structure shares basic features with standard document markup languages such as HTML, SGML, and XML. A document's logical structure is expressed as a hierarchy of *structure elements*, each represented by a dictionary object. Like their counterparts in other markup languages, PDF structure elements can have content and attributes. Their content can consist of references to document content, references to other structure elements, or both. In PDF, however, rendered document content takes over the role occupied by text in HTML, SGML, and XML.

A PDF document's logical structure is stored separately from its visible content, with pointers from each to the other. This separation allows the ordering and nesting of logical elements to be entirely independent of the order and location of graphics objects on the document's pages.

9.6.1 Structure Hierarchy

The logical structure of a document is described by a hierarchy of objects called the *structure hierarchy* or *structure tree*. At the root of the hierarchy is a dictionary object called the *structure tree root*, located via the **StructTreeRoot** entry in the document catalog (see Section 3.6.1, "Document Catalog"). The remainder of the hierarchy is formed of intermediate nodes called *structure elements*. At the leaves of the tree are individual *content items* associated with structure elements; these may be marked-content sequences, complete PDF objects, or other structure elements.

Tables 9.9 and 9.10 show the contents of the structure tree root and a structure element, respectively.

TABLE 9.9 Entries in the structure tree root

KEY	TYPE	VALUE
Type	name	<i>(Required)</i> The type of PDF object that this dictionary describes; must be StructTreeRoot for a structure tree root.
K	dictionary or array	<i>(Optional)</i> The immediate child or children of the structure tree root in the structure hierarchy. The value may be either a dictionary representing a single structure element or an array of such dictionaries.
IDTree	name tree	<i>(Required if any structure elements have element identifiers)</i> A name tree that maps element identifiers (see Table 9.10) to the structure elements they denote.
ParentTree	number tree	<p><i>(Required if any structure element contains PDF objects or marked-content sequences as content items)</i> A number tree (see Section 3.8.5, “Number Trees”) used in finding the structure elements to which content items belong. Each integer key in the number tree corresponds to a single page of the document or to an individual object (such as an annotation or an XObject) that is a content item in its own right. The integer key is given as the value of the StructParent or StructParents entry in that object (see “Finding Structure Elements from Content Items” on page 600). The form of the associated value depends on the nature of the object:</p> <ul style="list-style-type: none"> • For an object that is a content item in its own right, the value is an indirect reference to the object’s parent element (the structure element that contains it as a content item). • For a page object or content stream containing marked-content sequences that are content items, the value is an array of references to the parent elements of those marked-content sequences. <p>See “Finding Structure Elements from Content Items” on page 600 for further discussion.</p>
ParentTreeNextKey	integer	<i>(Optional)</i> An integer greater than any key in the parent tree, to be used as a key for the next entry added to the tree.
RoleMap	dictionary	<i>(Optional)</i> A dictionary mapping the names of structure types used in the document to their approximate equivalents in the set of standard structure types (see Section 9.7.4, “Standard Structure Types”).
ClassMap	dictionary	<i>(Optional)</i> A dictionary mapping name objects designating attribute classes to the corresponding attribute objects or arrays of attribute objects (see “Attribute Classes” on page 605).

TABLE 9.10 Entries in a structure element dictionary

KEY	TYPE	VALUE
Type	name	<i>(Optional)</i> The type of PDF object that this dictionary describes; if present, must be StructElem for a structure element.
S	name	<i>(Required)</i> The <i>structure type</i> , a name object identifying the nature of the structure element and its role within the document, such as a chapter, paragraph, or footnote (see Section 9.6.2, “Structure Types”). Names of structure types must conform to the guidelines described in Appendix E.
P	dictionary	<i>(Required; must be an indirect reference)</i> The structure element that is the immediate parent of this one in the structure hierarchy.
ID	string	<i>(Optional)</i> The <i>element identifier</i> , a string designating this structure element. The string must be unique among all elements in the document’s structure hierarchy. The IDTree entry in the structure tree root (see Table 9.9) defines the correspondence between element identifiers and the structure elements they denote.
Pg	dictionary	<i>(Optional; must be an indirect reference)</i> A page object representing a page on which some or all of the content items designated by the K entry are rendered.
K	(various)	<p><i>(Optional)</i> The contents of this structure element, which may consist of one or more marked-content sequences, PDF objects, and other structure elements. The value of this entry may be any of the following:</p> <ul style="list-style-type: none"> • An integer marked-content identifier denoting a marked-content sequence • A marked-content reference dictionary denoting a marked-content sequence • An object reference dictionary denoting a PDF object • A structure element dictionary denoting another structure element • An array, each of whose elements is one of the objects listed above <p>See Section 9.6.3, “Structure Content” for further discussion of each of these forms of representation.</p>
A	(various)	<i>(Optional)</i> The attribute object or objects, if any, associated with this structure element. Each attribute object is either a dictionary or a stream; the value of this entry may be either a single attribute object or an array of such objects together with their revision numbers (see Section 9.6.4, “Structure Attributes,” and “Attribute Revision Numbers” on page 606).

C	name or array	<p>(<i>Optional</i>) The attribute class or classes, if any, to which this structure element belongs. The value of this entry may be either a single class name or an array of class names together with their revision numbers (see “Attribute Classes” on page 605 and “Attribute Revision Numbers” on page 606).</p> <p>Note: If both the A and C entries are present and a given attribute is specified by both, the one specified by the A entry takes precedence.</p>
R	integer	<p>(<i>Optional</i>) The current revision number of this structure element (see “Attribute Revision Numbers” on page 606). The value must be a non-negative integer. Default value: 0.</p>
T	text string	<p>(<i>Optional</i>) The title of the structure element, a text string representing it in human-readable form. The title should characterize the specific structure element, such as Chapter 1, rather than merely a generic element type, such as Chapter.</p>
Lang	text string	<p>(<i>Optional; PDF 1.4</i>) A <i>language identifier</i> specifying the natural language for all text in the structure element except where overridden by language specifications for nested structure elements or marked content (see Section 9.8.1, “Natural Language Specification”). If this entry is absent, the language (if any) specified in the document catalog applies.</p>
Alt	text string	<p>(<i>Optional</i>) An alternate description of the structure element and its children in human-readable form, useful when extracting the document’s contents in support of accessibility to disabled users or for other purposes (see Section 9.8.2, “Alternate Descriptions”).</p>
ActualText	text string	<p>(<i>Optional; PDF 1.4</i>) Text that is an exact replacement for the structure element and its children. This replacement text (which should apply to as small a piece of content as possible) is useful when extracting the document’s contents in support of accessibility to disabled users or for other purposes (see Section 9.8.3, “Replacement Text”).</p>

9.6.2 Structure Types

Every structure element has a *structure type*, a name object that identifies the nature of the structure element and its role within the document (such as a chapter, paragraph, or footnote). To facilitate the interchange of content among PDF applications, Adobe has defined a set of standard structure types; see Section 9.7.4, “Standard Structure Types.” Applications are not required to adopt them, however, but may use any names they wish for their structure types.

Where names other than the standard ones are used, a *role map* may be provided in the structure tree root, mapping the structure types used in the document to their nearest equivalents in the standard set. For example, a structure type named `Section` used in the document might be mapped to the standard type `Sect`. The equivalence need not be exact; the role map merely indicates an approximate analogy between types, allowing applications other than the one creating a document to handle its nonstandard structure elements in a reasonable way.

Note: The same structure type may occur as both a key and a value in the role map, and circular chains of association are explicitly permitted. A single role map can thus define a bidirectional mapping. An application using the role map should follow the chain of associations until it either finds a structure type it recognizes or returns to one it has already encountered.

9.6.3 Structure Content

The content of a structure element may consist of one or more *content items* of any of the following kinds:

- Marked-content sequences embedded within content streams
- Complete PDF objects such as annotations and XObjects
- Other structure elements

The **K** entry in a structure element dictionary can have as its value either a single object denoting one of these items or an array of such objects. Items of any or all three kinds may be mixed in the same content array. The following sections describe how each kind of content item is denoted in a structure element's **K** entry.

Marked-Content Sequences as Content Items

For a sequence of graphics operators in a content stream to be included in the content of a structure element, they must be bracketed as a marked-content sequence between **BDC** and **EMC** operators (see Section 9.5, “Marked Content”). Furthermore, the marked-content sequence must have a property list (see Section 9.5.1, “Property Lists”) containing an **MCID** entry. This entry defines an integer *marked-content identifier* that uniquely identifies the marked-content sequence within its content stream. The structure element can then refer to the

sequence by specifying its content stream and its marked-content identifier within the stream.

Note: *Although the tag associated with a marked-content sequence is not directly related to the document’s logical structure, it should be the same as the structure type of the associated structure element.*

In the general case, a structure element refers to a marked-content sequence by means of a dictionary object called a *marked-content reference*. Table 9.11 shows the contents of this type of dictionary.

TABLE 9.11 Entries in a marked-content reference dictionary

KEY	TYPE	VALUE
Type	name	<i>(Required)</i> The type of PDF object that this dictionary describes; must be MCR for a marked-content reference.
Pg	dictionary	<i>(Optional; must be an indirect reference)</i> The page object representing the page on which the graphics objects in the marked-content sequence are rendered. This entry overrides any Pg entry in the structure element containing the marked-content reference; it is required if the structure element has no such entry.
Stm	stream	<i>(Optional; must be an indirect reference)</i> The content stream containing the marked-content sequence. This entry is needed only if the marked-content sequence resides in some other content stream associated with the page—for example, in a form XObject (see Section 4.9, “Form XObjects”) or an annotation’s appearance stream (Section 8.4.4, “Appearance Streams”). Default value: the content stream of the page identified by Pg .
StmOwn	(any)	<i>(Optional; must be an indirect reference)</i> The PDF object owning the stream identified by Stm —for example, the annotation to which an appearance stream belongs.
MCID	integer	<i>(Required)</i> The marked-content identifier of the marked-content sequence within its content stream.

Example 9.7 illustrates the use of a marked-content reference to refer to a marked-content sequence within the content stream of a page.

Note: *This and the following examples omit required **StructParents** entries in the objects used as content items (see “Finding Structure Elements from Content Items” on page 600).*

Example 9.7

```

1 0 obj                                % Structure element
  << /Type /StructElem
    /S /P                                % Structure type
    /P ...                                % Parent in structure hierarchy
    /K << /Type /MCR
      /Pg 2 0 R                            % Page containing marked-content sequence
      /MCID 0                              % Marked-content identifier
    >>
  >>
endobj

2 0 obj                                % Page object
  << /Type /Page
    /Contents 3 0 R                        % Content stream
  ...
  >>
endobj

3 0 obj                                % Page's content stream
  << /Length ... >>
stream
  ...
  /P << /MCID 0 >>                        % Start of marked-content sequence
  BDC
  ...
  (Here is some text) Tj
  ...
  EMC                                      % End of marked-content sequence
  ...
endstream
endobj

```

Marked-content sequences can be used to incorporate the content of a form XObject into a structure element in either of two ways:

- In the content stream that paints the form XObject with the **Do** operator, enclose the invocation of **Do** within a marked-content sequence and associate that sequence with the desired structure element (see Example 9.8). The entire content of the form XObject is considered to be part of the structure element's content, as if it were inserted into the marked-content sequence at the point of the **Do** operator. The form XObject has no internal substructure of its own and cannot in turn contain any marked-content sequences associated with this or other structure elements.

- In the content stream of the form XObject itself, enclose the desired content in a marked-content sequence associated with the structure element (see Example 9.9). The **Do** operator that paints the form XObject should *not* be part of any logical structure content item. However, in this case the form XObject itself can have arbitrary substructure, containing any number of marked-content sequences associated with logical structure elements.

Note: A form XObject that is painted with multiple invocations of the **Do** operator can be incorporated into the document's logical structure only by the first method, with each invocation of **Do** individually associated with a structure element.

Example 9.8

```

1 0 obj                                % Structure element
  << /Type /StructElem
    /S /P                                % Structure type
    /P ...                               % Parent in structure hierarchy
    /Pg 2 0 R                             % Page containing marked-content sequence
    /K 0                                   % Marked-content identifier
  >>
endobj

2 0 obj                                % Page object
  << /Type /Page
    /Resources << /XObject << /Fm4 4 0 R >> % Resource dictionary
                                     >> % containing form XObject
    /Contents 3 0 R                       % Content stream
    ...
  >>
endobj

3 0 obj                                % Page's content stream
  << /Length ... >>
stream
  ...
  /P << /MCID 0 >>                       % Start of marked-content sequence
  BDC
    /Fm4 Do                               % Paint form XObject
  EMC                                     % End of marked-content sequence
  ...
endstream
endobj

```

```

4 0 obj                                % FormXObject
  << /Type /XObject
    /Subtype /Form
    /Length ...
  >>
stream
...
(Here is some text) Tj
...
endstream
endobj

```

Example 9.9

```

1 0 obj                                % Structure element
  << /Type /StructElem
    /S /P                                % Structure type
    /P ...                                % Parent in structure hierarchy
    /K << /Type /MCR
      /Pg 2 0 R                            % Page containing marked-content sequence
      /Stm 4 0 R                            % Stream containing marked-content sequence
      /MCID 0                              % Marked-content identifier
    >>
  >>
endobj

2 0 obj                                % Page object
  << /Type /Page
    /Resources << /XObject << /Fm4 4 0 R >> % Resource dictionary
      >>                                    % containing form XObject
    /Contents 3 0 R                            % Content stream
  >>
endobj

3 0 obj                                % Page's content stream
  << /Length ... >>
stream
...
  /Fm4 Do                                    % Paint form XObject
...
endstream
endobj

```

```

4 0 obj                                % Form XObject
  << /Type /XObject
    /Subtype /Form
    /Length ...
  >>
stream
...
/P << /MCID 0 >>                       % Start of marked-content sequence
  BDC
  ...
  (Here is some text) Tj
  ...
  EMC                                    % End of marked-content sequence
...
endstream
endobj

```

In the common case where all or most of the marked-content sequences in a structure element's content are on the same page, the **Pg** entry identifying the page can be supplied in the structure element dictionary itself (see Table 9.10 on page 591), rather than in a separate marked-content reference for each sequence. Often, this allows the marked-content reference to be dispensed with entirely and the sequence identified in the structure element's **K** entry by simply giving its integer marked-content identifier directly, as in Example 9.8 above.

PDF Objects as Content Items

When a structure element's content includes an entire PDF object, such as an XObject or an annotation, that is associated with a page but not directly included in the page's content stream, the object is identified in the structure element's **K** entry by an *object reference dictionary* (see Table 9.12). Note that this form of reference is used only for entire objects; if the referenced content forms only part of the object's content stream, it is instead handled as a marked-content sequence, as described in the preceding section.

TABLE 9.12 Entries in an object reference dictionary

KEY	TYPE	VALUE
Type	name	<i>(Required)</i> The type of PDF object that this dictionary describes; must be OBJR for an object reference.
Pg	dictionary	<i>(Optional; must be an indirect reference)</i> The page object representing the page on which the object is rendered. This entry overrides any Pg entry in the structure element containing the object reference; it is required if the structure element has no such entry.
Obj	(any)	<i>(Required; must be an indirect reference)</i> The referenced object.

Note: *If the referenced object is rendered on multiple pages, each rendering requires a separate object reference; but if it is rendered multiple times on the same page, just a single object reference suffices to identify all of them. (If it is important to distinguish between multiple renditions of the same XObject on the same page, they should be accessed via marked-content sequences enclosing particular invocations of the **Do** operator, rather than via object references.)*

Structure Elements as Content Items

One structure element can contain another as a content item simply by referring to the other structure element in its **K** entry. Example 9.10 shows a structure element containing other structure elements as content items, along with a marked-content sequence from a page's content stream.

Example 9.10

```

1 0 obj                                % Containing structure element
  << /Type /StructElem
    /S /MixedContainer                 % Structure type
    /P ...                             % Parent in structure hierarchy
    /Pg 2 0 R                          % Page containing marked-content sequence
    /K [ 4 0 R                          % Three content items: a structure element
        0                               % a marked-content sequence
        5 0 R                          % another structure element
    ]
  >>
endobj

```

```

2 0 obj                                % Page object
  << /Type /Page
    /Contents 3 0 R                    % Content stream
    ...
  >>
endobj

3 0 obj                                % Page's content stream
  << /Length ... >>
stream
  ...
  /P << /MCID 0 >>                    % Start of marked-content sequence
  BDC
    (Here is some text) Tj
  ...
  EMC                                  % End of marked-content sequence
  ...
endstream
endobj

4 0 obj                                % First contained structure element
  << /Type /StructElem
  ...
  >>
endobj

5 0 obj                                % Second contained structure element
  << /Type /StructElem
  ...
  >>
endobj

```

Finding Structure Elements from Content Items

Because a stream cannot contain object references, there is no way for content items that are marked-content sequences to refer directly back to their parent structure elements (the ones to which they belong as content items). Instead, a different mechanism, the *structural parent tree*, is provided for this purpose. For consistency, content items that are entire PDF objects, such as XObjects, also use the parent tree to refer to their parent structure elements.

The parent tree is a number tree (see Section 3.8.5, “Number Trees”), accessed via the **ParentTree** entry in a document’s structure tree root (Table 9.9 on page 590). The tree contains an entry for each object that is a content item of at least one structure element and for each content stream containing at least one marked-content sequence that is a content item. The key for each entry is an integer given as the value of the **StructParent** or **StructParents** entry in the object (see below). The values of these entries are as follows:

- For an object identified as a content item by means of an object reference (see “PDF Objects as Content Items” on page 598), the value is an indirect reference to the parent structure element.
- For a content stream containing marked-content sequences that are content items, the value is an array of indirect references to the sequences’ parent structure elements. The array element corresponding to each sequence is found by using the sequence’s marked-content identifier as a zero-based index into the array.

***Note:** Because marked-content identifiers serve as indices into an array in the structural parent tree, their assigned values should be as small as possible to conserve space in the array.*

The **ParentTreeNextKey** entry in the structure tree root holds an integer value greater than any that is currently in use as a key in the structural parent tree. Whenever a new entry is added to the parent tree, it uses the current value of **ParentTreeNextKey** as its key; the value is then incremented to prepare for the next new entry to be added.

To locate the relevant parent tree entry, each object or content stream that is represented in the tree must contain a special dictionary entry, **StructParent** or **StructParents** (see Table 9.13). Depending on the type of content item, this entry may appear in the page object of a page containing marked-content sequences, in the stream dictionary of a form or image XObject, in an annotation dictionary, or in any other type of object dictionary that is included as a content item in a structure element. Its value is the integer key under which the entry corresponding to the object is to be found in the structural parent tree.

TABLE 9.13 Additional dictionary entries for structure element access

KEY	TYPE	VALUE
StructParent	integer	<i>(Required for all objects that are structural content items; PDF 1.3)</i> The integer key of this object's entry in the structural parent tree.
StructParents	integer	<i>(Required for all content streams containing marked-content sequences that are structural content items; PDF 1.3)</i> The integer key of this object's entry in the structural parent tree.

Note: *At most one of these two entries may be present in a given object. An object can be either a content item in its entirety or a container for marked-content sequences that are content items, but not both.*

For a content item identified by an object reference, the parent structure element can thus be found by using the value of the **StructParent** entry in the item's object dictionary as a retrieval key in the structural parent tree (found in the **ParentTree** entry of the structure tree root). The corresponding value retrieved from the parent tree is a reference to the parent structure element (see Example 9.11).

Example 9.11

```

1 0 obj                                % Parent structure element
  << /Type /StructElem
  ...
  /K << /Type /OBJR                    % Object reference
    /Pg 2 0 R                          % Page containing formXObject
    /Obj 4 0 R                          % Reference to formXObject
  >>
>>
endobj

2 0 obj                                % Page object
  << /Type /Page
    /Resources << /XObject << /Fm4 4 0 R >> % Resource dictionary
      >>                                % containing formXObject
    /Contents 3 0 R                    % Content stream
  ...
  >>
endobj

```

```

3 0 obj                                % Page's content stream
  << /Length ... >>
stream
  ...
  /Fm4 Do                               % Paint formXObject
  ...
endstream
endobj

4 0 obj                                % Form XObject
  << /Type /XObject
    /Subtype /Form
    /Length ...
    /StructParent 6                     % Parent tree key
  >>
stream
  ...
endstream
endobj

100 0 obj                               % Parent tree (accessed from structure tree root)
  << /Nums [ 0 101 0 R
            1 102 0 R
            ...
            6 1 0 R                     % Entry for page object 2; points back
            ...                          % to parent structure element
          ]
  >>
endobj

```

For a content item that is a marked-content sequence, the retrieval method is similar but slightly more complicated. Because a marked-content sequence is not an object in its own right, its parent tree key is found in the **StructParents** entry of the page object or other content stream in which the sequence resides. The value retrieved from the parent tree is not a reference to the parent structure element itself, but rather an array of such references—one for each marked-content sequence contained within that content stream. The parent structure element for the given sequence is found by using the sequence's marked-content identifier as an index into this array (see Example 9.12).

Example 9.12

```

1 0 obj                                % Parent structure element
  << /Type /StructElem
    ...
    /Pg 2 0 R                            % Page containing marked-content sequence
    /K 0                                  % Marked-content identifier
  >>
endobj

2 0 obj                                % Page object
  << /Type /Page
    /Contents 3 0 R                       % Content stream
    /StructParents 6                      % Parent tree key
    ...
  >>
endobj

3 0 obj                                % Page's content stream
  << /Length ... >>
stream
  ...
  /P << /MCID 0 >>                        % Start of marked-content sequence
  BDC
    (Here is some text) TJ
  ...
  EMC                                     % End of marked-content sequence
  ...
endstream
endobj

100 0 obj                               % Parent tree (accessed from structure tree root)
  << /Nums [ 0 101 0 R
    1 102 0 R
    ...
    6 [1 0 R]                             % Entry for page object 2; array element at index 0
    ...                                    % points back to parent structure element
  ]
  >>
endobj

```

9.6.4 Structure Attributes

An application or plug-in extension that processes logical structure can attach additional information, called *attributes*, to any structure element. The attribute information is held in an *attribute object* associated with the structure element. Any dictionary or stream can serve as an attribute object by including an **O** entry (see Table 9.14) identifying the application or plug-in that owns the attribute information; the owner can then add any additional entries it wishes to the object to hold the attributes. To facilitate the interchange of content among PDF applications, Adobe has defined a set of standard structure attributes, identified by specific standard owners; see Section 9.7.5, “Standard Structure Attributes.”

TABLE 9.14 Entry common to all attribute objects

KEY	TYPE	VALUE
O	name	<i>(Required)</i> The name of the application or plug-in extension owning the attribute data. The name must conform to the guidelines described in Appendix E.

Any application can attach attributes to any structure element, even one created by another application. Multiple applications can attach attributes to the same structure element; the **A** entry in the structure element dictionary (see Table 9.10 on page 591) can hold either a single attribute object or an array of such objects, together with *revision numbers* for coordinating attributes created by different owner (see “Attribute Revision Numbers” on page 606). An application creating or destroying the second attribute object for a structure element is responsible for converting the value of the **A** entry from a single object to an array or vice versa, as well as for maintaining the integrity of the revision numbers. No inherent order is defined for the attribute objects in an **A** array, but it is considered good practice to add new objects at the end of the array so that the first array element is the one belonging to the application that originally created the structure element.

Attribute Classes

If many structure elements share the same set of attribute values, they can be defined as an *attribute class* sharing the identical attribute object. Structure elements refer to the class by name; the association between class names and

attribute objects is defined by a dictionary called the *class map*, kept in the **ClassMap** entry of the structure tree root (see Table 9.9 on page 590). Each key in the class map is a name object denoting the name of a class; the corresponding value is an attribute object or an array of such objects.

Note: Despite the name, PDF attribute classes are unrelated to the concept of a class in object-oriented programming languages such as Java™ and C++. Attribute classes are strictly a mechanism for storing attribute information in a more compact form; they have no inheritance properties like those of true object-oriented classes.

The **C** entry in a structure element dictionary (see Table 9.10 on page 591) contains a class name or an array of class names (typically accompanied by revision numbers as well; see “Attribute Revision Numbers,” below). For each class named in the **C** entry, the corresponding attribute object or objects are considered to be attached to the given structure element along with those identified in the element’s **A** entry. (If both the **A** and **C** entries are present and a given attribute is specified by both, the one specified by the **A** entry takes precedence.)

Attribute Revision Numbers

When an application modifies a structure element or its contents, the change may affect the validity of attribute information attached to that structure element by other applications. A system of *revision numbers* allows applications to detect such changes made by other applications and update their own attribute information accordingly.

A structure element’s revision number is kept in the **R** entry in the structure element dictionary (see Table 9.10 on page 591). Initially, the revision number is 0 (the default value if no **R** entry is present); each time an application modifies the structure element or any of its content items, it must signal the change by incrementing the revision number. Note that the revision number is not the same thing as the generation number associated with an indirect object (see Section 3.2.9, “Indirect Objects”); the two are completely separate concepts.

As described below, each attribute object attached to a structure element carries an associated revision number. Each time an attribute object is created or modified, its revision number is set equal to the current value of the structure element’s **R** entry. By comparing the attribute object’s revision number with that of the structure element, an application can tell whether the contents of the

attribute object are still current or whether they have been outdated by more recent changes in the underlying structure element.

Because a single attribute object may be associated with more than one structure element (whose revision numbers may differ), the revision number is not stored directly in the attribute object itself, but rather in the array that associates the attribute object with the structure element. This allows the same attribute object to carry different revision numbers with respect to different structure elements. In the general case, each attribute object in a structure element's **A** array is represented by a pair of array elements, the first containing the attribute object itself and the second the integer revision number associated with it in this structure element. Similarly, the structure element's **C** array contains a pair of elements for each attribute class, the first containing the class name and the second the associated revision number. (Revision numbers equal to 0 can be omitted from both the **A** and **C** arrays; an attribute object or class name that is not followed by an integer array element is understood to have a revision number of 0.)

Note: A structure element's revision number changes only when the structure element itself or any of its content items is modified. Changes in an attached attribute object do not change the structure element's revision number (though they may cause the attribute object's revision number to be updated to match it).

Occasionally, an application may make such extensive changes to a structure element that they are likely to invalidate all previous attribute information associated with it. In this case, instead of incrementing the structure element's revision number, the application may choose to delete all unknown attribute objects from its **A** and **C** arrays. These two actions are mutually exclusive: the application should *either* increment the structure element's revision number *or* remove its attribute objects, but not both. Note that any application creating attribute objects must be prepared for the possibility that they may be deleted at any time by another application.

9.6.5 Example of Logical Structure

Example 9.13 shows portions of a PDF file with a simple document structure. The structure tree root (object 300) contains elements with structure types **Chap** (object 301) and **Para** (object 304). The **Chap** element, titled Chapter 1, contains elements with types **Head1** (object 302) and **Para** (object 303). The example also

illustrates the structure of a parent tree (object 400) mapping content items back to their parent structure elements, and an ID tree (object 403) mapping element identifiers to the structure elements they denote.

Example 9.13

```

1 0 obj                                % Document catalog
  << /Type /Catalog
    /Pages 100 0 R                      % Page tree
    /StructTreeRoot 300 0 R             % Structure tree root
  >>
endobj

100 0 obj                               % Page tree
  << /Type /Pages
    /Kids [ 101 1 R                    % First page object
            102 0 R                    % Second page object
          ]
    /Count 2                            % Page count
  >>
endobj

101 1 obj                               % First page object
  << /Type /Page
    /Parent 100 0 R                    % Parent is the page tree
    /Resources << /Font << /F1 6 0 R    % Font resources
                /F12 7 0 R
            >>
    /ProcSet [/PDF /Text] % Procedure sets
  >>
  /MediaBox [0 0 612 792]              % Media box
  /Contents 201 0 R                    % Content stream
  /StructParents 0                     % Parent tree key
  >>
endobj

201 0 obj                               % Content stream for first page
  << /Length ... >>
stream
  1 1 1 rg
  0 0 612 792 re f
  BT
  % Start of text object

```



```

/Para << /MCID 0 >>          % Start of marked-content sequence 0
  BDC
    0 0 0 rg
    /F12 1 Tf
    14 0 0 14 18 732 Tm
    (sentence.This is the very last sentence of the first paragraph.) Tj
  EMC                          % End of marked-content sequence 0

/Para << /MCID 1 >>          % Start of marked-content sequence 1
  BDC
    /F12 1 Tf
    14 0 0 14 18 570.8 Tm
    (This is the second paragraph.It has four fairly short and concise sentences.\
This is the next to last ) Tj
  EMC                          % End of marked-content sequence 1

/Para << /MCID 2 >>          % Start of marked-content sequence 2
  BDC
    1.1429 TL
    T*
    (sentence.This is the very last sentence of the second paragraph.) Tj
  EMC                          % End of marked-content sequence 2

ET                              % End of text object
endstream
endobj

300 0 obj                      % Structure tree root
<< /Type /StructTreeRoot
  /K [ 301 0 R
    304 0 R
  ]
  /RoleMap << /Chap /Chapter
    /Head1 /H
    /Para /P
  >>
  /ParentTree 400 0 R          % Number tree for parent elements
  /ParentTreeNextKey 2        % Next key to use in parent tree
  /IDTree 403 0 R            % Name tree for element identifiers
>>
endobj

```

```

301 0 obj                                % Structure element for a chapter
  << /Type /StructElem
    /S /Chap
    /ID (Chap1)                          % Element identifier
    /T (Chapter 1)                       % Human-readable title
    /P 300 0 R                            % Parent is the structure tree root
    /K [ 302 0 R                          % Two children: a section head
        303 0 R                            %   and a paragraph
      ]
  >>
endobj

302 0 obj                                % Structure element for a section head
  << /Type /StructElem
    /S /Head1
    /ID (Sec1.1)                          % Element identifier
    /T (Section 1.1)                     % Human-readable title
    /P 301 0 R                            % Parent is the chapter
    /Pg 101 1 R                           % Page containing content items
    /K 0                                   % Marked-content sequence 0
  >>
endobj

303 0 obj                                % Structure element for a paragraph
  << /Type /StructElem
    /S /Para
    /ID (Para1)                          % Element identifier
    /P 301 0 R                            % Parent is the chapter
    /Pg 101 1 R                           % Page containing first content item
    /K [ 1                                 % Marked-content sequence 1
      << /Type /MCR                        % Marked-content reference to second item
        /Pg 102 0 R                       % Page containing second item
        /MCID 0                           % Marked-content sequence 0
      >>
    ]
  >>
endobj

304 0 obj                                % Structure element for another paragraph
  << /Type /StructElem
    /S /Para
    /ID (Para2)                          % Element identifier
    /P 300 0 R                            % Parent is the structure tree root
    /Pg 102 0 R                           % Page containing content items
    /K [1 2]                              % Marked-content sequences 1 and 2
  >>
endobj

```

```

400 0 obj                                % Parent tree
  << /Nums [ 0 401 0 R                    % Parent elements for first page
            1 402 0 R                    % Parent elements for second page
          ]
  >>
endobj

401 0 obj                                % Array of parent elements for first page
  [ 302 0 R                                % Parent of marked-content sequence 0
    303 0 R                                % Parent of marked-content sequence 1
  ]
endobj

402 0 obj                                % Array of parent elements for second page
  [ 303 0 R                                % Parent of marked-content sequence 0
    304 0 R                                % Parent of marked-content sequence 1
    304 0 R                                % Parent of marked-content sequence 2
  ]
endobj

403 0 obj                                % ID tree root node
  << /Kids [404 0 R] >>                  % Reference to leaf node
endobj

404 0 obj                                % ID tree leaf node
  << /Limits [ (Chap1) (Sec1.3) ]         % Least and greatest keys in tree
    /Names [ (Chap1) 301 0 R             % Mapping from element identifiers
              (Sec1.1) 302 0 R           %   to structure elements
              (Sec1.2) 303 0 R
              (Sec1.3) 304 0 R
            ]
  >>
endobj

```

9.7 Tagged PDF

Tagged PDF (PDF 1.4) is a stylized use of PDF that builds on the logical structure framework described in Section 9.6, “Logical Structure,” by defining a set of standard structure types and attributes that allow page content (text, graphics, and images) to be extracted and reused for other purposes. In addition, it defines a set of rules for representing text in the page content so that characters, words, and

text order can be determined reliably. It is intended for use by tools that perform operations such as:

- Simple extraction of text and graphics for pasting into other applications
- Automatic reflow of text and associated graphics to fit a page of a different size than was assumed for the original layout
- Processing text for such purposes as searching, indexing, and spell-checking
- Conversion to other common file formats (such as HTML, XML, and RTF) with document structure and basic styling information preserved
- Making content accessible to the visually impaired

Tagged PDF encompasses three interrelated sets of conventions:

- *Page content.* All text is represented in a form that can be converted to Unicode. Word breaks are represented explicitly. Actual content is distinguished from artifacts of layout and pagination. Content is given in an order related to its appearance on the page, as determined by the authoring application.
- *Structure types.* A set of standard structure types define the meaning of structure elements such as paragraphs, headings, articles, and tables.
- *Structure attributes.* Standard structure attributes preserve styling information used by the authoring application in laying out content on the page.

Note: Tagged PDF is not intended to limit the use of logical structure to only the specific types and attributes described here, but rather to provide a set of standard fallback roles and minimum guaranteed attributes in order to enable consumer applications to perform operations such as those mentioned above. Producer applications are free to define additional structure types, so long as they also provide a role mapping to the nearest equivalent standard types, as described in Section 9.6.2, “Structure Types.” Likewise, producer applications are free to define additional structure attributes using any of the available extension mechanisms.

9.7.1 Mark Information Dictionary

The optional **MarkInfo** entry in the document catalog (see Section 3.6.1, “Document Catalog”) holds a *mark information dictionary* describing a document’s usage of Tagged PDF. As shown in Table 9.15, this dictionary contains a single entry, **Marked**, whose value is a boolean flag indicating whether the document

conforms to Tagged PDF conventions. For a document to be recognized as a Tagged PDF document, its catalog must contain such a dictionary with the value of the **Marked** flag set to **true**.

TABLE 9.15 Entry in the mark information dictionary

KEY	TYPE	VALUE
Marked	boolean	(<i>Optional</i>) A flag indicating whether the document conforms to Tagged PDF conventions. Default value: false .

9.7.2 Tagged PDF and Page Content

Like all PDF documents, a Tagged PDF document consists of a sequence of self-contained pages, each of which is described by one or more page content streams (including any subsidiary streams such as form XObjects and annotation appearances). Tagged PDF defines some further conventions for organizing and marking content streams so that additional information can be derived from them. These conventions include:

- Distinguishing between the author's original content and artifacts of the layout process
- Specifying a content order to guide the layout process if the page content must be reflowed
- Representing text in a form from which a Unicode representation and information about font characteristics can be unambiguously derived
- Representing word breaks unambiguously
- Marking text with information for making it accessible to the visually impaired

Real Content and Artifacts

The graphics objects in a page's content stream can be divided into two classes: *real content* and *artifacts*. The document's logical structure encompasses all graphics objects making up the real content and describes how those objects relate to one another; it does not include graphics objects that are mere artifacts of the layout and production process.

The *real content* of a document comprises objects representing material originally introduced by the document's author. This includes not only the page content stream and subsidiary form XObjects, but also any associated annotations meeting all of the following conditions:

- The annotation has an appearance stream (see Section 8.4.4, "Appearance Streams") containing a normal (**N**) appearance.
- The annotation's Hidden flag (see Section 8.4.2, "Annotation Flags") is not set.
- The annotation is included in the document's logical structure (see Section 9.6, "Logical Structure").

Artifacts are graphics objects that are not part of the author's original content but rather are generated by the PDF producer application in the course of pagination, layout, or other strictly mechanical processes. They include:

- *Pagination artifacts*. Ancillary page features such as running heads and folios (page numbers).
- *Layout artifacts*. Purely cosmetic typographical or design elements such as footnote rules or background screens.
- *Page artifacts*. Production aids extraneous to the document itself, such as cut marks and color bars.

Tagged PDF consumer applications may have their own ideas about what page content to consider relevant. A text-to-speech engine, for instance, probably should not speak running heads or page numbers when the page is turned. In general, consumer applications are allowed to:

- Disregard elements of page content (for example, specific types of artifact) that are not of interest
- Treat some page elements as *terminals* that are not to be examined further (for example, to treat an illustration as a unit for reflow purposes)
- Replace an element with alternate text (see Section 9.8.2, "Alternate Descriptions")

Depending on their specific goals, different consumer applications may make different decisions in this regard. The purpose of Tagged PDF is not to prescribe what the consumer application should do, but to provide sufficient declarative

and descriptive information to allow it to make its own appropriate choices about how to process the content.

Note: To support consumer applications in providing accessibility to disabled users, Tagged PDF documents should use the natural language specification (**Lang**), alternate description (**Alt**), replacement text (**ActualText**), and abbreviation expansion text (**E**) facilities described in Section 9.8, “Accessibility Support.”

Specification of Artifacts

An artifact is distinguished from real content by enclosing it in a marked-content sequence with the tag `Artifact`:

<code>/Artifact</code>		<code>/Artifact <i>propertyList</i></code>
<code>BMC</code>		<code>BDC</code>
<code>...</code>	or	<code>...</code>
<code>EMC</code>		<code>EMC</code>

The first form is used to identify a generic artifact, the second for those that have an associated property list. Table 9.16 shows the properties that can be included in such a property list.

TABLE 9.16 Property list entries for artifacts

KEY	TYPE	VALUE
Type	name	<i>(Optional)</i> The type of artifact that this property list describes; if present, must be one of the names Pagination , Layout , or Page .
BBox	rectangle	<i>(Optional)</i> An array of four numbers in default user space units giving the coordinates of the left, bottom, right, and top edges, respectively, of the artifact’s bounding box (the rectangle that completely encloses its visible extent).
Attached	array	<i>(Optional; pagination artifacts only)</i> An array of name objects containing one to four of the names Top , Bottom , Left , and Right , specifying the edges of the page, if any, to which the artifact is logically attached. Page edges are defined by the page’s crop box (see Section 9.10.1, “Page Boundaries”). The ordering of names within the array is immaterial. Including both Left and Right or both Top and Bottom indicates a full-width or full-height artifact, respectively.

Note: To aid in text reflow, it is recommended that artifacts be defined with property lists whenever possible. Artifacts lacking a specified bounding box are likely to be discarded during reflow.

Incidental Artifacts

In addition to objects that are explicitly marked as artifacts, the running text of a page may contain other elements and relationships that are not logically part of the document's real content, but merely incidental results of the process of laying out that content into a document. Since these incidental artifacts represent variations between the document's logical content and its visible representation, they are not reflected in its logical structure within the Tagged PDF file. They may include the following:

- *Hyphenation.* Among the artifacts introduced by text layout is the hyphen marking the incidental division of a word at the end of a line. In Tagged PDF, such an incidental word division must be represented by a *soft hyphen* character, which the Unicode mapping algorithm (see “Unicode Mapping” on page 620) translates to the Unicode value U+00AD. (This is distinct from an ordinary *hard hyphen*, whose Unicode value is U+002D.) The producer of a Tagged PDF document must distinguish explicitly between soft and hard hyphens, so that the consumer does not have to guess which type a given character represents.
- *Text discontinuities.* The running text of a page, as expressed in page content order (see “Page Content Order,” below), may contain places where the normal progression of text suffers a discontinuity. For example, the page may contain the beginnings of two separate articles (see Section 8.3.2, “Articles”), each of which is continued onto a later page of the document; the last words of the first article appearing on the page should not be run together with the first words of the second. Consumer applications can recognize such discontinuities by examining the document's logical structure.
- *Hidden page elements.* For a variety of reasons, elements of a document's logical content may be invisible on the page: they may be clipped, their color may match the background, or they may be obscured by other, overlapping objects. Consumer applications must still be able to recognize and process such hidden elements; for example, formerly invisible elements may become visible when a page is reflowed, or a text-to-speech engine may wish to speak text that is not visible to a sighted reader. For the purposes of Tagged PDF, page content is considered to include all text and illustrations in their entirety, whether or not they are visible when the document is displayed or printed.

Page Content Order

When dealing with material on a page-by-page basis, some Tagged PDF consumer applications may wish to process elements in *page content order*, determined by the sequencing of graphics objects within a page's content stream and of characters within a text object, rather than in the *logical structure order* defined by a depth-first traversal of the page's logical structure hierarchy. The two orderings are logically distinct and may or may not coincide; in particular, any artifacts the page may contain are included in the page content order but not in the logical structure order, since they are not considered part of the document's logical structure. The creator of a Tagged PDF document is responsible for establishing both an appropriate page content order for each page and an appropriate logical structure hierarchy for the entire document.

Because the primary requirement for page content order is to enable reflow to maintain elements in proper reading sequence, it should normally (for Western writing systems) proceed from top to bottom (and, in a multiple-column layout, from column to column), with artifacts in their correct relative places. In general, all parts of an article that appear on a given page should be kept together, even if it flows to scattered locations on the page. Illustrations or footnotes may be interspersed with the text of the associated article or may appear at the end of its content (or, in the case of footnotes, at the end of the entire page's logical content).

Sequencing of Annotations

Annotations associated with a page are not interleaved within the page's content stream, but are placed in the **Annots** array in its page object (see "Page Objects" on page 87). Consequently, the correct position of an annotation in the page content order is not readily apparent, but is determined from the document's logical structure.

Both page content (marked-content sequences) and annotations can be treated as content items that are referenced from structure elements (see Section 9.6.3, "Structure Content"). Structure elements of type Link or Form (see "Inline-Level Structure Elements" on page 632 and "Illustration Elements" on page 637) explicitly specify the association between a marked-content sequence and a corresponding annotation. In other cases, if the structure element corresponding to an annotation immediately precedes or follows (in the logical structure order) a structure element corresponding to a marked-content sequence, the annotation

is considered to precede or follow the marked-content sequence, respectively, in the page content order.

Note: If necessary, a Tagged PDF producer may introduce an empty marked-content sequence solely to serve as a structure element for the purpose of positioning adjacent annotations in the page content order.

Reverse-Order Show Strings

In writing systems that are read from right to left (such as Arabic or Hebrew), one might expect that the glyphs in a font would have their origins at the lower right and their widths (rightward horizontal displacements) specified as negative. For various technical and historical reasons, however, many such fonts follow the same conventions as those designed for Western writing systems, with glyph origins at the lower left and positive widths, as shown in Figure 5.4 on page 298. Consequently, showing text in such right-to-left writing systems requires either positioning each glyph individually (which is tedious and costly) or representing text with show strings whose characters are given in reverse order. When the latter method is used, the characters' correct page content order is the reverse of their order within the show string.

The marked-content tag `ReversedChars` informs the Tagged PDF consumer application that show strings within a marked-content sequence contain characters in the reverse of page content order. If the sequence encompasses multiple show strings, only the individual characters within each string are reversed; the strings themselves are in natural reading order. For example, the sequence

```
/ReversedChars
  BMC
    ( olleH) Tj
    -200 0 Td
    (.dlrow) Tj
  EMC
```

represents the text

```
Hello world.
```

The show strings may have a space character at the beginning or end to indicate a word break (see “Identifying Word Breaks” on page 622), but may not contain

interior spaces. This is not a serious limitation, since a space provides an opportunity to realign the typography without visible effect; and it serves the valuable purpose of limiting the scope of reversals for word-processing consumer applications.

Extraction of Character Properties

Character codes in a Tagged PDF document can be unambiguously converted into Unicode values representing the information content of the text, and some characteristics of the associated fonts can be deduced. These Unicode values and font characteristics can then be used for such operations as cut-and-paste editing, searching, text-to-speech conversion, and exporting to other applications or file formats.

Unicode Mapping

Every character code in a Tagged PDF document can be mapped to a corresponding Unicode value. Unicode defines 16-bit values for most of the characters used in the world's languages and writing systems, as well as providing a *vendor space* that applications can use as a private escape area. Information about Unicode can be found in the *Unicode Standard*, by the Unicode Consortium (see the Bibliography).

The methods for mapping a character code to a Unicode value are as follows:

- If the font dictionary contains a **ToUnicode** CMap (see Section 5.9, “ToUnicode CMaps”), use that CMap to convert the character code to Unicode.
- If the font is a simple font that uses one of the predefined encodings **MacRomanEncoding**, **MacExpertEncoding**, or **WinAnsiEncoding**, or that has an encoding whose **Differences** array includes only character names taken from the Adobe standard Latin character set and the set of named characters in the Symbol font (see Appendix D):
 1. Map the character code to a character name according to Table D.1 on page 710 and the font's **Differences** array.
 2. Look up the character name in the *Adobe Glyph List* (see the Bibliography) to obtain the corresponding Unicode value.

- If the font is a composite font that uses one of the predefined CMaps listed in Table 5.14 on page 343 (except Identity-H and Identity-V), or whose descendant CIDFont uses the Adobe-GB1, Adobe-CNS1, Adobe-Japan1, or Adobe-Korea1 character collection:
 1. Map the character code to a character identifier (CID) according to the font's CMap.
 2. Obtain the registry and ordering of the character collection used by the font's CMap (for example, Adobe and Japan1) from its **CIDSystemInfo** dictionary.
 3. Construct a second CMap name by concatenating the registry and ordering obtained in step 2 in the format *registry-ordering-UCS2* (for example, Adobe-Japan1-UCS2).
 4. Obtain the CMap with the name constructed in step 3 (available from the ASN Developer Program Web site; see the Bibliography).
 5. Map the CID obtained in step 1 according to the CMap obtained in step 4, producing a Unicode value.

If these methods fail to produce a Unicode value, there is no way to determine what the character code represents. Tagged PDF producers should take care to ensure that the PDF file contains enough information to map all characters to Unicode by one of the methods described above.

Note: An **Alt** or **ActualText** entry in a structure element dictionary (see Sections 9.8.2, “Alternate Descriptions,” and 9.8.3, “Replacement Text”) or an **E** (expansion text) property in a marked-content property list (Section 9.8.4, “Expansion of Abbreviations and Acronyms”) may affect the character stream that some Tagged PDF consumers actually use. For example, some consumers may choose to use the **Alt** or **ActualText** value and ignore all text and other content associated with the structure element and its descendants.

Note: Some uses of Tagged PDF require characters that may not be available in all fonts, such as the soft hyphen (see “Incidental Artifacts” on page 617). Such characters can be represented either by adding them to the font's encoding or CMap and using **ToUnicode** to map them to appropriate Unicode values, or by using an **ActualText** entry in the associated structure element to provide substitute characters.

Font Characteristics

In addition to a Unicode value, each text character in a content stream has an associated set of font characteristics. These characteristics are useful when exporting text to another application or file format that has a limited repertoire of available fonts.

Table 9.17 lists a common set of font characteristics corresponding to those used in CSS and XSL; more information can be found in the World Wide Web Consortium document *Extensible Stylesheet Language (XSL) 1.0* (see the Bibliography). Each of the characteristics can be derived from information available in the font descriptor's **Flags** entry (see Section 5.7.1, "Font Descriptor Flags").

TABLE 9.17 Derivation of font characteristics

CHARACTERISTIC	TYPE	DERIVATION
Serifed	boolean	The value of the Serif flag in the font descriptor's Flags entry.
Proportional	boolean	The complement of the FixedPitch flag in the font descriptor's Flags entry.
Italic	boolean	The value of the Italic flag in the font descriptor's Flags entry.
Smallcap	boolean	The value of the SmallCap flag in the font descriptor's Flags entry.

Note: The characteristics shown in the table apply only to characters contained in show strings within content streams; they do not exist for alternate description text (**Alt**), replacement text (**ActualText**), or abbreviation expansion text (**E**).

Note: For the standard 14 Type 1 fonts, the font descriptor may be missing; the well-known values for those fonts are used.

Identifying Word Breaks

A document's text stream defines not only the characters in a page's text, but also the words. Unlike a character, the notion of a word is not precisely defined, but depends on the purpose for which the text is being processed. A reflow tool needs to know where it can break the running text into lines; a text-to-speech engine needs to identify the words to be vocalized; spelling checkers and other applications all have their own ideas of what constitutes a word. It is not important for a Tagged PDF document to identify the words within the text stream according to a

single, unambiguous definition that will satisfy all of these clients; what is important is that there be enough information available for each client to make that determination for itself.

The consumer of a Tagged PDF document finds words by sequentially examining the Unicode character stream, perhaps augmented by replacement text specified with **ActualText** (see Section 9.8.3, “Replacement Text”). It does not need to guess about word breaks based on information such as glyph positioning on the page, font changes, or glyph sizes. The main consideration is to ensure that the spacing characters that would be present to separate words in a pure text representation are also present in the Tagged PDF.

Note that the identification of what constitutes a word is unrelated to how the text happens to be grouped into show strings. The division into show strings has no semantic significance; in particular, a space or other word-breaking character is still needed even if a word break happens to fall at the end of a show string.

***Note:** Some applications may identify words by simply separating them at every space character; others may be slightly more sophisticated and treat punctuation marks such as hyphens or em dashes as word separators as well. Still other applications may wish to identify possible line-break opportunities by using an algorithm similar to the one in Unicode Standard Annex #14, Line Breaking Properties, available from the Unicode Consortium (see the Bibliography). (Note that this specification implies that lines of text in Western writing systems usually end with a trailing space.)*

9.7.3 Basic Layout Model

Tagged PDF’s standard structure types and attributes are interpreted in the context of a basic layout model that determines how structure elements are to be arranged on the page. Although this model does not necessarily correspond to the one actually used for page layout by the application creating the document, it is designed to capture the general intent of the document’s underlying structure. (The PDF content stream itself specifies the exact appearance.) The goal is to provide sufficient information for Tagged PDF consumers to make their own layout decisions while preserving the authoring application’s intent as closely as their own layout models allow.

Note: The Tagged PDF layout model resembles the ones used in markup languages such as HTML, CSS, XSL, and RTF, but does not correspond exactly to any of them. The model is deliberately defined loosely in order to allow reasonable latitude in the interpretation of structure elements and attributes when converting to other document formats; some degree of variation in the resulting layout from one format to another is to be expected.

The basic layout model begins with the notion of a *reference area*. This is a rectangular region used by the layout application as a frame or guide in which to place the document's content. Some of the standard structure attributes, such as **StartIndent** and **EndIndent** (see "Layout Attributes for BLSEs" on page 642), are measured from the boundaries of the reference area. Reference areas are not specified explicitly, but are inferred from context; those of interest are generally the column area or areas in a general text layout, the outer bounding box of a table and those of its component cells, and the bounding box of an illustration or other floating element.

The standard structure types are divided into four main categories according to the roles they play in page layout:

- *Grouping elements* group other elements into sequences or hierarchies, but hold no content directly and have no direct effect on layout.
- *Block-level structure elements (BLSEs)* describe the overall layout of content on the page, proceeding in the *block-progression direction*.
- *Inline-level structure elements (ILSEs)* describe the layout of content within a BLSE, proceeding in the *inline-progression direction*.
- *Illustration elements* are compact sequences of content, in page content order, that are considered to be unitary objects with respect to page layout. An illustration can be treated as either a BLSE or an ILSE.

The meaning of the terms *block-progression direction* and *inline-progression direction* depends on the writing system in use, as specified by the standard attribute **WritingMode** (see "General Layout Attributes" on page 640). In Western writing systems, the block direction is from top to bottom and the inline direction from left to right; other writing systems use different directions for laying out content.

Because the progression directions can vary depending on the writing system, edges of areas and directions on the page must be identified by terms that are

neutral with respect to the progression order rather than by familiar terms such as *up*, *down*, *left*, and *right*. Block layout proceeds from *before* to *after*, inline from *start* to *end*. Thus, for example, in Western writing systems the before and after edges of a reference area are at the top and bottom, respectively, and the start and end edges are at the left and right. Another term, *shift direction* (the direction of shift for a superscript), refers to the direction opposite that for block progression—that is, from after to before (in Western writing systems, from bottom to top).

BLSEs are *stacked* within a reference area in block-progression order. In general, the first BLSE is placed against the before edge of the reference area; subsequent BLSEs are then stacked against preceding ones, progressing toward the after edge, until no more BLSEs will fit in the reference area. If the overflowing BLSE allows itself to be split—such as a paragraph that can be split between lines of text—a portion of it may be included in the current reference area and the remainder carried over to a subsequent reference area (either elsewhere on the same page or on another page of the document). Once the amount of content that fits in a reference area is determined, the placements of the individual BLSEs may be adjusted to bias the placement toward the before edge, the middle, or the after edge of the reference area, or the spacing within or between BLSEs may be adjusted to fill the full extent of the reference area.

Note: *BLSEs may be nested, with child BLSEs stacked within a parent BLSE in the same manner as BLSEs within a reference area. Except in a few instances noted below, such nesting of BLSEs does not result in the nesting of reference areas; a single reference area prevails for all levels of nested BLSEs.*

Within a BLSE, child ILSEs are *packed* into *lines*. (*Direct content items*—those that are immediate children of a BLSE rather than contained within a child ILSE—are implicitly treated as ILSEs for packing purposes.) Each line is treated as a synthesized BLSE and is stacked within the parent BLSE. Lines may be intermingled with other BLSEs within the parent area. This line-building process is analogous to the stacking of BLSEs within a reference area, except that it proceeds in the inline-progression rather than the block-progression direction: a line is packed with ILSEs beginning at the start edge of the containing BLSE and continuing until the end edge is reached and the line is full. The overflowing ILSE may allow itself to be broken at linguistically determined or explicitly marked break points (such as hyphenation points within a word), and the remaining fragment is then carried over to the next line.

Note: Certain values of an element's **Placement** attribute remove the element from the normal stacking or packing process and allow it instead to “float” to a specified edge of the enclosing reference area or parent BLSE; see “General Layout Attributes” on page 640 for further discussion.

Two enclosing rectangles are associated with each BLSE and ILSE (including direct content items that are treated implicitly as ILSEs):

- The *content rectangle* is derived from the shape of the enclosed content and defines the bounds used for the layout of any included child elements.
- The *allocation rectangle* includes any additional borders or spacing surrounding the element, affecting how it will be positioned with respect to adjacent elements and the enclosing content rectangle or reference area.

The definitions of these rectangles are determined by layout attributes associated with the structure element; see “Content and Allocation Rectangles” on page 648 for further discussion.

9.7.4 Standard Structure Types

Tagged PDF's *standard structure types* characterize the role of a content element within the document and, in conjunction with the standard structure attributes (described in Section 9.7.5, “Standard Structure Attributes”), how that content is laid out on the page. As discussed in Section 9.6.2, “Structure Types,” the structure type of a logical structure element is specified by the **S** entry in its structure element dictionary. To be considered a standard structure type, this value must be either:

- One of the standard structure type names described below
- An arbitrary name that is mapped to one of the standard names by the document's role map (see Section 9.6.2, “Structure Types”), possibly via multiple levels of mapping

Ordinarily, structure elements having standard structure types will be processed the same way whether the type is expressed directly or is determined indirectly via the role map. However, some consumer applications may ascribe additional semantics to nonstandard structure types, even though the role map associates them with standard ones. For instance, the actual values of the **S** entries may be

used when exporting to a tagged representation such as XML, while the corresponding role-mapped values are used when converting to presentation formats such as HTML or RTE, or for purposes such as reflow or accessibility to disabled users.

***Note:** Most of the standard element types are designed primarily for laying out text; the terminology reflects this usage. However, a layout can in fact include any type of content, such as path or image objects. The content items associated with a structure element are laid out on the page as if they were blocks of text (for a BLSE) or characters within a line of text (for an ILSE).*

Grouping Elements

Grouping elements are used solely to group other structure elements; they are not directly associated with content items. Table 9.18 describes the standard structure types for elements in this category.

TABLE 9.18 Standard structure types for grouping elements

STRUCTURE TYPE	DESCRIPTION
Document	(Document) A complete document. This is the root element of any structure tree containing multiple parts or multiple articles.
Part	(Part) A large-scale division of a document. This type of element is appropriate for grouping articles or sections.
Art	(Article) A relatively self-contained body of text constituting a single narrative or exposition. Articles should be disjoint; that is, they should not contain other articles as constituent elements.
Sect	(Section) A container for grouping related content elements. For example, a section might contain a heading, several introductory paragraphs, and two or more other sections nested within it as subsections.
Div	(Division) A generic block-level element or group of elements.
BlockQuote	(Block quotation) A portion of text consisting of one or more paragraphs attributed to someone other than the author of the surrounding text.
Caption	(Caption) A brief portion of text describing a table or figure.

TOC	<p>(Table of contents) A list made up of list items (structure types L and LI; see “List Elements” on page 630), each of which may contain either a table of contents item (structure type TOCI; see below) or another such list, along with a reference to the associated content (structure type Reference; see “Inline-Level Structure Elements” on page 632). A table of contents is thus potentially hierarchical; it is desirable for the hierarchy to correspond in structure to the structure hierarchy of the main body of the document.</p> <p><i>Note: Lists of figures and tables, as well as bibliographies, can be treated as tables of contents for purposes of the standard structure types.</i></p>
TOCI	<p>(Table of contents item) A piece of identifying text accompanied by one or more references to the associated content (structure type Reference; see “Inline-Level Structure Elements” on page 632). The reference elements can be used to show the label, title, and page number of the specified item in the main body of the document.</p>
Index	<p>(Index) A sequence of entries containing identifying text accompanied by reference elements (structure type Reference; see “Inline-Level Structure Elements” on page 632) that point out occurrences of the specified text in the main body of a document.</p>
NonStruct	<p>(Nonstructural element) A grouping element having no inherent structural significance; it serves solely for grouping purposes. This type of element differs from a division (structure type Div; see above) in that it is not interpreted or exported to other document formats; however, its descendants are to be processed normally.</p>
Private	<p>(Private element) A grouping element containing private content belonging to the application producing it. The structural significance of this type of element is unspecified, and is determined entirely by the producer application. Neither the Private element nor any of its descendants are to be interpreted or exported to other document formats.</p>

For most content extraction formats, the document must be a tree with a single top-level element; the structure tree root (identified by the **StructTreeRoot** entry in the document catalog) must have only one child in its **K** (kids) array. If the PDF file contains a complete document, the structure type Document is recommended for this top-level element in the logical structure hierarchy; if the file contains a well-formed document fragment, one of the structure types Part, Art, Sect, or Div may be used instead.

Block-Level Structure Elements

A *block-level structure element (BLSE)* is any region of text or other content that is laid out in the block-progression direction, such as a paragraph, heading, list item, or footnote. A structure element is a BLSE if its structure type (after role mapping, if any) is one of those listed in Table 9.19. All other standard structure types are treated as ILSEs, with the following exceptions:

- TR (Table row), TH (Table header), and TD (Table data), which are used to group elements within a table and are considered neither BLSEs nor ILSEs
- Elements with a **Placement** attribute (see “General Layout Attributes” on page 640) other than the default value of Inline

TABLE 9.19 Block-level structure elements

CATEGORY	STRUCTURE TYPES		
Paragraphlike elements	P	H1	H4
	H	H2	H5
		H3	H6
List elements	L	Lbl	
	LI	LBody	
Table element	Table		

In many cases, a BLSE appears as one compact, contiguous piece of page content; in other cases, it is discontinuous. Examples of the latter include a BLSE that extends across a page boundary or is interrupted in the page content order by another, nested BLSE or a directly included footnote. When necessary, Tagged PDF consumer applications can recognize such fragmented BLSEs from the logical structure and use this information to reassemble them and properly lay them out.

Paragraphlike Elements

Table 9.20 describes structure types for *paragraphlike elements* that consist of running text and other content laid out in the form of conventional paragraphs (as opposed to more specialized layouts such as lists and tables).

TABLE 9.20 Standard structure types for paragraphlike elements

STRUCTURE TYPE	DESCRIPTION
H	(Heading) A label for a subdivision of a document's content. It should be the first child of the division that it heads.
H1–H6	Headings with specific levels, for use in applications that cannot hierarchically nest their sections and thus cannot determine the level of a heading from its level of nesting.
P	(Paragraph) A low-level division of text.

List Elements

The structure types described in Table 9.21 are used for organizing the content of lists.

TABLE 9.21 Standard structure types for list elements

STRUCTURE TYPE	DESCRIPTION
L	(List) A sequence of items of like meaning and importance. Its immediate children should be an optional caption (structure type <i>Caption</i> ; see “Grouping Elements” on page 627) followed by one or more list items (structure type <i>LI</i> ; see below).
LI	(List item) An individual member of a list. Its children may be one or more labels, list bodies, or both (structure types <i>Lbl</i> or <i>LBody</i> ; see below).
Lbl	(Label) A name or number that distinguishes a given item from others in the same list or other group of like items. In a dictionary list, for example, it contains the term being defined; in a bulleted or numbered list, it contains the bullet character or the number of the list item and associated punctuation.
LBody	(List body) The descriptive content of a list item. In a dictionary list, for example, it contains the definition of the term. It can either contain the content directly or have other BLSEs, perhaps including nested lists, as children.

Table Elements

The structure types described in Table 9.22 are used for organizing the content of tables.

TABLE 9.22 Standard structure types for table elements

STRUCTURE TYPE	DESCRIPTION
Table	(Table) A two-dimensional layout of rectangular data cells, possibly having a complex substructure. It contains table rows (structure type TR; see below) as children, and may have a caption (structure type Caption; see “Grouping Elements” on page 627) as its first or last child.
TR	(Table row) A row of headings or data in a table. It may contain table header cells and table data cells (structure types TH and TD; see below).
TH	(Table header) A table cell containing header text describing one or more rows or columns of the table.
TD	(Table data) A table cell containing data that is part of the table’s content.

Note: The association of headers with rows and columns of data is to be determined heuristically by applications and is not defined here.

Usage Guidelines for Block-Level Structure

Because different consumer applications use PDF’s logical structure facilities in different ways, Tagged PDF does not enforce any strict rules regarding the order and nesting of elements using the standard structure types. Furthermore, each export format has its own conventions for logical structure. However, adhering to certain general guidelines will help to achieve the most consistent and predictable interpretation among different Tagged PDF consumers.

As described under “Grouping Elements” on page 627, a Tagged PDF document can have one or more levels of grouping elements, such as Document, Part, Art (Article), Sect (Section), and Div (Division). The descendants of these are BLSEs, such as H (Heading), P (Paragraph), and L (List), that hold the actual content. Their descendants, in turn, are the content items themselves or ILSEs that further describe the content.

Note: As noted earlier, elements with structure types that would ordinarily be treated as ILSEs can have a **Placement** attribute (see “General Layout Attributes” on page 640) that causes them to be treated as BLSEs instead. Such elements may be included as BLSEs in the same manner as headings and paragraphs.

The block-level structure can follow one of two principal paradigms:

- *Strongly structured.* The grouping elements nest to as many levels as necessary to reflect the organization of the material into articles, sections, subsections, and so on. At each level, the children of the grouping element consist of a heading (H), one or more paragraphs (P) for content at that level, and perhaps one or more additional grouping elements for nested subsections.
- *Weakly structured.* The document is relatively flat, having perhaps only one or two levels of grouping elements, with all the headings, paragraphs, and other BLSEs as their immediate children. In this case, the organization of the material is not reflected in the logical structure; however, it can be expressed by the use of headings with specific levels (H1–H6).

The strongly structured paradigm is used by some rich document models based on XML; the weakly structured paradigm is typical of documents represented in HTML.

Lists and tables should be organized using the specific structure types described under “List Elements” on page 630 and “Table Elements” on page 631. Likewise, tables of contents and indexes should be structured as described for the TOC and Index structure types under “Grouping Elements” on page 627.

Inline-Level Structure Elements

An *inline-level structure element (ILSE)* contains a portion of text or other content having specific styling characteristics or playing a specific role in the document. Within a paragraph or other block defined by a containing BLSE, consecutive ILSEs—possibly intermixed with other content items that are direct children of the parent BLSE—are laid out consecutively in the inline-progression direction (left to right in Western writing systems). The resulting content may then be broken into multiple *lines*, which in turn are stacked in the block-progression direction. It is possible for an ILSE in turn to contain a BLSE, which is then treated as a unitary item of layout in the inline direction. Table 9.23 lists the standard structure types for ILSEs.

TABLE 9.23 Standard structure types for inline-level structure elements

STRUCTURE TYPE	DESCRIPTION
Span	<p>(Span) A generic inline portion of text having no particular inherent characteristics. It can be used, for example, to delimit a range of text with a given set of styling attributes.</p> <p><i>Note:</i> Not all inline style changes need to be identified as a span. Text color and font changes (including modifiers such as bold, italic, and small caps) need not be so marked, since these can be derived from the PDF content (see “Font Characteristics” on page 622). However, it is necessary to use a span in order to apply explicit layout attributes such as LineHeight, BaselineShift, or TextDecorationType (see “Layout Attributes for ILSEs” on page 646).</p> <p><i>Note:</i> Marked-content sequences having the tag <i>Span</i> are also used to carry certain accessibility properties (Lang and E; see Sections 9.8.1, “Natural Language Specification,” and 9.8.4, “Expansion of Abbreviations and Acronyms”). Such sequences lack an MCID property and are not associated with any structure element. This use of the <i>Span</i> marked-content tag is distinct from its use as a structure type.</p>
Quote	<p>(Quotation) An inline portion of text attributed to someone other than the author of the surrounding text.</p> <p><i>Note:</i> The quoted text is contained inline within a single paragraph; this differs from the block-level element <i>BlockQuote</i> (see “Grouping Elements” on page 627), which consists of one or more complete paragraphs (or other elements presented as if they were complete paragraphs).</p>
Note	<p>(Note) An item of explanatory text, such as a footnote or an endnote, that is referred to from within the body of the document. It may have a label (structure type <i>Lbl</i>; see “List Elements” on page 630) as a child. The note itself may be included as a child of the structure element in the body text that refers to it, or it may be included elsewhere (such as in an endnotes section) and accessed via a reference (structure type <i>Reference</i>; see below).</p> <p><i>Note:</i> Tagged PDF does not prescribe the placement of footnotes in the page content order; they can be either inline or at the end of the page, at the discretion of the producer application.</p>
Reference	<p>(Reference) A citation to content elsewhere in the document.</p>

BibEntry	(Bibliography entry) A reference identifying the external source of some cited content. It may contain a label (structure type Lbl; see “List Elements” on page 630) as a child. <i>Note:</i> Although a bibliography entry is likely to include component parts identifying the cited content’s author, work, publisher, and so forth, at the time of publication no standard structure types are defined at this level of detail.
Code	(Code) A fragment of computer program text.
Link	(Link) An association between a portion of the ILSE’s content and a corresponding link annotation or annotations (see “Link Annotations” on page 501). Its children are one or more content items or child ILSEs and one or more object references (see “PDF Objects as Content Items” on page 598) identifying the associated link annotations. See “Link Elements,” below, for further discussion.

Link Elements

Link annotations (like all PDF annotations) are associated with a geometric region of the page rather than with a particular object in its content stream. Any connection between the link and the content is based solely on visual appearance rather than on an explicitly specified association. For this reason, link annotations in themselves are not useful to the visually impaired or to applications needing to determine which content can be activated to invoke a hypertext link.

Tagged PDF link elements (structure type Link) use PDF’s logical structure facilities to establish the association between content items and link annotations, providing functionality comparable to HTML hypertext links. The children of such an element consist of:

- One or more content items or other ILSEs (except other links)
- Object references (see “PDF Objects as Content Items” on page 598) to one or more link annotations associated with the content

A link element may contain several link annotations if the geometry of the content requires it; for instance, if a span of text wraps from the end of one line to the beginning of another, separate link annotations may be needed in order to cover the two portions of text. All of the child link annotations must have the

same target and action. In order to maintain a geometric association between the content and the annotation that is consistent with the logical association, all of the link element's content must be covered by the union of its child link annotations.

As an example, consider the following fragment of HTML code, which produces a line of text containing a hypertext link:

```
<html>
  <body>
    <p>
      Here is some text <a href=http://www.adobe.com>with a link</a> inside.
    </p>
  </body>
</html>
```

Example 9.14 shows an equivalent fragment of PDF using a link element, whose text it displays in blue and underlined. Example 9.15 shows an excerpt from the associated logical structure hierarchy.

Example 9.14

/P << /MCID 0 >>	% Marked-content sequence 0 (paragraph)
BDC	% Begin marked-content sequence
BT	% Begin text object
/T1_0 1 Tf	% Set text font and size
14 0 0 14 10.000 753.976 Tm	% Set text matrix
0.0 0.0 0.0 rg	% Set nonstroking color to black
(Here is some text) Tj	% Show text preceding link
ET	% End text object
EMC	% End marked-content sequence
/Link << /MCID 1 >>	% Marked-content sequence 1 (link)
BDC	% Begin marked-content sequence
0.7 w	% Set line width
[] 0 d	% Solid dash pattern
111.094 751.8587 m	% Move to beginning of underline
174.486 751.8587 l	% Draw underline
0.0 0.0 1.0 RG	% Set stroking color to blue
S	% Stroke underline

```

    BT                                     % Begin text object
      14 0 0 14 111.094 753.976 Tm       % Set text matrix
      0.0 0.0 1.0 rg                     % Set nonstroking color to blue
      (with a link) Tj                   % Show text of link
    ET                                     % End text object
  EMC                                     % End marked-content sequence
/P << /MCID 2 >>                         % Marked-content sequence 2 (paragraph)
  BDC                                     % Begin marked-content sequence
    BT                                     % Begin text object
      14 0 0 14 174.486 753.976 Tm       % Set text matrix
      0.0 0.0 0.0 rg                     % Set nonstroking color to black
      ( inside.) Tj                     % Show text following link
    ET                                     % End text object
  EMC                                     % End marked-content sequence

```

Example 9.15

```

501 0 obj                                % Structure element for paragraph
  << /Type /StructElem
    /S /P
    ...
    /K [ 0                                % Three children: marked-content sequence 0
      502 0 R                              % Link
      2                                    % Marked-content sequence 2
    ]
  >>
endobj

502 0 obj                                % Structure element for link
  << /Type /StructElem
    /S /Link
    ...
    /K [ 1                                % Two children: marked-content sequence 1
      503 0 R                              % Object reference to link annotation
    ]
  >>
endobj

503 0 obj                                % Object reference to link annotation
  << /Type /OBJR
    /Obj 600 0 R                          % Link annotation (not shown)
  >>
endobj

```

Illustration Elements

Tagged PDF defines an *illustration element* as any structure element whose structure type (after role mapping, if any) is one of those listed in Table 9.24. The illustration's content must consist of one or more complete graphics objects; it may not appear between the **BT** and **ET** operators delimiting a text object (see Section 5.3, "Text Objects"). It may include clipping only in the form of a contained marked clipping sequence, as defined in Section 9.5.2, "Marked Content and Clipping." In Tagged PDF, all such marked clipping sequences must carry the marked-content tag `Clip`.

TABLE 9.24 Standard structure types for illustration elements

STRUCTURE TYPE	DESCRIPTION
Figure	(Figure) An item of graphical content. Its placement may be specified with the Placement layout attribute (see "General Layout Attributes" on page 640).
Formula	(Formula) A mathematical formula. <i>Note: This structure type is useful only for identifying an entire content element as a formula; no standard structure types are defined for identifying individual components within the formula. From a formatting standpoint, the formula is treated similarly to a figure (structure type Figure; see above).</i>
Form	(Form) A widget annotation representing an interactive form field (see Section 8.6, "Interactive Forms"). Its only child is an object reference (see "PDF Objects as Content Items" on page 598) identifying the widget annotation; the annotation's appearance stream (see Section 8.4.4, "Appearance Streams") defines the rendering of the form element.

An illustration may have logical substructure, including other illustrations. For purposes of reflow, however, it is moved (and perhaps resized) as a unit, without examining its internal contents. To be useful for reflow, it must have a **BBox** attribute; it may also have **Placement**, **Width**, **Height**, and **BaselineShift** attributes (see "Layout Attributes" on page 639).

Often an illustration will be logically part of, or at least attached to, a paragraph or other element of a document. Any such containment or attachment is represented through the use of the **Figure** structure type; the **Figure** element indicates the point of attachment, while its **Placement** attribute describes the nature of the attachment. An illustration element without a **Placement** attribute is treated as an ILSE and laid out inline.

***Note:** For accessibility to disabled users and other text extraction purposes, an illustration element should always have an **Alt** entry or an **ActualText** entry (or both) in its structure element dictionary (see Sections 9.8.2, “Alternate Descriptions,” and 9.8.3, “Replacement Text”). **Alt** is a description of the illustration, whereas **ActualText** gives the exact text equivalent of a graphical illustration that has the appearance of text.*

9.7.5 Standard Structure Attributes

In addition to the standard structure types themselves, Tagged PDF also defines standard layout and styling attributes for structure elements of those types. These attributes enable predictable formatting to be applied during operations such as reflow and export of PDF content to other document formats.

***Note:** In addition to the standard structure attributes, several other optional entries—**Lang**, **Alt**, and **ActualText**—can appear directly in the structure element dictionary rather than in an attribute dictionary. These are described in Section 9.8, “Accessibility Support,” but are useful to other PDF consumers as well.*

Standard Attribute Owners

As discussed in Section 9.6.4, “Structure Attributes,” attributes are defined in *attribute objects*, which are dictionaries or streams attached to a structure element in either of two ways:

- The **A** entry in the structure element dictionary identifies an attribute object or an array of such objects.
- The **C** entry in the structure element dictionary gives the name of an *attribute class* or an array of such names. The class name is in turn looked up in the *class map*, a dictionary identified by the **ClassMap** entry in the structure tree root, yielding an attribute object corresponding to the class.

Each attribute object has an *owner*, specified by the object's **O** entry, which determines the interpretation of the attributes defined in the object's dictionary. Multiple owners may define like-named attributes with different value types or interpretations. Tagged PDF defines a set of standard attribute owners, shown in Table 9.25.

TABLE 9.25 Standard attribute owners

OWNER	DESCRIPTION
Layout	Attributes governing the layout of content
List	Attributes governing the numbering of lists
Table	Attributes governing the organization of cells in tables
XML-1.00	Additional attributes governing translation to XML, version 1.00
HTML-3.20	Additional attributes governing translation to HTML, version 3.20
HTML-4.01	Additional attributes governing translation to HTML, version 4.01
OEB-1.00	Additional attributes governing translation to OEB™, version 1.0
RTF-1.05	Additional attributes governing translation to Microsoft Rich Text Format, version 1.05
CSS-1.00	Additional attributes governing translation to a format using CSS, version 1.00
CSS-2.00	Additional attributes governing translation to a format using CSS, version 2.00

An attribute object owned by a specific export format, such as **XML-1.00**, is applied only when exporting PDF content to that format. Such format-specific attributes override any corresponding attributes owned by **Layout**, **List**, or **Table**. There may be also be additional format-specific attributes; the set of possible attributes is open-ended and is not explicitly specified or limited by Tagged PDF.

Layout Attributes

Layout attributes specify parameters of the layout process used to produce the appearance described by a document's PDF content. Attributes in this category

are defined in attribute objects whose **O** (owner) entry has the value **Layout** (or is one of the format-specific owner names listed in Table 9.25 on page 639). The intent is that these parameters can be used to reflow the content or export it to some other document format with at least basic styling preserved.

Table 9.26 summarizes the standard layout attributes and the structure elements to which they apply. The following sections describe the meaning and usage of these attributes.

TABLE 9.26 Standard layout attributes

STRUCTURE ELEMENTS	ATTRIBUTES
Any structure element	Placement WritingMode
Any BLSE ILSEs with Placement other than Inline	SpaceBefore SpaceAfter StartIndent EndIndent
BLSEs containing text	TextIndent TextAlign
Illustration elements (Figure, Formula, Form) Table	BBox Width Height
TH (Table header) TD (Table data)	Width Height BlockAlign InlineAlign
Any ILSE BLSEs containing ILSEs or containing direct or nested content items	LineHeight BaselineShift TextDecorationType

General Layout Attributes

The layout attributes described in Table 9.27 can apply to structure elements of any of the standard types, whether at the block level (BLSEs) or the inline level (ILSEs).

TABLE 9.27 Standard layout attributes common to all standard structure types

KEY	TYPE	VALUE	
Placement	name	<i>(Optional)</i> The positioning of the element with respect to the enclosing reference area and other content:	
		Block	Stacked in the block-progression direction within an enclosing reference area or parent BLSE.
		Inline	Packed in the inline-progression direction within an enclosing BLSE.
		Before	Placed so that the before edge of the element's allocation rectangle (see "Content and Allocation Rectangles" on page 648) coincides with that of the nearest enclosing reference area. The element may float, if necessary, to achieve the specified placement (see note below). The element is treated as a block occupying the full extent of the enclosing reference area in the inline direction; other content is stacked so as to begin at the after edge of the element's allocation rectangle.
		Start	Placed so that the start edge of the element's allocation rectangle (see "Content and Allocation Rectangles" on page 648) coincides with that of the nearest enclosing reference area. The element may float, if necessary, to achieve the specified placement (see note below). Other content that would intrude into the element's allocation rectangle is laid out as a runaround.
End	Placed so that the end edge of the element's allocation rectangle (see "Content and Allocation Rectangles" on page 648) coincides with that of the nearest enclosing reference area. The element may float, if necessary, to achieve the specified placement (see note below). Other content that would intrude into the element's allocation rectangle is laid out as a runaround.		

When applied to an ILSE, any value except `Inline` causes the element to be treated as a BLSE instead. Default value: `Inline`.

Note: Elements with **Placement** values of `Before`, `Start`, or `End` are removed from the normal stacking or packing process and allowed to "float" to the specified edge of the enclosing reference area or parent BLSE. Multiple such floating elements may be positioned adjacent to one another against the specified edge of the reference area, or placed serially against the edge, in the order encountered.

Complex cases such as floating elements that interfere with each other or do not fit on the same page may be handled differently by different layout applications; Tagged PDF merely identifies the elements as floating and indicates their desired placement.

WritingMode	name	(Optional) The directions of layout progression for packing of ILSEs (inline progression) and stacking of BLSEs (block progression):
	LrTb	Inline progression from left to right; block progression from top to bottom. This is the typical writing mode for Western writing systems.
	RlTb	Inline progression from right to left; block progression from top to bottom. This is the typical writing mode for Arabic and Hebrew writing systems.
	TbRl	Inline progression from top to bottom; block progression from right to left. This is the typical writing mode for Chinese and Japanese writing systems.

The specified layout directions apply to the given structure element and all of its descendants to any level of nesting. Default value: LrTb.

For elements that produce multiple columns, the writing mode defines the direction of column progression within the reference area: the inline direction determines the stacking direction for columns and the default flow order of text from column to column. For tables, the writing mode controls the layout of rows and columns: table rows (structure type TR) are stacked in the block direction, cells within a row (structure type TD) in the inline direction.

Note: *The inline-progression direction specified by the writing mode is subject to local override within the text being laid out, as described in Unicode Standard Annex #9, The Bidirectional Algorithm, available from the Unicode Consortium (see the Bibliography).*

Layout Attributes for BLSEs

Table 9.28 describes layout attributes that apply only to block-level structure elements (BLSEs).

*Note: Inline-level structure elements (ILSEs) with a **Placement** attribute other than the default value of *Inline* are treated as BSLEs, and hence are also subject to the attributes described here.*

TABLE 9.28 Additional standard layout attributes specific to block-level structure elements

KEY	TYPE	VALUE
SpaceBefore	number	<p>(Optional) The amount of extra space preceding the before edge of the BLSE, measured in default user space units in the block-progression direction. This value is added to any adjustments induced by the LineHeight attributes of ILSEs within the first line of the BLSE (see “Layout Attributes for ILSEs” on page 646). If the preceding BLSE has a SpaceAfter attribute, the greater of the two attribute values is used. Default value: 0.</p> <p><i>Note: This attribute is disregarded for the first BLSE placed in a given reference area.</i></p>
SpaceAfter	number	<p>(Optional) The amount of extra space following the after edge of the BLSE, measured in default user space units in the block-progression direction. This value is added to any adjustments induced by the LineHeight attributes of ILSEs within the last line of the BLSE (see “Layout Attributes for ILSEs” on page 646). If the following BLSE has a SpaceBefore attribute, the greater of the two attribute values is used. Default value: 0.</p> <p><i>Note: This attribute is disregarded for the last BLSE placed in a given reference area.</i></p>
StartIndent	number	<p>(Optional) The distance from the start edge of the reference area to that of the BLSE, measured in default user space units in the inline-progression direction. This attribute applies only to structure elements with a Placement attribute of Block or Start (see “General Layout Attributes” on page 640); it is disregarded for those with other Placement values. Default value: 0.</p> <p><i>Note: A negative value for this attribute places the start edge of the BLSE outside that of the reference area. The results are implementation-dependent and may not be supported by all Tagged PDF consumer applications or export formats.</i></p> <p><i>Note: If a structure element with a StartIndent attribute is placed adjacent to a floating element with a Placement attribute of Start, the actual value used for the element’s starting indent will be its own StartIndent attribute or the inline extent of the adjacent floating element, whichever is greater. This value may then be further adjusted by the element’s TextIndent attribute, if any.</i></p>

EndIndent	number	<p>(<i>Optional</i>) The distance from the end edge of the BLSE to that of the reference area, measured in default user space units in the inline-progression direction. This attribute applies only to structure elements with a Placement attribute of Block or End (see “General Layout Attributes” on page 640); it is disregarded for those with other Placement values. Default value: 0.</p> <p>Note: A negative value for this attribute places the end edge of the BLSE outside that of the reference area. The results are implementation-dependent and may not be supported by all Tagged PDF consumer applications or export formats.</p> <p>Note: If a structure element with an EndIndent attribute is placed adjacent to a floating element with a Placement attribute of End, the actual value used for the element’s ending indent will be its own EndIndent attribute or the inline extent of the adjacent floating element, whichever is greater.</p>								
TextIndent	number	<p>(<i>Optional; applies only to some BLSEs, as described below</i>) The additional distance, measured in default user space units in the inline-progression direction, from the start edge of the BLSE, as specified by StartIndent (above), to that of the first line of text. A negative value indicates a hanging indent. Default value: 0.</p> <p>This attribute applies only to paragraphlike BLSEs and those of structure types Lbl (Label), LBody (List body), TH (Table header), and TD (Table data), provided that they contain content other than nested BLSEs.</p>								
TextAlign	name	<p>(<i>Optional; applies only to BLSEs containing text</i>) The alignment, in the inline-progression direction, of text and other content within lines of the BLSE:</p> <table border="0" style="margin-left: 2em;"> <tr> <td style="padding-right: 1em;">Start</td> <td>Aligned with the start edge.</td> </tr> <tr> <td>Center</td> <td>Centered between the start and end edges.</td> </tr> <tr> <td>End</td> <td>Aligned with the end edge.</td> </tr> <tr> <td>Justify</td> <td>Aligned with both the start and end edges, with internal spacing within each line expanded, if necessary, to achieve such alignment. The last (or only) line is aligned with the start edge only, as for Start (above).</td> </tr> </table> <p>Default value: Start.</p>	Start	Aligned with the start edge.	Center	Centered between the start and end edges.	End	Aligned with the end edge.	Justify	Aligned with both the start and end edges, with internal spacing within each line expanded, if necessary, to achieve such alignment. The last (or only) line is aligned with the start edge only, as for Start (above).
Start	Aligned with the start edge.									
Center	Centered between the start and end edges.									
End	Aligned with the end edge.									
Justify	Aligned with both the start and end edges, with internal spacing within each line expanded, if necessary, to achieve such alignment. The last (or only) line is aligned with the start edge only, as for Start (above).									
BBox	rectangle	<p>(<i>Illustrations and tables only; required if the element appears in its entirety on a single page</i>) An array of four numbers in default user space units giving the coordinates of the left, bottom, right, and top edges, respectively, of the element’s bounding box (the rectangle that completely encloses its visible content). This attribute applies only to elements of structure type Figure, Formula, Form, or Table.</p>								

Width number
 or name *(Optional; illustrations, tables, table headers, and table cells only; strongly recommended for table cells)* The desired width of the element’s content rectangle (see “Content and Allocation Rectangles” on page 648), measured in default user space units in the inline-progression direction. This attribute applies only to elements of structure type Figure, Formula, Form, Table, TH (Table header), or TD (Table data).

The name Auto in place of a numeric value indicates that no specific width constraint is to be imposed; the element’s width is determined by the intrinsic width of its content. Default value: Auto.

Height number
 or name *(Optional; illustrations, tables, table headers, and table cells only)* The desired height of the element’s content rectangle (see “Content and Allocation Rectangles” on page 648), measured in default user space units in the block-progression direction. This attribute applies only to elements of structure type Figure, Formula, Form, Table, TH (Table header), or TD (Table data).

The name Auto in place of a numeric value indicates that no specific height constraint is to be imposed; the element’s height is determined by the intrinsic height of its content. Default value: Auto.

BlockAlign name *(Optional; table cells only)* The alignment, in the block-progression direction, of content within the table cell:

- Before Before edge of the first child’s allocation rectangle aligned with that of the table cell’s content rectangle.
- Middle Children centered within the table cell, so that the distance between the before edge of the first child’s allocation rectangle and that of the table cell’s content rectangle is the same as the distance between the after edge of the last child’s allocation rectangle and that of the table cell’s content rectangle.
- After After edge of the last child’s allocation rectangle aligned with that of the table cell’s content rectangle.
- Justify Children aligned with both the before and after edges of the table cell’s content rectangle. The first child is placed as described above for Before and the last child as described for After, with equal spacing between the children. If there is only one child, it is aligned with the before edge only, as for Before.

This attribute applies only to elements of structure type TH (Table header) or TD (Table data), and controls the placement of all BLSEs that are children of the given element. The table cell’s content rectangle (see “Content and Allocation Rectangles” on page 648) becomes the reference area for all of its descendants. Default value: Before.

InlineAlign	name	<p>(<i>Optional; table cells only</i>) The alignment, in the inline-progression direction, of content within the table cell:</p> <table border="0"> <tr> <td style="padding-left: 2em;">Start</td> <td>Start edge of each child's allocation rectangle aligned with that of the table cell's content rectangle</td> </tr> <tr> <td style="padding-left: 2em;">Center</td> <td>Each child centered within the table cell, so that the distance between the start edges of the child's allocation rectangle and the table cell's content rectangle is the same as the distance between their end edges</td> </tr> <tr> <td style="padding-left: 2em;">End</td> <td>End edge of each child's allocation rectangle aligned with that of the table cell's content rectangle</td> </tr> </table> <p>This attribute applies only to elements of structure type TH (Table header) or TD (Table data), and controls the placement of all BLSEs that are children of the given element. The table cell's content rectangle (see "Content and Allocation Rectangles" on page 648) becomes the reference area for all of its descendants. Default value: Start.</p>	Start	Start edge of each child's allocation rectangle aligned with that of the table cell's content rectangle	Center	Each child centered within the table cell, so that the distance between the start edges of the child's allocation rectangle and the table cell's content rectangle is the same as the distance between their end edges	End	End edge of each child's allocation rectangle aligned with that of the table cell's content rectangle
Start	Start edge of each child's allocation rectangle aligned with that of the table cell's content rectangle							
Center	Each child centered within the table cell, so that the distance between the start edges of the child's allocation rectangle and the table cell's content rectangle is the same as the distance between their end edges							
End	End edge of each child's allocation rectangle aligned with that of the table cell's content rectangle							

Layout Attributes for ILSEs

The attributes described in Table 9.29 apply to inline-level structure elements (ILSEs). They may also be specified for a block-level element (BLSE) and will apply to any content items that are its immediate children.

TABLE 9.29 Standard layout attributes specific to inline-level structure elements

KEY	TYPE	VALUE				
LineHeight	number or name	<p>(<i>Optional</i>) The element's preferred height, measured in default user space units in the block-progression direction. The height of a line is determined by the largest LineHeight value for any complete or partial ILSE that it contains.</p> <p>The name Normal or Auto in place of a numeric value indicates that no specific height constraint is to be imposed; the element's height is set to a reasonable value based on the content's font size:</p> <table border="0"> <tr> <td style="padding-left: 2em;">Normal</td> <td>Adjust the line height to include any nonzero value specified for BaselineShift (see below).</td> </tr> <tr> <td style="padding-left: 2em;">Auto</td> <td>Do not adjust for the value of BaselineShift.</td> </tr> </table> <p>Default value: Normal.</p>	Normal	Adjust the line height to include any nonzero value specified for BaselineShift (see below).	Auto	Do not adjust for the value of BaselineShift .
Normal	Adjust the line height to include any nonzero value specified for BaselineShift (see below).					
Auto	Do not adjust for the value of BaselineShift .					

This attribute applies to all ILSEs (including implicit ones) that are children of this element or of its nested ILSEs, if any; it does not apply to nested BLSEs.

Note: When translating to a specific export format, the values *Normal* and *Auto*, if specified, are used directly if they are available in the target format. The meaning of the term “reasonable value,” used above, is left to the consumer application to determine; it can be assumed to be approximately 1.2 times the font size, but this value may vary depending on the export format. In the absence of a numeric value for **LineHeight** or an explicit value for the font size, a reasonable method of calculating the line height from the information in a Tagged PDF file is to find the difference between the associated font’s **Ascent** and **Descent** values (see Section 5.7, “Font Descriptors”), map it from glyph space to default user space (see Section 5.3.3, “Text Space Details”), and use the maximum resulting value for any character in the line.

BaselineShift number

(Optional) The distance, in default user space units, by which the element’s baseline is shifted relative to that of its parent element. The shift direction is the opposite of the block-progression direction specified by the prevailing **WritingMode** attribute (see “General Layout Attributes” on page 640); thus positive values shift the baseline toward the before edge and negative values toward the after edge of the reference area (upward and downward, respectively, in Western writing systems). Default value: 0.

The shifted element might be a superscript, a subscript, or an inline graphic. The shift applies to the element itself, its content, and all of its descendants. Any further baseline shift applied to a child of this element is measured relative to the shifted baseline of this (parent) element.

TextDecorationType name

(Optional) The *text decoration*, if any, to be applied to the element’s text.

None	No text decoration
Underline	A line below the text
Overline	A line above the text
LineThrough	A line through the middle of the text

Default value: None.

This attribute applies to all text content items that are children of this element or of its nested ILSEs, if any; it does not apply to nested BLSEs or to content items other than text.

Note: The color, position, and thickness of the decoration should be uniform across all children, in spite of changes in color, font size, or other variations in the content’s text characteristics.

Content and Allocation Rectangles

As defined in Section 9.7.3, “Basic Layout Model,” an element’s *content rectangle* is an enclosing rectangle derived from the shape of the element’s content, which defines the bounds used for the layout of any included child elements; the *allocation rectangle* includes any additional borders or spacing surrounding the element, affecting how it will be positioned with respect to adjacent elements and the enclosing content rectangle or reference area.

The exact definition of the content rectangle depends on the element’s structure type:

- For a table cell (structure type TH or TD), the content rectangle is determined from the bounding box of all graphics objects in the cell’s content, taking into account any explicit bounding boxes (such as the **BBox** entry in a form XObject). This implied size can be explicitly overridden by the cell’s **Width** and **Height** attributes. The cell’s height is then further adjusted to equal the maximum height of any cell in its row, and its width to the maximum width of any cell in its column.
- For any other BLSE, the height of the content rectangle is the sum of the heights of all BLSEs it contains, plus any additional spacing adjustments between these elements.
- For an ILSE that contains text, the height of the content rectangle is set by the **LineHeight** attribute. The width is determined by summing the widths of the contained characters, adjusted for any indents, letter spacing, word spacing, or line-end conditions.
- For an ILSE that contains an illustration or table, the content rectangle is determined from the bounding box of all graphics objects in the content, taking into account any explicit bounding boxes (such as the **BBox** entry in a form XObject). This implied size can be explicitly overridden by the element’s **Width** and **Height** attributes.
- For an ILSE that contains a mixture of elements, the height of the content rectangle is determined by aligning the child objects relative to one another based on their text baseline (for text ILSEs) or end edge (for nontext ILSEs), along with any applicable **BaselineShift** attribute (for all ILSEs), and then finding the extreme top and bottom for all elements.

***Note:** Some applications may apply this process to all elements within the block; others may apply it on a line-by-line basis.*

The allocation rectangle is derived from the content rectangle in a way that also depends on the structure type:

- For a BLSE, the allocation rectangle is equal to the content rectangle with its before and after edges adjusted by the element's **SpaceBefore** and **SpaceAfter** attributes, if any, but with no changes to the start and end edges.
- For an ILSE, the allocation rectangle is the same as the content rectangle.

Note: Future versions of Tagged PDF are likely to include additional attributes that can adjust all four edges of the allocation rectangle for both BLSEs and ILSEs.

Illustration Attributes

Certain additional restrictions arise in connection with particular uses of illustration elements (structure types Figure, Formula, or Form):

- When an illustration element has a **Placement** attribute of Block, it must have a **Height** attribute with an explicitly specified numerical value (not Auto). This value is the sole source of information about the illustration's extent in the block-progression direction.
- When an illustration element has a **Placement** attribute of Inline, it must have a **Width** attribute with an explicitly specified numerical value (not Auto). This value is the sole source of information about the illustration's extent in the inline-progression direction.
- When an illustration element has a **Placement** attribute of Inline, Start, or End, the value of its **BaselineShift** attribute is used to determine the position of its after edge relative to the text baseline; **BaselineShift** is ignored for all other values of **Placement**. (An illustration element with a **Placement** value of Start can be used to create a dropped capital; one with a **Placement** value of Inline can be used to create a raised capital.)

List Attribute

The **ListNumbering** attribute, described in Table 9.30, is carried by an L (List) element, but controls the interpretation of the Lbl (Label) elements within the list's Ll (List item) elements (see "List Elements" on page 630). This attribute is

defined in attribute objects whose **O** (owner) entry has the value **List** (or is one of the format-specific owner names listed in Table 9.25 on page 639).

TABLE 9.30 Standard list attribute

KEY	TYPE	VALUE																		
ListNumbering	name	<p>(Optional) The numbering system used to generate the content of the Lbl (Label) elements in an autonumbered list, or the symbol used to identify each item in an unnumbered list:</p> <table> <tr> <td>None</td> <td>No autonumbering; Lbl elements (if present) contain arbitrary text not subject to any numbering scheme</td> </tr> <tr> <td>Disc</td> <td>Solid circular bullet</td> </tr> <tr> <td>Circle</td> <td>Open circular bullet</td> </tr> <tr> <td>Square</td> <td>Solid square bullet</td> </tr> <tr> <td>Decimal</td> <td>Decimal arabic numerals (1–9, 10–99, ...)</td> </tr> <tr> <td>UpperRoman</td> <td>Uppercase roman numerals (I, II, III, IV, ...)</td> </tr> <tr> <td>LowerRoman</td> <td>Lowercase roman numerals (i, ii, iii, iv, ...)</td> </tr> <tr> <td>UpperAlpha</td> <td>Uppercase letters (A, B, C, ...)</td> </tr> <tr> <td>LowerAlpha</td> <td>Lowercase letters (a, b, c, ...)</td> </tr> </table> <p>Default value: None.</p> <p>Note: The alphabet used for <i>UpperAlpha</i> and <i>LowerAlpha</i> is determined by the prevailing Lang entry (see Section 9.8.1, “Natural Language Specification”).</p> <p>Note: The set of possible values may be expanded as Unicode identifies additional numbering systems.</p>	None	No autonumbering; Lbl elements (if present) contain arbitrary text not subject to any numbering scheme	Disc	Solid circular bullet	Circle	Open circular bullet	Square	Solid square bullet	Decimal	Decimal arabic numerals (1–9, 10–99, ...)	UpperRoman	Uppercase roman numerals (I, II, III, IV, ...)	LowerRoman	Lowercase roman numerals (i, ii, iii, iv, ...)	UpperAlpha	Uppercase letters (A, B, C, ...)	LowerAlpha	Lowercase letters (a, b, c, ...)
None	No autonumbering; Lbl elements (if present) contain arbitrary text not subject to any numbering scheme																			
Disc	Solid circular bullet																			
Circle	Open circular bullet																			
Square	Solid square bullet																			
Decimal	Decimal arabic numerals (1–9, 10–99, ...)																			
UpperRoman	Uppercase roman numerals (I, II, III, IV, ...)																			
LowerRoman	Lowercase roman numerals (i, ii, iii, iv, ...)																			
UpperAlpha	Uppercase letters (A, B, C, ...)																			
LowerAlpha	Lowercase letters (a, b, c, ...)																			

Note: This attribute is used to allow a content extraction tool to autonumber a list. However, the Lbl elements within the table should nevertheless contain the resulting numbers explicitly, so that the document can be reflowed or printed without the need for autonumbering.

Table Attributes

The attributes described in Table 9.31 apply only to table cells (structure types TH and TD; see “Table Elements” on page 631). Attributes in this category are defined

in attribute objects whose **O** (owner) entry has the value **Table** (or is one of the format-specific owner names listed in Table 9.25 on page 639).

TABLE 9.31 Standard table attributes

KEY	TYPE	VALUE
RowSpan	integer	<i>(Optional)</i> The number of rows in the enclosing table that are spanned by the cell. The cell expands by adding rows in the block-progression direction specified by the table's WritingMode attribute. Default value: 1.
ColSpan	integer	<i>(Optional)</i> The number of columns in the enclosing table that are spanned by the cell. The cell expands by adding columns in the inline-progression direction specified by the table's WritingMode attribute. Default value: 1.

9.8 Accessibility Support

PDF includes several facilities in support of accessibility of documents to disabled users. In particular, many visually impaired computer users use screen readers to read documents aloud. To enable proper vocalization, either through a screen reader or by some more direct invocation of a text-to-speech engine, PDF supports the following features:

- Specifying the natural language used for text in a PDF document—for example, as English or Spanish (see Section 9.8.1, “Natural Language Specification”)
- Providing textual descriptions for images or other items that do not translate naturally into text (Section 9.8.2, “Alternate Descriptions”), or replacement text for content that does translate into text but is represented in a nonstandard way (such as with a ligature or illuminated character; see Section 9.8.3, “Replacement Text”)
- Specifying the expansion of abbreviations or acronyms (Section 9.8.4, “Expansion of Abbreviations and Acronyms”)

The core of this support lies in the ability to determine the logical order of content in a PDF document, independently of the content's appearance or layout, through logical structure and Tagged PDF, as described under “Page Content Order” on page 618. An accessibility application can extract the content of a document for presentation to a disabled user by traversing the structure hierarchy and presenting the contents of each node. For this reason, producers of PDF files

must ensure that all information in a document is reachable via the structure hierarchy, and they are strongly encouraged to use the facilities described in this section.

Note: Text can be extracted and examined or reused for purposes other than accessibility; see Section 9.7, “Tagged PDF.”

Additional guidelines for accessibility support of content published on the World Wide Web can be found in the World Wide Web Consortium document *Web Content Accessibility Guidelines* and the documents it points to (see the Bibliography).

9.8.1 Natural Language Specification

The natural language used for text in a document is determined in a hierarchical fashion, based on whether an optional **Lang** entry (*PDF 1.4*) is present in any of several possible locations. At the highest level, the document’s default language (which applies to both text strings and text within content streams) can be specified by a **Lang** entry in the document catalog (see Section 3.6.1, “Document Catalog”). Below this, the language can be specified for the following:

- Structure elements of any type (see Section 9.6.1, “Structure Hierarchy”), through a **Lang** entry in the structure element dictionary.
- Marked-content sequences that are not in the structure hierarchy (see Section 9.5, “Marked Content”), through a **Lang** entry in a property list attached to the marked-content sequence with a **Span** tag. (Note that although **Span** is also a standard structure type, as described under “Inline-Level Structure Elements” on page 632, its use here is entirely independent of logical structure.)

The following sections provide details on the value of the **Lang** entry and the hierarchical manner in which the language for text in a document is determined.

Note: Text strings encoded in Unicode may include an escape sequence indicating the language of the text, overriding the prevailing **Lang** entry (see Section 3.8.1, “Text Strings”).

Language Identifiers

The value of the **Lang** entry in the document catalog, structure element dictionary, or property list is a text string that specifies the language with a *language*

identifier having the syntax defined in Internet RFC 1766, *Tags for the Identification of Languages*. This syntax, which is summarized below, is also used to identify languages in XML, according to the World Wide Web Consortium document *Extensible Markup Language (XML) 1.0*; see the Bibliography for more information about these documents. An empty string indicates that the language is unknown.

Language identifiers can be based on codes defined by the International Organization for Standardization in ISO 639 and ISO 3166 (see the Bibliography) or registered with the Internet Assigned Numbers Authority (IANA, whose Web site is located at <<http://iana.org/>>), or they can include codes created for private use. A language identifier consists of a primary code optionally followed by one or more subcodes (each preceded by a hyphen). The primary code can be any of the following:

- A 2-character ISO 639 language code—for example, en for English or es for Spanish
- The letter i, designating an IANA-registered identifier
- The letter x, for private use

The first subcode can be a 2-character ISO 3166 country code, as in en-US, or a 3- to 8-character subcode registered with IANA, as in en-cockney or i-cherokee (except in private identifiers, for which subcodes are not registered). Subcodes beyond the first can be any that have been registered with IANA.

Language Specification Hierarchy

The **Lang** entry in the document catalog specifies the natural language for all text in the document except where overridden by language specifications for structure elements or for marked-content sequences that are not in the structure hierarchy (for example, within an entirely unstructured document). Examples in this section illustrate the hierarchical manner in which the language for text in a document is determined.

Example 9.16 shows how a language specified for the document as a whole could be overridden by one specified for a marked-content sequence within a page's content stream, independent of any logical structure. In this case, the **Lang** entry in the document catalog (not shown) has the value en-US, meaning U.S. English, and it is overridden by the **Lang** property attached (with the **Span** tag) to the

marked-content sequence *Hasta la vista*. The **Lang** property identifies the language for this marked content sequence with the value *es-MX*, meaning Mexican Spanish.

Example 9.16

```

2 0 obj                                % Page object
  << /Type /Page
    /Contents 3 0 R                    % Content stream
  ...
  >>
endobj

3 0 obj                                % Page's content stream
  << /Length ... >>
stream
BT
  (See you later, or as Arnold would say,) Tj
  /Span << /Lang (es-MX) >>          % Start of marked-content sequence
  BDC
  (Hasta la vista.) Tj
  EMC                                  % End of marked-content sequence
ET
endstream
endobj

```

Where logical structure is described (by a structure hierarchy) within a document, the **Lang** entry in the document catalog sets the default for the document, as usual; below that, any language specifications within the structure hierarchy apply in this order:

- A structure element's language specification
- Within a structure element, a language specification for a nested structure element or marked-content sequence

In Example 9.17, the **Lang** entry in the structure element dictionary (specifying English) applies to the marked-content sequence having an **MCID** (marked-content identifier) value of 0 within the indicated page's content stream. However, nested within that marked-content sequence is another one in which the **Lang** property attached with the **Span** tag (specifying Spanish) overrides the structure element's language specification.

Note: This example and the next one below omit required **StructParents** entries in the objects used as content items (see “Finding Structure Elements from Content Items” on page 600).

Example 9.17

```

1 0 obj                                % Structure element
  << /Type /StructElem
    /S /P                                % Structure type
    /P ...                                % Parent in structure hierarchy
    /K << /Type /MCR
      /Pg 2 0 R                            % Page containing marked-content sequence
      /MCID 0                              % Marked-content identifier
    >>
    /Lang (en-US)                          % Language specification for this element
  >>
endobj

2 0 obj                                % Page object
  << /Type /Page
    /Contents 3 0 R                        % Content stream
    ...
  >>
endobj

3 0 obj                                % Page's content stream
  << /Length ... >>
stream
  BT
    /P << /MCID 0 >>                       % Start of marked-content sequence
    BDC
      (See you later, or as Arnold would say,) Tj
    /Span << /Lang (es-MX) >> % Start of nested marked-content sequence
    BDC
      (Hasta la vista.) Tj
    EMC                                     % End of nested marked-content sequence
  EMC                                       % End of marked-content sequence
  ET
endstream
endobj

```

If only part of the page content is contained in the structure hierarchy, and the structured content is nested within nonstructured content to which a different language specification applies, the structure element's language specification takes precedence. In Example 9.18, the page's content stream consists of a marked-content sequence that specifies (via the `Span` tag with a **Lang** property) Spanish as its language, but nested within it is content that is part of a structure element (indicated by the **MCID** entry in that property list), and the language specification that applies to the latter content is that of the structure element, English.

Example 9.18

```

1 0 obj                                % Structure element
  << /Type /StructElem
    /S /P                                % Structure type
    /P ...                               % Parent in structure hierarchy
    /K << /Type /MCR
      /Pg 2 0 R                          % Page containing marked-content sequence
      /MCID 0                            % Marked-content identifier
    >>
    /Lang (en-US)                        % Language specification for this element
  >>
endobj

2 0 obj                                % Page object
  << /Type /Page
    /Contents 3 0 R                      % Content stream
    ...
  >>
endobj

3 0 obj                                % Page's content stream
  << /Length ... >>
stream
  /Span << /Lang (es-MX) >>             % Start of marked-content sequence
  BDC
  (Hasta la vista, ) Tj
  /P << /MCID 0 >>                      % Start of structured marked-content sequence,
  BDC                                    % to which structure element's language applies
  (as Arnold would say.) Tj
  EMC                                    % End of structured marked-content sequence
  EMC                                    % End of marked-content sequence
endstream
endobj

```


In other words, a language identifier attached to a marked-content sequence with the `Span` tag specifies the language for all text in the sequence except for nested marked content that is contained in the structure hierarchy (in which case the structure element's language applies) and except where overridden by language specifications for other nested marked content.

9.8.2 Alternate Descriptions

PDF documents can be enhanced by providing alternate descriptions (as human-readable text) for images, formulas, or other items that do not translate naturally into text. A text-to-speech engine, for example, could vocalize the text of the alternate description for visually impaired users.

An alternate description can be specified for the following:

- A structure element (see Section 9.6.1, “Structure Hierarchy”), through an **Alt** entry in the structure element dictionary
- Any type of annotation (see Section 8.4, “Annotations”) that does not already have a text representation, through a **Contents** entry in the annotation dictionary

For annotation types that normally display text, that text (specified in the **Contents** entry of the annotation dictionary) is the natural source for vocalization purposes. For annotation types that do not display text, a **Contents** entry (*PDF 1.4*) may optionally be included to specify an alternate description. Sound annotations, which are vocalized by default and so need no alternate description for that purpose, may include a **Contents** entry specifying a description that will be displayed in a pop-up window for the benefit of hearing-impaired users.

In addition, an alternate name can be specified for an interactive form field (see Section 8.6, “Interactive Forms”), to be used in place of the actual field name wherever the field must be identified in the user interface (such as in error or status messages referring to the field). This alternate name, specified in the optional **TU** entry of the field dictionary, can be useful for vocalization purposes.

The following are true of all alternate descriptions (or field names):

- They are text strings, which may be encoded in either **PDFDocEncoding** or Unicode character encoding. As described in Section 3.8.1, “Text Strings,” Unicode defines an escape sequence for indicating the language of the text; this mecha-

nism enables the alternate description to change from the language specified by the prevailing **Lang** entry (as described in the preceding section).

- The text is considered to be a word or phrase substitution for the current structure element. For example, if each of two (or more) elements in a sequence has an **Alt** entry in its dictionary, they should be treated as if a word break is present between them.

9.8.3 Replacement Text

Just as alternate descriptions can be provided for images and other items that do not translate naturally into text (as described in the preceding section), replacement text can be specified for content that does translate into text but that is represented in a nonstandard way. These nonstandard representations might include, for example, glyphs for ligatures or custom characters, or inline graphics corresponding to letters in an illuminated manuscript or to dropped capitals.

Replacement text can be specified for a structure element (see Section 9.6.1, “Structure Hierarchy”), in the optional **ActualText** entry (*PDF 1.4*) of the structure element dictionary. The **ActualText** value is not a description but a replacement for the content, providing text that is equivalent to what a sighted reader would see when viewing the content. In contrast to the value of **Alt**, which is considered to be a word or phrase substitution for the structure element, the value of **ActualText** is considered to be a character substitution for the structure element. Thus, if each of two (or more) elements in a sequence has an **ActualText** entry in its dictionary, they should be treated as if no word break is present between them.

Like alternate descriptions (and other text strings), replacement text, if encoded in Unicode, may include an escape sequence for indicating the language of the text, overriding the prevailing **Lang** entry (see Section 3.8.1, “Text Strings”).

9.8.4 Expansion of Abbreviations and Acronyms

Abbreviations and acronyms can pose a problem for text-to-speech engines. Sometimes the full pronunciation for an abbreviation can be divined without aid; for instance, a search through a dictionary will probably reveal that “Blvd.” is pronounced “boulevard” and that “Ave.” is pronounced “avenue.” However, some abbreviations are difficult to resolve, as in the sentence “Dr. Healwell works at 123 Industrial Dr.” For this reason, the expansion of an abbreviation or acronym can

be supplied through an **E** property (*PDF 1.4*) attached to a marked-content sequence with a `Span` tag. For example:

```
BT
  /Span <</E (Doctor) >>
  BDC
    (Dr.) Tj
  EMC
  (Healwell works at 123 Industrial ) Tj
  /Span <</E (Drive) >>
  BDC
    (Dr.) Tj
  EMC
ET
```

The **E** value (a text string) is considered to be a word or phrase substitution for the tagged text and so should be treated as if a word break separates it from any surrounding text. Like other text strings, the expansion text, if encoded in Unicode, may include an escape sequence for indicating the language of the text (see Section 3.8.1, “Text Strings”).

Some abbreviations or acronyms are conventionally not expanded into words. For the text “CBS,” for example, either no expansion should be supplied (leaving its pronunciation up to the text-to-speech engine) or, to be safe, the expansion “C B S” should be specified.

9.9 Web Capture

Web Capture is a PDF 1.3 feature that allows information from Internet-based or locally resident HTML, PDF, GIF, JPEG, and ASCII text files to be imported into a PDF file. This feature is implemented in Acrobat 4.0 and later viewers by a Web Capture plug-in extension (sometimes called AcroSpider). The information in the Web Capture data structures enables viewer applications to perform the following operations:

- Save locally and preserve the visual appearance of material from the World Wide Web
- Retrieve additional material from the Web and add it to an existing PDF file
- Update or modify existing material previously captured from the Web

- Find source information for material captured from the Web, such as the URL (if any) from which it was captured
- Find all material in a PDF file that was generated from a given URL
- Find all material in a PDF file that matches a given digital identifier (MD5 hash)

The information needed to perform these operations is recorded in two data structures in the PDF file:

- The *Web Capture information dictionary* holds document-level information related to Web Capture.
- The Web Capture *content database* keeps track of the material retrieved by Web Capture and where it came from, enabling Web Capture to avoid needlessly downloading material that is already present in the file.

The following sections provide a detailed overview of these structures. See Appendix C for information about implementation limits in Web Capture.

Note: The following discussion centers on HTML and GIF files, although Web Capture handles other file types as well.

9.9.1 Web Capture Information Dictionary

The optional **SpiderInfo** entry in the document catalog (see Section 3.6.1, “Document Catalog”) holds an optional *Web Capture information dictionary* containing document-level information related to Web Capture. Table 9.32 shows the contents of this dictionary.

TABLE 9.32 Entries in the Web Capture information dictionary

KEY	TYPE	VALUE
V	number	(Required) The Web Capture version number. For PDF 1.3, the version number is 1.0. <i>Note:</i> This value is a single real number, not a major and minor version number. Thus, for example, a version number of 1.2 would be considered greater than 1.15.
C	array	(Optional) An array of indirect references to Web Capture command dictionaries (see “Command Dictionaries” on page 672) describing commands that were used in building the PDF file. The commands appear in the array in the order in which they were executed in building the file.

9.9.2 Content Database

Web Capture retrieves HTML files from URLs and converts them into PDF. The resulting PDF file may contain the contents of multiple HTML pages. Conversely, since HTML pages do not have a fixed size, a single HTML page may give rise to multiple PDF pages. To keep track of the correspondences, Web Capture maintains a *content database* mapping URLs and digital identifiers to PDF objects such as pages and XObjects. By looking up digital identifiers in the database, Web Capture can determine whether newly downloaded content is identical to content already retrieved from a different URL. This allows it to perform optimizations such as storing only one copy of an image that is referenced by multiple HTML pages.

Web Capture's content database is organized into *content sets*. Each content set is a dictionary holding information about a group of related PDF objects generated from the same source data. Content sets are of two subtypes: *page sets* and *image sets*. When Web Capture converts an HTML file into PDF pages, for example, it creates a page set to hold information about the pages. Similarly, when it converts a GIF image into one or more image XObjects, it creates an image set describing those XObjects.

The content set corresponding to a given data source can be accessed in either of two ways:

- By the URLs from which it was retrieved
- By a digital identifier generated from the source data itself (see “Digital Identifiers” on page 664)

The **URLS** and **IDS** entries in a PDF document's name dictionary (see Section 3.6.3, “Name Dictionary”) contain name trees mapping URLs and digital identifiers, respectively, to Web Capture content sets. Figure 9.1 shows a simple example. An HTML file retrieved from the URL `<http://www.adobe.com/>` has been converted into three pages in the PDF file. The entry for that URL in the **URLS** name tree points to a page set containing the three pages. Similarly, the **IDS** name tree contains an entry pointing to the same page set, associated with the digital identifier calculated from the HTML source (the string shown in the figure as 904B...1EA2).




FIGURE 9.1 Simple Web Capture file structure

Entries in the **URLS** and **IDS** name trees may refer to an array of content sets instead of just a single content set. The content sets need not have the same subtype, but may include both page sets and image sets. In Figure 9.2, for example, a GIF file has been retrieved from a URL (`<http://www.adobe.com/getacro.gif>`) and converted into a single PDF page. As in Figure 9.1, a page set has been created to hold information about the new page. However, since the retrieval also resulted in a new image XObject, an image set has also been created. Instead of

pointing directly to a single content set, the **URLS** and **IDS** entries point to an array containing both the page set and the image set.




FIGURE 9.2 Complex Web Capture file structure

URL Strings

URLs associated with Web Capture content sets must be reduced to a predictable, canonical form before being used as keys in the **URLS** name tree. The following steps describe how to perform this reduction, using terminology from Internet RFCs 1738, *Uniform Resource Locators*, and 1808, *Relative Uniform Resource Locators* (see the Bibliography). This algorithm is relevant for HTTP, FTP, and file URLs:

1. If the URL is relative, make it absolute.
2. If the URL contains one or more number sign characters (#), strip the leftmost number sign and any characters after it.
3. Convert the scheme section to lowercase ASCII.
4. If there is a host section, convert it to lowercase ASCII.
5. If the scheme is file and the host is localhost, strip the host section.
6. If there is a port section and the port is the default port for the given protocol (80 for HTTP or 21 for FTP), strip the port section.
7. If the path section contains dot (.) or double-dot (..) subsequences, transform the path as described in section 4 of RFC 1808.

***Note:** Because the percent character (%) is “unsafe” according to RFC 1738 and is also the escape character for encoded characters, it is not possible in general to distinguish a URL with unencoded characters from one with encoded characters. For example, it is impossible to decide whether the sequence %00 represents a single encoded null character or a sequence of three unencoded characters. Hence no number of encoding or decoding passes on a URL will ever cause it to reach a stable state. Empirically, URLs embedded in HTML files have unsafe characters encoded with one encoding pass, and Web servers perform one decoding pass on received paths (though CGI scripts are free to make their own decisions). Canonical URLs are thus assumed to have undergone one and only one encoding pass. A URL whose initial encoding state is known can be safely transformed into a URL that has undergone only one encoding pass.*

Digital Identifiers

Digital identifiers associated with Web Capture content sets by the **IDS** name tree are generated using the MD5 message-digest algorithm (described in Internet

RFC 1321, *The MD5 Message-Digest Algorithm*; see the Bibliography). The exact data passed to the algorithm depends on the type of content set and the nature of the identifier being calculated.

For a page set, the source data itself is passed to the MD5 algorithm first, followed by strings representing the digital identifiers of any auxiliary data files (such as images) referenced in the source data, in the order in which they are first referenced. (If an auxiliary file is referenced more than once, its identifier is passed only the first time.) This produces a composite identifier representing the visual appearance of the pages in the page set. Two HTML source files that are identical, but for which the referenced images contain different data—for example, if they have been generated by a script or are pointed to by relative URLs—will not produce the same identifier.

Note: *When the source data is taken from a PDF file, the identifier will be generated solely from the contents of that file; there is no auxiliary data. (See also implementation note 107 in Appendix H.)*

A page set can also have a *text identifier*, calculated by applying the MD5 algorithm to just the rendered text present in the source data. For an HTML file, for example, the text identifier is based solely on the text between markup tags; no images are used in the calculation.

For an image set, the digital identifier is calculated by passing the source data for the original image to the MD5 algorithm. For example, the identifier for an image set created from a GIF image is calculated from the contents of the GIF itself.

Unique Name Generation

In generating PDF pages from a data source, Web Capture converts items such as hypertext links and HTML form fields into corresponding named destinations and interactive form fields. These items must have names that do not conflict with those of existing items in the file. Also, when updating the file, Web Capture may need to locate all destinations and fields constructed for a given page set. Accordingly, each destination or field is given a unique name, derived from its original name but constructed so as to avoid conflicts with similarly named items in other page sets.

Note: As used here, the term *name* refers to a *string*, not a *name object*.

The unique name is formed by appending an encoded form of the page set's digital identifier string to the original name of the destination or field. The identifier string must be encoded to remove characters that have special meaning in destinations and fields. For example, since the period character (.) is used as the field separator in interactive form field names, it must not appear in the identifier portion of the unique name; it is therefore encoded internally as two bytes, 92 and 112, corresponding to the ASCII characters \p. Note that since the backslash character (\) has special meaning for the syntax of string objects, it must be preceded by another backslash when written in the PDF file. For example, if the original digital identifier string were

```
alpha.beta
```

this would be encoded internally as

```
alpha\pbeta
```

and written in the PDF file as

```
(alpha\\pbeta)
```

Similarly, the null character (character code 0) is encoded internally as the two bytes 92 and 48, corresponding to the ASCII characters \0. If the original digital identifier string were

```
alphaØbeta
```

(where Ø denotes the null character), it would be encoded internally as

```
alpha\0beta
```

and written in the PDF file as

```
(alpha\\0beta)
```

Finally, the backslash character itself is encoded internally as the two bytes 92 and 92, corresponding to the characters \\. In written form, each of these in turn requires a preceding backslash. Thus the digital identifier string

```
alpha\beta
```

would be encoded internally as

```
alpha\\beta
```

and written in the PDF file as

```
(alpha\\\\\\beta)
```

If the name is used for an interactive form field, there is an additional encoding to ensure uniqueness and compatibility with interactive forms. Each byte in the source string, encoded as described above, is replaced by two bytes in the destination string. The first byte in each pair is 65 (corresponding to the ASCII character A) plus the high-order 4 bits of the source byte; the second byte is 65 plus the low-order 4 bits of the source byte.

9.9.3 Content Sets

A Web Capture *content set* is a dictionary describing a set of PDF objects generated from the same source data. It may include information common to all the objects in the set as well as about the set itself. Table 9.33 shows the contents of this type of dictionary.

Page Sets

A *page set* is a content set containing a group of PDF page objects generated from a common source, such as an HTML file. The pages are listed in the **O** array (see Table 9.33) in the same order in which they were initially added to the file. A single page object may not belong to more than one page set. Table 9.34 shows the content set dictionary entries specific to this type of content set.

The optional **TID** (text identifier) entry may be used to store an identifier generated from the text of the pages belonging to the page set (see “Digital Identifiers” on page 664). This identifier may be used, for example, to determine whether the text of a document has changed. A text identifier may not be appropriate for some page sets (such as those with no text), and should be omitted in these cases.

TABLE 9.33 Entries common to all Web Capture content sets

KEY	TYPE	VALUE
Type	name	<i>(Optional)</i> The type of PDF object that this dictionary describes; if present, must be SpiderContentSet for a Web Capture content set.
S	name	<i>(Required)</i> The subtype of content set that this dictionary describes: SPS (“Spider page set”) A page set SIS (“Spider image set”) An image set
ID	string	<i>(Required)</i> The digital identifier of the content set (see “Digital Identifiers” on page 664). If the content set has been located via the URLS name tree, this allows its related entry in the IDS name tree to be found.
O	array	<i>(Required)</i> An array of indirect references to the objects belonging to the content set. The order of objects in the array is undefined in general, but may be restricted by specific content set subtypes.
SI	dictionary or array	<i>(Required)</i> A source information dictionary (see Section 9.9.4, “Source Information”), or an array of such dictionaries, describing the sources from which the objects belonging to the content set were created.
CT	string	<i>(Optional)</i> The <i>content type</i> , a string characterizing the source from which the objects belonging to the content set were created. The string should conform to the content type specification described in Internet RFC 2045, <i>Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies</i> (see the Bibliography). For example, for a page set consisting of a group of PDF pages created from an HTML file, the content type would be text/html.
TS	date	<i>(Optional)</i> A time stamp giving the date and time at which the content set was created.

TABLE 9.34 Additional entries specific to a Web Capture page set

KEY	TYPE	VALUE
S	name	<i>(Required)</i> The subtype of content set that this dictionary describes; must be SPS (“Spider page set”) for a page set.
T	text string	<i>(Optional)</i> The title of the page set, a text string representing it in human-readable form.
TID	string	<i>(Optional)</i> A text identifier generated from the text of the page set, as described in “Digital Identifiers” on page 664.

Image Sets

An *image set* is a content set containing a group of image XObjects generated from a common source, such as multiple frames of an animated GIF image. (Web Capture 4.0 will always generate a single image XObject for a given image.) A single XObject may not belong to more than one image set. Table 9.35 shows the content set dictionary entries specific to this type of content set.

TABLE 9.35 Additional entries specific to a Web Capture image set

KEY	TYPE	VALUE
S	name	<i>(Required)</i> The subtype of content set that this dictionary describes; must be SIS (“Spider image set”) for an image set.
R	integer or array	<i>(Required)</i> The reference counts (see below) for the image XObjects belonging to the image set. For an image set containing a single XObject, the value is simply the integer reference count for that XObject. If the image set contains multiple XObjects, the value is an array of reference counts parallel to the O array (see Table 9.33 on page 668); that is, each element in the R array holds the reference count for the image XObject at the corresponding position in the O array.

Each image XObject in an image set has a reference count indicating the number of PDF pages referring to that XObject. The reference count is incremented whenever Web Capture creates a new page referring to the XObject (including copies of already existing pages) and decremented whenever such a page is destroyed. (The reference count is incremented or decremented only once per page, regardless of the number of times the XObject may be referenced by that same page.) When the reference count reaches 0, it is assumed that there are no remaining pages referring to the XObject, and that it can be removed from the image set’s **O** array. (See implementation note 108 in Appendix H.)

9.9.4 Source Information

The **SI** entry in a content set dictionary (see Table 9.33 on page 668) identifies one or more *source information dictionaries* containing information about the locations from which the source data for the content set was retrieved. Table 9.36 shows the contents of this type of dictionary.

TABLE 9.36 Entries in a source information dictionary

KEY	TYPE	VALUE
AU	string or dictionary	<i>(Required)</i> A string or URL alias dictionary (see “URL Alias Dictionaries,” below) identifying the URLs from which the source data was retrieved.
TS	date	<i>(Optional)</i> A time stamp giving the most recent date and time at which the content set’s contents were known to be up to date with the source data.
E	date	<i>(Optional)</i> An expiration stamp giving the date and time at which the content set’s contents should be considered out of date with the source data.
S	integer	<p><i>(Optional)</i> A code indicating the type of form submission, if any, by which the source data was accessed (see “Submit-Form Actions” on page 550):</p> <ul style="list-style-type: none"> 0 Not accessed via a form submission 1 Accessed via an HTTP GET request 2 Accessed via an HTTP POST request <p>This entry should be present only in source information dictionaries associated with page sets. Default value: 0.</p>
C	dictionary	<i>(Optional; must be an indirect reference)</i> A command dictionary (see “Command Dictionaries” on page 672) describing the command that caused the source data to be retrieved. This entry should be present only in source information dictionaries associated with page sets.

In the simplest case, the content set’s **SI** entry just contains a single source information dictionary. However, it is not uncommon for the same source data to be accessible via two or more unrelated URLs. When Web Capture detects such a condition (by comparing digital identifiers), it generates a single content set from the source data, containing just one copy of the relevant PDF pages or image XObjects, but creates multiple source information dictionaries describing the separate ways in which the original source data can be accessed. It then stores an array containing these multiple source information dictionaries as the value of the **SI** entry in the content set dictionary.

A source information dictionary’s **AU** (aliased URLs) entry identifies the URLs from which the source data was retrieved. If there is only one such URL, a simple string suffices as the value of this entry. If multiple URLs map to the same location through redirection, the **AU** value is a URL alias dictionary representing them (see “URL Alias Dictionaries,” below).

Note: For file size efficiency, it is recommended that the entire URL alias dictionary (excluding the URL strings) be represented as a direct object, as its internal structure should never be shared or externally referenced.

The **TS** (time stamp) entry allows each source location associated with a content set to have its own time stamp. This is necessary because the time stamp in the content set dictionary itself (see Table 9.33 on page 668) merely refers to the creation date of the content set. A hypothetical “Update Content Set” command might reset the time stamp in the source information dictionary to the current time if it found that the source data had not changed since the time stamp was last set.

The **E** (expiration) entry specifies an expiration date for each source location associated with a content set. If the current date and time are later than those specified, the contents of the content set should be considered out of date with the original source.

URL Alias Dictionaries

When a URL is accessed via HTTP, a response header may be returned indicating that the requested data is to be found at a different URL. This *redirection* process may be repeated in turn at the new URL, and can potentially continue indefinitely. It is not uncommon to find multiple URLs that all lead eventually to the same destination through one or more redirections. A *URL alias dictionary* represents such a set of URL chains leading to a common destination. Table 9.37 shows the contents of this type of dictionary.

TABLE 9.37 Entries in a URL alias dictionary

KEY	TYPE	VALUE
U	string	<i>(Required)</i> The destination URL to which all of the chains specified by the C entry lead.
C	array	<i>(Optional)</i> An array of one or more arrays of strings, each representing a chain of URLs leading to the common destination specified by U .

The **C** (chains) entry should be omitted if the URL alias dictionary contains only one URL. If **C** is present, its value is an array of arrays, each representing a chain of URLs leading to the common destination. Within each chain, the URLs are

stored as strings in the order in which they occur in the redirection sequence. The common destination (the last URL in a chain) may be omitted, since it is already identified by the **U** entry. (See implementation note 109 in Appendix H.)

Command Dictionaries

A Web Capture *command dictionary* represents a command executed by Web Capture to retrieve one or more pieces of source data that were used to create new pages or modify existing pages. The entries in this dictionary represent parameters that were originally specified interactively by the user who requested that the Web content be captured. This information is recorded so that the command can subsequently be repeated to update the captured content. Table 9.38 shows the contents of this type of dictionary.

TABLE 9.38 Entries in a Web Capture command dictionary

KEY	TYPE	VALUE
URL	string	<i>(Required)</i> The initial URL from which source data was requested.
L	integer	<i>(Optional)</i> The number of levels of pages retrieved from the initial URL. Default value: 1.
F	integer	<i>(Optional)</i> A set of flags specifying various characteristics of the command (see Table 9.39). Default value: 0.
P	string or stream	<i>(Optional)</i> Data that was posted to the URL.
CT	string	<i>(Optional)</i> A content type describing the data posted to the URL. Default value: application/x-www-form-urlencoded.
H	string	<i>(Optional)</i> Additional HTTP request headers sent to the URL.
S	dictionary	<i>(Optional)</i> A command settings dictionary containing settings used in the conversion process (see “Command Settings” on page 674).

The **URL** entry specifies the initial URL for the retrieval command. The **L** (levels) entry specifies the number of levels of pages requested to be retrieved from this URL. If the **L** entry is omitted, its value is assumed to be 1, denoting retrieval of the initial URL only.

The value of the command dictionary’s **F** entry is an unsigned 32-bit integer containing flags specifying various characteristics of the command. Bit positions

within the flag word are numbered from 1 (low-order) to 32 (high-order). Table 9.39 shows the meanings of the flags; all undefined flag bits are reserved and must be set to 0.

TABLE 9.39 Web Capture command flags

BIT POSITION	NAME	MEANING
1	SameSite	If set, pages were retrieved only from the host specified in the initial URL.
2	SamePath	If set, pages were retrieved only from the path specified in the initial URL (see below).
3	Submit	If set, the command represents a form submission (see below).

The SamePath flag, if set, indicates that pages were retrieved only if they were in the same path specified in the initial URL. A page is considered to be in the same path if its scheme and network location components (as defined in Internet RFC 1808, *Relative Uniform Resource Locators*) match those of the initial URL and its path component matches up to and including the last forward slash (/) character in the initial URL. For example, the URL

`http://www.adobe.com/fiddle/faddle/foo.html`

is considered to be in the same path as the initial URL

`http://www.adobe.com/fiddle/initial.html`

The comparison is case-insensitive for the scheme and network location components and case-sensitive for the path component.

If the Submit flag is set, the command represents a form submission. If no **P** (posted data) entry is present, the submitted data is encoded in the URL (an HTTP GET request). If **P** is present, the command represents an HTTP POST request. In this case, the value of the Submit flag is ignored. If the posted data is small enough, it may be represented by a string; for large amounts of data, a stream is recommended, as it can offer compression.

The **CT** (content type) entry is relevant only for POST requests. It describes the content type of the posted data, as described in Internet RFC 2045, *Multipurpose*

Internet Mail Extensions (MIME), Part One: Format of Internet Message Bodies (see the Bibliography).

The **H** (headers) entry specifies additional HTTP request headers that were sent in the request for the URL. Each header line in the string is terminated with a carriage return and a line feed. For example:

```
(Referer: http://frumble.com\015\012From:veeble@frotz.com\015\012)
```

The HTTP request header format is specified in Internet RFC 2068, *Hypertext Transfer Protocol—HTTP/1.1* (see the Bibliography).

The **S** (settings) entry specifies a command settings dictionary (see the next section), holding settings specific to the conversion engines. If this entry is omitted, default values are assumed. It is recommended that command settings dictionaries be shared by any command dictionaries that use the same settings.

Command Settings

The **S** (settings) entry in a command dictionary contains a *command settings dictionary*, which holds settings for conversion engines used in converting the results of the command to PDF. Table 9.40 shows the contents of this type of dictionary.

TABLE 9.40 Entries in a Web Capture command settings dictionary

KEY	TYPE	VALUE
G	dictionary	<i>(Optional)</i> A dictionary containing global conversion engine settings relevant to all conversion engines. If this key is absent, default settings will be used.
C	dictionary	<i>(Optional)</i> Settings for specific conversion engines. Each key in this dictionary is the internal name of a conversion engine (see below). The associated value is a dictionary containing the settings associated with that conversion engine. If the settings for a particular conversion engine are not found in the dictionary, default settings will be used.

Each key in the **C** dictionary is the internal name of a conversion engine, which should be a name object of the following form:

```
/company:product:version:contentType
```

where

company is the name (or abbreviation) of the company that created the conversion engine.

product is the name of the conversion engine. This field may be left blank, but the trailing colon character (:) is still required.

version is the version of the conversion engine.

contentType is an identifier for the content type that the settings are associated with. This is required because some converters may handle multiple content types.

For example:

```
/ADBE:H2PDF:1.0:HTML
```

Note that all fields in the internal name are case-sensitive. The *company* field must conform to the naming guidelines described in Appendix E; the values of the other fields are unrestricted, except that they must not contain a colon.

Note: *It must be possible to make a deep copy of a command settings dictionary without explicit knowledge of the settings it may contain. To facilitate this operation, the directed graph of PDF objects rooted by the command settings dictionary must be entirely self-contained; that is, it must not contain any object referred to from elsewhere in the PDF file.*

9.9.5 Object Attributes Related to Web Capture

A given page object or image XObject can belong to at most one Web Capture content set, called its *parent content set*. However, the object has no direct pointer to its parent content set; such a pointer might present problems for an application that traces all pointers from an object to determine, for example, what resources the object depends on. Instead, the object's **ID** entry (see Tables 3.18 on page 88 and 4.35 on page 267) contains the digital identifier of the parent content set, which can be used to locate the parent content set via the **IDS** name tree in the document's name dictionary. (If the **IDS** entry for the identifier contains an array of content sets, the parent can be found by searching the array for the content set whose **O** entry includes the child object.)

In the course of creating PDF pages from HTML files, Web Capture frequently scales the contents down to fit on fixed-sized pages. The **PZ** (preferred zoom) entry in a page object (see “Page Objects” on page 87) specifies a magnification factor by which the page can be scaled to undo the downscaling and view the page at its original size. That is, when the page is viewed at the preferred magnification factor, one unit in default user space will correspond to one original source pixel.

9.10 Prepress Support

This section describes features of PDF that support prepress production workflows:

- The specification of *page boundaries* governing various aspects of the prepress process, such as cropping, bleed, and trimming (Section 9.10.1, “Page Boundaries”)
- Facilities for including *printer’s marks* such as registration targets, gray ramps, color bars, and cut marks to assist in the production process (Section 9.10.2, “Printer’s Marks”)
- Information for generating *color separations* for pages in a document (Section 9.10.3, “Separation Dictionaries”)
- *Output intents* for matching the color characteristics of a document with those of a target output device or production environment in which it will be printed (Section 9.10.4, “Output Intents”)
- Support for the generation of *traps* to minimize the visual effects of misregistration between multiple colorants (Section 9.10.5, “Trapping Support”)
- The *Open Prepress Interface (OPI)* for creating low-resolution proxies for high-resolution images (Section 9.10.6, “Open Prepress Interface (OPI)”)

9.10.1 Page Boundaries

A PDF page may be prepared either for a finished medium, such as a sheet of paper, or as part of a prepress process in which the content of the page is placed on an intermediate medium, such as film or an imposed reproduction plate. In the latter case, it is important to distinguish between the intermediate page and the finished page. The intermediate page may often include additional production-related content, such as bleeds or printer marks, that falls outside the

boundaries of the finished page. To handle such cases, a PDF page can define as many as five separate boundaries to control various aspects of the imaging process:

- The *media box* defines the boundaries of the physical medium on which the page is to be printed. It may include any extended area surrounding the finished page for bleed, printing marks, or other such purposes. It may also include areas close to the edges of the medium that cannot be marked because of physical limitations of the output device. Content falling outside this boundary can safely be discarded without affecting the meaning of the PDF file.
- The *crop box* defines the region to which the contents of the page are to be clipped (cropped) when displayed or printed. Unlike the other boxes, the crop box has no defined meaning in terms of physical page geometry or intended use; it merely imposes clipping on the page contents. However, in the absence of additional information (such as imposition instructions specified in a JDF or PJTF job ticket), the crop box will determine how the page's contents are to be positioned on the output medium. The default value is the page's media box.
- The *bleed box* (*PDF 1.3*) defines the region to which the contents of the page should be clipped when output in a production environment. This may include any extra "bleed area" needed to accommodate the physical limitations of cutting, folding, and trimming equipment. The actual printed page may include printing marks that fall outside the bleed box. The default value is the page's crop box.
- The *trim box* (*PDF 1.3*) defines the intended dimensions of the finished page after trimming. It may be smaller than the media box, to allow for production-related content such as printing instructions, cut marks, or color bars. The default value is the page's crop box.
- The *art box* (*PDF 1.3*) defines the extent of the page's meaningful content (including potential white space) as intended by the page's creator. The default value is the page's crop box.

These boundaries are specified by the **MediaBox**, **CropBox**, **BleedBox**, **TrimBox**, and **ArtBox** entries, respectively, in the page object dictionary (see Table 3.18 on page 88). All of them are rectangles expressed in default user space units. The crop, bleed, trim, and art boxes should not ordinarily extend beyond the boundaries of the media box; if they do, they will be effectively reduced to their inter-

section with the media box. Figure 9.3 illustrates the relationships among these boundaries. (The crop box is not shown in the figure because it has no defined relationship with any of the other boundaries.)




FIGURE 9.3 Page boundaries

How the various boundaries are used depends on the purpose to which the page is being put. Typical purposes might include the following:

- *Placing the content of a page in another application.* The art box determines the boundary of the content that is to be placed in the application. Depending on the applicable usage conventions, the placed content may be clipped to either the art box or the bleed box. (For example, a quarter-page advertisement to be placed on a magazine page might be clipped to the art box on the two sides of the ad that face into the middle of the page and to the bleed box on the two sides that bleed over the edge of the page.) The media box and trim box are ignored.
- *Printing a finished page.* This case is typical of desktop or shared page printers, in which the page content is positioned directly on the final output medium. The art box and bleed box are ignored. The media box may be used as advice for selecting media of the appropriate size. The crop box and trim box, if present, should be the same as the media box. (See implementation note 110 in Appendix H.)
- *Printing an intermediate page for use in a prepress process.* The art box is ignored. The bleed box defines the boundary of the content to be imaged. The trim box specifies the positioning of the content on the medium; it may also be used to generate cut or fold marks outside the bleed box. Content falling within the media box but outside the bleed box may or may not be imaged, depending on the specific production process being used.
- *Building an imposition of multiple pages on a press sheet.* The art box is ignored. The bleed box defines the clipping boundary of the content to be imaged; content outside the bleed box will be ignored. The trim box specifies the positioning of the page's content within the imposition. Cut and fold marks are typically generated for the imposition as a whole.

In the scenarios above, an application that interprets the bleed, trim, and art boxes for some purpose will typically alter the crop box so as to impose the clipping that those boxes prescribe.

Display of Page Boundaries

For the user's convenience, viewer applications may wish to offer the ability to display guidelines on the screen for the various page boundaries. The optional **BoxColorInfo** entry in a page object (see "Page Objects" on page 87) holds a *box*

color information dictionary (PDF 1.4) specifying the colors and other visual characteristics to be used for such display. Viewer applications will typically provide a user interface to allow the user to set these characteristics interactively. Note that this information is page-specific and can vary from one page to another.

As shown in Table 9.41, the box color information dictionary contains an optional entry for each of the possible page boundaries other than the media box; the value of each entry is in turn a *box style dictionary*, whose contents are shown in Table 9.42. If a given entry is absent, the viewer application should use its own current default settings instead.

9.10.2 Printer's Marks

Printer's marks are graphic symbols or text added to a page to assist production personnel in identifying components of a multiple-plate job and maintaining consistent output during production. Examples commonly used in the printing industry include:

- Registration targets for aligning plates
- Gray ramps and color bars for measuring colors and ink densities
- Cut marks showing where the output medium is to be trimmed

Although PDF producer applications traditionally include such marks in the content stream of a document, they are logically separate from the content of the page itself and typically appear outside the boundaries (the crop box, trim box, and art box) defining the extent of that content (see Section 9.10.1, “Page Boundaries”).

Printer's mark annotations (PDF 1.4) provide a mechanism for incorporating printer's marks into the PDF representation of a page, while keeping them separate from the actual page content. Each page in a PDF document may contain any number of such annotations, each of which represents a single printer's mark.

Note: *Because printer's marks typically fall outside the page's content boundaries, each mark must be represented as a separate annotation. Otherwise—if, for example, the cut marks at the four corners of the page were defined in a single annotation—the annotation rectangle would encompass the entire contents of the page and could interfere with the user's ability to select content or interact with other annotations on the page. Defining printer's marks in separate annotations also facilitates the implementation of a drag-and-drop user interface for specifying them.*

TABLE 9.41 Entries in a box color information dictionary

KEY	TYPE	VALUE
CropBox	dictionary	<i>(Optional)</i> A box style dictionary (see Table 9.42) specifying the visual characteristics for displaying guidelines for the page's crop box. This entry is ignored if no crop box is defined in the page object.
BleedBox	dictionary	<i>(Optional)</i> A box style dictionary (see Table 9.42) specifying the visual characteristics for displaying guidelines for the page's bleed box. This entry is ignored if no bleed box is defined in the page object.
TrimBox	dictionary	<i>(Optional)</i> A box style dictionary (see Table 9.42) specifying the visual characteristics for displaying guidelines for the page's trim box. This entry is ignored if no trim box is defined in the page object.
ArtBox	dictionary	<i>(Optional)</i> A box style dictionary (see Table 9.42) specifying the visual characteristics for displaying guidelines for the page's art box. This entry is ignored if no art box is defined in the page object.

TABLE 9.42 Entries in a box style dictionary

KEY	TYPE	VALUE
C	array	<i>(Required)</i> An array of three numbers in the range 0.0 to 1.0, representing the components in the DeviceRGB color space of the color to be used for displaying the guidelines. Default value: [0.0 0.0 0.0].
W	number	<i>(Optional)</i> The guideline width in default user space units. Default value: 1.
S	name	<i>(Optional)</i> The guideline style: <ul style="list-style-type: none"> S (Solid) A solid rectangle. D (Dashed) A dashed rectangle. The dash pattern is specified by the D entry (see below). Other guideline styles may be defined in the future. Default value: S.
D	array	<i>(Optional)</i> A <i>dash array</i> defining a pattern of dashes and gaps to be used in drawing dashed guidelines (guideline style D above). The dash array is specified in default user space units, in the same format as in the line dash pattern parameter of the graphics state (see "Line Dash Pattern" on page 155). The dash phase is not specified and is assumed to be 0. For example, a D entry of [3 2] specifies guidelines drawn with 3-point dashes alternating with 2-point gaps. Default value: [3].

The visual presentation of a printer's mark is defined by a form XObject specified as an appearance stream in the **N** (normal) entry of the printer's mark annotation's appearance dictionary (see Section 8.4.4, "Appearance Streams"). More than one appearance may be defined for the same printer's mark, to meet the requirements of different regions or production facilities; in this case, the appearance dictionary's **N** entry holds a subdictionary containing the alternate appearances, each identified by an arbitrary key. The **AS** (appearance state) entry in the annotation dictionary designates one of them to be displayed or printed.

Note: The printer's mark annotation's appearance dictionary may include **R** (roll-over) or **D** (down) entries, but appearances defined in either of these entries will never be displayed or printed.

Like all annotations, a printer's mark annotation is defined by an annotation dictionary (see Section 8.4.1, "Annotation Dictionaries"); its annotation type is **PrinterMark**. The **AP** (appearances) and **F** (flags) entries (which ordinarily are optional) must be present, as must the **AS** (appearance state) entry if the appearance dictionary **AP** contains more than one appearance stream. The Print and ReadOnly flags in the **F** entry must be set and all others clear (see Section 8.4.2, "Annotation Flags"). Table 9.43 shows an additional annotation dictionary entry specific to this type of annotation.

TABLE 9.43 Additional entries specific to a printer's mark annotation

KEY	TYPE	VALUE
Subtype	name	<i>(Required)</i> The type of annotation that this dictionary describes; must be PrinterMark for a printer's mark annotation.
MN	name	<i>(Optional)</i> An arbitrary name identifying the type of printer's mark, such as ColorBar or RegistrationTarget.

The form dictionary defining a printer's mark can contain the optional entries shown in Table 9.44 in addition to the standard ones common to all form dictionaries (see Section 4.9.1, "Form Dictionaries").

TABLE 9.44 Additional entries specific to a printer's mark form dictionary

KEY	TYPE	VALUE
MarkStyle	text string	<i>(Optional; PDF 1.4)</i> A text string representing the printer's mark in human-readable form, suitable for presentation to the user on the screen.
Colorants	dictionary	<i>(Optional; PDF 1.4)</i> A dictionary identifying the individual colorants associated with a printer's mark such as a color bar. For each entry in this dictionary, the key is a colorant name and the value is an array defining a Separation color space for that colorant (see "Separation Color Spaces" on page 201). The key must match the colorant name given in that color space.

9.10.3 Separation Dictionaries

In high-end printing workflows, pages are ultimately produced as sets of *separations*, one per colorant (see "Separation Color Spaces" on page 201). Ordinarily, each page in a PDF file is treated as a composite page that paints graphics objects using all the process colorants and perhaps some spot colorants as well. In other words, all separations for a page are generated from a single PDF description of that page.

In some workflows, however, pages are *pre-separated* prior to the generation of the PDF file. In a pre-separated PDF file, the separations for a page are described as separate page objects, each painting only a single colorant (usually specified in the **DeviceGray** color space). When this is done, additional information is needed to identify the actual colorant associated with each separation and to group together the page objects representing all the separations for a given page. This information is contained in a *separation dictionary* (PDF 1.3) in the **Separation-Info** entry of each page object (see "Page Objects" on page 87). Table 9.45 shows the contents of this type of dictionary.

TABLE 9.45 Entries in a separation dictionary

KEY	TYPE	VALUE
Pages	array	<i>(Required)</i> An array of indirect references to page objects representing separations of the same document page. One of the page objects in the array must be the one with which this separation dictionary is associated, and all of them must have separation dictionaries (SeparationInfo entries) containing Pages arrays identical to this one.
DeviceColorant	name or string	<i>(Required)</i> The name of the device colorant to be used in rendering this separation, such as Cyan or PANTONE 35 CV.
ColorSpace	array	<i>(Optional)</i> An array defining a Separation or DeviceN color space (see “Separation Color Spaces” on page 201 and “DeviceN Color Spaces” on page 205). This provides additional information about the color specified by DeviceColorant —in particular, the alternate color space and tint transformation function that would be used to represent the colorant as a process color. This information enables a viewer application to preview the separation in a color that approximates the device colorant. The value of DeviceColorant must match the space’s colorant name (if it is a Separation space) or be one of the space’s colorant names (if it is a DeviceN space).

9.10.4 Output Intents

Output intents (PDF 1.4) provide a means for matching the color characteristics of a PDF document with those of a target output device or production environment in which the document will be printed. The optional **OutputIntents** entry in the document catalog (see Section 3.6.1, “Document Catalog”) holds an array of *output intent dictionaries*, each describing the color reproduction characteristics of a possible output device or production condition. The contents of these dictionaries can vary for different devices and conditions; the dictionary’s **S** entry specifies an *output intent subtype* that determines the format and meaning of the remaining entries.

This use of multiple output intents allows the production process to be customized to the expected workflow and the specific tools available. For example, one production facility might process files conforming to a recognized format standard such as PDF/X-1, while another uses custom Acrobat plug-in extensions to produce *RGB* output for document distribution on the World Wide Web; each of

these workflows would require different sets of output intent information. Multiple output intents also allow the same PDF file to be distributed unmodified to multiple production facilities. The choice of which output intent to use in a given production environment is a matter for agreement between the purchaser and provider of production services; PDF intentionally does not include a selector for choosing a particular output intent from within the PDF file itself.

At the time of publication, only one output intent subtype, **GTS_PDFX**, has been defined, corresponding to the PDF/X format standard (available on the World Wide Web at <<http://www.npes.org/standards/>>); Table 9.46 shows the contents of this type of output intent dictionary. Other subtypes may be added in the future; the names of any such additional subtypes must conform to the naming guidelines described in Appendix E.

TABLE 9.46 Entries in a PDF/X output intent dictionary

KEY	TYPE	VALUE
Type	name	<i>(Optional)</i> The type of PDF object that this dictionary describes; if present, must be OutputIntent for an output intent dictionary.
S	name	<i>(Required)</i> The output intent subtype; must be GTS_PDFX for a PDF/X output intent.
OutputCondition	text string	<i>(Optional)</i> A text string concisely identifying the intended output device or production condition in human-readable form. This is the preferred method of defining such a string for presentation to the user.
OutputConditionIdentifier	string	<i>(Required)</i> A string identifying the intended output device or production condition in human- or machine-readable form. If human-readable, this string may be used in lieu of an OutputCondition string for presentation to the user.

A typical value for this entry would be the name of a production condition maintained in an industry-standard registry such as the *ICC Characterization Data Registry* (see the Bibliography). If the designated condition matches that in effect at production time, it is the responsibility of the production software to provide the corresponding ICC profile as defined in the registry.

If the intended production condition is not a recognized standard, the value **Custom** is recommended for this entry; the **DestOutputProfile** entry defines the ICC profile and the **Info** entry is used for further human-readable identification.

RegistryName	string	<i>(Optional)</i> A string (conventionally a uniform resource identifier, or URI) identifying the registry in which the condition designated by OutputConditionIdentifier is defined.
Info	text string	<i>(Required if OutputConditionIdentifier does not specify a standard production condition; optional otherwise)</i> A human-readable text string containing additional information or comments about the intended target device or production condition.
DestOutputProfile	stream	<i>(Required if OutputConditionIdentifier does not specify a standard production condition; optional otherwise)</i> An ICC profile stream defining the transformation from the PDF document's source colors to output device colorants. The format of the profile stream is the same as that used in specifying an ICCBased color space (see "ICCBased Color Spaces" on page 189). The output transformation uses the profile's "from CIE" information (<i>BToA</i> in ICC terminology); the "to CIE" (<i>AToB</i>) information can optionally be used to remap source color values to some other destination color space, such as for screen preview or hardcopy proofing. (See implementation note 111 in Appendix H.)

Note: PDF/X is actually a family of standards representing varying levels of conformance. The standard for a given conformance level may prescribe further restrictions on the usage and meaning of entries in the output intent dictionary. Any such restrictions take precedence over the descriptions given in Table 9.46.

The ICC profile information in an output intent dictionary supplements rather than replaces that in an **ICCBased** or default color space (see "ICCBased Color Spaces" on page 189 and "Default Color Spaces" on page 194). Those mechanisms are specifically intended for describing the characteristics of source color component values; an output intent can be used in conjunction with them to convert source colors to those required for a specific production condition or to enable the display or proofing of the intended output.

The data in an output intent dictionary is provided for informational purposes only, and PDF consumer applications are free to disregard it. In particular, there is no expectation that PDF production tools will automatically convert colors expressed in the same source color space to the specified target space before generating output. (In some workflows, such conversion may, in fact, be undesirable. For example, when working with *CMYK* source colors tagged with a source ICC

profile solely for purposes of characterization, converting such colors from four components to three and back is unnecessary and will result in a loss of fidelity in the values of the black component; see “Implicit Conversion of CIE-Based Color Spaces” on page 195 for further discussion.) On the other hand, when source colors are expressed in different base color spaces—for example, when combining separately generated images on the same PDF page—it is possible (though not required) to use the destination profile specified in the output intent dictionary to convert source colors to the same target color space. (See implementation note 112 in Appendix H.)

Example 9.19 shows a PDF/X output intent dictionary based on an industry-standard production condition (CGATS TR 001) from the *ICC Characterization Data Registry*; Example 9.20 shows one for a custom production condition.

Example 9.19

```
<< /Type /OutputIntent                % Output intent dictionary
    /S /GTS_PDFX
    /OutputCondition (CGATS TR 001 (SWOP))
    /OutputConditionIdentifier (CGATS TR 001)
    /RegistryName (http://www.color.org)
    /DestOutputProfile 100 0 R
>>

100 0 obj                               % ICC profile stream
<< /N 4
    /Length 1605
    /Filter /ASCIIHexDecode
>>
stream
00 00 02 0C 61 70 ... >
endstream
endobj
```

Example 9.20

```
<< /Type /OutputIntent                % Output intent dictionary
    /S /GTS_PDFX
    /OutputCondition (Coated)
    /OutputConditionIdentifier (Custom)
    /Info (Coated 150lpi)
    /DestOutputProfile 100 0 R
>>
```

```
100 0 obj                                % ICC profile stream
  << /N 4
    /Length 1605
    /Filter /ASCIIHexDecode
  >>
  stream
  00 00 02 0C 61 70 ... >
  endstream
endobj
```

9.10.5 Trapping Support

On devices such as offset printing presses, which mark multiple colorants on a single sheet of physical medium, mechanical limitations of the device can cause imprecise alignment, or *misregistration*, between colorants. This can produce unwanted visual artifacts such as brightly colored gaps or bands around the edges of printed objects. In high-quality reproduction of color documents, such artifacts are commonly avoided by creating an overlap, called a *trap*, between areas of adjacent color.

Figure 9.4 shows an example of trapping. The light and medium grays represent two different colorants, which are used to paint the background and the glyph denoting the letter A. The first figure shows the intended result, with the two colorants properly registered. The second figure shows what happens when the colorants are misregistered. In the third figure, traps have been overprinted along the boundaries, obscuring the artifacts caused by the misregistration. (For emphasis, the traps are shown here in dark gray; in actual practice, their color would be similar to one of the adjoining colors.)




FIGURE 9.4 *Trapping example*

Trapping can be implemented by the application generating a PDF file, by some intermediate application that adds traps to a PDF document, or by the raster image processor (RIP) that produces final output. In the last two cases, the trapping process is controlled by a set of *trapping instructions*, which define two kinds of information:

- *Trapping zones* within which traps should be created
- *Trapping parameters* specifying the nature of the traps within each zone

Trapping zones and trapping parameters are discussed fully in Sections 6.3.2 and 6.3.3, respectively, of the *PostScript Language Reference*, Third Edition. Trapping instructions are not directly specified in a PDF file (as they are in a PostScript file); instead, they are specified in a *job ticket* that accompanies the PDF file or can be embedded within it. Various standards exist for the format of job tickets; two of them, JDF (Job Definition Format) and PJTF (Portable Job Ticket Format), are described in the CIP4 document *JDF Specification* and in Adobe Technical Note #5620, *Portable Job Ticket Format* (see the Bibliography).

When trapping is performed before the production of final output, the resulting traps are placed in the PDF file for subsequent use. The traps themselves are described as a content stream in a trap network annotation (see below). The stream dictionary can include additional entries describing the method that was used to produce the traps and other information about their appearance.

Trap Network Annotations

A complete set of traps generated for a given page under a specified set of trapping instructions is called a *trap network* (*PDF 1.3*). It is a form XObject containing graphics objects for painting the required traps on the page. A page may have more than one trap network based on different trapping instructions, presumably intended for different output devices. All of the trap networks for a given page are contained in a single *trap network annotation* (see Section 8.4, “Annotations”). There can be at most one trap network annotation per page, which must be the last element in the page’s **Annots** array (see “Page Objects” on page 87). This ensures that the trap network is printed after all of the page’s other contents. (See implementation note 113 in Appendix H.)

The form XObject defining a trap network is specified as an appearance stream in the **N** (normal) entry of the trap network annotation’s appearance dictionary (see

Section 8.4.4, “Appearance Streams”). If more than one trap network is defined for the same page, the **N** entry holds a subdictionary containing the alternate trap networks, each identified by an arbitrary key. The **AS** (appearance state) entry in the annotation dictionary designates one of them as the *current trap network* to be displayed or printed.

*Note: The trap network annotation’s appearance dictionary may include **R** (rollover) or **D** (down) entries, but appearances defined in either of these entries will never be printed.*

Like all annotations, a trap network annotation is defined by an annotation dictionary (see Section 8.4.1, “Annotation Dictionaries”); its annotation type is **TrapNet**. The **AP** (appearances), **AS** (appearance state), and **F** (flags) entries (which ordinarily are optional) must be present, with the Print and ReadOnly flags set and all others clear (see Section 8.4.2, “Annotation Flags”). Table 9.47 shows the additional annotation dictionary entries specific to this type of annotation.

The **Version** and **AnnotStates** entries, if present, are used to detect changes in the content of a page that might require regenerating its trap networks. The **Version** array identifies elements of the page’s content that might be changed by an editing application and thus invalidate its trap networks. Because there is at most one **Version** array per trap network annotation (and thus per page), any application generating a new trap network must also verify the validity of existing trap networks by enumerating the objects identified in the array and verifying that the results exactly match the array’s current contents. Any trap networks found to be invalid must be regenerated. (See implementation notes 114 and 115 in Appendix H.)

Beginning with PDF 1.4, the **LastModified** entry can be used in place of the **Version** array to track changes to a page’s trap network. (The trap network annotation must include either a **LastModified** entry or the combination of **Version** and **AnnotStates**, but not all three.) If the modification date in the **LastModified** entry of the page object (see “Page Objects” on page 87) is more recent than the one in the trap network annotation dictionary, then the page’s trap networks are invalid and must be regenerated. Note, however, that not all editing applications and plug-in extensions correctly maintain these modification dates; this method of tracking trap network modifications can be used reliably only in a controlled workflow environment where the integrity of the modification dates is assured.

TABLE 9.47 Additional entries specific to a trap network annotation

KEY	TYPE	VALUE
Subtype	name	<i>(Required)</i> The type of annotation that this dictionary describes; must be TrapNet for a trap network annotation.
Contents	text string	<i>(Optional; PDF 1.4)</i> An alternate description of the annotation's contents in human-readable form, useful when extracting the document's contents in support of accessibility to disabled users or for other purposes (see Section 9.8.2, "Alternate Descriptions").
LastModified	date	<i>(Required if Version and AnnotStates are absent; must be absent if Version and AnnotStates are present; PDF 1.4)</i> The date and time (see Section 3.8.2, "Dates") when the trap network was most recently modified.
Version	array	<p><i>(Required if AnnotStates is present; must be absent if LastModified is present)</i> An unordered array of all objects present in the page description at the time the trap networks were generated and that, if changed, could affect the appearance of the page. If present, the array must include the following objects:</p> <ul style="list-style-type: none"> • All content streams identified in the page object's Contents entry (see "Page Objects" on page 87) • All resource objects (other than procedure sets) in the page's resource dictionary (see Section 3.7.2, "Resource Dictionaries") • All resource objects (other than procedure sets) in the resource dictionaries of any form XObjects on the page (see Section 4.9, "Form XObjects") • All OPI dictionaries associated with XObjects on the page (see Section 9.10.6, "Open Prepress Interface (OPI)")
AnnotStates	array	<i>(Required if Version is present; must be absent if LastModified is present)</i> An array of name objects representing the appearance states (value of the AS entry) for annotations associated with the page. The appearance states must be listed in the same order as the annotations in the page's Annots array (see "Page Objects" on page 87). For an annotation with no AS entry, the corresponding array element should be null . No appearance state should be included for the trap network annotation itself.
FontFauxing	array	<i>(Optional)</i> An array of font dictionaries representing fonts that were "fauxed" (replaced by substitute fonts) during the generation of trap networks for the page.

Trap Network Appearances

Each entry in the **N** (normal) subdictionary of a trap network annotation's appearance dictionary holds an appearance stream defining a trap network associated with the given page. Like all appearances, a trap network is a stream object defining a form XObject (see Section 4.9, "Form XObjects"). The body of the stream contains the graphics objects needed to paint the traps making up the trap network. Its dictionary entries include, besides the standard entries for a form dictionary, the additional entries shown in Table 9.48.

TABLE 9.48 Additional entries specific to a trap network appearance stream

KEY	TYPE	VALUE
PCM	name	<i>(Required)</i> The name of the process color model that was assumed when this trap network was created; equivalent to the PostScript page device parameter ProcessColorModel (see Section 6.2.5 of the <i>PostScript Language Reference</i> , Third Edition). Valid values are DeviceGray , DeviceRGB , DeviceCMYK , DeviceCMY , DeviceRGBK , and DeviceN .
SeparationColorNames	array	<i>(Optional)</i> An array of names identifying the colorants that were assumed when this network was created; equivalent to the PostScript page device parameter of the same name (see Section 6.2.5 of the <i>PostScript Language Reference</i> , Third Edition). Colorants implied by the process color model PCM are available automatically and need not be explicitly declared. If this entry is absent, the colorants implied by PCM are assumed.
TrapRegions	array	<i>(Optional)</i> An array of indirect references to TrapRegion objects defining the page's trapping zones and the associated trapping parameters, as described in Adobe Technical Note #5620, <i>Portable Job Ticket Format</i> . These references are to objects comprising portions of a PJTF job ticket that is embedded in the PDF file. When the trapping zones and parameters are defined by an external job ticket (or by some other means, such as with JDF), this entry is absent.
TrapStyles	text string	<i>(Optional)</i> A human-readable text string that applications can use to describe this trap network to the user (for example, to allow switching between trap networks).

Note: Pre-separated PDF files (see Section 9.10.3, “Separation Dictionaries”) cannot be trapped, because traps are defined along the borders between different colors and a pre-separated file uses only one color. Pre-separation must therefore occur after trapping, not before. An application pre-separating a trapped PDF file is responsible for calculating new **Version** arrays for the separated trap networks.

9.10.6 Open Prepress Interface (OPI)

The workflow in a prepress environment often involves multiple applications in areas such as graphic design, page layout, word processing, photo manipulation, and document construction. As pieces of the final document are moved from one application to another, it is useful to separate the data of high-resolution images, which can be quite large—in some cases, many times the size of the rest of the document combined—from that of the document itself. The *Open Prepress Interface (OPI)* is a mechanism, originally developed by Aldus Corporation, for creating low-resolution placeholders, or *proxies*, for such high-resolution images. The proxy typically consists of a downsampled version of the full-resolution image, to be used for screen display and proofing. Before the document is printed, it passes through a filter known as an *OPI server*, which replaces the proxies with the original full-resolution images.

In PostScript programs, OPI proxies are defined by PostScript code surrounded by special *OPI comments*, which specify such information as the placement and cropping of the image and adjustments to its size, rotation, color, and other attributes. In PDF, proxies are embedded in a document as image or form XObjects with an associated *OPI dictionary (PDF 1.2)* containing the same information conveyed in PostScript by the OPI comments. Two versions of OPI are supported, versions 1.3 and 2.0. In OPI 1.3, a proxy consisting of a single image, with no changes in the graphics state, may be represented as an image XObject; otherwise it must be a form XObject. In OPI 2.0, the proxy always entails changes in the graphics state and hence must be represented as a form XObject. (See implementation notes 116 and 117 in Appendix H.)

An XObject representing an OPI proxy must contain an **OPI** entry in its image or form dictionary (see Tables 4.35 on page 267 and 4.41 on page 284). The value of this entry is an *OPI version dictionary* (Table 9.49) identifying the version of OPI to which the proxy corresponds. This dictionary consists of a single entry, whose key is the name **1.3** or **2.0** and whose value is the OPI dictionary defining the proxy’s OPI attributes.

TABLE 9.49 Entry in an OPI version dictionary

KEY	TYPE	VALUE
<i>version number</i>	dictionary	<i>(Required; PDF 1.2)</i> An OPI dictionary specifying the attributes of this proxy (see Tables 9.50 and 9.51). The key for this entry must be the name 1.3 or 2.0 , identifying the version of OPI to which the proxy corresponds.

Note: As in any other PDF dictionary, the key in an OPI version dictionary must be a name object. The OPI version dictionary would thus be written in the PDF file in either the form

```
<< /1.3 d0R >>                                % OPI 1.3 dictionary
```

or

```
<< /2.0 d0R >>                                % OPI 2.0 dictionary
```

where *d* is the object number of the corresponding OPI dictionary.

Tables 9.50 and 9.51 describe the contents of the OPI dictionaries for OPI 1.3 and OPI 2.0, respectively, along with the corresponding PostScript OPI comments. The dictionary entries are listed in the order in which the corresponding OPI comments should appear in a PostScript program. Complete details on the meanings of these entries and their effects on OPI servers can be found in *OPI: Open Prepress Interface Specification 1.3* (available from Adobe) and Adobe Technical Note #5660, *Open Prepress Interface (OPI) Specification, Version 2.0*.

TABLE 9.50 Entries in a version 1.3 OPI dictionary

KEY	TYPE	OPI COMMENT	VALUE
Type	name		<i>(Optional)</i> The type of PDF object that this dictionary describes; if present, must be OPI for an OPI dictionary.
Version	number		<i>(Required)</i> The version of OPI to which this dictionary refers; must be the number 1.3 (not the name 1.3, as in an OPI version dictionary).
F	file specification	%ALDImageFilename	<i>(Required)</i> The external file containing the image corresponding to this proxy. (See implementation note 118 in Appendix H.)

ID	string	%ALDImageID	<i>(Optional)</i> An identifying string denoting the image.
Comments	text string	%ALDObjectComments	<i>(Optional)</i> A human-readable comment, typically containing instructions or suggestions to the operator of the OPI server on how to handle the image.
Size	array	%ALDImageDimensions	<i>(Required)</i> An array of two integers of the form $[pixelsWide\ pixelsHigh]$ specifying the dimensions of the image in pixels.
CropRect	rectangle	%ALDImageCropRect	<i>(Required)</i> An array of four integers of the form $[left\ top\ right\ bottom]$ specifying the portion of the image to be used.
CropFixed	array	%ALDImageCropFixed	<i>(Optional)</i> An array with the same form and meaning as CropRect , but expressed in real numbers instead of integers. Default value: the value of CropRect .
Position	array	%ALDImagePosition	<i>(Required)</i> An array of eight numbers of the form $[ll_x\ ll_y\ ul_x\ ul_y\ ur_x\ ur_y\ lr_x\ lr_y]$ specifying the location on the page of the cropped image, where (ll_x, ll_y) are the user space coordinates of the lower-left corner, (ul_x, ul_y) those of the upper-left corner, (ur_x, ur_y) those of the upper-right corner, and (lr_x, lr_y) those of the lower-right corner. The specified coordinates must define a parallelogram; that is, they must satisfy the conditions $ul_x - ll_x = ur_x - lr_x$ and $ul_y - ll_y = ur_y - lr_y$ The combination of Position and CropRect determines the image's scaling, rotation, reflection, and skew.

Resolution	array	%ALDImageResolution	<p>(<i>Optional</i>) An array of two numbers of the form <code>[horizRes vertRes]</code> specifying the resolution of the image in samples per inch.</p>
ColorType	name	%ALDImageColorType	<p>(<i>Optional</i>) The type of color specified by the Color entry. Valid values are Process, Spot, and Separation. Default value: Spot.</p>
Color	array	%ALDImageColor	<p>(<i>Optional</i>) An array of four numbers and a string of the form <code>[C M Y K colorName]</code> specifying the value and name of the color in which the image is to be rendered. The values of <i>C</i>, <i>M</i>, <i>Y</i>, and <i>K</i> must all be in the range 0.0 to 1.0. Default value: [0.0 0.0 0.0 1.0 (Black)].</p>
Tint	number	%ALDImageTint	<p>(<i>Optional</i>) A number in the range 0.0 to 1.0 specifying the concentration of the color specified by Color in which the image is to be rendered. Default value: 1.0.</p>
Overprint	boolean	%ALDImageOverprint	<p>(<i>Optional</i>) A flag specifying whether the image is to overprint (true) or knock out (false) underlying marks on other separations. Default value: false.</p>
ImageType	array	%ALDImageType	<p>(<i>Optional</i>) An array of two integers of the form <code>[samples bits]</code> specifying the number of samples per pixel and bits per sample in the image.</p>
GrayMap	array	%ALDImageGrayMap	<p>(<i>Optional</i>) An array of 2^n integers in the range 0 to 65,535 (where n is the number of bits per sample) recording changes made to the brightness or contrast of the image.</p>
Transparency	boolean	%ALDImageTransparency	<p>(<i>Optional</i>) A flag specifying whether white pixels in the image are to be treated as transparent. Default value: true.</p>

Tags	array	%ALDImageAsciiTag<NNN>	<p>(Optional) An array of pairs of the form</p> <p style="text-align: center;">[tagNum₁ tagText₁ ... tagNum_n tagText_n]</p> <p>where each <i>tagNum</i> is an integer representing a TIFF tag number and each <i>tagText</i> is a string representing the corresponding ASCII tag value.</p>
-------------	-------	------------------------	---

TABLE 9.51 Entries in a version 2.0 OPI dictionary

KEY	TYPE	OPI COMMENT	VALUE
Type	name		(Optional) The type of PDF object that this dictionary describes; if present, must be OPI for an OPI dictionary.
Version	number		(Required) The version of OPI to which this dictionary refers; must be the number 2 or 2.0 (not the name 2.0, as in an OPI version dictionary).
F	file specification	%%ImageFilename	(Required) The external file containing the low-resolution proxy image. (See implementation note 118 in Appendix H.)
MainImage	string	%%MainImage	(Optional) The pathname of the file containing the full-resolution image corresponding to this proxy, or any other identifying string that uniquely identifies the full-resolution image.
Tags	array	%%TIFFASCIITag	<p>(Optional) An array of pairs of the form</p> <p style="text-align: center;">[tagNum₁ tagText₁ ... tagNum_n tagText_n]</p> <p>where each <i>tagNum</i> is an integer representing a TIFF tag number and each <i>tagText</i> is a string or an array of strings representing the corresponding ASCII tag value.</p>
Size	array	%%ImageDimensions	<p>(Optional; see note below) An array of two numbers of the form</p> <p style="text-align: center;">[width height]</p> <p>specifying the dimensions of the image in pixels.</p>

CropRect	rectangle	%%ImageCropRect	<p>(Optional; see note below) An array of four numbers of the form</p> <p style="padding-left: 40px;">[left top right bottom]</p> <p>specifying the portion of the image to be used.</p> <p><i>Note: The Size and CropRect entries should either both be present or both absent. If present, they must satisfy the conditions</i></p> <p style="padding-left: 40px;">$0 \leq \text{left} < \text{right} \leq \text{width}$</p> <p style="padding-left: 40px;">and</p> <p style="padding-left: 40px;">$0 \leq \text{top} < \text{bottom} \leq \text{height}$</p> <p>(Note that in this coordinate space, the positive y axis extends vertically downward; hence the requirement that $\text{top} < \text{bottom}$.)</p>
Overprint	boolean	%%ImageOverprint	<p>(Optional) A flag specifying whether the image is to overprint (true) or knock out (false) underlying marks on other separations. Default value: false.</p>
Inks	name or array	%%ImageInks	<p>(Optional) A name object or array specifying the colorants to be applied to the image. The value may be the name <code>full_color</code> or <code>registration</code> or an array of the form</p> <p style="padding-left: 40px;">[/monochrome <i>name</i>₁ <i>tint</i>₁ ... <i>name</i>_{<i>n</i>} <i>tint</i>_{<i>n</i>}]</p> <p>where each <i>name</i> is a string representing the name of a colorant and each <i>tint</i> is a real number in the range 0.0 to 1.0 specifying the concentration of that colorant to be applied.</p>
IncludedImageDimensions	array	%%IncludedImageDimensions	<p>(Optional) An array of two integers of the form</p> <p style="padding-left: 40px;">[<i>pixelsWide pixelsHigh</i>]</p> <p>specifying the dimensions of the included image in pixels.</p>
IncludedImageQuality	number	%%IncludedImageQuality	<p>(Optional) A number indicating the quality of the included image. Valid values are 1, 2, and 3.</p>

APPENDIX A

Operator Summary

THIS APPENDIX LISTS all the operators used in PDF content streams, in alphabetical order. Corresponding PostScript language operators are given in Table A.1 only when they are exact or near-exact equivalents of the PDF operators. Table and page references are to the place in the text where each operator is introduced.

TABLE A.1 PDF content stream operators

OPERATOR	POSTSCRIPT EQUIVALENT	DESCRIPTION	TABLE	PAGE
b	closepath, fill, stroke	Close, fill, and stroke path using nonzero winding number rule	4.10	167
B	fill, stroke	Fill and stroke path using nonzero winding number rule	4.10	167
b*	closepath, eofill, stroke	Close, fill, and stroke path using even-odd rule	4.10	167
B*	eofill, stroke	Fill and stroke path using even-odd rule	4.10	167
BDC		(PDF 1.2) Begin marked-content sequence with property list	9.8	584
BI		Begin inline image object	4.38	278
BMC		(PDF 1.2) Begin marked-content sequence	9.8	584
BT		Begin text object	5.4	308
BX		(PDF 1.1) Begin compatibility section	3.20	95
c	curveto	Append curved segment to path (three control points)	4.9	163
cm	concat	Concatenate matrix to current transformation matrix	4.7	156
CS	setcolorspace	(PDF 1.1) Set color space for stroking operations	4.21	216
cs	setcolorspace	(PDF 1.1) Set color space for nonstroking operations	4.21	217

d	setdash	Set line dash pattern	4.7	156
d0	setcharwidth	Set glyph width in Type 3 font	5.10	326
d1	setcachedevice	Set glyph width and bounding box in Type 3 font	5.10	326
Do		Invoke named XObject	4.34	261
DP		(PDF 1.2) Define marked-content point with property list	9.8	584
EI		End inline image object	4.38	278
EMC		(PDF 1.2) End marked-content sequence	9.8	584
ET		End text object	5.4	308
EX		(PDF 1.1) End compatibility section	3.20	95
f	fill	Fill path using nonzero winding number rule	4.10	167
F	fill	Fill path using nonzero winding number rule (obsolete)	4.10	167
f*	eofill	Fill path using even-odd rule	4.10	167
G	setgray	Set gray level for stroking operations	4.21	217
g	setgray	Set gray level for nonstroking operations	4.21	217
gs		(PDF 1.2) Set parameters from graphics state parameter dictionary	4.7	156
h	closepath	Close subpath	4.9	163
i	setflat	Set flatness tolerance	4.7	156
ID		Begin inline image data	4.38	278
j	setlinejoin	Set line join style	4.7	156
J	setlinecap	Set line cap style	4.7	156
K	setcmykcolor	Set <i>CMYK</i> color for stroking operations	4.21	218
k	setcmykcolor	Set <i>CMYK</i> color for nonstroking operations	4.21	218
l	lineto	Append straight line segment to path	4.9	163
m	moveto	Begin new subpath	4.9	163
M	setmiterlimit	Set miter limit	4.7	156
MP		(PDF 1.2) Define marked-content point	9.8	584
n		End path without filling or stroking	4.10	167

q	gsave	Save graphics state	4.7	156
Q	grestore	Restore graphics state	4.7	156
re		Append rectangle to path	4.9	164
RG	setrgbcolor	Set <i>RGB</i> color for stroking operations	4.21	217
rg	setrgbcolor	Set <i>RGB</i> color for nonstroking operations	4.21	217
ri		Set color rendering intent	4.7	156
s	closepath, stroke	Close and stroke path	4.10	167
S	stroke	Stroke path	4.10	167
SC	setcolor	(<i>PDF 1.1</i>) Set color for stroking operations	4.21	217
sc	setcolor	(<i>PDF 1.1</i>) Set color for nonstroking operations	4.21	217
SCN	setcolor	(<i>PDF 1.2</i>) Set color for stroking operations (ICCBased and special color spaces)	4.21	217
scn	setcolor	(<i>PDF 1.2</i>) Set color for nonstroking operations (ICCBased and special color spaces)	4.21	217
sh	shfill	(<i>PDF 1.3</i>) Paint area defined by shading pattern	4.24	232
T*		Move to start of next text line	5.5	310
Tc		Set character spacing	5.2	302
Td		Move text position	5.5	310
TD		Move text position and set leading	5.5	310
Tf	selectfont	Set text font and size	5.2	302
Tj	show	Show text	5.6	311
TJ		Show text, allowing individual glyph positioning	5.6	311
TL		Set text leading	5.2	302
Tm		Set text matrix and text line matrix	5.5	310
Tr		Set text rendering mode	5.2	302
Ts		Set text rise	5.2	302
Tw		Set word spacing	5.2	302
Tz		Set horizontal text scaling	5.2	302
v	curveto	Append curved segment to path (initial point replicated)	4.9	163

w	setlinewidth	Set line width	4.7	156
W	clip	Set clipping path using nonzero winding number rule	4.11	172
W*	eoclip	Set clipping path using even-odd rule	4.11	172
y	curveto	Append curved segment to path (final point replicated)	4.9	163
'		Move to next line and show text	5.6	311
"		Set word and character spacing, move to next line, and show text	5.6	311

APPENDIX B

Operators in Type 4 Functions

THIS APPENDIX SUMMARIZES the PostScript operators that can appear in a type 4 function, as discussed in Section 3.9.4, “Type 4 (PostScript Calculator) Functions.” For details on these operators, see the *PostScript Language Reference*, Third Edition.

B.1 Arithmetic Operators

num_1	num_2	add	<i>sum</i>	Return num_1 plus num_2
num_1	num_2	sub	<i>difference</i>	Return num_1 minus num_2
num_1	num_2	mul	<i>product</i>	Return num_1 times num_2
num_1	num_2	div	<i>quotient</i>	Return num_1 divided by num_2
int_1	int_2	idiv	<i>quotient</i>	Return int_1 divided by int_2 as an integer
int_1	int_2	mod	<i>remainder</i>	Return remainder after dividing int_1 by int_2
num_1		neg	num_2	Return negative of num_1
num_1		abs	num_2	Return absolute value of num_1
num_1		ceiling	num_2	Return ceiling of num_1
num_1		floor	num_2	Return floor of num_1
num_1		round	num_2	Round num_1 to nearest integer
num_1		truncate	num_2	Remove fractional part of num_1
num		sqrt	<i>real</i>	Return square root of num
$angle$		sin	<i>real</i>	Return sine of $angle$ degrees
$angle$		cos	<i>real</i>	Return cosine of $angle$ degrees
num	den	atan	<i>angle</i>	Return arc tangent of num/den in degrees
$base$	$exponent$	exp	<i>real</i>	Raise $base$ to $exponent$ power
num		ln	<i>real</i>	Return natural logarithm (base e)
num		log	<i>real</i>	Return common logarithm (base 10)
num		cvi	<i>int</i>	Convert to integer
num		cvr	<i>real</i>	Convert to real

B.2 Relational, Boolean, and Bitwise Operators

$any_1 any_2$ eq <i>bool</i>	Test equal
$any_1 any_2$ ne <i>bool</i>	Test not equal
$num_1 num_2$ gt <i>bool</i>	Test greater than
$num_1 num_2$ ge <i>bool</i>	Test greater than or equal
$num_1 num_2$ lt <i>bool</i>	Test less than
$num_1 num_2$ le <i>bool</i>	Test less than or equal
$bool_1 int_1$ $bool_2 int_2$ and $bool_3 int_3$	Perform logical bitwise and
$bool_1 int_1$ $bool_2 int_2$ or $bool_3 int_3$	Perform logical bitwise inclusive or
$bool_1 int_1$ $bool_2 int_2$ xor $bool_3 int_3$	Perform logical bitwise exclusive or
$bool_1 int_1$ not $bool_2 int_2$	Perform logical bitwise not
int_1 <i>shift</i> bitshift int_2	Perform bitwise shift of int_1 (positive is left)
- true <i>true</i>	Return boolean value <i>true</i>
- false <i>false</i>	Return boolean value <i>false</i>

B.3 Conditional Operators

<i>bool</i> { <i>expr</i> } if -	Execute <i>expr</i> if <i>bool</i> is <i>true</i>
<i>bool</i> { <i>expr</i> ₁ } { <i>expr</i> ₂ } ifelse -	Execute <i>expr</i> ₁ if <i>bool</i> is <i>true</i> , <i>expr</i> ₂ if <i>false</i>

B.4 Stack Operators

<i>any</i> pop -	Discard top element
$any_1 any_2$ exch $any_2 any_1$	Exchange top two elements
<i>any</i> dup <i>any any</i>	Duplicate top element
$any_1 \dots any_n n$ copy $any_1 \dots any_n any_1 \dots any_n$	Duplicate top <i>n</i> elements
$any_n \dots any_0 n$ index $any_n \dots any_0 any_n$	Duplicate arbitrary element
$any_{n-1} \dots any_0 n j$ roll $any_{(j-1) \bmod n} \dots any_0 any_{n-1} \dots any_{j \bmod n}$	Roll <i>n</i> elements up <i>j</i> times

APPENDIX C

Implementation Limits

IN GENERAL, PDF does not restrict the size or quantity of things described in the file format, such as numbers, arrays, images, and so on. However, a PDF viewer application running on a particular processor and in a particular operating environment does have such limits. If a viewer application attempts to perform an action that exceeds one of the limits, it will display an error.

PostScript interpreters also have implementation limits, listed in Appendix B of the *PostScript Language Reference*, Third Edition. It is possible to construct a PDF file that does not violate viewer application limits but will not print on a PostScript printer. Keep in mind that these limits vary according to the PostScript language level, interpreter version, and the amount of memory available to the interpreter.

This appendix describes typical limits for Acrobat. These limits fall into two main classes:

- *Architectural limits.* The hardware on which a viewer application executes imposes certain constraints. For example, an integer is usually represented in 32 bits, limiting the range of allowed integers. In addition, the design of the software imposes other constraints, such as a limit to the number of elements in an array or string.
- *Memory limits.* The amount of memory available to a viewer application limits the number of memory-consuming objects that can be held simultaneously.

PDF itself has one architectural limit: Because ten digits are allocated to byte offsets, the size of a file is limited to 10^{10} bytes (approximately 10 gigabytes).

C.1 General Implementation Limits

Table C.1 describes the architectural limits for Acrobat viewer applications running on 32-bit machines. Because Acrobat implementations are subject to these limits, applications producing PDF files are strongly advised to remain within them. Note, however, that memory limits will often be exceeded before architectural limits (such as the limit on the number of indirect objects) are reached.

TABLE C.1 Architectural limits

QUANTITY	LIMIT	DESCRIPTION
integer	2,147,483,647	Largest integer value; equal to $2^{31} - 1$.
	-2,147,483,648	Smallest integer value; equal to -2^{31} .
real	$\pm 32,767$	Largest and smallest real values (approximate).
	$\pm 1/65,536$	Nonzero real values closest to 0 (approximate); equal to $\pm 10^{-38}$. Values closer than these are automatically converted to 0.
	5	Number of significant decimal digits of precision in fractional part (approximate).
string	65,535	Maximum length of a string, in bytes.
name	127	Maximum length of a name, in bytes.
array	8191	Maximum capacity of an array, in elements.
dictionary	4095	Maximum capacity of a dictionary, in entries.
indirect object	8,388,607	Maximum number of indirect objects in a PDF file.
q/Q nesting	28	Maximum depth of graphics state nesting by q and Q operators. (This is not a limit of Acrobat as such, but arises from the fact that q and Q are implemented by the PostScript gsave and grestore operators when generating PostScript output; see implementation note 119 in Appendix H.)
DeviceN components	8	Maximum number of colorants or tint components in a DeviceN color space.
CID	65,535	Maximum value of a CID (character identifier).

Acrobat has some additional architectural limits:

- Thumbnail images may be no larger than 106 by 106 samples, and should be created at one-eighth scale for 8.5-by-11-inch and A4-size pages.
- The minimum allowed page size in Acrobat 4.0 is 3 by 3 units in default user space (approximately 0.04 by 0.04 inch); the maximum is 14,400 by 14,400 units (200 by 200 inches). (See implementation note 120 in Appendix H.)
- The magnification factor of a view is constrained to be between approximately 8 percent and 3200 percent. These limits are not fixed; they vary with the size of the page being displayed, as well as with the size of the pages previously viewed within the file.
- When Acrobat reads a PDF file with a damaged or missing cross-reference table, it attempts to rebuild the table by scanning all the objects in the file. However, the generation numbers of deleted entries are lost if the cross-reference table is missing or severely damaged. Reconstruction fails if any object identifiers do not appear at the start of a line or if the **endobj** keyword does not appear at the start of a line. Also, reconstruction fails if a stream contains a line beginning with the word **endstream**, aside from the required **endstream** that delimits the end of the stream.

Memory limits cannot be characterized as precisely as architectural limits can, because the amount of available memory and the ways in which it is allocated vary from one product to another. Memory is automatically reallocated from one use to another when necessary: when more memory is needed for a particular purpose, it can be taken away from memory allocated to another purpose if that memory is currently unused or its use is nonessential (a cache, for example). Also, data is often saved to a temporary file when memory is limited. Because of this behavior, it is not possible to state limits for such items as the number of pages in a document, number of text annotations or hypertext links on a page, number of graphics objects on a page, or number of fonts on a page or in a document.

C.2 Implementation Limits Affecting Web Capture

The data structures constructed by the Web Capture plug-in extension (*PDF 1.3*; see Section 9.9) depend on the maximum length of an array, k , which is 8191 elements in the Acrobat 4 implementation.

- A content set array can associate at most k content sets with a given name.
- A content set can reference at most k objects.
- There can be at most k source information dictionaries associated with a single content set.
- A URL alias dictionary can contain at most k chains, and each chain can contain at most k URLs.
- A maximum of k command dictionaries can be stored in the **C** array of the Web Capture information dictionary.
- There can be at most $k + 2$ entries in the **C** dictionary of a Web Capture command settings dictionary.

APPENDIX D

Character Sets and Encodings

THIS APPENDIX LISTS the character sets and encodings that are assumed to be predefined in any PDF viewer application. Only simple fonts, encompassing Latin text and some symbols, are described here. See “Predefined CMaps” on page 343 for a list of predefined CMaps for CID-keyed fonts.

Section D.1, “Latin Character Set and Encodings,” describes the entire character set for Adobe’s standard Latin-text fonts. This is the character set supported by the Times, Helvetica, and Courier font families, which are among the standard 14 predefined fonts (see “Standard Type 1 Fonts” on page 319). For each named character, an octal character code is given in four different encodings: **StandardEncoding**, **MacRomanEncoding**, **WinAnsiEncoding**, and **PDFDocEncoding** (see Table D.1). Unencoded characters are indicated by a dash (—).

Section D.2, “Expert Set and MacExpertEncoding,” describes the so-called “expert” character set, which contains additional characters useful for sophisticated typography, such as small capitals, ligatures, and fractions. For each named character, an octal character code is given in **MacExpertEncoding**. Note that the built-in encoding in an expert font program is usually different from **MacExpertEncoding**.

Sections D.3, “Symbol Set and Encoding,” and D.4, “ZapfDingbats Set and Encoding,” describe the character sets and built-in encodings for the **Symbol** and **ZapfDingbats** (ITC Zapf Dingbats) font programs, which are among the standard 14 predefined fonts. These fonts have built-in encodings that are unique to each font. (The characters for ZapfDingbats are ordered by code instead of by name, since the names in that font are meaningless.)

TABLE D.1 Latin-text encodings

ENCODING	DESCRIPTION
StandardEncoding	Adobe standard Latin-text encoding. This is the built-in encoding defined in Type 1 Latin-text font programs (but generally not in TrueType font programs). PDF does not have a predefined encoding named StandardEncoding . However, it is useful to describe this encoding, since a font's built-in encoding can be used as the base encoding from which differences are specified in an encoding dictionary.
MacRomanEncoding	Mac OS standard encoding for Latin text in Western writing systems. PDF has a predefined encoding named MacRomanEncoding that can be used with both Type 1 and TrueType fonts.
WinAnsiEncoding	Windows Code Page 1252, often called the “Windows ANSI” encoding. This is the standard Windows encoding for Latin text in Western writing systems. PDF has a predefined encoding named WinAnsiEncoding that can be used with both Type 1 and TrueType fonts.
PDFDocEncoding	Encoding for text strings in a PDF document <i>outside</i> the document's content streams. This is one of two encodings (the other being Unicode) that can be used to represent text strings; see Section 3.8.1, “Text Strings.” PDF does not have a predefined encoding named PDFDocEncoding ; it is not customary to use this encoding to show text from fonts.
MacExpertEncoding	An encoding for use with expert fonts—ones containing the expert character set. PDF has a predefined encoding named MacExpertEncoding . Despite its name, it is not a platform-specific encoding; however, only certain fonts have the appropriate character set for use with this encoding. No such fonts are among the standard 14 predefined fonts.

D.1 Latin Character Set and Encodings

CHAR	NAME	CHAR CODE (OCTAL)				CHAR	NAME	CHAR CODE (OCTAL)			
		STD	MAC	WIN	PDF			STD	MAC	WIN	PDF
A	A	101	101	101	101	Œ	OE	352	316	214	226
Æ	AE	341	256	306	306	Ó	Oacute	—	356	323	323
Á	Aacute	—	347	301	301	Ô	Ocircumflex	—	357	324	324
Â	Acircumflex	—	345	302	302	Ö	Odieresis	—	205	326	326
Ä	Adieresis	—	200	304	304	Ò	Ograve	—	361	322	322
À	Agrave	—	313	300	300	Ø	Oslash	351	257	330	330
Å	Aring	—	201	305	305	Õ	Otilde	—	315	325	325
Ã	Atilde	—	314	303	303	P	P	120	120	120	120
B	B	102	102	102	102	Q	Q	121	121	121	121
C	C	103	103	103	103	R	R	122	122	122	122
Ç	Ccedilla	—	202	307	307	S	S	123	123	123	123
D	D	104	104	104	104	Š	Scaron	—	—	212	227
E	E	105	105	105	105	T	T	124	124	124	124
É	Eacute	—	203	311	311	Þ	Thorn	—	—	336	336
Ê	Ecircumflex	—	346	312	312	U	U	125	125	125	125
Ë	Edieresis	—	350	313	313	Ú	Uacute	—	362	332	332
È	Egrave	—	351	310	310	Û	Ucircumflex	—	363	333	333
Ð	Eth	—	—	320	320	Ü	Udieresis	—	206	334	334
€	Euro ¹	—	—	200	240	Û	Ugrave	—	364	331	331
F	F	106	106	106	106	V	V	126	126	126	126
G	G	107	107	107	107	W	W	127	127	127	127
H	H	110	110	110	110	X	X	130	130	130	130
I	I	111	111	111	111	Y	Y	131	131	131	131
Í	Iacute	—	352	315	315	Ý	Yacute	—	—	335	335
Î	Icircumflex	—	353	316	316	ÿ	Ydieresis	—	331	237	230
Ï	Idieresis	—	354	317	317	Z	Z	132	132	132	132
Ì	Igrave	—	355	314	314	Ž	Zcaron ²	—	—	216	231
J	J	112	112	112	112	a	a	141	141	141	141
K	K	113	113	113	113	á	aacute	—	207	341	341
L	L	114	114	114	114	â	acircumflex	—	211	342	342
Ł	Lslash	350	—	—	225	´	acute	302	253	264	264
M	M	115	115	115	115	ä	adieresis	—	212	344	344
N	N	116	116	116	116	æ	ae	361	276	346	346
Ñ	Ntilde	—	204	321	321	à	agrave	—	210	340	340
O	O	117	117	117	117	&	ampersand	046	046	046	046

CHAR	NAME	CHAR CODE (OCTAL)				CHAR	NAME	CHAR CODE (OCTAL)			
		STD	MAC	WIN	PDF			STD	MAC	WIN	PDF
å	aring	—	214	345	345	è	ecircumflex	—	220	352	352
^	asciicircum	136	136	136	136	ë	edieresis	—	221	353	353
~	asciitilde	176	176	176	176	è	egrave	—	217	350	350
*	asterisk	052	052	052	052	8	eight	070	070	070	070
@	at	100	100	100	100	...	ellipsis	274	311	205	203
ã	atilde	—	213	343	343	—	emdash	320	321	227	204
b	b	142	142	142	142	–	endash	261	320	226	205
\	backslash	134	134	134	134	=	equal	075	075	075	075
	bar	174	174	174	174	ð	eth	—	—	360	360
{	braceleft	173	173	173	173	!	exclam	041	041	041	041
}	braceright	175	175	175	175	¡	exclamdown	241	301	241	241
[bracketleft	133	133	133	133	f	f	146	146	146	146
]	bracketright	135	135	135	135	fi	fi	256	336	—	223
˘	breve	306	371	—	030	5	five	065	065	065	065
‡	brokenbar	—	—	246	246	fl	fl	257	337	—	224
•	bullet ³	267	245	225	200	f	florin	246	304	203	206
c	c	143	143	143	143	4	four	064	064	064	064
ˇ	caron	317	377	—	031	/	fraction	244	332	—	207
ç	ccedilla	—	215	347	347	g	g	147	147	147	147
¸	cedilla	313	374	270	270	ß	germandbls	373	247	337	337
¢	cent	242	242	242	242	`	grave	301	140	140	140
^	circumflex	303	366	210	032	>	greater	076	076	076	076
:	colon	072	072	072	072	«	guillemotleft ⁴	253	307	253	253
,	comma	054	054	054	054	»	guillemotright ⁴	273	310	273	273
©	copyright	—	251	251	251	<	guilsinglleft	254	334	213	210
¤	currency ¹	250	333	244	244	>	guilsinglright	255	335	233	211
d	d	144	144	144	144	h	h	150	150	150	150
†	dagger	262	240	206	201	˘	hungarumlaut	315	375	—	034
‡	daggerdbl	263	340	207	202	-	hyphen ⁵	055	055	055	055
°	degree	—	241	260	260	i	i	151	151	151	151
¨	dieresis	310	254	250	250	í	iacute	—	222	355	355
÷	divide	—	326	367	367	î	icircumflex	—	224	356	356
\$	dollar	044	044	044	044	ï	idieresis	—	225	357	357
·	dotaccent	307	372	—	033	ì	igrave	—	223	354	354
ı	dotlessi	365	365	—	232	j	j	152	152	152	152
e	e	145	145	145	145	k	k	153	153	153	153
é	eacute	—	216	351	351	l	l	154	154	154	154

CHAR	NAME	CHAR CODE (OCTAL)				CHAR	NAME	CHAR CODE (OCTAL)			
		STD	MAC	WIN	PDF			STD	MAC	WIN	PDF
<	less	074	074	074	074	q	q	161	161	161	161
¬	logicalnot	—	302	254	254	?	question	077	077	077	077
ł	lslash	370	—	—	233	¿	questiondown	277	300	277	277
m	m	155	155	155	155	"	quotedbl	042	042	042	042
˘	macron	305	370	257	257	„	quotedblbase	271	343	204	214
−	minus	—	—	—	212	“	quotedblleft	252	322	223	215
μ	mu	—	265	265	265	”	quotedblright	272	323	224	216
×	multiply	—	—	327	327	‘	quoteleft	140	324	221	217
n	n	156	156	156	156	’	quoteright	047	325	222	220
9	nine	071	071	071	071	,	quotesinglbase	270	342	202	221
ñ	ntilde	—	226	361	361	'	quotesingle	251	047	047	047
#	numbersign	043	043	043	043	r	r	162	162	162	162
o	o	157	157	157	157	®	registered	—	250	256	256
ó	oacute	—	227	363	363	°	ring	312	373	—	036
ô	ocircumflex	—	231	364	364	s	s	163	163	163	163
ö	odieresis	—	232	366	366	š	scaron	—	—	232	235
œ	oe	372	317	234	234	§	section	247	244	247	247
ˆ	ogonek	316	376	—	035	;	semicolon	073	073	073	073
ò	ograve	—	230	362	362	7	seven	067	067	067	067
1	one	061	061	061	061	6	six	066	066	066	066
½	onehalf	—	—	275	275	/	slash	057	057	057	057
¼	onequarter	—	—	274	274		space ⁶	040	040	040	040
¹	onesuperior	—	—	271	271	£	sterling	243	243	243	243
^a	ordfeminine	343	273	252	252	t	t	164	164	164	164
^o	ordmasculine	353	274	272	272	þ	thorn	—	—	376	376
ø	oslash	371	277	370	370	3	three	063	063	063	063
õ	otilde	—	233	365	365	¾	threequarters	—	—	276	276
p	p	160	160	160	160	³	threesuperior	—	—	263	263
¶	paragraph	266	246	266	266	˜	tilde	304	367	230	037
(parenleft	050	050	050	050	™	trademark	—	252	231	222
)	parenright	051	051	051	051	2	two	062	062	062	062
%	percent	045	045	045	045	²	twosuperior	—	—	262	262
.	period	056	056	056	056	u	u	165	165	165	165
·	periodcentered	264	341	267	267	ú	uacute	—	234	372	372
‰	perthousand	275	344	211	213	û	ucircumflex	—	236	373	373
+	plus	053	053	053	053	ü	udieresis	—	237	374	374
±	plusminus	—	261	261	261	ù	ugrave	—	235	371	371

CHAR	NAME	CHAR CODE (OCTAL)				CHAR	NAME	CHAR CODE (OCTAL)			
		STD	MAC	WIN	PDF			STD	MAC	WIN	PDF
_	underscore	137	137	137	137	ÿ	ydieresis	—	330	377	377
v	v	166	166	166	166	¥	yen	245	264	245	245
w	w	167	167	167	167	z	z	172	172	172	172
x	x	170	170	170	170	ž	zcaron ²	—	—	236	236
y	y	171	171	171	171	0	zero	060	060	060	060
ý	yacute	—	—	375	375						

1. In PDF 1.3, the euro character was added to the Adobe standard Latin character set. It is encoded as 200 in **WinAnsiEncoding** and 240 in **PDFDocEncoding**, assigning codes that were previously unused. Apple changed the Mac OS Latin-text encoding for code 333 from the currency character to the euro character. However, this incompatible change has *not* been reflected in PDF's **MacRomanEncoding**, which continues to map code 333 to currency. If the euro character is desired, an encoding dictionary can be used to specify this single difference from **MacRomanEncoding**.
2. In PDF 1.3, the existing Zcaron and zcaron characters were added to **WinAnsiEncoding** as the previously unused codes 216 and 236.
3. In **WinAnsiEncoding**, all unused codes greater than 40 map to the bullet character. However, only code 225 is specifically assigned to the bullet character; other codes are subject to future reassignment.
4. The character names guillemotleft and guillemotright are misspelled. The correct spelling for this punctuation character is *guillemet*. However, the misspelled names are the ones actually used in the fonts and encodings containing these characters.
5. The hyphen character is also encoded as 255 in **WinAnsiEncoding**. The meaning of this duplicate code is “soft hyphen,” but it is typographically the same as hyphen.
6. The space character is also encoded as 312 in **MacRomanEncoding** and as 240 in **WinAnsiEncoding**. The meaning of this duplicate code is “nonbreaking space,” but it is typographically the same as space.

D.2 Expert Set and MacExpertEncoding

CHAR	NAME	CODE	CHAR	NAME	CODE
Æ	AEsmall	276	J	Jsmall	152
Á	Aacutesmall	207	K	Ksmall	153
Â	Acircumflexsmall	211	Ł	Lslashsmall	302
´	Acutesmall	047	L	Lsmall	154
Ä	Adieresissmall	212	˘	Macronsmall	364
À	Agravesmall	210	M	Msmall	155
Å	Aringsmall	214	N	Nsmall	156
A	Asmall	141	Ñ	Ntildesmall	226
Ã	Atildesmall	213	Œ	OEmall	317
ˆ	Brevesmall	363	Ó	Oacutesmall	227
B	Bsmall	142	Ô	Ocircumflexsmall	231
˘	Caronsmall	256	Ö	Odieresissmall	232
Ç	Ccedillasmall	215	˙	Ogoneksmall	362
¸	Cedillasmall	311	Ò	Ogravesmall	230
ˆ	Circumflexsmall	136	Ø	Oslashsmall	277
C	Csmall	143	O	Osmall	157
¨	Dieresissmall	254	Ö	Otildesmall	233
˙	Dotaccentsmall	372	P	Psmall	160
D	Dsmall	144	Q	Qsmall	161
É	Eacutesmall	216	°	Ringsmall	373
Ê	Ecircumflexsmall	220	R	Rsmall	162
Ë	Edieresissmall	221	Š	Scaronsmall	247
È	Egravesmall	217	S	Ssmall	163
E	Esmall	145	Þ	Thornsmall	271
Ð	Ethsmall	104	˜	Tildesmall	176
F	Fsmall	146	T	Tsmall	164
`	Gravesmall	140	Ú	Uacutesmall	234
G	Gsmall	147	Û	Ucircumflexsmall	236
H	Hsmall	150	Ü	Udieresissmall	237
˘	Hungarumlautsmall	042	Ù	Ugravesmall	235
Í	Iacutesmall	222	U	Usmall	165
Î	Icircumflexsmall	224	V	Vsmall	166
Ï	Idieresissmall	225	W	Wsmall	167
Ì	Igravesmall	223	X	Xsmall	170
I	Ismall	151	Ý	Yacutesmall	264

CHAR	NAME	CODE	CHAR	NAME	CODE
ÿ	Ydieresis <small></small>	330	⁄	fouroidstyle	064
ŷ	Y <small></small>	171	⁴	foursuperior	335
ž	Zcaron <small></small>	275	/	fraction	057
z	Z <small></small>	172	-	hyphen	055
&	ampersand <small></small>	046	-	hypheninferior	137
ˆ	asuperior	201	-	hyphensuperior	321
ˆ	bsuperior	365	ı	isuperior	351
¢	centinferior	251	ı	lsuperior	361
¢	centoldstyle	043	ˆ	msuperior	367
¢	centsuperior	202	9	nineinferior	273
:	colon	072	9	nineoldstyle	071
₯	colonmonetary	173	9	ninesuperior	341
,	comma	054	ˆ	nsuperior	366
,	commainferior	262	.	onedotenleader	053
,	commasuperior	370	⅛	oneeighth	112
\$	dollarinferior	266	ı	onefitted	174
\$	dollaroldstyle	044	½	onehalf	110
\$	dollarsuperior	045	ı	oneinferior	301
ˆ	dsuperior	353	ı	oneoldstyle	061
8	eightinferior	245	¼	onequarter	107
8	eightoldstyle	070	ı	onesuperior	332
8	eightsuperior	241	⅓	onethird	116
ˆ	esuperior	344	ˆ	osuperior	257
ı	exclamdown <small></small>	326	(parenleftinferior	133
!	exclam <small></small>	041	(parenleftsuperior	050
ff	ff	126)	parenrightinferior	135
ffi	ffi	131)	parenrightsuperior	051
ffl	ffl	132	.	period	056
fi	fi	127	.	periodinferior	263
–	figuredash	320	·	periodsuperior	371
⅝	fiveeighths	114	¿	questiondown <small></small>	300
5	fiveinferior	260	?	questionsmall	077
5	fiveoldstyle	065	ˆ	rsuperior	345
5	fivesuperior	336	₨	rupiah	175
fl	fl	130	;	semicolon	073
4	fourinferior	242	⅞	seveneighths	115

CHAR	NAME	CODE	CHAR	NAME	CODE
7	seveninferior	246	—	threequartersemdash	075
7	sevenoldstyle	067	³	threesuperior	334
7	sevensuperior	340	ˆ	tsuperior	346
6	sixinferior	244	..	twodotenleader	052
6	sixoldstyle	066	₂	twoinferior	252
6	sixsuperior	337	₂	twooldstyle	062
	space	040	₂	twosuperior	333
ˢ	ssuperior	352	⅔	twothirds	117
⅜	threeeighths	113	₀	zeroinferior	274
₃	threeinferior	243	₀	zerooldstyle	060
₃	threeoldstyle	063	₀	zerosuperior	342
¾	threequarters	111			

D.3 Symbol Set and Encoding

CHAR	NAME	CODE	CHAR	NAME	CODE
A	Alpha	101	↔	arrowboth	253
B	Beta	102	↔	arrowdblboth	333
X	Chi	103	⇩	arrowdbldown	337
Δ	Delta	104	⇐	arrowdblleft	334
E	Epsilon	105	⇒	arrowdblright	336
H	Eta	110	⇧	arrowdblup	335
€	Euro	240	↓	arrowdown	257
Γ	Gamma	107	—	arrowhorizex	276
Œ	Ifraktur	301	←	arrowleft	254
I	Iota	111	→	arrowright	256
K	Kappa	113	↑	arrowup	255
Λ	Lambda	114		arrowvertex	275
M	Mu	115	*	asteriskmath	052
N	Nu	116		bar	174
Ω	Omega	127	β	beta	142
O	Omicron	117	{	braceleft	173
Φ	Phi	106	}	braceright	175
Π	Pi	120	{	bracelefttp	354
Ψ	Psi	131	}	braceleftmid	355
℞	Rfraktur	302	{	braceleftbt	356
P	Rho	122	}	bracerighttp	374
Σ	Sigma	123	}	bracerightmid	375
T	Tau	124	}	bracerightbt	376
Θ	Theta	121		braceex	357
Υ	Upsilon	125	[bracketleft	133
Υ	Upsilon1	241]	bracketright	135
Ξ	Xi	130	[bracketlefttp	351
Z	Zeta	132		bracketleftex	352
ℵ	aleph	300	[bracketleftbt	353
α	alpha	141]	bracketrighttp	371
&	ampersand	046		bracketrightex	372
∠	angle	320]	bracketrightbt	373
⟨	angleleft	341	•	bullet	267
⟩	angleright	361	↵	carriagereturn	277
≈	aproxequal	273	χ	chi	143

CHAR	NAME	CODE	CHAR	NAME	CODE
⊗	circlemultiply	304	∫	integralbt	365
⊕	circleplus	305	∩	intersection	307
♣	club	247	ι	iota	151
:	colon	072	κ	kappa	153
,	comma	054	λ	lambda	154
≡	congruent	100	<	less	074
©	copyrightsans	343	≤	lessequal	243
©	copyrightserif	323	∧	logicaland	331
°	degree	260	¬	logicalnot	330
δ	delta	144	∨	logicalor	332
◆	diamond	250	◇	lozenge	340
÷	divide	270	-	minus	055
·	dotmath	327	'	minute	242
8	eight	070	μ	mu	155
∈	element	316	×	multiply	264
...	ellipsis	274	9	nine	071
∅	emptyset	306	∉	notelement	317
ε	epsilon	145	≠	notequal	271
=	equal	075	⊄	notsubset	313
≡	equivalence	272	ν	nu	156
η	eta	150	#	numbersign	043
!	exclam	041	ω	omega	167
∃	existential	044	ω	omega1	166
5	five	065	ο	omicron	157
f	florin	246	1	one	061
4	four	064	(parenleft	050
/	fraction	244)	parenright	051
γ	gamma	147	/	parenlefttp	346
∇	gradient	321		parenleftex	347
>	greater	076	\	parenleftbt	350
≥	greaterequal	263	\	parenrighttp	366
♥	heart	251		parenrightex	367
∞	infinity	245)	parenrightbt	370
∫	integral	362	∂	partialdiff	266
∫	integraltp	363	%	percent	045
	integralex	364	.	period	056

CHAR	NAME	CODE	CHAR	NAME	CODE
\perp	perpendicular	136	\sim	similar	176
ϕ	phi	146	6	six	066
φ	phi1	152	/	slash	057
π	pi	160		space	040
+	plus	053	\spadesuit	spade	252
\pm	plusminus	261	\ni	suchthat	047
\prod	product	325	\sum	summation	345
\subset	probersubset	314	τ	tau	164
\supset	probersuperset	311	\therefore	therefore	134
\propto	proportional	265	θ	theta	161
ψ	psi	171	ϑ	theta1	112
?	question	077	3	three	063
$\sqrt{\quad}$	radical	326	TM	trademarksans	344
$_$	radicalex	140	TM	trademarkserif	324
\subseteq	reflexsubset	315	2	two	062
\supseteq	reflexsuperset	312	$_$	underscore	137
®	registersans	342	\cup	union	310
®	registerserif	322	\forall	universal	042
ρ	rho	162	υ	upsilon	165
"	second	262	\wp	weierstrass	303
;	semicolon	073	ξ	xi	170
7	seven	067	0	zero	060
σ	sigma	163	ζ	zeta	172
ς	sigma1	126			

D.4 ZapfDingbats Set and Encoding

CHAR	NAME	CODE	CHAR	NAME	CODE	CHAR	NAME	CODE	CHAR	NAME	CODE
	space	040	☞	a30	103	✱	a65	146	♠	a109	253
✂	a1	041	☞	a31	104	✱	a66	147	①	a120	254
✂	a2	042	☞	a32	105	✱	a67	150	②	a121	255
✂	a202	043	◆	a33	106	✱	a68	151	③	a122	256
✂	a3	044	◇	a34	107	✱	a69	152	④	a123	257
☞	a4	045	★	a35	110	✱	a70	153	⑤	a124	260
☞	a5	046	☆	a36	111	●	a71	154	⑥	a125	261
☞	a119	047	☼	a37	112	○	a72	155	⑦	a126	262
☞	a118	050	☆	a38	113	■	a73	156	⑧	a127	263
☞	a117	051	☆	a39	114	□	a74	157	⑨	a128	264
☞	a11	052	☆	a40	115	□	a203	160	⑩	a129	265
☞	a12	053	☆	a41	116	□	a75	161	❶	a130	266
☞	a13	054	☆	a42	117	□	a204	162	❷	a131	267
☞	a14	055	☆	a43	120	▲	a76	163	❸	a132	270
☞	a15	056	★	a44	121	▼	a77	164	❹	a133	271
☞	a16	057	☆	a45	122	◆	a78	165	❺	a134	272
☞	a105	060	✱	a46	123	❖	a79	166	❻	a135	273
☞	a17	061	✱	a47	124	◐	a81	167	❼	a136	274
☞	a18	062	☞	a48	125		a82	170	❽	a137	275
☞	a19	063	★	a49	126		a83	171	❾	a138	276
☞	a20	064	✱	a50	127	■	a84	172	❿	a139	277
✕	a21	065	✱	a51	130	‘	a97	173	①	a140	300
✕	a22	066	✱	a52	131	’	a98	174	②	a141	301
✕	a23	067	✱	a53	132	“	a99	175	③	a142	302
✕	a24	070	✱	a54	133	”	a100	176	④	a143	303
✚	a25	071	✱	a55	134	☞	a101	241	⑤	a144	304
✚	a26	072	✱	a56	135	☞	a102	242	⑥	a145	305
✚	a27	073	☞	a57	136	☞	a103	243	⑦	a146	306
✚	a28	074	☞	a58	137	♥	a104	244	⑧	a147	307
†	a6	075	☞	a59	140	☞	a106	245	⑨	a148	310
†	a7	076	☞	a60	141	☞	a107	246	⑩	a149	311
†	a8	077	✱	a61	142	☞	a108	247	❶	a150	312
✚	a9	100	✱	a62	143	♣	a112	250	❷	a151	313
☆	a10	101	✱	a63	144	♦	a111	251	❸	a152	314
☞	a29	102	✱	a64	145	♥	a110	252	❹	a153	315

CHAR	NAME	CODE	CHAR	NAME	CODE	CHAR	NAME	CODE	CHAR	NAME	CODE
⑤	a154	316	➤	a192	332	➡	a176	346	⇒	a184	363
⑥	a155	317	→	a166	333	➡	a177	347	➡	a197	364
⑦	a156	320	➡	a167	334	➡	a178	350	➡	a185	365
⑧	a157	321	→	a168	335	➡	a179	351	➡	a194	366
⑨	a158	322	➡	a169	336	➡	a193	352	➡	a198	367
⑩	a159	323	➡	a170	337	➡	a180	353	➡	a186	370
➡	a160	324	➡	a171	340	➡	a199	354	➡	a195	371
→	a161	325	➡	a172	341	➡	a181	355	→	a187	372
↔	a163	326	➤	a173	342	➡	a200	356	➡	a188	373
↕	a164	327	➤	a162	343	➡	a182	357	➡	a189	374
➤	a196	330	➤	a174	344	➡	a201	361	➡	a190	375
➡	a165	331	➡	a175	345	➡	a183	362	➤	a191	376

APPENDIX E

PDF Name Registry

THIS APPENDIX DISCUSSES a registry, maintained for developers by Adobe Systems, that contains private names and formats used by PDF producers or Acrobat plug-in extensions.

Acrobat enables third parties to add private data to PDF documents and to add plug-in extensions that change viewer behavior based on this data. However, Acrobat users have certain expectations when opening a PDF document, no matter what plug-ins are available. PDF enforces certain restrictions on private data in order to meet these expectations.

A PDF producer or Acrobat viewer plug-in extension may define new types of action, destination, annotation, security, and file system handlers. If a user opens a PDF document and the plug-in that implements the new type of object is unavailable, the viewer will behave as described in Appendix H.

A PDF producer or Acrobat plug-in extension may also add keys to any PDF object that is implemented as a dictionary, except the file trailer dictionary (see Section 3.4.4, “File Trailer”). In addition, a PDF producer or Acrobat plug-in may create tags that indicate the role of marked-content operators (*PDF 1.2*), as described in Section 9.5, “Marked Content.”

To avoid conflicts with third-party names and with future versions of PDF, Adobe maintains a registry for certain private names and formats. Developers must only add private data that conforms to the registry rules. The registry includes three classes:

- *First class*. Names and data formats that are of value to a wide range of developers. All names defined in any version of the PDF specification are first-class names. Plug-in extensions that are publicly available should often use

first-class names for their private data. First-class names and data formats must be registered with Adobe and will be made available for all developers to use. To submit a private name and format for consideration as first-class, contact Adobe at either of the following addresses:

Adobe Solutions Network
Adobe Systems Incorporated
345 Park Avenue
San Jose, CA 95110-2704

<acrodevsup@adobe.com>

- *Second class.* Names that are applicable to a specific developer. (Adobe does not register second-class data formats.) Adobe distributes second-class names by registering developer-specific prefixes, which must be used as the first characters in the names of all private data added by the developer. Adobe will not register the same prefix to two different developers, thereby ensuring that different developers' second-class names will not conflict. It is the responsibility of the developer to ensure that it does not itself use the same name in conflicting ways. Contact Adobe (at either of the addresses listed above) to request a prefix for second-class names.
- *Third class.* Names that can be used only in files that will never be seen by other third parties, because they may conflict with third-class names defined by others. Third-class names all begin with a specific prefix reserved by Adobe for private plug-in extensions. This prefix, which is *XX*, must be used as the first characters in the names of all private data added by the developer. It is not necessary to contact Adobe to register third-class names.

Note: *New keys for the document information dictionary (see Section 9.2.1, “Document Information Dictionary”) or a thread information dictionary (in the I entry of a thread dictionary; see Section 8.3.2, “Articles”) need not be registered.*

APPENDIX F

Linearized PDF

A LINEARIZED PDF FILE is one that has been organized in a special way to enable efficient incremental access in a network environment. The file is valid PDF in all respects, and is compatible with all existing viewers and other PDF applications. Enhanced viewer applications can recognize that a PDF file has been linearized and can take advantage of that organization (as well as added “hint” information) to enhance viewing performance.

The Linearized PDF file organization is an optional feature available beginning in PDF 1.2. Its primary goal is to achieve the following behavior:

- When a document is opened, display the first page as quickly as possible. The first page to be viewed can be an arbitrary page of the document, not necessarily page 0 (though opening at page 0 is most common).
- When the user requests another page of an open document (for example, by going to the next page or by following a link to an arbitrary page), display that page as quickly as possible.
- When data for a page is delivered over a slow channel, display the page incrementally as it arrives. To the extent possible, display the most useful data first.
- Permit user interaction, such as following a link, to be performed even before the entire page has been received and displayed.

This behavior should be achieved for documents of arbitrary size. The total number of pages in the document should have little or no effect on the user-perceived performance of viewing any particular page.

The primary focus of Linearized PDF is optimized viewing of read-only PDF documents. It is intended that the Linearized PDF will be generated once and read many times. Incremental update is still permitted, but the resulting PDF is

no longer linearized and subsequently will be treated as ordinary PDF. Linearizing it again may require reprocessing the entire file; see Section F.4.6, “Accessing an Updated File,” for details.

Linearized PDF requires two additions to the PDF specification:

- Rules for the ordering of objects in the PDF file
- Additional data structures, called *hint tables*, that enable efficient navigation within the document

Both of these additions are relatively simple to describe; however, using them effectively requires a deeper understanding of their purpose. Consequently, this appendix goes considerably beyond a simple specification of these PDF extensions, to include background, motivation, and strategies.

- Section F.1, “Background and Assumptions,” provides background information about the properties of the World Wide Web that are relevant to the design of Linearized PDF.
- Section F.2, “Linearized PDF Document Structure,” specifies the file format and object-ordering requirements of Linearized PDF.
- Section F.3, “Hint Tables,” specifies the detailed representation of the hint tables.
- Section F.4, “Access Strategies,” outlines strategies for accessing Linearized PDF over a network, which in turn determine the optimal way to organize the PDF file itself.

The reader is assumed to be familiar with the basic architecture of the Web, including terms such as URL, HTTP, and MIME.

F.1 Background and Assumptions

The principal problem addressed by the Linearized PDF design is the access of PDF documents through the World Wide Web. This environment has the following important properties:

- The access protocol (HTTP) is a transaction consisting of a request and a response. The client presents a request in the form of a URL, and the server sends a response consisting of one or more MIME-tagged data blocks.

- After a transaction has completed, obtaining more data requires a new request-response transaction. The connection between client and server does not ordinarily persist beyond the end of a transaction, although some implementations may attempt to cache the open connection in order to expedite subsequent transactions with the same server.
- Round-trip delay can be significant. A request-response transaction can take up to several seconds, independent of the amount of data requested.
- The data rate may be limited. A typical bottleneck is a slow modem link between the client and the Internet service provider.

These properties are generally shared by other wide-area network architectures besides the Web. Also, CD-ROMs share some of these properties, since they have relatively slow seek times and limited data rates compared to magnetic media. The remainder of this appendix focuses on the Web.

There are some additional properties of the HTTP protocol that are relevant to the problem of accessing PDF files efficiently. These properties may not all be shared by other protocols or network environments.

- When a PDF file is initially accessed (such as by following a URL hyperlink from some other document), the file type is not known to the client. Therefore, the client initiates a transaction to retrieve the entire document and then inspects the MIME tag of the response as it arrives. Only at that point is the document known to be PDF. Additionally, with a properly configured server environment, the length of the document becomes known at that time.
- The client can abort a response while the transaction is still in progress, if it decides that the remainder of the data is not of immediate interest. In HTTP, aborting the transaction requires closing the connection, which will interfere with the strategy of caching the open connection between transactions.
- The client can request retrieval of portions of a document by specifying one or more byte ranges (by offset and count) in the HTTP request headers. Each range can be relative to either the beginning or the end of the file. The client can specify as many ranges as it wants in the request, and the response will consist of multiple blocks, each properly tagged.
- The client can initiate multiple concurrent transactions in an attempt to obtain multiple responses in parallel. This is commonly done, for instance, to retrieve inline images referenced from an HTML document. This strategy is

not always reliable and may backfire if the transactions interfere with each other by competing for scarce resources in the server or the communication channel.

***Note:** Extensive experimentation has determined that having multiple concurrent transactions does not work very well for PDF in some important environments. Therefore, Linearized PDF is designed to enable good performance to be achieved using only one transaction at a time. In particular, this means that the client must have sufficient information to determine the byte ranges for all the objects required to display a given page of the PDF file, so that it can specify all those byte ranges in a single request.*

Finally, the following additional assumptions are made about the PDF viewer application and its local environment:

- The viewer application has plenty of local temporary storage available. It should rarely need to retrieve a given portion of a PDF document more than once from the server.
- The viewer application is able to display PDF data quickly once it has been received. The performance bottleneck is assumed to be in the transport system (throughput or round-trip delay), not in the processing of data after it arrives.

The consequence of these assumptions is that it may be advantageous for the client to do considerable extra work in order to minimize delays due to communications. Such work includes maintaining local caches and reordering actions according to when the needed data becomes available.

F.2 Linearized PDF Document Structure

Except as noted below, all elements of a Linearized PDF file are as specified in Section 3.4, “File Structure,” and all indirect objects in the file are numbered sequentially in two groups, based on their order of appearance in the file.

- The first group consists of the document catalog, certain other document-level objects, and all objects belonging to the first page of the document. These are numbered sequentially starting at the first object number after the last number of the second group. (The stream containing the hint tables, called a *hint*

stream, may be numbered out of sequence; see Section F.2.5, “Hint Streams (Parts 5 and 10).”)

- The second group consists of all remaining objects in the document, including all pages after the first, all shared objects (objects referenced from more than one page, not counting objects referenced from the first page), and so forth. These are numbered sequentially starting at 1.

These groups of objects are indexed by exactly two cross-reference table sections, located as shown in Example F.1. The composition of these groups is discussed in more detail in the sections that follow (ordered by the part number as shown in this example, with one section for parts 5 and 10). All objects have a generation number of 0.

Example F.1

Part 1: Header

```
%PDF-1.1
% ... Binary characters ...
```

Part 2: Linearization parameter dictionary

```
43 0 obj
  << /Linearized 1.0      % Version
     /L 54567            % File length
     /H [475 598]        % Primary hint stream offset and length (part 5)
     /O 45               % Object number of first page's page object (part 6)
     /E 5437             % Offset of end of first page
     /N 11               % Number of pages in document
     /T 52786            % Offset of first entry in main cross-reference table (part 11)
  >>
endobj
```

Part 3: First-page cross-reference table and trailer

```
xref
43 14
0000000052 00000 n
0000000392 00000 n
0000001073 00000 n
... Cross-reference entries for remaining objects in the first page ...
0000000475 00000 n
```

```

trailer
  << /Size 57          % Total number of cross-reference table entries in document
     /Prev 52776       % Offset of main cross-reference table (part 11)
     /Root 44 0 R      % Indirect reference to catalog (part 4)
     ... Any other entries, such as Info and Encrypt ... % (part 9)
  >>
startxref
0          % Dummy cross-reference table offset
%%EOF

```

Part 4: Document catalog and other required document-level objects

```

44 0 obj
  << /Type /Catalog
     /Pages 42 0 R
  >>
endobj
... Other objects ...

```

Part 5: Primary hint stream (may precede or follow part 6)

```

56 0 obj
  << /Length 457
     ... Possibly other stream attributes, such as Filter ...
     /S 221          % Position of shared object hint table
     ... Possibly entries for other hint tables ...
  >>
stream
  ... Page offset hint table ...
  ... Shared object hint table ...
  ... Possibly other hint tables ...
endstream
endobj

```

Part 6: First-page section (may precede or follow part 5)

```

45 0 obj
  << /Type /Page
     ...
  >>
endobj
... Outline hierarchy (if the PageMode value in the document catalog is UseOutlines) ...
... Objects for first page, including both shared and nonshared objects ...

```

Part 7: Remaining pages

```

1 0 obj
  << /Type /Page
    ... Other page attributes, such as MediaBox, Parent, and Contents ...
  >>
endobj

```

... Nonshared objects for this page ...

... Each successive page followed by its nonshared objects ...

... Last page followed by its nonshared objects ...

Part 8: Shared objects for all pages except the first

... Shared objects ...

Part 9: Objects not associated with pages, if any

... Other objects ...

Part 10: Overflow hint stream (optional)

... Overflow hint stream ...

Part 11: Main cross-reference table and trailer

```

xref
0 43
0000000000 65535 f
... Cross-reference entries for all except first page's objects ...
trailer
  << /Size 43 >>          % Trailer need not contain other entries; in particular,
startxref                 % it should not have a Prev entry
257                       % Offset of first-page cross-reference table (part 3)
%%EOF

```

F.2.1 Header (Part 1)

The Linearized PDF file begins with the standard header line (see Section 3.4.1, “File Header”). Linearization is independent of PDF version number and can be applied to any PDF file of version 1.1 or greater.

The “binary characters” following the percent sign on the second line are characters with codes 128 or greater, as recommended in Section 3.4.1, “File Header.”

F.2.2 Linearization Parameter Dictionary (Part 2)

Following the header, the first object in the body of the file (part 2) must be an indirect dictionary object, the *linearization parameter dictionary*, containing the parameters listed in Table F.1. All values in this dictionary must be direct objects. Note that there are no references to this dictionary anywhere in the document. (However, there is a normal entry for it in the first-page cross-reference table, part 3.)

The linearization parameter dictionary must be entirely contained within the first 1024 bytes of the PDF file. This limits the amount of data a viewer application must read before deciding whether the file is linearized.

F.2.3 First-Page Cross-Reference Table and Trailer (Part 3)

Part 3 contains the cross-reference table for all the first page's objects (discussed in Section F.2.6, "First-Page Section (Part 6)") as well as for the document catalog and document-level objects appearing before the first page (discussed in Section F.2.4, "Document Catalog and Document-Level Objects (Part 4)"). Additionally, it contains entries for the linearization parameter dictionary (at the beginning) and the primary hint stream (at the end). This table is a valid cross-reference table as defined in Section 3.4.3, "Cross-Reference Table," although its position in the file is unconventional. It consists of a single cross-reference subsection, with no free entries.

Below the table is the first-page trailer. The **startxref** line at the end of the trailer gives the offset of the first-page cross-reference table. The trailer's **Prev** entry gives the offset of the main cross-reference table near the end of the file. Again, this is valid PDF syntax, although the trailers are linked in an unusual order. A PDF viewer application that is unaware of linearization interprets the first-page cross-reference table as an update to an original document that is indexed by the main cross-reference table.

The first-page trailer must contain valid **Size** and **Root** entries, as well as any other entries needed to display the document. The **Size** value must be the combined number of entries in both the first-page cross-reference table and the main cross-reference table.

The first-page trailer may optionally end with **startxref**, an integer, and %%EOF, just as in an ordinary trailer. This information is ignored.

TABLE F.1 Entries in the linearization parameter dictionary

PARAMETER	TYPE	VALUE
Linearized	number	<i>(Required)</i> A version identification for the linearized format. As usual, a change in the integer part indicates an incompatible change in the linearized format, while a change in the fractional part indicates a backward-compatible change. The current version is 1.0.
L	integer	<i>(Required)</i> The length of the entire file in bytes. This must be exactly equal to the actual length of the PDF file. A mismatch indicates that the file is not linearized and must be treated as ordinary PDF, ignoring linearization information. (If the mismatch resulted from appending an update, the linearization information may still be correct but requires validation; see Section F.4.6, “Accessing an Updated File,” for details.)
H	array	<i>(Required)</i> An array of either two or four integers, $[offset_1\ length_1]$ or $[offset_1\ length_1\ offset_2\ length_2]$. $offset_1$ is the offset of the primary hint stream from the beginning of the file. (This is the beginning of the stream object, not the beginning of the stream data.) $length_1$ is the length of this stream, including stream object overhead. If the length of the primary hint stream is given by an indirect reference, the stream length object immediately follows the stream object, and $length_1$ also includes the length of the stream length object, including object overhead. (See implementation note 121 in Appendix H.) If there is an overflow hint stream, $offset_2$ and $length_2$ specify its offset and length. (See implementation note 122 in Appendix H.)
O	integer	<i>(Required)</i> The object number of the first page’s page object.
E	integer	<i>(Required)</i> The offset of the end of the first page (the end of part 6 in Example F.1), relative to the beginning of the file. (See implementation note 123 in Appendix H.)
N	integer	<i>(Required)</i> The number of pages in the document.
T	integer	<i>(Required)</i> The offset of the white-space character preceding the first entry of the main cross-reference table (the entry for object number 0), relative to the beginning of the file. Note that this differs from the Prev entry in the first-page trailer, which gives the location of the xref line that precedes the table.
P	integer	<i>(Optional)</i> The page number of the first page (see Section F.2.6, “First-Page Section (Part 6)”). Default value: 0.

F.2.4 Document Catalog and Document-Level Objects (Part 4)

Following the first-page cross-reference table and trailer are the catalog dictionary and other objects that are required when the document is opened. These additional objects (constituting part 4) include the values of the following entries, if they are present and are indirect objects:

- The **ViewerPreferences** entry in the catalog.
- The **PageMode** entry in the catalog. (Note that if the value of **PageMode** is `UseOutlines`, the outline hierarchy is located in part 6; otherwise, the outline hierarchy, if any, is located in part 9. See Section F.2.9, “Other Objects (Part 9)” for details.)
- The **Threads** entry in the catalog, along with all thread dictionaries it refers to. This does not include the threads’ information dictionaries or the individual bead dictionaries belonging to the threads.
- The **OpenAction** entry in the catalog.
- The **AcroForm** entry in the catalog. Only the top-level interactive form dictionary is needed, not the objects that it refers to.
- The **Encrypt** entry in the first-page trailer dictionary. All values in the encryption dictionary must be located here also.

Objects that are not ordinarily needed when the document is opened should not be located here but instead should be at the end of the file; see Section F.2.9, “Other Objects (Part 9).” This includes objects such as page tree nodes, the document information dictionary, and the definitions for named destinations.

Note that the objects located here are indexed by the first-page cross-reference table, even though they are not logically part of the first page.

F.2.5 Hint Streams (Parts 5 and 10)

The core of the linearization information is stored in data structures known as *hint tables*, whose format is described in Section F.3, “Hint Tables.” They provide indexing information that enables the client to construct a single request for all the objects that are needed to display any page of the document or to retrieve certain other information efficiently. The hint tables may contain additional information to optimize access by plug-in extensions to application-specific data.

The hint tables are not logically part of the information content of the document; they can be derived from the document. Any action that changes the document—for instance, appending an incremental update—will invalidate the hint tables. The document will remain a valid PDF file but will no longer be linearized; see Section F.4.6, “Accessing an Updated File,” for details.

The hint tables are binary data structures that are enclosed in a stream object. Syntactically, this stream is a normal PDF indirect object. However, there are no references to the stream anywhere in the document, so it is not logically part of the document; an operation that regenerates the document may remove the stream.

Usually, all the hint tables are contained in a single stream, known as the *primary hint stream*. Optionally, there may be an additional stream containing more hints, known as the *overflow hint stream*. The contents of the two hint streams are to be concatenated and treated as if they were a single unbroken stream.

The primary hint stream, which is required, is shown as part 5 in Example F.1. The order of this part and the first-page section, shown as part 6, may be reversed; see Section F.4, “Access Strategies,” for considerations on the choice of placement. The overflow hint stream, part 10, is optional. (See implementation note 122 in Appendix H.)

The location and length of the primary hint stream, and of the overflow hint stream if present, are given in the linearization parameter dictionary at the beginning of the file.

The hint streams are assigned the last object numbers in the file—that is, after the object number for the last object in the first page. Their cross-reference table entries are at the end of the first-page cross-reference table. This object number assignment is independent of the physical locations of the hint streams in the file. (This convention keeps their object numbers out of the way of the numbering of the linearized objects.)

The values of all entries in the hint streams’ dictionaries must be direct objects, and may contain no indirect object references.

In addition to the standard stream attributes, the dictionary of the primary hint stream contains entries giving the position of the beginning of each hint table in the stream. These positions are given in bytes relative to the beginning of the

stream data (after decoding filters, if any, are applied) and with the overflow hint stream concatenated if present. The dictionary of the overflow hint stream should not contain these entries. The keys designating the standard hint tables in the primary hint stream's dictionary are listed in Table F.2; Section F.3, "Hint Tables," documents the format of these hint tables. Additionally, there is a required page offset hint table, which must be the first table in the stream and must start at offset 0.

TABLE F.2 Standard hint tables

KEY	HINT TABLE
S	<i>(Required)</i> Shared object hint table
T	<i>(Present only if thumbnail images exist)</i> Thumbnail hint table
O	<i>(Present only if a document outline exists)</i> Outline hint table
A	<i>(Present only if article threads exist)</i> Thread information hint table
E	<i>(Present only if named destinations exist)</i> Named destination hint table
V	<i>(Present only if an interactive form dictionary exists)</i> Interactive form hint table
I	<i>(Present only if a document information dictionary exists)</i> Information dictionary hint table
C	<i>(Present only if a logical structure hierarchy exists; PDF 1.3)</i> Logical structure hint table
L	<i>(PDF 1.3)</i> Page label hint table

New keys may be registered for additional hint tables required for new PDF features or for application-specific data accessed by plug-in extensions. See Appendix E for further information.

F.2.6 First-Page Section (Part 6)

As mentioned earlier, the section containing objects belonging to the first page of the document may either precede or follow the primary hint stream. The starting file offset and length of this section can be determined from the hint tables. In addition, the **E** entry in the linearization parameter dictionary specifies the end of the first page (as an offset relative to the beginning of the file), and the **O** entry gives the object number of the first page's page object.

This part of the file contains all the objects needed to display the first page of the document. Ordinarily, the “first page” is page 0—that is, the leftmost leaf page node in the page tree. However, if the document catalog contains an **OpenAction** entry that specifies opening at some page other than page 0, then that page is the “first page” and should be located here. The page number of the first page is given in the **P** entry of the linearization parameter dictionary. (See also implementation note 124 in Appendix H.)

The objects contained here should include the following:

- The page object for the first page. This must be the first object in this part of the file. Its object number is given in the linearization parameter dictionary. This page object must explicitly specify all required attributes, such as **Resources** and **MediaBox**; the attributes cannot be inherited from ancestor page tree nodes.
- The entire outline hierarchy, if the value of the **PageMode** entry in the catalog is **UseOutlines**. (If the **PageMode** entry is omitted or has some other value and the document has an outline hierarchy, the outline hierarchy appears in part 9; see Section F.2.9, “Other Objects (Part 9)” for details.)
- All objects that the page object refers to, to an arbitrary depth, except page tree nodes or other page objects. This includes objects referred to by its **Contents**, **Resources**, **Annots**, and **B** entries, but not **Thumb**.

The order of objects referenced from the page object should facilitate early user interaction and incremental display of the page data as it arrives. The following order is recommended:

1. The **Annots** array and all annotation dictionaries, to a depth sufficient to allow those annotations to be activated. Information required to draw the annotation can be deferred until later, since annotations are always drawn on top of (hence after) the contents.
2. The **B** (beads) array and all bead dictionaries, if any, for this page. If any beads exist for this page, the **B** array is required to be present in the page dictionary. Additionally, each bead in the thread (not just the first bead) must contain a **T** entry referring to the associated thread dictionary.
3. The resource dictionary, but not the resource objects contained in the dictionary.

4. Resource objects, other than the types listed below, in the order that they are first referenced (directly or indirectly) from the content stream. If the contents are represented as an array of streams, each resource object should precede the stream in which it is first referenced. Note that **Font**, **FontDescriptor**, and **Encoding** resources should be included here, but not substitutable font files referenced from font descriptors (see the last item below).
5. The page contents (**Contents**). If large, this should be represented as an array of indirect references to content streams, which in turn are interleaved with the resources they require. If small, the entire contents should be a single content stream preceding the resources.
6. Image XObjects, in the order that they are first referenced. Images are assumed to be large and slow to transfer, so the viewer application defers rendering images until all the other contents have been displayed.
7. **FontFile** streams, which contain the actual definitions of embedded fonts. These are assumed to be large and slow to transfer, so the viewer application uses substitute fonts until the real ones have arrived. Only those fonts for which substitution is possible can be deferred in this way. (Currently, this includes any Type 1 or TrueType font that has a font descriptor with the Nonsymbolic flag set, indicating the Adobe standard Latin character set).

See Section F.4, “Access Strategies,” for additional discussion about object order and incremental drawing strategies.

F.2.7 Remaining Pages (Part 7)

Part 7 of the Linearized PDF file contains the page objects and nonshared objects for all remaining pages of the file, with the objects for each page grouped together. The pages are contiguous and are ordered by page number. If the first page of the file is not page 0, this section starts with page 0 and skips over the first page when its position in the sequence is reached.

For each page, the objects required to display that page are grouped together, except for resources and other objects that are shared with other pages. Shared objects are located in the shared objects section (part 8). The starting file offset and length of any page can be determined from the hint tables.

The recommended order of objects within a page is essentially the same as in the first page. In particular, the page object must be the first object in each section.

In most cases, unlike for the first page, there will be little benefit from interleaving contents with resources. This is because most resources other than images—fonts in particular—are shared among multiple pages and therefore reside in the shared objects section. Image XObjects usually are not shared, but they should appear at the end of the page’s section of the file, since rendering of images is deferred.

F.2.8 Shared Objects (Part 8)

Part 8 of the file contains objects, primarily named resources, that are referenced from more than one page but that are not referenced (directly or indirectly) from the first page. The hint tables contain an index of these objects. For more information on named resources, see Section 3.7.2, “Resource Dictionaries.”

The order of these objects is essentially arbitrary. However, wherever a resource consists of a multiple-level structure, all components of the structure should be grouped together. If only the top-level object is referenced from outside the group, the entire group can be described by a single entry in the shared object hint table. This helps to minimize the size of the shared object hint table and the number of individual references from entries in the page offset hint table. (See also implementation note 125 in Appendix H.)

F.2.9 Other Objects (Part 9)

Following the shared objects are any other objects that are part of the document but are not required for displaying pages. These objects are divided into functional categories. Objects within each of these categories should be grouped together; the relative order of the categories is unimportant.

- *The page tree.* This can be located here, since the viewer application never needs to consult it. Note that all **Resources** attributes and other inheritable attributes of the page objects must be pushed down and replicated in each of the leaf page objects (but they may contain indirect references to shared objects).
- *Thumbnail images.* These should simply be ordered by page number. Note that the thumbnail image for page 0 should be first, even if the first page of the document is some page other than 0. Each thumbnail image consists of one or more objects, which may refer to objects in the thumbnail shared objects section (see the next item).

- *Thumbnail shared objects.* These are objects that are shared among some or all thumbnail images and are not referenced from any other objects.
- *The outline hierarchy,* if not located in part 6. The order of objects should be the same as the order in which they are displayed by the viewer application. This is a preorder traversal of the outline tree, skipping over any subtree that is closed (that is, whose parent's **Count** value is negative); following that should be the subtrees that were skipped over, in the order in which they would have appeared if they were all open.
- *Thread information dictionaries,* referenced from the **I** entries of thread dictionaries. Note that the thread dictionaries themselves are located with the document catalog, and the bead dictionaries with the individual pages.
- *Named destinations.* These objects include the value of the **Dests** or **Names** entry in the document catalog and all the destination objects that it refers to. See Section F.4.2, "Opening at an Arbitrary Page."
- *The document information dictionary* and the objects contained within it.
- *The interactive form field hierarchy.* This does not include the top-level interactive form dictionary, which is located with the document catalog.
- *Other entries* in the document catalog that are not referenced from any page.
- *(PDF 1.3) The logical structure hierarchy.*

F.2.10 Main Cross-Reference and Trailer (Part 11)

Part 11 is the cross-reference table for all objects in the PDF file except those listed in the first-page cross-reference table (part 3). As indicated earlier, this cross-reference table plays the role of the original cross-reference table for the file (before any updates are appended). It must conform to the following rules:

- It consists of a single cross-reference subsection, beginning at object number 0.
- The first entry (for object number 0) must be a free entry.
- The remaining entries are for in-use objects, which are numbered consecutively starting at 1.

As indicated earlier, the **startxref** line gives the offset of the first-page cross-reference table. The **Prev** entry of the first-page trailer gives the offset of the main

cross-reference table. The main trailer has no **Prev** entry, and in fact does not need to contain any entries other than **Size**.

F.3 Hint Tables

The core of the linearization information is stored in two or more hint tables, as indicated by the attributes of the primary hint stream (see Section F.2.5, “Hint Streams (Parts 5 and 10)”). The format of the standard hint tables is described in this section.

There can be additional hint tables for application-specific data that is accessed by plug-in extensions. A generic format for such hint tables is defined; see Section F.3.4, “Generic Hint Tables.” Alternatively, the format of a hint table can be private to the application; see Appendix E for further information.

Each hint table consists of a portion of the stream, beginning at the position in the stream indicated by the corresponding stream attribute. Additionally, there is a required page offset hint table, which must be the first table in the stream and must start at offset 0. (If there is an overflow hint stream, its contents are to be appended seamlessly to the primary hint stream; hint table positions are relative to the beginning of this combined stream.) In general, this byte stream is treated as a bit stream, high-order bit first, which is then subdivided into fields of arbitrary width without regard to byte boundaries. However, each hint table begins at a byte boundary.

The hint tables are designed to encode the required information as compactly as possible. Interpreting the hint tables requires reading them sequentially; they are not designed for random access. The client is expected to read and decode the tables once and retain the information for as long as the document remains open.

A hint table encodes the positions of various objects in the file. The representation is either explicit (an offset from the beginning of the file) or implicit (accumulated lengths of preceding objects). Regardless of the representation, the resulting positions must be interpreted as if the primary hint stream itself were not present. That is, a position greater than the *hint stream offset* must have the *hint stream length* added to it in order to determine the actual offset relative to the beginning of the file. (The hint stream offset and hint stream length are the values $offset_1$ and $length_1$ in the **H** array in the linearization parameter dictionary at the beginning of the file.)

The reason for this rule is that the length of the primary hint stream depends on the information contained within the hint tables, and this is not known until after they have been generated. Any information contained in the hint tables must not depend on knowing the primary hint stream's length in advance.

Note that this rule applies only to offsets given in the hint tables and not to offsets given in the cross-reference tables or linearization parameter dictionary. Also, the offset and length of the overflow hint stream, if present, need not be taken into account, since this object follows all other objects in the file.

F.3.1 Page Offset Hint Table

The page offset hint table provides information required for locating each page. Additionally, for each page except the first, it also enumerates all shared objects that the page references, directly or indirectly.

This table begins with a header section, described in Table F.3, followed by one or more per-page entries, described in Table F.4. Note that the items making up each per-page entry are not contiguous; they are broken up with items from entries for other pages. The order of items making up the per-page entries is as follows:

1. Item 1 for all pages, in page order starting with the first page
2. Item 2 for all pages, in page order starting with the first page
3. Item 3 for all pages, in page order starting with the first page
4. Item 4 for all shared objects in the second page, followed by item 4 for all shared objects in the third page, and so on
5. Item 5 for all shared objects in the second page, followed by item 5 for all shared objects in the third page, and so on
6. Item 6 for all pages, in page order starting with the first page
7. Item 7 for all pages, in page order starting with the first page

Note: All the “bits needed” items in Table F.3, such as item 3, may have values in the range 0 through 32. Although that range requires only 6 bits, 16-bit numbers are used.

TABLE F.3 Page offset hint table, header section

ITEM	SIZE (BITS)	DESCRIPTION
1	32	The least number of objects in a page (including the page object itself).
2	32	The location of the first page's page object.
3	16	The number of bits needed to represent the difference between the greatest and least number of objects in a page.
4	32	The least length of a page in bytes. This is the least length from the beginning of a page object to the last byte of the last object used by that page.
5	16	The number of bits needed to represent the difference between the greatest and least length of a page, in bytes.
6	32	The least offset of the start of any content stream, relative to the beginning of its page. (See implementation note 126 in Appendix H.)
7	16	The number of bits needed to represent the difference between the greatest and least offset to the start of the content stream. (See implementation note 126 in Appendix H.)
8	32	The least content stream length. (See implementation note 127 in Appendix H.)
9	16	The number of bits needed to represent the difference between the greatest and least content stream length. (See implementation note 127 in Appendix H.)
10	16	The number of bits needed to represent the greatest number of shared object references.
11	16	The number of bits needed to represent the numerically greatest shared object identifier used by the pages (discussed further in Table F.4, item 4).
12	16	The number of bits needed to represent the numerator of the fractional position for each shared object reference. For each shared object referenced from a page, there is an indication of where in the page's content stream the object is first referenced. That position is given as the numerator of a fraction, whose denominator is specified once for the entire document (in the next item in this table). The fraction is explained in more detail in Table F.4, item 5.
13	16	The denominator of the fractional position for each shared object reference.

TABLE F.4 Page offset hint table, per-page entry

ITEM	SIZE (BITS)	DESCRIPTION
1	See Table F.3, item 3	A number that, when added to the least number of objects in a page (Table F.3, item 1), gives the number of objects in the page. The first object of the first page has an object number that is the value of the O entry in the linearization parameter dictionary at the beginning of the file. The first object of the second page has an object number of 1. Object numbers for subsequent pages can be determined by accumulating the number of objects in all previous pages.
2	See Table F.3, item 5	A number that, when added to the least page length (Table F.3, item 4), gives the length of the page in bytes. The location of the first object of the first page can be determined from its object number (the O entry in the linearization parameter dictionary) and the cross-reference table entry for that object (see Section F.2.3, “First-Page Cross-Reference Table and Trailer (Part 3)”). The locations of subsequent pages can be determined by accumulating the lengths of all previous pages. Note that it is necessary to skip over the primary hint stream, wherever it is located.
3	See Table F.3, item 10	The number of shared objects referenced from the page. Note that this must be 0 for the first page, and that the next two items start with the second page.
4	See Table F.3, item 11	<i>(One item for each shared object referenced from the page)</i> A <i>shared object identifier</i> —that is, an index into the shared object hint table (described in Section F.3.2, “Shared Object Hint Table”). Note that a single entry in the shared object hint table can designate a group of shared objects, only one of which is referenced from outside the group. That is, shared object identifiers are not directly related to object numbers. This identifier combines with the numerators provided in item 5 to form a <i>shared object reference</i> .
5	See Table F.3, item 12	<i>(One item for each shared object referenced from the page)</i> The numerator of the fractional position for each shared object reference, in the same order as the preceding item. The fraction indicates where in the page’s content stream the shared object is first referenced. This item is interpreted as the numerator of a fraction whose denominator is specified once for the entire document (Table F.3, item 13). If the denominator is d , a numerator ranging from 0 to $d - 1$ indicates the corresponding portion of the page’s content stream. For example, if the denominator is 4, a numerator of 0, 1, 2, or 3 indicates that the first reference lies in the first, second, third, or fourth quarter of the content stream, respectively.

There are two (or more) other possible values for the numerator, which indicate that the shared object is not referenced from the content stream but is needed by annotations or other objects that are drawn after the contents. The value d indicates that the shared object is needed before image XObjects and other nonshared objects that are at the end of the page. A value of $d + 1$ or greater indicates that the shared object is needed after those objects.

This method of dividing the page into fractions is only approximate. Determining the first reference to a shared object entails inspecting the unencoded content stream. The relationship between positions in the unencoded and encoded streams is not necessarily linear.

- | | | |
|---|-----------------------|--|
| 6 | See Table F.3, item 7 | A number that, when added to the least offset to the start of the content stream (Table F.3, item 6), gives the offset in bytes of the start of the page's content stream, relative to the beginning of the page. This is the offset of the stream object, not the stream data. (See implementation note 126 in Appendix H.) |
| 7 | See Table F.3, item 9 | A number that, when added to the least content stream length (Table F.3, item 8), gives the length of the page's content stream in bytes. This includes object overhead preceding and following the stream data. (See implementation note 127 in Appendix H.) |
-

F.3.2 Shared Object Hint Table

The shared object hint table gives information required to locate shared objects (see Section F.2.8, “Shared Objects (Part 8)”). Shared objects can be physically located in either of two places: objects that are referenced from the first page are located with the first-page objects (part 6); all other shared objects are located in the shared objects section (part 8).

A single entry in the shared object hint table can actually describe a group of adjacent objects, under the following condition: Only the first object in the group is referenced from outside the group; the remaining objects in the group are referenced only from other objects in the same group. The objects in a group must have adjacent object numbers.

The page offset hint table, interactive form hint table, and logical structure hint table refer to an entry in the shared object hint table by a simple index that is its sequential position in the table, counting from 0.

The shared object hint table consists of a header section (Table F.5), followed by one or more shared object group entries (Table F.6). There are two sequences of shared object group entries: the ones for objects located in the first page, followed by the ones for objects located in the shared objects section. The entries have the same format in both cases. Note that the items making up each shared object group entry are not contiguous; they are broken up with items from entries for other shared object groups. The order of items in each sequence is as follows:

1. Item 1 for the first group, item 1 for the second group, and so on
2. Item 2 for the first group, item 2 for the second group, and so on
3. Item 3 for the first group, item 3 for the second group, and so on
4. Item 4 for the first group, item 4 for the second group, and so on

All objects associated with the first page (part 6) have entries in the shared object hint table, whether or not they are actually shared. The first entry refers to the beginning of the first page and has an object count and length that span all the initial nonshared objects. The next entry refers to a group of shared objects. Subsequent entries span additional groups of either shared or nonshared objects consecutively, until all shared objects in the first page have been enumerated. (The entries that refer to nonshared objects will never be used.)

TABLE F.5 Shared object hint table, header section

ITEM	SIZE (BITS)	DESCRIPTION
1	32	The object number of the first object in the shared objects section (part 8).
2	32	The location of the first object in the shared objects section.
3	32	The number of shared object entries for the first page (including nonshared objects, as noted above).
4	32	The number of shared object entries for the shared objects section. This includes the number of shared object entries for the first page (that is, the value of item 3).
5	16	The number of bits needed to represent the greatest number of objects in a shared object group. (See also implementation note 128 in Appendix H.)
6	32	The least length of a shared object group in bytes.
7	16	The number of bits needed to represent the difference between the greatest and least length of a shared object group, in bytes.

TABLE F.6 Shared object hint table, shared object group entry

ITEM	SIZE (BITS)	DESCRIPTION
1	See Table F.5, item 7	A number that, when added to the least shared object group length (Table F.5, item 6), gives the length of the object group in bytes. The location of the first object of the first page is given in the page offset hint table, header section (Table F.3, item 4). The locations of subsequent object groups can be determined by accumulating the lengths of all previous object groups until all shared objects in the first page have been enumerated. Following that, the location of the first object in the shared objects section can be obtained from the header section of the shared object hint table (Table F.5, item 2).
2	1	A flag indicating whether the shared object signature (item 3) is present; its value is 1 if the signature is present and 0 if it is absent. (See also implementation note 129 in Appendix H.)
3	128	<i>(Only if item 2 is 1)</i> The <i>shared object signature</i> , a 16-byte MD5 hash that uniquely identifies the resource that the group of objects represents. This is intended to enable the client to substitute a locally cached copy of the resource instead of reading it from the PDF file. Note that this signature is unrelated to signature fields in interactive forms, as defined in the section “Signature Fields” on page 547.
4	See Table F.5, item 5	A number equal to 1 less than the number of objects in the group. The first object of the first page is the one whose object number is given by the O entry in the linearization parameter dictionary at the beginning of the file. Object numbers for subsequent entries can be determined by accumulating the number of objects in all previous entries, until all shared objects in the first page have been enumerated. Following that, the first object in the shared objects section has a number that can be obtained from the header section of the shared object hint table (Table F.5, item 1). (See also implementation note 130 in Appendix H.)

Note: *In a document consisting of only one page, all of that page’s objects are nevertheless treated as if they were shared; the shared object hint table reflects this. (See implementation note 131 in Appendix H.)*

F.3.3 Thumbnail Hint Table

The thumbnail hint table consists of a header section (Table F.7), followed by the thumbnails section, which includes one or more per-page entries (Table F.8),

each of which describes the thumbnail image for a single page. The entries are in page number order starting with page 0, even if the document catalog contains an **OpenAction** entry that specifies opening at some page other than page 0. Thumbnail images may exist for some pages and not for others.

TABLE F.7 Thumbnail hint table, header section

ITEM	SIZE (BITS)	DESCRIPTION
1	32	The object number of the first thumbnail image (that is, the thumbnail image that is described by the first entry in the thumbnails section).
2	32	The location of the first thumbnail image, specified as an offset from the beginning of the file to that thumbnail image, minus the length of the thumbnail hint table.
3	32	The number of pages that have thumbnail images.
4	16	The number of bits needed to represent the greatest number of consecutive pages that do not have a thumbnail image.
5	32	The least length of a thumbnail image in bytes.
6	16	The number of bits needed to represent the difference between the greatest and least length of a thumbnail image.
7	32	The least number of objects in a thumbnail image.
8	16	The number of bits needed to represent the difference between the greatest and least number of objects in a thumbnail image.
9	32	The object number of the first object in the thumbnail shared objects section (a subsection of part 8). These are objects (color spaces, for example) that are referenced from some or all thumbnail objects and are not referenced from any other objects. The thumbnail shared objects are undifferentiated; there is no indication of which shared objects are referenced from any given page's thumbnail image.
10	32	The location of the first object in the thumbnail shared objects section, specified as the offset from the beginning of the file to that object, minus the length of the thumbnail hint table.
11	32	The number of thumbnail shared objects.
12	32	The length of the thumbnail shared objects section in bytes.

TABLE F.8 Thumbnail hint table, per-page entry

ITEM	SIZE (BITS)	DESCRIPTION
1	See Table F.7, item 4	<i>(Optional)</i> The number of preceding pages lacking a thumbnail image. This indicates how many pages without a thumbnail image lie between the previous entry's page and this one.
2	See Table F.7, item 8	A number that, when added to the least number of objects in a thumbnail image (Table F.7, item 7), gives the number of objects in this page's thumbnail image.
3	See Table F.7, item 6	A number that, when added to the least length of a thumbnail image (Table F.7, item 5), gives the length of this page's thumbnail image in bytes.

The order of items in Table F.8 is as follows:

1. Item 1 for all pages, in page order starting with the first page
2. Item 2 for all pages, in page order starting with the first page
3. Item 3 for all pages, in page order starting with the first page

F.3.4 Generic Hint Tables

Certain categories of objects are associated with the document as a whole rather than with individual pages (see Section F.2.9, “Other Objects (Part 9)”), and it is sometimes useful to provide hints for accessing those objects efficiently. For each category of hints, there is a separate entry in the primary hint stream giving the starting position of the table within the stream (see Section F.2.5, “Hint Streams (Parts 5 and 10)”).

There is a generic representation for such hints, specified below. This representation is useful for some standard categories of objects, such as outline items, article threads, and named destinations. It may also be useful for application-specific objects accessed by plug-in extensions. It is considerably more convenient for a plug-in to use the generic hint representation than to specify custom hints.

A generic hint table describes a single group of objects that are located together in the PDF file. The entries in this table are listed in Table F.9.

TABLE F.9 Generic hint table

ITEM	SIZE (BITS)	DESCRIPTION
1	32	The object number of the first object in the group.
2	32	The location of the first object in the group.
3	32	The number of objects in the group.
4	32	The length of the object group in bytes.

F.3.5 Interactive Form and Logical Structure Hint Tables

If an interactive form dictionary or structure tree root is present, the interactive form and logical structure hint tables refer to the contents of those dictionaries; see Section F.2.9, “Other Objects (Part 9).” Interactive forms and logical structure hierarchies can refer to objects that are shared with other parts of the document, in which case the hint table lists those shared objects.

An interactive form or logical structure hint table begins with the same entries as in a generic hint table, followed by three additional entries, as shown in Table F.10.

TABLE F.10 Interactive form or logical structure hint table

ITEM	SIZE (BITS)	DESCRIPTION
1	32	The object number of the first object in the group.
2	32	The location of the first object in the group.
3	32	The number of objects in the group.
4	32	The length of the object group in bytes.
5	32	The number of shared object references.
6	16	The number of bits needed to represent the numerically greatest shared object identifier used by the interactive form or the logical structure hierarchy.
7...	See Table F.3, item 11	Starting with item 7, each of the remaining items in this table is a shared object identifier—that is, an index into the shared object hint table (described in Section F.3.2, “Shared Object Hint Table”).

F.3.6 Other Hint Tables

The outline, thread information, named destination, information dictionary, and page label hint tables use the generic hint table representation; see Section F.3.4, “Generic Hint Tables.” The objects that they refer to are grouped together as described in Section F.2.9, “Other Objects (Part 9).”

F.4 Access Strategies

This section outlines how the client can take advantage of the structure of a Linearized PDF file in order to retrieve and display it efficiently. This material is not formally a part of the Linearized PDF specification, but it may help explain the rationale for the organization.

F.4.1 Opening at the First Page

As described earlier, when a document is initially accessed, a request is issued to retrieve the entire file, starting at the beginning. Consequently, Linearized PDF is organized so that all the data required to display the first page is at the beginning of the file. This includes all resources that are referenced from the first page, whether or not they are also referenced from other pages.

The first page is usually but not necessarily page 0. If the document catalog contains an **OpenAction** entry that specifies opening at some page other than page 0, that page will be the one physically located at the beginning of the document. Thus, opening a document at the default place (rather than a specific destination) requires simply waiting for the first-page data to arrive; no additional transactions are required.

In an ordinary PDF viewer application, opening a document requires first positioning to the end to obtain the **startxref** line. Since a Linearized PDF file has the first page’s cross-reference table at the beginning, reading the **startxref** line is not necessary. All that is required is to verify that the file length given in the linearization parameter dictionary at the beginning of the file matches the actual length of the file, indicating that no updates have been appended to the PDF file.

The primary hint stream is located either before or after the first-page section. This means that it will also be retrieved as part of the initial sequential read of the file. The client is expected to interpret and retain all the information in the hint

tables. They are reasonably compact and are not designed to be obtained from the file in random pieces.

The client must now decide whether to continue reading the remainder of the document sequentially or to abort the initial transaction and access subsequent pages using separate transactions requesting byte ranges. This decision is a function of the size of the file, the data rate of the channel, and the overhead cost of a transaction.

F.4.2 Opening at an Arbitrary Page

The viewer application may be requested to open a PDF file at an arbitrary page. The page can be specified in one of three ways:

- By page number (remote go-to action, integer page specifier)
- By named destination (remote go-to action, name or string page specifier)
- By article thread (thread action)

Additionally, an indexed search results in opening a document by page number. Handling this case efficiently is especially important.

As indicated above, when the document is initially opened, it is retrieved sequentially starting at the beginning. As soon as the hint tables have been received, the client has sufficient information to request retrieval of any page of the document given its page number. Therefore, it can abort the initial transaction and issue a new transaction for the target page, as described in Section F.4.3, “Going to Another Page of an Open Document.”

The position of the primary hint stream (part 5) with respect to the first-page section (part 6) determines how quickly this can be done. If the primary hint stream precedes the first-page section, the initial transaction can be aborted very quickly; however, this is at the cost of increased delay when opening the document at the first page. On the other hand, if the primary hint stream follows the first-page section, displaying the first page is quicker (since the hint tables are not needed for that), but opening at an arbitrary page is delayed by the time required to receive the first page. The decision whether to favor opening at the first page or opening at an arbitrary page must be made at the time a PDF file is linearized.

If an overflow hint stream exists, obtaining it requires issuing an additional transaction. For this reason, inclusion of an overflow hint stream in Linearized PDF, although permitted, is not recommended. The feature exists to allow the linearizer to write the PDF file with space reserved for a primary hint stream of an estimated size, and then go back and fill in the hint tables. If the estimate is too small, the linearizer can append an overflow stream containing the remaining hint table data. This enables the PDF file to be written in one pass, which may be an advantage if the performance of writing PDF is considered important.

Opening at a named destination requires the viewer application first to read the entire **Dests** or **Names** dictionary, for which a hint is present. Using this information, it is possible to determine the page containing the specific destination identified by the name.

Opening to an article requires the viewer application first to read the entire **Threads** array, which is located with the document catalog at the beginning of the document. Using this information, it is possible to determine the page containing the first bead of any thread. Opening at other than the first bead of a thread requires chaining through all the beads until the desired one is reached; there are no hints to accelerate this.

F.4.3 Going to Another Page of an Open Document

Given a page number and the information in the hint tables, it is now straightforward for the client to construct a single request to retrieve any arbitrary page of the document. The request should include:

- The objects of the page itself, whose byte range can be determined from the entry in the page offset hint table.
- The portion of the main cross-reference table referring to those objects. This can be computed from main cross-reference table location (the **T** entry in the linearization parameter dictionary) and the cumulative object number in the page offset hint table.
- The shared objects referenced from the page, whose byte ranges can be determined from information in the shared object hint table.
- The portion or portions of the main cross-reference table referring to those objects, as described above.

The purpose of the fractions in the page offset hint table is to enable the client to schedule retrieval of the page in a way that allows incremental display of the data as it arrives. It accomplishes this by constructing a request that interleaves pieces of the page contents with the shared resources that the contents refer to. This serves much the same purpose as the physical interleaving that is done for the first page.

F.4.4 Drawing a Page Incrementally

The ordering of objects in pages and the organization of the hint tables are intended to allow progressive update of the display and early opportunities for user interaction when the data is arriving slowly. The viewer application must recognize instances in which the targets of indirect object references have not yet arrived and, where possible, rearrange the order in which it acts on the objects in the page.

The following sequence of actions is recommended:

1. Activate the annotations, but do not draw them yet. Also activate the cursor feedback for any article threads in the page.
2. Begin drawing the contents. Whenever there is a reference to an image XObject that has not yet arrived, skip over it. Whenever there is a reference to a font whose definition is an embedded font file that has not yet arrived, draw the text using a substitute font (if that is possible).
3. Draw the annotations.
4. Draw the images as they arrive, together with anything that overlaps them.
5. Once the embedded font definitions have arrived, redraw the text using the correct fonts, together with anything that overlaps the text.

The last two steps should be done using an off-screen buffer if possible, to avoid objectionable flashing during the redraw process.

On encountering a reference XObject (see Section 4.9.3, “Reference XObjects”), the viewer application may choose to initially display the object itself as a proxy and defer the retrieval and rendering of the imported content. Note that, since all XObjects in a Linearized PDF file follow the content stream of the page on which they appear, their retrieval is already deferred; the use of a reference XObject will result in an additional level of deferral.

F.4.5 Following an Article Thread

As indicated earlier, the bead dictionaries for any article thread that visits a given page are located with that page. This enables the bead rectangles to be activated and proper cursor feedback to be shown.

If the user follows a thread, the viewer application can obtain the object number from the **N** or **P** entry of the bead dictionary. This identifies a target bead, which is located with the page to which it belongs. Given this object number, the viewer application can then go to that page, as discussed in Section F.4.3, “Going to Another Page of an Open Document.”

F.4.6 Accessing an Updated File

As stated earlier, if a Linearized PDF file subsequently has an incremental update appended to it, the linearization and hints are no longer valid. Actually, this is not necessarily true, but the viewer application must do some additional work to validate them.

When the viewer application sees that the file is longer than the length given in the linearization parameter dictionary, it must issue an additional transaction to read everything that was appended. It must then analyze the objects in that update to see whether any of them modify objects that are in the first page or that are the targets of hints. If so, it must augment its internal data structures as necessary to take the updates into account.

For a PDF file that has received only a small update, this approach may be worthwhile. Accessing the file this way will be quicker than accessing it without hints or retrieving the entire file before displaying any of it.

APPENDIX G

Example PDF Files

THIS APPENDIX PRESENTS several examples showing the structure of actual PDF files:

- A minimal file that can serve as a starting point for creating other PDF files (and that is the basis of later examples)
- A simple example that shows a text string—the classic “Hello World”—and a simple graphics example that draws lines and shapes
- A fragment of a PDF file that illustrates the structure of the page tree for a large document and, similarly, two fragments that illustrate the structure of an outline hierarchy
- Finally, an example showing the structure of a PDF file as it is updated several times, illustrating multiple body sections, cross-reference sections, and trailers

*Note: The **Length** values of stream objects in the examples and the byte addresses in cross-reference tables are not necessarily accurate.*

G.1 Minimal PDF File

Example G.1 is a PDF file that does not draw anything; it is almost the minimum acceptable PDF file. It is not strictly the minimum acceptable because it contains an outline dictionary (**Outlines** in the document catalog) with a zero count (in which case this object would normally be omitted); a page content stream (**Contents** in the page object); and a resource dictionary (**Resources** in the page object) containing a **ProcSet** array. These objects were included to make this file useful as a starting point for creating other, more realistic PDF files.

Table G.1 lists the objects present in this example.

TABLE G.1 Objects in minimal example

OBJECT NUMBER	OBJECT TYPE
1	Catalog (document catalog)
2	Outlines (outline dictionary)
3	Pages (page tree node)
4	Page (page object)
5	Content stream
6	Procedure set array

Note: When using Example G.1 as a starting point for creating other files, remember to update the **ProcSet** array as needed (see Section 9.1, “Procedure Sets”). Also, remember that the cross-reference table entries may need to have a trailing space (see Section 3.4.3, “Cross-Reference Table”).

Example G.1

```
%PDF-1.4
1 0 obj
  << /Type /Catalog
    /Outlines 2 0 R
    /Pages 3 0 R
  >>
endobj

2 0 obj
  << /Type Outlines
    /Count 0
  >>
endobj

3 0 obj
  << /Type /Pages
    /Kids [4 0 R]
    /Count 1
  >>
endobj
```

```
4 0 obj
  << /Type /Page
    /Parent 3 0 R
    /MediaBox [0 0 612 792]
    /Contents 5 0 R
    /Resources << /ProcSet 6 0 R >>
  >>
```

```
>>
endobj
```

```
5 0 obj
  << /Length 35 >>
stream
...Page-marking operators...
endstream
endobj
```

```
6 0 obj
  [/PDF]
endobj
```

```
xref
0 7
0000000000 65535 f
0000000009 00000 n
0000000074 00000 n
0000000120 00000 n
0000000179 00000 n
0000000300 00000 n
0000000384 00000 n
```

```
trailer
  << /Size 7
    /Root 1 0 R
  >>
```

```
startxref
408
%%EOF
```

G.2 Simple Text String Example

Example G.2 is the classic “Hello World” example built from the preceding example. It shows a single line of text consisting of the string Hello World, illustrating the use of fonts and several text-related PDF operators. The string is displayed in 24-point Helvetica; because Helvetica is one of the standard 14 fonts, no font descriptor is needed.

Table G.2 lists the objects present in this example.

TABLE G.2 Objects in simple text string example	
OBJECT NUMBER	OBJECT TYPE
1	Catalog (document catalog)
2	Outlines (outline dictionary)
3	Pages (page tree node)
4	Page (page object)
5	Content stream
6	Procedure set array
7	Font (Type 1 font)

Example G.2

```
%PDF-1.4
1 0 obj
  << /Type /Catalog
    /Outlines 2 0 R
    /Pages 3 0 R
  >>
endobj

2 0 obj
  << /Type /Outlines
    /Count 0
  >>
endobj
```



```
3 0 obj
  << /Type /Pages
    /Kids [4 0 R]
    /Count 1
  >>
endobj

4 0 obj
  << /Type /Page
    /Parent 3 0 R
    /MediaBox [0 0 612 792]
    /Contents 5 0 R
    /Resources << /ProcSet 6 0 R
      /Font << /F1 7 0 R >>
    >>
  >>
endobj

5 0 obj
  << /Length 73 >>
stream
  BT
    /F1 24 Tf
    100 100 Td
    (Hello World) Tj
  ET
endstream
endobj

6 0 obj
  [/PDF /Text]
endobj

7 0 obj
  << /Type /Font
    /Subtype /Type1
    /Name /F1
    /BaseFont /Helvetica
    /Encoding /MacRomanEncoding
  >>
endobj
```

```

xref
0 8
0000000000 65535 f
0000000009 00000 n
0000000074 00000 n
0000000120 00000 n
0000000179 00000 n
0000000364 00000 n
0000000466 00000 n
0000000496 00000 n

trailer
<< /Size 8
    /Root 1 0 R
>>
startxref
625
%%EOF

```

G.3 Simple Graphics Example

Example G.3 draws a thin black line segment, a thick black dashed line segment, a filled and stroked rectangle, and a filled and stroked cubic Bézier curve. Table G.3 lists the objects present in this example, and Figure G.1 shows the resulting output. (Each shape has a red border, and the rectangle is filled with light blue.)

TABLE G.3 Objects in simple graphics example

OBJECT NUMBER	OBJECT TYPE
1	Catalog (document catalog)
2	Outlines (outline dictionary)
3	Pages (page tree node)
4	Page (page object)
5	Content stream
6	Procedure set array




FIGURE G.1 *Output of Example G.3*

Example G.3

```
%PDF-1.4
1 0 obj
  << /Type /Catalog
    /Outlines 2 0 R
    /Pages 3 0 R
  >>
endobj

2 0 obj
  << /Type /Outlines
    /Count 0
  >>
endobj

3 0 obj
  << /Type /Pages
    /Kids [4 0 R]
    /Count 1
  >>
endobj
```

```
4 0 obj
  << /Type /Page
    /Parent 3 0 R
    /MediaBox [0 0 612 792]
    /Contents 5 0 R
    /Resources << /ProcSet 6 0 R >>
  >>
endobj

5 0 obj
  << /Length 883 >>
stream
  % Draw a black line segment, using the default line width.
  150 250 m
  150 350 l
  S

  % Draw a thicker, dashed line segment.
  4 w                               % Set line width to 4 points
  [4 6] 0 d                          % Set dash pattern to 4 units on, 6 units off
  150 250 m
  400 250 l
  S

  [] 0 d                              % Reset dash pattern to a solid line
  1 w                                  % Reset line width to 1 unit

  % Draw a rectangle with a 1-unit red border, filled with light blue.
  1.0 0.0 0.0 RG                      % Red for stroke color
  0.5 0.75 1.0 rg                     % Light blue for fill color
  200 300 50 75 re
  B

  % Draw a curve filled with gray and with a colored border.
  0.5 0.1 0.2 RG
  0.7 g
  300 300 m
  300 400 400 400 400 300 c
  b
endstream
endobj

6 0 obj
  [/PDF]
endobj
```

```

xref
0 7
0000000000 65535 f
0000000009 00000 n
0000000074 00000 n
0000000120 00000 n
0000000179 00000 n
0000000300 00000 n
0000001532 00000 n

trailer
<< /Size 7
    /Root 1 0 R
>>
startxref
1556
%%EOF

```

G.4 Page Tree Example

Example G.4 is a fragment of a PDF file illustrating the structure of the page tree for a large document. It contains the page tree nodes for a 62-page document; Figure G.2 shows the structure of this page tree. Numbers in the figure are object numbers corresponding to the objects in the example.




FIGURE G.2 Page tree for Example G.4

Example G.4

```
337 0 obj
  << /Type /Pages
    /Kids [ 335 0 R
            336 0 R
            ]
    /Count 62
  >>
endobj

335 0 obj
  << /Type /Pages
    /Parent 337 0 R
    /Kids [ 4 0 R
            43 0 R
            77 0 R
            108 0 R
            139 0 R
            170 0 R
            ]
    /Count 36
  >>
endobj

336 0 obj
  << /Type /Pages
    /Parent 337 0 R
    /Kids [ 201 0 R
            232 0 R
            263 0 R
            294 0 R
            325 0 R
            ]
    /Count 26
  >>
endobj
```

```
4 0 obj
  << /Type /Pages
    /Parent 335 0 R
    /Kids [ 3 0 R
            16 0 R
            21 0 R
            26 0 R
            31 0 R
            37 0 R
          ]
    /Count 6
  >>
endobj

43 0 obj
  << /Type /Pages
    /Parent 335 0 R
    /Kids [ 42 0 R
            48 0 R
            53 0 R
            58 0 R
            63 0 R
            70 0 R
          ]
    /Count 6
  >>
endobj

77 0 obj
  << /Type /Pages
    /Parent 335 0 R
    /Kids [ 76 0 R
            82 0 R
            87 0 R
            92 0 R
            97 0 R
            102 0 R
          ]
    /Count 6
  >>
endobj
```

```
108 0 obj
  << /Type /Pages
    /Parent 335 0 R
    /Kids [ 107 0 R
            113 0 R
            118 0 R
            123 0 R
            128 0 R
            133 0 R
          ]
    /Count 6
  >>
endobj

139 0 obj
  << /Type /Pages
    /Parent 335 0 R
    /Kids [ 138 0 R
            144 0 R
            149 0 R
            154 0 R
            159 0 R
            164 0 R
          ]
    /Count 6
  >>
endobj

170 0 obj
  << /Type /Pages
    /Parent 335 0 R
    /Kids [ 169 0 R
            175 0 R
            180 0 R
            185 0 R
            190 0 R
            195 0 R
          ]
    /Count 6
  >>
endobj
```



```
201 0 obj
  << /Type /Pages
    /Parent 336 0 R
    /Kids [ 200 0 R
            206 0 R
            211 0 R
            216 0 R
            221 0 R
            226 0 R
          ]
    /Count 6
  >>
endobj

232 0 obj
  << /Type /Pages
    /Parent 336 0 R
    /Kids [ 231 0 R
            237 0 R
            242 0 R
            247 0 R
            252 0 R
            257 0 R
          ]
    /Count 6
  >>
endobj

263 0 obj
  << /Type /Pages
    /Parent 336 0 R
    /Kids [ 262 0 R
            268 0 R
            273 0 R
            278 0 R
            283 0 R
            288 0 R
          ]
    /Count 6
  >>
endobj
```

```

294 0 obj
  << /Type /Pages
    /Parent 336 0 R
    /Kids [ 293 0 R
            299 0 R
            304 0 R
            309 0 R
            314 0 R
            319 0 R
          ]
    /Count 6
  >>
endobj

325 0 obj
  << /Type /Pages
    /Parent 336 0 R
    /Kids [ 324 0 R
            330 0 R
          ]
    /Count 2
  >>
endobj

```

G.5 Outline Hierarchy Example

This section from a PDF file illustrates the structure of an outline hierarchy with six items. Example G.5 shows the outline with all items open, as illustrated in Figure G.3.

On-screen appearance	Object number	Count
Document	21	6
Section 1	22	4
Section 2	25	0
Subsection 1	26	1
Section 3	27	0
Section 3	28	0
Summary	29	0

FIGURE G.3 Document outline as displayed in Example G.5

Example G.5

```
21 0 obj
  << /Type /Outlines
    /First 22 0 R
    /Last 29 0 R
    /Count 6
  >>
endobj

22 0 obj
  << /Title (Document)
    /Parent 21 0 R
    /Next 29 0 R
    /First 25 0 R
    /Last 28 0 R
    /Count 4
    /Dest [3 0 R /XYZ 0 792 0]
  >>
endobj

25 0 obj
  << /Title (Section 1)
    /Parent 22 0 R
    /Next 26 0 R
    /Dest [3 0 R /XYZ null 701 null]
  >>
endobj

26 0 obj
  << /Title (Section 2)
    /Parent 22 0 R
    /Prev 25 0 R
    /Next 28 0 R
    /First 27 0 R
    /Last 27 0 R
    /Count 1
    /Dest [3 0 R /XYZ null 680 null]
  >>
endobj
```

```

27 0 obj
  << /Title (Subsection 1)
    /Parent 26 0 R
    /Dest [3 0 R /XYZ null 670 null]
  >>
endobj

28 0 obj
  << /Title (Section 3)
    /Parent 22 0 R
    /Prev 26 0 R
    /Dest [7 0 R /XYZ null 500 null]
  >>
endobj

29 0 obj
  << /Title (Summary)
    /Parent 21 0 R
    /Prev 22 0 R
    /Dest [8 0 R /XYZ null 199 null]
  >>
endobj

```

Example G.6 is the same as Example G.5, except that one of the outline items has been closed in the display. The outline appears as shown in Figure G.4.

On-screen appearance	Object number	Count
Document	21	5
Section 1	22	3
Section 2	25	0
Section 3	26	-1
Section 3	28	0
Summary	29	0

FIGURE G.4 Document outline as displayed in Example G.6

Example G.6

```
21 0 obj
  << /Type /Outlines
    /First 22 0 R
    /Last 29 0 R
    /Count 5
  >>
endobj

22 0 obj
  << /Title (Document)
    /Parent 21 0 R
    /Next 29 0 R
    /First 25 0 R
    /Last 28 0 R
    /Count 3
    /Dest [3 0 R /XYZ 0 792 0]
  >>
endobj

25 0 obj
  << /Title (Section 1)
    /Parent 22 0 R
    /Next 26 0 R
    /Dest [3 0 R /XYZ null 701 null]
  >>
endobj

26 0 obj
  << /Title (Section 2)
    /Parent 22 0 R
    /Prev 25 0 R
    /Next 28 0 R
    /First 27 0 R
    /Last 27 0 R
    /Count -1
    /Dest [3 0 R /XYZ null 680 null]
  >>
endobj
```

```
27 0 obj
  << /Title (Subsection 1)
    /Parent 26 0 R
    /Dest [3 0 R /XYZ null 670 null]
  >>
endobj

28 0 obj
  << /Title (Section 3)
    /Parent 22 0 R
    /Prev 26 0 R
    /Dest [7 0 R /XYZ null 500 null]
  >>
endobj

29 0 obj
  << /Title (Summary)
    /Parent 21 0 R
    /Prev 22 0 R
    /Dest [8 0 R /XYZ null 199 null]
  >>
endobj
```

G.6 Updating Example

This example shows the structure of a PDF file as it is updated several times; it illustrates multiple body sections, cross-reference sections, and trailers. In addition, it shows that once an object has been assigned an object identifier, it keeps that identifier until the object is deleted, even if the object is altered. Finally, the example illustrates the reuse of cross-reference entries for objects that have been deleted, along with the incrementing of the generation number after an object has been deleted.

The original file is that shown in Example G.1 on page 758. The updates are divided into four stages, with the file saved after each:

1. Four text annotations are added.
2. The text of one of the annotations is altered.
3. Two of the text annotations are deleted.
4. Three text annotations are added.

The sections following show the segments added to the file at each stage. Throughout this example, objects are referred to by their object identifiers, which are made up of the object number and the generation number, rather than simply by their object numbers as in earlier examples. This is necessary because the example reuses object numbers, so the objects they denote are not unique.

Note: The tables in these sections show only those objects that are modified during the updating process. Objects from Example G.1 that are not altered during the update are not shown.

G.6.1 Stage 1: Add Four Text Annotations

Four text annotations are added to the initial file and the file is saved. Table G.4 lists the objects involved in this update.

OBJECT IDENTIFIER	OBJECT TYPE
4 0	Page (page object)
7 0	Annotation array
8 0	Annot (annotation dictionary)
9 0	Annot (annotation dictionary)
10 0	Annot (annotation dictionary)
11 0	Annot (annotation dictionary)

Example G.7 shows the lines added to the file by this update. The page object is updated because an **Annots** entry has been added to it. Note that the file's trailer now contains a **Prev** entry, which points to the original cross-reference section in the file, while the **startxref** value at the end of the trailer points to the cross-reference section added by the update.

Example G.7

```
4 0 obj
  << /Type /Page
    /Parent 3 0 R
    /MediaBox [0 0 612 792]
    /Contents 5 0 R
    /Resources << /ProcSet 6 0 R >>
    /Annots 7 0 R
  >>
endobj

7 0 obj
  [ 8 0 R
    9 0 R
    10 0 R
    11 0 R
  ]
endobj

8 0 obj
  << /Type /Annot
    /Subtype /Text
    /Rect [44 616 162 735]
    /Contents (Text #1)
    /Open true
  >>
endobj

9 0 obj
  << /Type /Annot
    /Subtype /Text
    /Rect [224 668 457 735]
    /Contents (Text #2)
    /Open false
  >>
endobj

10 0 obj
  << /Type /Annot
    /Subtype /Text
    /Rect [239 393 328 622]
    /Contents (Text #3)
    /Open true
  >>
endobj
```



```
11 0 obj
  << /Type /Annot
    /Subtype /Text
    /Rect [34 398 225 575]
    /Contents (Text #4)
    /Open false
  >>
endobj

xref
0 1
0000000000 65535 f
4 1
0000000632 00000 n
7 5
0000000810 00000 n
0000000883 00000 n
0000001024 00000 n
0000001167 00000 n
0000001309 00000 n

trailer
  << /Size 12
    /Root 1 0 R
    /Prev 408
  >>
startxref
1452
%%EOF
```

G.6.2 Stage 2: Modify Text of One Annotation

One text annotation is modified and the file is saved. Example G.8 shows the lines added to the file by this update. Note that the file now contains two copies of the object with identifier 10 0 (the text annotation that was modified) and that the added cross-reference section points to the more recent version of the object. This added cross-reference section contains one subsection, containing only an entry for the object that was modified. In addition, the **Prev** entry in the file's trailer has been updated to point to the cross-reference section added in the previous stage, while the **startxref** value at the end of the trailer points to the newly added cross-reference section.

Example G.8

```

10 0 obj
  << /Type /Annot
    /Subtype /Text
    /Rect [239 393 328 622]
    /Contents (Modified Text #3)
    /Open true
  >>
endobj

xref
0 1
0000000000 65535 f
10 1
000001703 00000 n

trailer
<< /Size 12
    /Root 1 0 R
    /Prev 1452
  >>
startxref
1855
%%EOF

```

G.6.3 Stage 3: Delete Two Annotations

Two text annotations are deleted and the file is saved. Table G.5 lists the objects updated.

TABLE G.5 Object usage after deleting two text annotations

OBJECT IDENTIFIER	OBJECT TYPE
7 0	Annotation array
8 0	Free
9 0	Free

The **Annots** array is the only object that is written in this update. It is updated because it now contains two annotations fewer.

Example G.9 shows the lines added when the file was saved. Note that objects with identifiers 8 0 and 9 0 have been deleted, as can be seen from the fact that their entries in the cross-reference section end with the keyword **f**.

Example G.9

```
7 0 obj
  [ 10 0 R
    11 0 R
  ]
endobj

xref
0 1
0000000008 65535 f
7 3
0000001978 00000 n
0000000009 00001 f
0000000000 00001 f

trailer
  << /Size 12
      /Root 1 0 R
      /Prev 1855
  >>
startxref
2027
%%EOF
```

The cross-reference section added at this stage contains four entries, representing object number 0, the **Annots** array, and the two deleted text annotations.

- The cross-reference entry for object number 0 is updated because it is the head of the linked list of free entries and must now point to the entry for the newly freed object number 8. The entry for object number 8 points to the entry for object number 9 (the next free entry), while the entry for object number 9 is the last free entry in the cross-reference table, indicated by the fact that it points back to object number 0.
- The entries for the two deleted text annotations are marked as free and as having generation numbers of 1, which will be used for any objects that reuse these cross-reference entries. Keep in mind that, although the two objects have been deleted, they are still present in the file. It is the cross-reference table that records the fact that they have been deleted.

The **Prev** entry in the trailer has again been updated, so that it points to the cross-reference section added at the previous stage, and the **startxref** value points to the newly added cross-reference section.

G.6.4 Stage 4: Add Three Annotations

Finally, three new text annotations are added to the file. Table G.6 lists the objects involved in this update.

TABLE G.6 Object usage after adding three text annotations

OBJECT IDENTIFIER	OBJECT TYPE
7 0	Annotation array
8 1	Annot (annotation dictionary)
9 1	Annot (annotation dictionary)
12 0	Annot (annotation dictionary)

Object numbers 8 and 9, which were used for the two annotations deleted in the previous stage, have been reused; however, the new objects have been given a generation number of 1. In addition, the third text annotation added has been assigned the previously unused object identifier of 12 0.

Example G.10 shows the lines added to the file by this update. The added cross-reference section contains five entries, corresponding to object number 0, the **Annots** array, and the three annotations added. The entry for object number 0 is updated because the previously free entries for object numbers 8 and 9 have been reused. The entry for object number 0 now shows that there are no free entries in the cross-reference table. The **Annots** array is updated to reflect the addition of the three text annotations.

Example G.10

```

7 0 obj
  [ 10 0 R
    11 0 R
    8 1 R
    9 1 R
    12 0 R
  ]
endobj

```

```
8 1 obj
  << /Type /Annot
    /Subtype /Text
    /Rect [58 657 172 742]
    /Contents (New Text #1)
    /Open true
  >>
endobj

9 1 obj
  << /Type /Annot
    /Subtype /Text
    /Rect [389 459 570 537]
    /Contents (New Text #2)
    /Open false
  >>
endobj

12 0 obj
  << /Type /Annot
    /Subtype /Text
    /Rect [44 253 473 337]
    /Contents (New Text #3\203a longer text annotation which we will continue \
onto a second line)
    /Open true
  >>
endobj

xref
0 1
0000000000 65535 f
7 3
000002216 00000 n
000002302 00001 n
000002447 00001 n
12 1
000002594 00000 n

trailer
  << /Size 13
    /Root 1 0 R
    /Prev 2027
  >>
startxref
2814
%%EOF
```

The annotation with object identifier 12 0 illustrates splitting a long text string across multiple lines, as well as the technique for including nonstandard characters in a string. In this case, the character is an ellipsis (...), which is character code 203 (octal) in **PDFDocEncoding**, the encoding used for text annotations.

As in previous updates, the trailer's **Prev** entry and **startxref** value have been updated.

APPENDIX H

Compatibility and Implementation Notes

THE GOAL OF the Adobe Acrobat family of products is to enable people to exchange and view electronic documents easily and reliably. Ideally, this means that any Acrobat viewer application should be able to display the contents of any PDF file, even if the PDF file was created long before or long after the viewer application itself. Of course, new versions of viewer applications are introduced to provide additional capabilities not present before. Furthermore, beginning with Acrobat 2.0, viewer applications may accept plug-in extensions, making some Acrobat 2.0 (and later) viewers more capable than others, depending on what extensions are present.

Both the viewer applications and PDF itself have been designed to enable users to view everything in the document that the viewer application understands and to ignore or inform the user about objects not understood. The decision whether to ignore or inform the user is made on a feature-by-feature basis.

The original PDF specification did not define how a viewer application should behave when it reads a file that does not conform to the specification. This appendix provides that information. The PDF version associated with a file determines how it should be treated when a viewer application encounters a problem.

In addition, this appendix includes notes on the Adobe Acrobat implementation for details that are not strictly defined by the PDF specifications.

H.1 PDF Version Numbers

PDF version numbers take the form $M.m$, where M is the major and m the minor version number. Adobe increments the major version number when the PDF specification changes in such a way that existing viewer applications will be un-

likely to read a document without serious errors that prevent pages from being viewed. The minor version number is incremented if the changes do not prevent existing viewer applications from continuing to work, such as the addition of new page description operators. The version number does not change at all if PDF changes in a way that existing viewer applications are unlikely to detect. Such changes might include the addition of private data, such as additional entries in the document catalog, that can be gracefully ignored by applications that do not understand it.

The header in the first line of a PDF file specifies a PDF version (see Section 3.4.1, “File Header”). In PDF 1.4, a PDF version can also be specified in the **Version** entry of the document catalog, essentially updating the version associated with the file by overriding the one specified in the file header (see Section 3.6.1, “Document Catalog”). As described in the following paragraphs, the viewer application’s behavior upon opening or saving a document depends on what it perceives to be the document’s PDF version (compared to the viewer’s native file format—for example, PDF 1.3 for Acrobat 4.0—which is also referred to as the viewer’s PDF version). Viewers that are not PDF 1.4-aware may perceive the document’s version incorrectly, because they will look for it only in the PDF file’s header and will not see the version (if any) specified in the document catalog.

An Acrobat viewer will attempt to read any PDF file, even if the file’s version is more recent than that of the viewer itself. It will read without errors any file that does not require a plug-in extension, even if the file’s version is older than the viewer’s. Some documents may require a plug-in to display an annotation, follow a link, or execute an action. Viewer behavior in this situation is described below in Section H.3, “Implementation Notes.” However, a plug-in is never required in order to display the contents of a page.

If a viewer application opens a document with a major version number newer than it expects, it warns the user that it is unlikely to be able to read the document successfully and that the user will not be able to change or save the document. At the first error related to document processing, the viewer notifies the user that an error has occurred but that no further errors will be reported. (Some errors are always reported, including file I/O errors, extension loading errors, out-of-memory errors, and notifications that a command has failed.) Processing continues if possible. Acrobat does not permit a document with a newer than expected major version number to be inserted into another document.

If a viewer application opens a document with a minor version number newer than it expects, it notifies the user that the document may contain information the viewer does not understand. (This describes Acrobat 5.0's behavior; versions prior to that do not so notify the user.) If the viewer encounters an error, it notifies the user that the document version is newer than expected, that an error has occurred, and that no further errors will be reported. Acrobat permits a document with a newer minor version to be inserted into another document.

Whether and how the version of a document changes when the document is modified and saved depends on several factors. If the document has a newer version than expected, the viewer will not alter the version—that is, a document's version will never be changed to an older version. If the document has an older version than expected, the viewer will update the document's version to match the viewer's version. If a user modifies a document by inserting another document into it, the saved document's version will be whichever is the most recent of the viewer's version, the document's original version, and the inserted document's version.

If the version of a document changes, viewers that are not PDF 1.4-aware will not be able to save the document using an incremental update, because updating the header requires rewriting the entire file. Among other disadvantages, this can cause existing digital signatures to become invalid. Since viewers that are PDF 1.4-aware can use the **Version** entry in the document catalog to update the document's version, they can incrementally save the document (and will do so if necessary to preserve existing signatures). For example, if an Acrobat 5.0 user modifies a document having a PDF version earlier than 1.4, the document may be updated incrementally when saved (with the updated version of 1.4 in the document catalog); however, if an Acrobat 4.0 user modifies a document having a PDF version earlier than 1.3, the entire file will be rewritten when saved (with a new header indicating version 1.3).

Again, the discussion of viewer behavior above applies to what the viewer perceives to be a document's PDF version, which may be different from the document's actual version if the viewer does not look for the **Version** entry in the document's catalog (a PDF 1.4 feature). One consequence of this is that a file may be rewritten when it could have been incrementally updated. For example, suppose an Acrobat 4.0 user opens a document having a version of 1.4 (newer than expected), specified in the catalog's **Version** entry. Acrobat 4.0 will deter-

mine the version by looking only at the document's header. There are two cases to consider:

- The header specifies version 1.2 or earlier. If the user alters and saves the document, the viewer will update the document's version to match its own, by rewriting the file with a new header indicating version 1.3.
- The header specifies version 1.3 or later. If the user alters and saves the document, the viewer will allow the file to be incrementally updated, since it does not believe the version needs updating.

In both cases, the version number in the document catalog will be maintained at 1.4, so later versions of Acrobat will recognize the correct version number.

H.2 Feature Compatibility

Many PDF features are introduced simply by adding new entries to existing dictionaries. Earlier versions of viewer applications will not notice the existence of such entries and will behave as if they were not there. Such new features are therefore both forward- and backward-compatible. Likewise, adding entries not described in the PDF specification to dictionary objects does not affect the viewers' behavior. (See Appendix E for information on how to choose key names that are compatible with future versions of PDF.)

In some cases, a new feature is impossible to ignore, because doing so would preclude some vital operation such as viewing or printing a page. For instance, if a page's content stream is encoded with some new type of filter, then there is no way to view or print the page, even though the content stream (if decoded) would be perfectly understood by the viewer. There is little choice but to give an error in cases like these. Such new features are forward-compatible, but not backward-compatible.

There are a few cases in which new features are defined in a way that earlier viewer versions will ignore, but the output will be degraded in some way without any error indication. The most significant example of this is transparency. All of the transparency features introduced in PDF 1.4 are defined as new entries in existing dictionaries (including the graphics state parameter dictionary). A viewer that does not understand transparency will treat transparency group XObjects as if they were opaque form XObjects. This is a significant enough deviation from the

intended behavior that it is worth pointing out as a compatibility issue (and so is covered in implementation notes in this appendix).

If a PDF document undergoes editing by an application that does not understand some of the features that the document uses, the occurrences of those features may or may not survive. If a dictionary object such as an annotation is copied into another document during a page insertion (or, beginning with Acrobat 2.0, during a page extraction), all entries are copied. If a value is an indirect reference to another object, that object may be copied as well, depending on the entry.

H.3 Implementation Notes

This section gives notes on the implementation of Adobe Acrobat and on compatibility between different versions of PDF. The notes are listed in the order of the sections to which they refer in the main text.

1.2, "Introduction to PDF 1.4 Features"

1. The native file formats of Adobe Acrobat products are PDF 1.2 for Acrobat 3.0, PDF 1.3 for Acrobat 4.0, and PDF 1.4 for Acrobat 5.0.

3.1.2, "Comments"

2. Acrobat viewers do not preserve comments when saving a file.

3.2.4, "Name Objects"

3. In PDF 1.1, the number sign character (#) could be used as part of a name (for example, /A#B), and the specifications did not specifically prohibit embedded spaces (although Adobe producer applications did not provide a way to write names containing them). In PDF 1.2, the number sign became an escape character, preceding two hexadecimal digits. Thus a 3-character name A-space-B can now be written as /A#20B (since 20 is the hexadecimal code for the space character). This means that the name /A#B is no longer valid, since the number sign is not followed by two hexadecimal digits. A name object with this value must be written as /A#23B, since 23 is the hexadecimal code for the character #.
4. In cases where a PostScript name must be preserved, or where a string is permitted in PostScript but not in PDF, the Acrobat Distiller application

uses the # convention as necessary. When an Acrobat viewer generates PostScript, it “inverts” the convention by writing a string, where that is permitted, or a name otherwise. For example, if the string (Adobe Green) were used as a key in a dictionary, the Distiller program would use the name /Adobe#20Green and the viewer would generate (Adobe Green).

5. In Acrobat 4.0 and earlier versions, a name object being treated as text will typically be interpreted in a host platform encoding, which depends on the operating system and the local language. For Asian languages, this encoding may be something like Shift-JIS or Big Five. Consequently, it will be necessary to distinguish between names encoded this way and ones encoded as UTF-8. Fortunately, UTF-8 encoding is very stylized and its use can usually be recognized. A name that is found not to conform to UTF-8 encoding rules can instead be interpreted according to host platform encoding.

3.2.7, “Stream Objects”

6. When a stream specifies an external file, PDF 1.1 parsers ignore the file and always use the bytes between **stream** and **endstream**.
7. Acrobat viewers accept the name **DP** as an abbreviation for the **DecodeParms** key in any stream dictionary. If both **DP** and **DecodeParms** entries are present, **DecodeParms** takes precedence.

3.3, “Filters”

8. Acrobat viewers accept the abbreviated filter names shown in Table H.1 in addition to the standard ones. Although the abbreviated names are intended for use only in the context of inline images (see Section 4.8.6, “Inline Images”), they are also accepted as filter names in any stream object.

TABLE H.1 Abbreviations for standard filter names

STANDARD FILTER NAME	ABBREVIATION
ASCIHexDecode	AHx
ASCII85Decode	A85
LZWDecode	LZW
FlateDecode (<i>PDF 1.2</i>)	Fl (uppercase F, lowercase L)
RunLengthDecode	RL
CCITTFaxDecode	CCF
DCTDecode	DCT

9. If an unrecognized filter is encountered, Acrobat viewers report the context in which the filter was found. If errors occur while a page is being displayed, only the first error is reported. The subsequent behavior depends on the context, as described in Table H.2. Acrobat operations that process pages, such as the Find command and the Create Thumbnails command, stop as soon as an error occurs.

TABLE H.2 Acrobat behavior with unknown filters

CONTEXT	BEHAVIOR
Content stream	Page processing stops.
Indexed color space	The image does not appear, but page processing continues.
Image resource	The image does not appear, but page processing continues.
Inline image	Page processing stops.
Thumbnail image	An error is reported and no more thumbnail images are displayed, but the thumbnails can be deleted and created again.
Form XObject	The form does not appear, but page processing continues.
Type 3 glyph description	The glyph does not appear, but page processing continues. The text position is adjusted based on the glyph width.
Embedded font	The viewer behaves as if the font is not embedded.

3.3.7, “DCTDecode Filter”

10. Acrobat 4.0 and later viewers do not support the combination of the **DCTDecode** filter with any other filter if the encoded data uses the progressive JPEG format. If a version of the Acrobat viewer earlier than 4.0 encounters **DCTDecode** data encoded in progressive JPEG format, an error occurs that will be handled according to Table H.2.

3.4, “File Structure”

11. The restriction on line length is not enforced by any Acrobat viewer.

3.4.1, “File Header”

12. Acrobat viewers require only that the header appear somewhere within the first 1024 bytes of the file.
13. Acrobat viewers will also accept a header of the form

```
%!PS-Adobe-N.n PDF-M.m
```

3.4.4, “File Trailer”

14. Acrobat viewers require only that the %%EOF marker appear somewhere within the last 1024 bytes of the file.

3.5, “Encryption”

15. An Acrobat viewer will fail to open a document encrypted with a **V** value defined in a version of PDF that the viewer does not understand.

3.5.2, “Standard Security Handler” (Standard Encryption Dictionary)

16. Acrobat viewers implement this limited mode of printing as “Print As Image,” except on UNIX systems, where this feature is not available.

3.5.2, “Standard Security Handler” (Password Algorithms)

17. In Acrobat 2.0 and 2.1 viewers, the standard security handler uses the empty string if there is no owner password in step 1 of Algorithm 3.3.

3.6.1, “Document Catalog”

18. Acrobat 5.0 will avoid adding a **Version** entry to the document catalog; it will do so only if it must. Once it has done so, however, it will never remove the **Version** entry. For documents containing a **Version** entry, Acrobat 5.0 will attempt to ensure that the version specified in the header matches the version specified in the **Version** entry; if this is not possible, it will at least ensure that the latter is later than (and therefore overrides) the version specified in the header.
19. An earlier version of this specification documented the **PageLayout** entry as being in the viewer preferences dictionary (see Section 8.1, “Viewer Preferences”); it is actually implemented in the document catalog instead.
20. In PDF 1.2, an additional entry in the document catalog, named **AA**, was defined but was never implemented. The **AA** entry that is newly introduced in PDF 1.4 is entirely different from the one that was contemplated for PDF 1.2.

3.6.2, “Page Tree” (Page Objects)

21. In PDF 1.2, an additional entry in the page object, named **Hid**, was defined but was never implemented. Beginning with PDF 1.3, this entry is obsolete and should be ignored.
22. Acrobat 5.0 does not accept a **Contents** array containing no elements.
23. In a document containing articles, if the first page with an article bead does not have a **B** entry, Acrobat viewers rebuild the **B** array for all pages of the document.
24. In PDF 1.2, additional-actions dictionaries were inheritable; beginning with PDF 1.3, they no longer are.

3.7.1, “Content Streams”

25. Acrobat viewers report an error the first time they find an unknown operator or an operator with too few operands, but continue processing the content stream. No further errors are reported.

3.9.1, “Type 0 (Sampled) Functions”

26. When printing, Acrobat performs only linear interpolation, regardless of the value of the **Order** entry.

3.9.2, “Type 2 (Exponential Interpolation) Functions”

27. Since Type 2 functions are not defined in PDF 1.2 or earlier versions, Acrobat 3.0 (whose native file format is PDF 1.2) will report an error, “Invalid Function Resource,” if it encounters a function of this type.

3.9.3, “Type 3 (Stitching) Functions”

28. Since Type 3 functions are not defined in PDF 1.2 or earlier versions, Acrobat 3.0 (whose native file format is PDF 1.2) will report an error, “Invalid Function Resource,” if it encounters a function of this type.

3.9.4, “Type 4 (PostScript Calculator) Functions”

29. Since Type 4 functions are not defined in PDF 1.2 or earlier versions, Acrobat 3.0 (whose native file format is PDF 1.2) will report an error, “Invalid Function Resource,” if it encounters a function of this type.
30. Acrobat uses single-precision floating-point numbers for all real-number operations in a type 4 function.

4.5.2, “Color Space Families”

31. If an Acrobat viewer encounters an unknown color space family name, it displays an error specifying the name, but reports no further errors thereafter.

4.5.5, “Special Color Spaces” (Multitone Examples)

32. This method of representing multitones is used by Adobe Photoshop® 5.0.2 and subsequent versions when exporting EPS files. Beginning with version 4.0, Acrobat exports Level 3 EPS files using this method, and can also export Level 1 EPS files that use the “Level 1 separation” conventions of Adobe Technical Note #5044, *Color Separation Conventions for PostScript Language Programs*. These conventions are used to emit multitone images as calls to “customcolorimage” with overprinting, which can then

be placed in page layout applications such as Adobe PageMaker[®], Adobe InDesign[™], and QuarkXPress[®].

4.6, "Patterns"

33. Acrobat viewers prior to version 4.0 do not display patterns on the screen, although they do print them to PostScript output devices.

4.7, "External Objects"

34. If an Acrobat viewer encounters an XObject of an unknown type, it displays an error specifying the type of XObject, but reports no further errors thereafter.

4.8.4, "Image Dictionaries"

35. Image XObjects in PDF 1.2 and earlier versions are all implicitly unmasked images. A PDF consumer that does not recognize the **Mask** entry will treat the image as unmasked without raising an error.
36. All Acrobat viewers ignore the **Name** entry in an image dictionary.

4.8.5, "Masked Images"

37. Explicit masking and color key masking are features of PostScript LanguageLevel 3. Acrobat 4.0 and later versions do not attempt to emulate the effect of masked images when printing to LanguageLevel 1 or LanguageLevel 2 output devices; they print the base image without the mask.

The Acrobat 4.0 viewer will display masked images, but only when the amount of data in the mask is below a certain limit. Above that, the viewer will display the base image without the mask.

4.9.1, "Form Dictionaries"

38. All Acrobat viewers ignore the **Name** entry in a form dictionary.

5.2.5, "Text Rendering Mode"

39. In Acrobat 4.05 and earlier versions, text-showing operators such as **Tj** first perform the fills for all the glyphs in the string being shown, followed

by the strokes for all the glyphs. This produces incorrect results if glyphs overlap.

5.3.2, “Text-Showing Operators”

40. In versions of Acrobat prior to 3.0, the horizontal coordinate of the text position after the **TJ** operator paints a character glyph and moves by any specified offset must not be less than it was before the glyph was painted.
41. In Acrobat 4.0 and earlier viewers, position adjustments specified by numbers in a **TJ** array are performed incorrectly if the horizontal scaling parameter, T_h , is different from its default value of 100.

5.5.1, “Type 1 Fonts”

42. All Acrobat viewers ignore the **Name** entry in a font dictionary.
43. Acrobat 5.0 uses the glyph widths stored in the font dictionary to override the widths of glyphs in the font program itself, which improves the consistency of the display and printing of the document. This addresses the situation in which the font program used by the viewer application is different from the one used by the application that produced the document.

The font program with the altered glyph widths might be embedded or not. If it is embedded, its widths should exactly match the widths in the font dictionary. If the font program is not embedded, Acrobat will override the widths in the font program on the viewer application’s system with the widths specified in the font dictionary.

It is important that the widths in the font dictionary match the actual glyph widths of the font program that was used to produce the document. Consumers of PDF files depend on these widths in many different contexts, including viewing, printing, “fauxing” (font substitution), reflow, and word search. These operations may malfunction if arbitrary adjustments are made to the widths so that they do not represent the glyph widths intended by the PDF producer.

It is recommended that diagnostic and preflight tools check the glyph widths in the font dictionary against those in an embedded font program and flag any inconsistencies. It would also be helpful if the tools could optionally check for consistency with the widths in font programs that are not embedded; this is useful for checking a PDF file immediately after it is produced, when the original font programs are still available.

Note: This implementation note is also referred to in Section 5.6.3, “CIDFonts” (Glyph Metrics in CIDFonts).

5.5.1, “Type 1 Fonts” (Standard Type 1 Fonts)

44. Acrobat 3.0 and earlier viewers may ignore attempts to override the standard fonts. Also, Acrobat 4.0 and earlier viewers incorrectly allow substitution fonts, such as TimesNewRoman and ArialMT, to be specified without **FirstChar**, **LastChar**, **Widths**, and **FontDescriptor** entries.

Table H.3 shows the complete list of font names that are accepted as the names of standard fonts. In each group, the first name (for example, Helvetica) is the proper one; the others (Arial, ArialMT) are alternatives.

TABLE H.3 Names of standard fonts

Times–Roman	Helvetica	Courier
TimesNewRoman	Arial	CourierNew
TimesNewRomanPS	ArialMT	CourierNewPSMT
TimesNewRomanPSMT		
Times–Bold	Helvetica–Bold	Courier–Bold
TimesNewRoman–Bold	Helvetica,Bold	Courier,Bold
TimesNewRoman,Bold	Arial–Bold	CourierNew–Bold
TimesNewRomanPS–Bold	Arial,Bold	CourierNew,Bold
TimesNewRomanPS–BoldMT	Arial–BoldMT	CourierNewPS–BoldMT
Times–Italic	Helvetica–Oblique	Courier–Oblique
TimesNewRoman–Italic	Helvetica–Italic	Courier,Italic
TimesNewRoman,Italic	Helvetica,Italic	CourierNew–Italic
TimesNewRomanPS–Italic	Arial–Italic	CourierNew,Italic
TimesNewRomanPS–ItalicMT	Arial,Italic	CourierNewPS–ItalicMT
	Arial–ItalicMT	
Times–BoldItalic	Helvetica–BoldOblique	Courier–BoldOblique
TimesNewRoman–BoldItalic	Helvetica–BoldItalic	Courier,BoldItalic
TimesNewRoman,BoldItalic	Helvetica,BoldItalic	CourierNew–BoldItalic
TimesNewRomanPS–BoldItalic	Arial–BoldItalic	CourierNew,BoldItalic
TimesNewRomanPS–BoldItalicMT	Arial,BoldItalic	CourierNewPS–BoldItalicMT
	Arial–BoldItalicMT	
Symbol		ZapfDingbats

5.5.3, “Font Subsets”

45. For Acrobat 3.0 and earlier viewers, all font subsets whose **BaseFont** names differ only in their tags should have the same font descriptor values and should map character names to glyphs in the same way; otherwise, glyphs may be shown unpredictably. This restriction is eliminated in Acrobat 4.0.

5.5.4, “Type 3 Fonts”

46. In principle, the value of the **Encoding** entry could also be the name of a predefined encoding or an encoding dictionary whose **BaseEncoding** entry is a predefined encoding. However, Acrobat 4.0 and earlier viewers do not implement this correctly.
47. For compatibility with Acrobat 2.0 and 2.1, the names of resources in a Type 3 font’s resource dictionary must match those in the page object’s resource dictionary for all pages in which the font is referenced. If backward compatibility is not required, any valid names may be used.

5.6.4, “CMaps”

48. Embedded CMap files, other than **ToUnicode** CMaps, do not work properly in Acrobat 4.0 viewers; this has been corrected in Acrobat 4.05.

5.7, “Font Descriptors”

49. Acrobat viewers prior to version 3.0 ignore the **FontFile3** entry. If a font uses the Adobe standard Latin character set (as defined in Section D.1, “Latin Character Set and Encodings”), Acrobat creates a substitute font. Otherwise, Acrobat displays an error message (once per document) and substitutes any characters in the font with the bullet character.

5.8, “Embedded Font Programs”

50. For simple fonts, font substitution is performed using multiple master Type 1 fonts. This substitution can be performed only for fonts that use the Adobe standard Latin character set (as defined in Section D.1, “Latin Character Set and Encodings”). In Acrobat 3.0.1 and later, Type 0 fonts that use a CMap whose **CIDSystemInfo** dictionary defines the Adobe-GB1, Adobe-CNS1, Adobe-Japan1, or Adobe-Korea1 character collection can

also be substituted. To make a document portable, it is necessary to embed fonts that cannot be substituted. The only exceptions are the Symbol and ZapfDingbats fonts, which are assumed to be present.

6.4.2, “Spot Functions”

51. When the Acrobat Distiller encounters a call to the PostScript **setscreen** or **sethalftone** operator that includes a spot function, it compares the PostScript code defining the spot function with that of the predefined spot functions shown in Table 6.1. If the code matches one of the predefined functions, Distiller puts the name of that function into the halftone dictionary; Acrobat will then use that function when printing the PDF document to a PostScript output device. If the code does not match any of the predefined spot functions, Distiller samples the specified spot function and generates a function for the halftone dictionary; when printing to a PostScript device, Acrobat will generate a spot function that interpolates values from that function.

When producing PDF version 1.3 or later, Distiller represents the spot function using a Type 4 (PostScript calculator) function whenever possible (see Section 3.9.4, “Type 4 (PostScript Calculator) Functions”). In this case, Acrobat will use this function directly when printing the document.

7.5.2, “Specifying Blending Color Space and Blend Mode”

52. PDF 1.3 or earlier viewers will ignore all transparency-related graphics state parameters (blend mode, soft mask, alpha constant, and alpha source). All graphics objects will be painted opaquely.

Note: This implementation note is also referred to in Sections 7.5.3, “Specifying Shape and Opacity” (Mask Shape and Opacity, Constant Shape and Opacity) and 7.5.4, “Specifying Soft Masks” (Soft-Mask Dictionaries).

7.5.3, “Specifying Shape and Opacity” (Mask Shape and Opacity)

53. PDF 1.3 or earlier viewers will ignore the **SMask** entry in an image dictionary. All images will be painted opaquely.

Note: This implementation note is also referred to in Section 7.5.4, “Specifying Soft Masks” (Soft-Mask Images).

8.1, “Viewer Preferences”

54. Earlier versions of the PDF specification erroneously described an additional entry, **PageLayout**, as being in the viewer preferences dictionary; it is actually implemented in the document catalog instead (see Section 3.6.1, “Document Catalog”).

8.2.2, “Document Outline”

55. In PDF 1.2, an additional entry in the outline item dictionary, named **AA**, was defined but was never implemented. Beginning with PDF 1.3, this entry is obsolete and should be ignored.
56. Acrobat viewers report an error when a user activates an outline item whose destination is of an unknown type.

8.3.1, “Page Labels”

57. Acrobat viewers up to version 3.0 ignore the **PageLabels** entry and label pages with decimal numbers starting at 1.

8.4.1, “Annotation Dictionaries”

58. Acrobat 4.05 and earlier versions assume that a given annotation dictionary is referenced from only one page. Attempting to share an annotation dictionary among multiple pages will produce unpredictable behavior.
59. Acrobat viewers update the annotation dictionary’s **M** entry only for text annotations.
60. Acrobat 2.0 and 2.1 viewers ignore the annotation dictionary’s **BS**, **AP**, and **AS** entries.
61. Acrobat viewers ignore the horizontal and vertical corner radii in the annotation dictionary’s **Border** entry; the border is always drawn with square corners.
62. Acrobat viewers support a maximum of ten elements in the dash array of the annotation dictionary’s **Border** entry.

8.4.2, "Annotation Flags"

63. Acrobat viewers prior to version 3.0 ignore an annotation's Hidden and Print flags. Annotations that should be hidden are shown; annotations that should be printed are not printed. Acrobat 3.0 ignores the Print flag for text and link annotations.

8.4.3, "Border Styles"

64. If an Acrobat viewer encounters a border style it does not recognize, the border style defaults to S (Solid).

8.4.5, "Annotation Types"

65. Acrobat viewers display annotations whose types they do not recognize in closed form, with an icon containing a question mark. Such an annotation can be selected, moved, or deleted, but if the user attempts to activate it, an alert appears giving the annotation type and reporting that a required plug-in is unavailable.

8.4.5, "Annotation Types" (Link Annotations)

66. Acrobat viewers report an error when a user activates a link annotation whose destination is of an unknown type.

8.4.5, "Annotation Types" (Markup Annotations)

67. In Acrobat 4.0 and later versions, the text is oriented with respect to the vertex with the smallest *y* value (or the leftmost of those, if there are two such vertices) and the next vertex in a counterclockwise direction, regardless of whether these are the first two points in the **QuadPoints** array.

8.4.5, "Annotation Types" (Ink Annotations)

68. Acrobat viewers always connect the points along each path with straight lines.

8.4.5, "Annotation Types" (Movie Annotations)

69. Acrobat viewers report the following error when they encounter an annotation of type **Movie**: "The Plug-in required by this 'Movie' annotation is

unavailable.” The annotation is displayed as a gray rectangle with a question mark.

8.5.2, “Trigger Events”

70. In PDF 1.2, the additional-actions dictionary could contain entries named **NP** (next page), **PP** (previous page), **FP** (first page), and **LP** (last page). The actions associated with these entries were never implemented; beginning with PDF 1.3, these entries are obsolete and should be ignored.
71. In PDF 1.2, additional-actions dictionaries were inheritable; beginning with PDF 1.3, they no longer are.
72. In Acrobat 3.0, the **O** and **C** events in a page object’s additional-actions dictionary are ignored if the document is not being displayed in a page-oriented layout mode. Beginning with Acrobat 4.0, the actions associated with these events are executed if the document is in a page-oriented or single-column layout; they are ignored if it is in a multiple-column layout.

8.5.3, “Action Types” (Launch Actions)

73. The Acrobat viewer for the Windows platform uses the Windows function ShellExecute to launch an application. The **Win** dictionary entries correspond to the parameters of ShellExecute.

8.5.3, “Action Types” (URI Actions)

74. URI actions are resolved by the Acrobat WebLink plug-in extension.
75. If the appropriate plug-in extension (WebLink) is not present, Acrobat viewers report the following error when a link annotation that uses a URI action is activated: “The plug-in required by this URI action is not available.”

8.5.3, “Action Types” (Sound Actions)

76. In PDF 1.2, the value of the **Sound** entry was allowed to be a file specification. Beginning with PDF 1.3, this is no longer supported, but the same effect can be achieved by using an external stream.

8.5.3, “Action Types” (Movie Actions)

77. Acrobat viewers prior to version 3.0 report an error when they encounter an action of type **Movie**.

8.5.3, “Action Types” (Hide Actions)

78. Acrobat viewers prior to version 3.0 report the following error when encountering an action of type **Hide**: “The plug-in needed for this **Hide** action is not available.”
79. In Acrobat viewers, the change in an annotation’s Hidden flag as a result of a hide action is temporary, in the sense that the user can subsequently close the document without being prompted to save changes, and the effect of the hide action will be lost. However, if the user does explicitly save the document before closing, such changes *will* be saved and will thus become permanent.

8.5.3, “Action Types” (Named Actions)

80. Acrobat viewers prior to version 3.0 report the following error when encountering an action of type **Named**: “The plug-in needed for this **Named** action is not available.”
81. Acrobat viewers extend the list of named actions in Table 8.45 to include most of the menu item names available in the viewer.

8.6.2, “Field Dictionaries” (Field Names)

82. Beginning in Acrobat 3.0, partial field names may not contain a period.
83. Acrobat versions 3.0 and later do not support Unicode encoding of field names.

8.6.2, “Field Dictionaries” (Variable Text)

84. Acrobat viewers may insert additional entries in the **DR** resource dictionary, such as **Encoding**, as a convenience for keeping track of objects being used to construct form fields. Such objects are not actually resources and are not referenced from the appearance stream.

8.6.3, “Field Types” (Choice Fields)

85. In Acrobat 3.0, the **Opt** array must be homogenous: its elements must be either all text strings or all arrays.

8.6.4, “Form Actions” (Submit-Form Actions)

86. In Acrobat viewers, if the response to a submit-form action uses Forms Data Format (FDF), then the URL must end in #FDF so that it will be recognized as such by the Acrobat software and handled properly. Conversely, if the response is in any other format, the URL should not end in #FDF.

8.6.4, “Form Actions” (Import-Data Actions)

87. Acrobat viewers set the **F** entry to a relative file specification locating the FDF file with respect to the current PDF document file. If the designated FDF file is not found when the import-data action is performed, Acrobat tries to locate the file in a few “well-known” locations depending on the host platform. On the Windows platform, for example, it looks in the directory from which Acrobat was loaded, the current directory, the System directory, the Windows directory, and any directories listed in the PATH environment variable; on Mac OS, it looks in the Preferences folder and the Acrobat folder.
88. When performing an import-data action, Acrobat viewers import the contents of the FDF file into the current document’s interactive form, ignoring the **F** and **ID** entries in the FDF dictionary of the FDF file itself.

8.6.4, “Form Actions” (JavaScript Actions)

89. Because JavaScript 1.2 is not Unicode-compatible, **PDFDocEncoding** and the Unicode encoding are translated to a platform-specific encoding prior to interpretation by the JavaScript engine.

8.6.6, “Forms Data Format” (FDF Header)

90. Because a bug in versions of Acrobat prior to 5.0 prevents them from accepting any other version number, the FDF file header is permanently frozen at version 1.2. All further updates to the version number will be made via the **Version** entry in the FDF catalog dictionary instead.

8.6.6, "Forms Data Format" (FDF Catalog)

91. The Acrobat implementation of interactive forms displays the value of the **Status** entry, if any, in an alert note when importing an FDF file.
92. The only **Encoding** value supported by Acrobat 4.0 is Shift-JIS. Acrobat 5 supports Shift-JIS, UHC, GBK, and BigFive. If any other value is specified, the default, **PDFDocEncoding**, will be used.

8.6.6, "Forms Data Format" (FDF Fields)

93. Of all the possible entries shown in Table 8.72 on page 564, Acrobat 3.0 will export only the **V** entry when generating FDF, and Acrobat 4.0 and later versions will export only the **V** and **AP** entries. It will, however, import FDF files containing fields using any of the described entries.
94. If the FDF dictionary in an FDF file received as a result of a submit-form action contains an **F** entry specifying a form other than the one currently being displayed, Acrobat fetches the specified form before importing the FDF file.
95. When exporting a form to an FDF file, Acrobat sets the **F** entry in the FDF dictionary to a relative file specification giving the location of the FDF file relative to that of the file from which it was exported.
96. If an FDF file being imported contains fields whose fully qualified names are not present in the form, Acrobat will discard those fields. This feature can be useful, for example, if an FDF file containing commonly used fields (such as name and address) is used to populate various types of form, not all of which necessarily include all of the fields available in the FDF file.
97. As shown in Table 8.72 on page 564, the only required entry in the field dictionary is **T**. One possible use for exporting FDF with fields containing **T** entries but no **V** entries is to indicate to a server which fields are desired in the FDF files returned in response. For example, a server accessing a database might use this information to decide whether to transmit all fields in a record or just some selected ones. As noted in implementation note 96 above, the Acrobat implementation of interactive forms will ignore fields in the imported FDF file that do not exist in the form.
98. The Acrobat implementation of forms allows the option of submitting the data in a submit-form action in HTML Form format. This is for the benefit of existing server scripts written to process such forms. Note, however,

that any such existing scripts that generate new HTML forms in response will need to be modified to generate FDF instead.

8.6.6, “Forms Data Format” (FDF Pages)

99. Acrobat renames fields by prepending a page number, a template name, and an ordinal number to the field name. The ordinal number corresponds to the order in which the template is applied to a page, with 0 being the first template specified for the page. For example, if the first template used on the fifth page has the name `Template` and has the **Rename** flag set to **true**, fields defined in that template will be renamed by prepending the character string `P5.Template_0.` to their field names.
100. Adobe Extreme™ printing systems require that the **Rename** flag be **true**.

8.7, “Sounds”

101. Acrobat supports a maximum of two sound channels.

9.1, “Procedure Sets”

102. Acrobat viewers prior to version 5.0 respond to requests for unknown procedure sets by warning the user that a required procedure set is unavailable and canceling the printing operation. Acrobat 5.0 ignores procedure sets.

9.2, “Metadata”

103. Acrobat viewers display the document’s metadata in the Document Properties dialog box and impose a limit of 255 bytes on any string representing one of those values.

9.2.2, “Metadata Streams”

104. For backward compatibility, applications that create PDF 1.4 documents should include the metadata for a document in the document information dictionary as well as in the document’s metadata stream. Applications that support PDF 1.4 should check for the existence of a metadata stream and synchronize the information in it with that in the document information dictionary. The Adobe metadata framework provides a date stamp for metadata expressed in the framework. If this date stamp is equal to or

later than the document modification date recorded in the document information dictionary, the metadata stream can be taken as authoritative. If, however, the document modification date recorded in the document information dictionary is later than the metadata stream's date stamp, it is likely that the document has been saved by an application that is not aware of PDF 1.4 metadata streams. In this case, information stored in the document information dictionary should be taken to override any semantically equivalent items in the metadata stream.

9.3, "File Identifiers"

105. Although the **ID** entry is not required, all Adobe applications that produce PDF files include this entry. Acrobat adds this entry when saving a file if it is not already present.
106. Adobe applications pass the suggested information to the MD5 message digest algorithm to calculate file identifiers. Note that the calculation of the file identifier need not be reproducible; all that matters is that the identifier is likely to be unique. For example, two implementations of this algorithm might use different formats for the current time; this will cause them to produce different file identifiers for the same file created at the same time, but does not affect the uniqueness of the identifier.

9.9.2, "Content Database" (Digital Identifiers)

107. The Acrobat Web Capture plug-in treats external streams referenced within a PDF file as auxiliary data. Such streams are not used in generating the digital identifier.

9.9.3, "Content Sets" (Image Sets)

108. In Acrobat 4.0 and later versions, if the indirect reference to an image XObject is not removed from the **O** array when its reference count reaches 0, the XObject will never be garbage-collected during a save operation. The image set's reference to the XObject may thus be considered a weak one that is relevant only for caching purposes; when the last strong reference goes away, so does the weak one.

9.9.4, “Source Information” (URL Alias Dictionaries)

109. Acrobat viewers use an indirect object reference to a shared string for each URL in a URL alias dictionary; these strings can then be shared among the chains and with other data structures. It is recommended that other PDF viewer applications adopt this same implementation.

9.10.1, “Page Boundaries”

110. Acrobat provides various user-specified options for determining how the region specified by the crop box is to be imposed on the output medium during printing. Although these options have varied from one Acrobat version to another, the default behavior is as follows:
 1. Select the media size and orientation according to the operating system’s Print Setup dialog. (Acrobat itself has no direct control over this process.)
 2. Compute an effective crop box by first clipping it with the media box, then rotating the page according to the page object’s **Rotate** entry, if specified.
 3. Center the crop box on the medium, rotating it if necessary to enable it to fit in both dimensions.
 4. Optionally, scale the page up or down so that the crop box coincides with the edges of the medium in the more restrictive dimension.

The description above applies only in simple printing workflows that lack any other information about how PDF pages are to be imposed on the output medium. In some workflows, there will be additional information, either in the PDF file itself (**BleedBox**, **TrimBox**, or **ArtBox**) or in a separate job ticket (such as JDF or PJTF). In these circumstances, other rules will apply, which depend on the details of the workflow.

Consequently, it is recommended that PDF files initially be created with the crop box the same as the media box (or equivalently, with the crop box omitted). This ensures that if the page is printed on that size medium, the crop box will coincide with the edges of the medium, producing predictable and dependable positioning of the page contents. On the other hand, if the page is printed on a different size medium, the page may be repositioned or scaled in implementation-defined or user-specified ways.

9.10.4, “Output Intents”

111. Acrobat viewers do not make use of the “to CIE” (*AToB*) information in an output intent’s ICC profile.
112. Acrobat 5.0 does not make direct use of the destination profile in the output intent dictionary, but third-party plug-in extensions might do so.

9.10.5, “Trapping Support” (Trap Network Annotations)

113. Older viewers may fail to maintain the trap network annotation’s required position at the end of the **Annots** array.
114. Older viewers may fail to validate trap networks before printing.
115. In Acrobat 4.0, saving a PDF file with the Optimize option selected would cause a page’s trap networks to be incorrectly invalidated even if the contents of the page had not been changed. This occurred because the new, optimized content stream generated for the page differed from the original content stream still referenced by the trap network annotation’s **Version** array. This problem has been corrected in later versions of Acrobat.

9.10.6, “Open Prepress Interface (OPI)”

116. The Acrobat 3.0 Distiller application converts OPI comments into OPI dictionaries; when the Acrobat 3.0 viewer prints a PDF file to a PostScript file or printer, it converts the OPI dictionary back to OPI comments. However, the OPI information has no effect on the displayed image or form XObject.
117. Acrobat viewer and Distiller applications prior to version 4.0 do not support OPI 2.0.
118. In Acrobat 3.0, the value of the **F** entry in an OPI dictionary must be a string.

C.1, “General Implementation Limits”

119. Acrobat viewers prior to 5.0 use the PostScript **save** and **restore** operators, rather than **gsave** and **grestore**, to implement **q** and **Q** and are subject to a nesting limit of 12 levels.

120. In Acrobat viewers prior to version 4.0, the minimum allowed page size is 72 by 72 units in default user space (1 inch by 1 inch); the maximum is 3240 by 3240 units (45 by 45 inches).

F.2.2, “Linearization Parameter Dictionary (Part 2)”

121. Acrobat requires a white-space character to follow the left bracket ([) character that begins the **H** array.
122. Acrobat does not currently support reading or writing files that have an overflow hint stream.

Note: This implementation note is also referred to in Section F.2.5, “Hint Streams (Parts 5 and 10).”

123. Acrobat generates a value for the **E** parameter that incorrectly includes an object beyond the end of the first page as if it were part of the first page.

F.2.6, “First-Page Section (Part 6)”

124. Acrobat always treats page 0 as the first page for linearization, regardless of the value of **OpenAction**.

F.2.8, “Shared Objects (Part 8)”

125. Acrobat does not generate shared object groups containing more than one object.

F.3.1, “Page Offset Hint Table”

126. In Acrobat, items 6 and 7 in the header section of the page offset hint table are set to 0. As a result, item 6 of the per-page entry effectively does not exist; its value is taken to be 0. That is, the sequence of bytes constituting the content stream for a page is described as if the content stream were the first object in the page, even though it is not.
127. Acrobat 4.0 and later versions always set item 8 equal to 0. They also set item 9 equal to the value of item 5, and set item 7 of each per-page hint table entry (Table F.4) to be the same as item 2 of the per-page entry. Acrobat ignores all of these entries when reading the file.

F.3.2, “Shared Object Hint Table”

128. In Acrobat, item 5 in the header section of the shared objects hint table is unused and is always set to 0.
129. MD5 signatures are not implemented in Acrobat; item 2 in a shared object group entry must be 0.
130. Acrobat does not support more than one shared object in a group; item 4 in a shared object group entry should always be 0.
131. In a document consisting of only one page, items 1 and 2 in the shared object hint table are not meaningful; Acrobat writes unpredictable values for these items.

Bibliography

THIS BIBLIOGRAPHY PROVIDES details on books and documents, from both Adobe Systems and other sources, that are referred to in this book.

Resources from Adobe Systems Incorporated

All of these resources from Adobe Systems are available on the Adobe Solutions Network (ASN) Developer Program site on the World Wide Web, located at

<<http://partners.adobe.com/asn/developer/>>

Document version numbers and dates given in this Bibliography are the latest at the time of publication; more recent versions may be found on the Web site.

The ASN can also be contacted as follows:

Adobe Solutions Network
Adobe Systems Incorporated
345 Park Avenue
San Jose, CA 95110-2704

(800) 685-3510 (from North America)
(206) 675-6145 (from other areas)

<acrodevsup@adobe.com>

Adobe Glyph List, Version 1.2. Available through the document *Unicode and Glyph Names* in the Type Technology Forum on the ASN Developer Program Web site.

Adobe Patent Clarification Notice. Available on the Technical Notes page of the ASN Developer Program Web site.

Adobe Type 1 Font Format. Explains the internal organization of a PostScript Type 1 font program. Also see Adobe Technical Note #5015, *Type 1 Font Format Supplement*.

OPI: Open Prepress Interface Specification 1.3. Also see Adobe Technical Note #5660, *Open Prepress Interface (OPI) Specification, Version 2.0.*

PDF Public-Key Digital Signature and Encryption Specification. Available in the Adobe Acrobat Software Development Kit (SDK), on either the ASN Developer Program Web site or the Acrobat SDK CD.

PostScript Language Reference, Third Edition, Addison-Wesley, Reading, MA, 1999.

XMP: Extensible Metadata Platform. To be available on the Technical Notes page of the ASN Developer Program Web site. (Not yet available at the time of publication.)

Technical Notes:

- Technical Note #5001, *PostScript Language Document Structuring Conventions Specification, Version 3.0*
- Technical Note #5004, *Adobe Font Metrics File Format Specification, Version 4.1*
Adobe font metrics (AFM) files are available through the Type Technology Forum on the ASN Developer Program Web site.
- Technical Note #5014, *Adobe CMap and CID Font Files Specification, Version 1.0*
- Technical Note #5015, *Type 1 Font Format Supplement*
- Technical Note #5044, *Color Separation Conventions for PostScript Language Programs*
- Technical Note #5078, *Adobe-Japan1-4 Character Collection for CID-Keyed Fonts*
- Technical Note #5079, *Adobe-GB1-4 Character Collection for CID-Keyed Fonts*
- Technical Note #5080, *Adobe-CNS1-3 Character Collection for CID-Keyed Fonts*
- Technical Note #5088, *Font Naming Issues*
- Technical Note #5092, *CID-Keyed Font Technology Overview*
- Technical Note #5093, *Adobe-Korea1-0 Character Collection for CID-Keyed Fonts*
- Technical Note #5094, *Adobe CJK Character Collections and CMaps for CID-Keyed Fonts*

- Technical Note #5097, *Adobe-Japan2-0 Character Collection for CID-Keyed Fonts*
- Technical Note #5116, *Supporting the DCT Filters in PostScript Level 2*
- Technical Note #5176, *The Compact Font Format Specification*
- Technical Note #5177, *The Type 2 Charstring Format*
- Technical Note #5186, *Acrobat Forms JavaScript Object Specification, Version 4.05*
- Technical Note #5411, *ToUnicode Mapping File Tutorial*
- Technical Note #5620, *Portable Job Ticket Format, Version 1.1*
- Technical Note #5660, *Open Prepress Interface (OPI) Specification, Version 2.0*

Other Resources

Aho, A. V., Hopcroft, J. E., and Ullman, J. D., *Data Structures and Algorithms*, Addison-Wesley, Reading, MA, 1983. Includes a discussion of balanced trees.

Apple Computer, Inc., *TrueType Reference Manual*. Available on Apple's Web site at <http://fonts.apple.com/TTRefMan/index.html>.

Arvo, J. (ed.), *Graphics Gems II*, Academic Press, 1994. The section "Geometrically Continuous Cubic Bézier Curves" by Hans-Peter Seidel describes the mathematics used to smoothly join two cubic Bézier curves.

CIP4. See International Cooperation for the Integration of Processes in Prepress, Press and Postpress.

Fairchild, M. D., *Color Appearance Models*, Addison-Wesley, Reading, MA, 1997. Covers color vision, basic colorimetry, color appearance models, cross-media color reproduction, and the current CIE standards activities. Updates, software, and color appearance data are available at <http://www.cis.rit.edu/people/faculty/fairchild/CAM.html>.

Foley, J. D. et al., *Computer Graphics: Principles and Practice*, Addison-Wesley, Reading, MA, 1996. (First edition was Foley, J. D. and van Dam, A., *Fundamentals of Interactive Computer Graphics*, Addison-Wesley, Reading, MA, 1982.) Covers many graphics-related topics, including a thorough treatment of the mathematics of Bézier cubics and Gouraud shadings.

Glassner, A. S. (ed.), *Graphics Gems*, Academic Press, 1993. The section “An Algorithm for Automatically Fitting Digitized Curves” by Philip J. Schneider describes an algorithm for determining the set of Bézier curves approximating an arbitrary set of user-provided points. Appendix 2 contains an implementation of the algorithm, written in the C programming language. Other sections relevant to the mathematics of Bézier curves include “Solving the Nearest-Point-On-Curve Problem” and “A Bézier Curve-Based Root-Finder,” both by Philip J. Schneider, and “Some Properties of Bézier Curves” by Ronald Goldman. The source code appearing in the appendix is available via anonymous FTP, as described in the preface to *Graphics Gems III* (edited by D. Kirk; see its entry below).

Hewlett-Packard Corporation, *PANOSE Classification Metrics Guide*. Available on the Agfa Monotype Web site at <http://www.agfamonotype.com/print_manu/pan1.htm/>.

Hunt, R. W. G., *The Reproduction of Colour*, 5th ed., Fisher Books, England, 1996. A comprehensive general reference on color reproduction; includes an introduction to the CIE system.

International Color Consortium (ICC). The following are available with related documents at <<http://www.color.org>>:

- Specification ICC.1:1998-09, *File Format for Color Profiles*, and Document ICC.1A:1999-04, *Addendum 2 to Specification ICC.1:1998-09*
- *ICC Characterization Data Registry*

International Cooperation for the Integration of Processes in Prepress, Press and Postpress (CIP4), *JDF Specification, Version 1.0*. Available through the CIP4 Web site at <<http://www.cip4.org>>.

International Electrotechnical Commission (IEC), IEC/3WD 61966-2.1, *Colour Measurement and Management in Multimedia Systems and Equipment, Part 2.1: Default RGB Colour Space—sRGB*. Available through Hewlett-Packard’s sRGB Web site at <<http://www.srgb.com>>.

International Organization for Standardization (ISO). The following standards are available through <<http://www.iso.ch>>:

- ISO 639, *Codes for the Representation of Names of Languages*
- ISO 3166, *Codes for the Representation of Names of Countries and Their Subdivisions*

- ISO/IEC 8824-1, *Abstract Syntax Notation One (ASN.1): Specification of Basic Notation*
- ISO/IEC 10918-1, *Digital Compression and Coding of Continuous-Tone Still Images* (informally known as the JPEG standard, for the Joint Photographic Experts Group, the ISO group that developed the standard)

International Telecommunication Union (ITU), Recommendations T.4 and T.6. These standards for Group 3 and Group 4 facsimile encoding (which replace those formerly provided in the CCITT *Blue Book*, Vol. VII.3) can be ordered from ITU at <<http://www.itu.int>>.

Internet Engineering Task Force (IETF) Requests for Comments (RFCs). The following RFCs are available through <<http://www.rfc-editor.org>>:

- RFC 1321, *The MD5 Message-Digest Algorithm*
- RFC 1738, *Uniform Resource Locators*
- RFC 1766, *Tags for the Identification of Languages*
- RFC 1808, *Relative Uniform Resource Locators*
- RFC 1866, *Hypertext Markup Language 2.0 Proposed Standard*
- RFC 1950, *ZLIB Compressed Data Format Specification, Version 3.3*
- RFC 1951, *DEFLATE Compressed Data Format Specification, Version 1.3*
- RFC 2045, *Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies*
- RFC 2046, *Multipurpose Internet Mail Extensions (MIME) Part Two: Media Types*
- RFC 2068, *Hypertext Transfer Protocol—HTTP/1.1*
- RFC 2083, *PNG (Portable Network Graphics) Specification, Version 1.0*

Kirk, D. (ed.), *Graphics Gems III*, Academic Press, 1994. The section “Interpolation Using Bézier Curves” by Gershon Elber contains an algorithm for calculating a Bézier curve that passes through a user-specified set of points. The algorithm uses not only cubic Bézier curves, which are supported in PDF, but also higher-order Bézier curves. The appendix contains an implementation of the algorithm, written in the C programming language. The source code appearing in the appendix is available via anonymous FTP, as described in the book’s preface.

Lunde, K., *CJKV Information Processing*, O'Reilly & Associates, Sebastopol, CA, 1999. Excellent background material on CMaps, character sets, encodings, and the like.

Microsoft Corporation, *TrueType 1.0 Font Files Technical Specification*. Available at <<http://www.microsoft.com/typography/tt/tt.htm>>.

Netscape Communications Corporation, *Client-Side JavaScript Reference*. Available through Netscape's developer site at <<http://developer.netscape.com>>.

Pennebaker, W. B. and Mitchell, J. L., *JPEG: Still Image Data Compression Standard*, Van Nostrand Reinhold, New York, 1992.

Porter, T. and Duff, T., "Compositing Digital Images," *Computer Graphics*, Vol. 18 No. 3, July 1984. *Computer Graphics* is the newsletter of the ACM's special interest group SIGGRAPH; for more information, see <<http://www.acm.org>>.

Unicode Consortium publications:

- *The Unicode Standard, Version 2.0*, Addison-Wesley, Reading, MA, 1996. The latest information is available at <<http://www.unicode.org>>.
- Unicode Standard Annex #9, *The Bidirectional Algorithm, Version 3.1.0*, and Unicode Standard Annex #14, *Line Breaking Properties, Version 3.1.0*. These technical reports are available at <<http://www.unicode.org>>.

World Wide Web Consortium (W3C). The following publications are available through the W3C Web site at <<http://www.w3.org/>>:

- *Extensible Markup Language (XML) 1.0*
- *Extensible Stylesheet Language (XSL) 1.0*
- *Web Content Accessibility Guidelines 1.0*

Index

Page references in **boldface** mark principal or defining occurrences of a topic.

- ' (apostrophe) character
 - in dates 100
 - as text-showing operator 134
 - as text-showing operator. *See* ' (apostrophe) operator
- ' (apostrophe) operator 134, 302, 305, **311**, 702
- \ (backslash) character 29
 - as DOS (Windows) file name delimiter 120, 521
 - as escape character **30–31**, 119, 122, 312
 - escape sequence for **30**, 312
 - in unique names (Web Capture) 666–667
- : (colon) character
 - in conversion engine names 675
 - as DOS file name delimiter 120
 - as Mac OS file name delimiter 120
- \$ (dollar sign) character 343
- ... (ellipsis) character 782
- ! (exclamation point) character
 - in ASCII base-85 encoding **44**, **45**
- < (left angle bracket) character 26
 - double, as dictionary delimiter **35**, 67, 560
 - as hexadecimal string delimiter 29, **32**, 35, 122
- { (left brace) character 26
 - as delimiter in PostScript calculator functions 116
- [(left bracket) character 26
 - as array delimiter **34**, 808
- ((left parenthesis) character 26
 - escape sequence for **30**, 312
 - as literal string delimiter 29
- (minus sign) character
 - in dates 100
- # (number sign) character
 - as hexadecimal escape character in names **33**, 34, 322, 787–788
 - in uniform resource locators (URLs) 664
- % (percent sign) character 26
 - as comment delimiter 27
 - in uniform resource locators, “unsafe” 664
- . (period) character
 - double, in relative file specifications 120
 - double, in uniform resource locators (URLs) 664
 - in field names 533, 801
 - in file names 121
 - in handler names 548
 - . (period) character (*continued*)
 - in uniform resource locators (URLs) 664
 - in unique names (Web Capture) 666
- + (plus sign) character
 - in dates 100
 - in font subset names 323
- " (quotation mark) character
 - as text-showing operator. *See* " (quotation mark) operator
- " (quotation mark) operator 134, 302, 305, **311**, 702
- > (right angle bracket) character 26
 - double, as dictionary delimiter **35**, 67, 560
 - as EOD marker 44
 - as hexadecimal string delimiter 29, **32**, 35, 122
- } (right brace) character 26
 - as delimiter in PostScript calculator functions 116
-] (right bracket) character 26
 - as array delimiter 34
-) (right parenthesis) character 26
 - escape sequence for **30**, 312
 - as literal string delimiter 29
- / (slash) character 26, 94
 - as file specification delimiter **118**, 122
 - as name delimiter **32**, 34, 83, 357, 561
 - in uniform resource locators (URLs) 673
 - as UNIX file name delimiter 121
- ~> (tilde, right angle bracket) character sequence
 - as EOD marker **44**, 45
- _ (underscore) character
 - in file specifications 121
 - in multiple master font names 320
- ¥ (yuan symbol) character 343
- 1.3** entry (OPI version dictionary) 693, **694**
- 2.0** entry (OPI version dictionary) 693, **694**
- 83pv–RKSJ–H predefined CMap **344**, 346
- 90ms–RKSJ–H predefined CMap **344**, 346
- 90ms–RKSJ–V predefined CMap **344**, 346
- 90msp–RKSJ–H predefined CMap **344**, 346
- 90msp–RKSJ–V predefined CMap **344**, 346
- 90pv–RKSJ–H predefined CMap **344**, 346

A**A** entry

annotation dictionary 492, 513, 515, 519
 FDF field dictionary 565
 hint stream dictionary 736
 icon fit dictionary 566
 movie annotation dictionary 492, 511, 570
 outline item dictionary 478, 513, 519
 structure element dictionary 591, 592, 605, 606, 607, 638

A85 filter abbreviation 280, 789

AA entry

annotation dictionary 492, 514, 532
 document catalog 85, 515, 791
 document catalog (obsolete) 791
 FDF field dictionary 565
 field dictionary 514, 532
 outline item dictionary (obsolete) 798
 page object 90, 514

abbreviations and acronyms, expansion of 651, 658–659
 font characteristics unavailable for 622
 in Tagged PDF 616
 and Unicode natural language escape 659

abs operator (PostScript) 116, 703

absolute file specifications 119–120

AbsoluteColorimetric rendering intent 198

Abstract Syntax Notation One (ASN.1): Specification of Basic Notation (ISO/IEC 8824-1) 100, 815

AC entry (appearance characteristics dictionary) 537

accented characters 329, 356

access flags 76–77

access permissions xx, 74

copyright 6
 flags 76–77
 operations 74–75

accessibility to disabled users xx, 5, 573, 651–659

access permissions 74, 75, 77
 alternate descriptions 490, 592, 691
 annotation contents 501, 509, 510, 511, 512
 content extraction for 651
 field contents 531
 replacement text 592
 Span marked-content tag 633
 standard structure types 627
 Tagged PDF and 613, 614, 616, 638
See also

abbreviations and acronyms, expansion of
 alternate descriptions
 natural language specification
 replacement text

accurate screens algorithm 393–394

AccurateScreens entry (type 1 halftone dictionary) 393, 394

ACFM (Adobe Composite Font Metrics) file format 300

achromatic highlight, diffuse 185

achromatic shadow, diffuse 185

Acrobat® PDF viewer application xix, 1

Create Thumbnails command 789, 807
 Document Properties dialog box 804
 error reporting 784–785, 786, 789, 790, 791, 792, 793, 796, 798, 799, 800, 801, 804

Find command 789

implementation limits 1, 3, 705–707

implementation notes 3, 787–809

indirect processing of PDF 20, 21

JPEG implementation 61

LZW compression 48

native file formats 787, 792

plug-in extensions 684, 723

Print As Image feature 75, 790

RC4 encryption algorithm 73

scan conversion 404

Software Development Kit (SDK) 812

trademark 7

TrueType font encodings, treatment of 333–334

Type 0 fonts, naming of 353

version compatibility 2, 3, 353, 783–809

Web Capture plug-in extension 85

WebLink plug-in extension 800

Acrobat Distiller® PDF producer application 21, 576

balanced trees 86

OPI comments 807

PostScript names, compatibility with 787–788

spot functions 797

Acrobat Forms JavaScript Object Specification (Adobe Technical Note #5186) 556, 813

Acrobat Reader® PDF viewer application xix, 1, 7

AcroForm entry (document catalog) 85, 529, 734

AcroForms

See interactive forms

acronyms

See abbreviations and acronyms, expansion of

AcroSpider

See Web Capture plug-in extension

action dictionaries 84, 513–514

metadata inapplicable to 579

Next entry 514

S entry 514

Type entry 514

See also

go-to action dictionaries
 hide action dictionaries
 import-data action dictionaries

action dictionaries (*continued*)*See also*

- JavaScript action dictionaries
- launch action dictionaries
- movie action dictionaries
- named-action dictionaries
- remote go-to action dictionaries
- reset-form action dictionaries
- sound action dictionaries
- submit-form action dictionaries
- thread action dictionaries
- URI action dictionaries

action handlers 723

Action object type 514

action types 513, 514, 518–528, 550–557

FirstPage 527**GoTo** 518, 519**GoToR** 518, 520**Hide** 518, 527, 801**ImportData** 518, 556**JavaScript** 518, 556**LastPage** 527**Launch** 518, 521**Movie** 518, 526, 801**Named** 518, 528, 801**NextPage** 527**PrevPage** 527**ResetForm** 518, 555**Sound** 518, 525**SubmitForm** 518, 551**Thread** 518, 522**URI** 518, 523

actions 9, 84, 513–528

for annotations 492, 501, 513

chaining of 514

destinations for 474

for FDF fields 565

handlers 723

and named destinations 476

for outline items 477, 478, 513

plug-in extensions for 723, 784

type. *See* action types*See also*

- action dictionaries
- additional-actions dictionaries
- FirstPage** named action
- go-to actions
- hide actions
- import-data actions
- JavaScript actions
- LastPage** named action
- launch actions
- movie actions

actions (*continued*)*See also*

- named actions
- NextPage** named action
- PrevPage** named action
- remote go-to actions
- reset-form actions
- sound actions
- submit-form actions
- thread actions
- trigger events
- URI actions

activating

- annotations 488, 492, 501, 508, 510, 513, 565, 568, 737, 754, 799

- outline items 477, 478, 513, 798

ActualText entry (structure element dictionary) 592, 616, 638, 658**Alt** entry, compared with 658

and font characteristics 622

for illustrations 638

and Unicode mapping 621

and word breaks 623

add operator (PostScript) 116, 703

Add-RKSJ-H predefined CMap 344, 346

Add-RKSJ-V predefined CMap 344, 346

additional-actions dictionaries 85, 90, 514–516, 526, 791

for annotations 492

BI entry (widget annotation) 515**C** entry

- form field 516, 529

- page object 515, 800

D entry (annotation) 515**DC** entry (document catalog) 516**DP** entry (document catalog) 516**DS** entry (document catalog) 516**E** entry (annotation) 515**F** entry (form field) 516

for FDF fields 565

Fo entry (widget annotation) 515

for form fields 532

FP entry (obsolete) 800

inheritability of 800

K entry (form field) 516**LP** entry (obsolete) 800**NP** entry (obsolete) 800**O** entry (page object) 515, 800**PP** entry (obsolete) 800**U** entry (annotation) 515**V** entry (form field) 516**WP** entry (document catalog) 516**WS** entry (document catalog) 516**X** entry (annotation) 515, 517

- additive color representation **178**
 - for blend modes 416, 465
 - in blending color space 415, 450
 - and default color spaces 195
 - DeviceRGB** color space 179
 - and halftones 383
 - overprinting, not typically subject to 465
 - primary color components 178, 180
 - in soft-mask images 448
 - transfer functions, input to 381
 - transfer functions, output from 381, 383
- additive colorants 201
 - See also*
 - blue colorant
 - green colorant
 - red colorant
- additive output devices 201, 203, 383
- Adobe Acrobat Software Development Kit (SDK) 812
- Adobe CJK Character Collections and CMaps for CID-Keyed Fonts* (Adobe Technical Note #5094) **812**
- Adobe CMap and CID Font Files Specification* (Adobe Technical Note #5014) 5, 335, 336, 348, 352, 369, **812**
- Adobe-CNS1 character collection 344, 346, 362, 369, 621, 796
- Adobe-CNS1-3 Character Collection for CID-Keyed Fonts* (Adobe Technical Note #5080) **812**
- Adobe Composite Font Metrics (ACFM) file format 300
- Adobe Font Metrics (AFM) file format 300, 319, 812
- Adobe Font Metrics File Format Specification* (Adobe Technical Note #5004) 300, 319, **812**
- Adobe Garamond™ typeface 318
- Adobe-GB1 character collection 344, 346, 362, 369, 621, 796
- Adobe-GB1-4 Character Collection for CID-Keyed Fonts* (Adobe Technical Note #5079) **812**
- Adobe Glyph List* 333, 620, **811**
- Adobe imaging model xix, 2, 10, **11–12**
 - and indirect generation of PDF 19, 21
 - memory representation, independent of 14
 - and PostScript 5, 290
 - rendering 373
 - and transparent annotations 492, 496
- Adobe-Japan1 character collection 345, 346–347, 362, 363, 369, 621, 796
- Adobe-Japan1-4 Character Collection for CID-Keyed Fonts* (Adobe Technical Note #5078) **812**
- Adobe-Japan2 character collection 362, 363
- Adobe-Japan2-0 Character Collection for CID-Keyed Fonts* (Adobe Technical Note #5097) **813**
- Adobe-Korea1 character collection 345, 347, 362, 363, 369, 621, 796
- Adobe-Korea1-0 Character Collection for CID-Keyed Fonts* (Adobe Technical Note #5093) **812**
- Adobe Patent Clarification Notice 7, **811**
- Adobe.PPKLite** signature handler 549
- Adobe products
 - See*
 - Acrobat® PDF viewer application
 - Acrobat Distiller® PDF producer application
 - Acrobat Reader® PDF viewer application
 - Adobe Garamond™ typeface
 - ePaper® Solutions network publishing software
 - Extreme™ printing systems
 - FrameMaker® document publishing software
 - Illustrator® graphics software
 - InDesign™ page layout software
 - Minion® typeface
 - Myriad® typeface
 - PageMaker® page layout software
 - Photoshop® image editing software
 - Poetica® typeface
 - XMP™ (Extensible Metadata Platform) framework
- Adobe Solutions Network (ASN) 337, **811**
 - contact addresses 724, 811
 - Developer Program Web site 300, 319, 337, 347, 621, **811**, 812
 - telephone numbers 811
- Adobe standard encoding
 - See* **StandardEncoding** standard character encoding
- Adobe Technical Notes 348, **812**
 - #5001 (*PostScript Language Document Structuring Conventions Specification*) **812**
 - #5004 (*Adobe Font Metrics File Format Specification*) 300, 319, **812**
 - #5014 (*Adobe CMap and CID Font Files Specification*) 5, 335, 336, 348, 352, 369, **812**
 - #5015 (*Type 1 Font Format Supplement*) 5, 320, 811, **812**
 - #5044 (*Color Separation Conventions for PostScript Language Programs*) 792, **812**
 - #5078 (*Adobe-Japan1-4 Character Collection for CID-Keyed Fonts*) **812**
 - #5079 (*Adobe-GB1-4 Character Collection for CID-Keyed Fonts*) **812**
 - #5080 (*Adobe-CNS1-3 Character Collection for CID-Keyed Fonts*) **812**
 - #5088 (*Font Naming Issues*) 320, **812**
 - #5092 (*CID-Keyed Font Technology Overview*) 335, **812**
 - #5093 (*Adobe-Korea1-0 Character Collection for CID-Keyed Fonts*) **812**
 - #5094 (*Adobe CJK Character Collections and CMaps for CID-Keyed Fonts*) **812**
 - #5097 (*Adobe-Japan2-0 Character Collection for CID-Keyed Fonts*) **813**

- Adobe Technical Notes (*continued*)
- #5116 (*Supporting the DCT Filters in PostScript Level 2*) **813**
 - #5176 (*The Compact Font Format Specification*) 5, 316, 365, **813**
 - #5177 (*The Type 2 Charstring Format*) 5, **813**
 - #5186 (*Acrobat Forms JavaScript Object Specification*) 556, **813**
 - #5411 (*ToUnicode Mapping File Tutorial*) 369, **813**
 - #5620 (*Portable Job Ticket Format*) 24, 374, 689, 692, **813**
 - #5660 (*Open Prepress Interface (OPI) Specification*) 694, 812, **813**
- Adobe Type 1 Font Format 5, 315, 316, 365, 366, 367, **811**
- advance timing
See display duration
- AFM (Adobe Font Metrics) file format 300, 319, 812
- After block alignment **645**
- after edge **625**
of allocation rectangle 649
in layout 625, 641, 643, 645, 647, 649
- After** entry (JavaScript dictionary) **563**
- AHx** filter abbreviation **280**, 789
- AIFF (Audio Interchange File Format) 569
- AIFF-C (Audio Interchange File Format, Compressed) 569
- AIS** entry (graphics state parameter dictionary) **159**, 443
- ALaw sound encoding format **569**
- %ALDImageAsciiTag OPI comment (PostScript) **697**
- %ALDImageColor OPI comment (PostScript) **696**
- %ALDImageColorType OPI comment (PostScript) **696**
- %ALDImageCropFixed OPI comment (PostScript) **695**
- %ALDImageCropRect OPI comment (PostScript) **695**
- %ALDImageDimensions OPI comment (PostScript) **695**
- %ALDImageFilename OPI comment (PostScript) **694**
- %ALDImageGrayMap OPI comment (PostScript) **696**
- %ALDImageID OPI comment (PostScript) **695**
- %ALDImageOverprint OPI comment (PostScript) **696**
- %ALDImagePosition OPI comment (PostScript) **695**
- %ALDImageResolution OPI comment (PostScript) **696**
- %ALDImageTint OPI comment (PostScript) **696**
- %ALDImageTransparency OPI comment (PostScript) **696**
- %ALDImageType OPI comment (PostScript) **696**
- %ALDObjectComments OPI comment (PostScript) **695**
- Aldus Corporation 693
- “Algorithm for Automatically Fitting Digitized Curves, An” (Schneider) **814**
- algorithm tags (PNG predictor functions) 51
- all-cap fonts 358
- All** colorant name
DeviceN color spaces, prohibited in 206
in **Separation** color spaces **203**
- AllCap font flag **358**
- allocation rectangle **626**, **648–649**
in layout 641, 645, 646
- alpha **411**
alpha source parameter **149**, 159, 268, **443**, **444**
backdrop. *See* backdrop alpha
in basic compositing formula 413, 414
current alpha constant **149**, 159, 268
group backdrop 429
group. *See* group alpha
interpretation of 419–420
notation for 413
object 429, 430, 431, 433
premultiplied. *See* preblending of soft-mask image data result. *See* result alpha
shape and opacity, product of 413, 420, 424, 429, 435
source. *See* source alpha
- alpha constant, current
See current alpha constant
- alpha mask
See soft masks
- Alpha** soft-mask subtype 445, **446**
- alpha source parameter **149**
AIS entry (graphics state parameter dictionary) 159
constant opacity **444**
constant shape **444**
ignored by older viewer applications 797
mask opacity **443**
mask shape **443**
soft-mask images 268
- Alphabetic** character class 362, 363
- AlphaNum** character class 363
- Alt** entry (structure element dictionary) **592**, 616, 638, 657, 658
ActualText entry, compared with 658
and font characteristics 622
for illustrations 638
and Unicode mapping 621
- alternate color space **204**
and color separations 684
for **DeviceN** color spaces 115, 195, 206–207, 216, 236
and flattening of transparent content 470
for **ICCBased** color spaces 190, 191
and overprinting 204
for **Separation** color spaces 195, 203, 204, 207, 236
in soft masks 447
and spot color components 457, 458
and transparent imaging model 204, 415
and transparent overprinting 462

- alternate descriptions 651, **657–658**
 - for annotations 490, 657, 691
 - font characteristics unavailable for 622
 - for sound annotations 657
 - for structure elements 592, 615, 657, 658
 - in Tagged PDF 616
 - and Unicode natural language escape 657–658
- Alternate** entry (ICC profile stream dictionary) **190**, 191
- alternate field names 531, 657
- alternate image dictionaries **273–274**
 - DefaultForPrinting** entry **274**
 - Image** entry **274**
- alternate images 262, 267, 269, **273–275**
 - printing 274
 - See also*
 - alternate image dictionaries
- Alternates** entry (image dictionary) **269**, 448
- anamorphic scaling **566**
- and** operator (PostScript) **116**, **704**
- angle (halftone screen) **384**
 - type 1 halftones 391, 392, 393
 - type 10 halftones 391, 394, 395, 397
 - type 16 halftones 391, 399
- angle brackets (< >) 26
 - double, as dictionary delimiters 35, 67, 560
 - as hexadecimal string delimiters 29, **32**, 35, 122
- Angle** entry (type 1 halftone dictionary) **393**
- Annot** entry (movie action dictionary) **526**
- Annot** object type **490**, 775, 780
- annotation dictionaries 4, 90, **489–492**, 775, 780, 798
 - A** entry **492**, 513, 515, 519
 - AA** entry **492**, 514, 532
 - AP** entry **491**, 497, 500, 502, 503, 505, 506, 507, 508, 509, 510, 534, 548, 549, 554, 682, 690, 798
 - AS** entry **491**, 498, 682, 690, 691, 798
 - Border** entry **490**, 495, 496, 798
 - BS** entry **490**, 491, 495, 798
 - C** entry **491**, 509
 - CA** entry **491**, 496
 - Contents** entry **490**, 509, 657
 - F** entry **490**, 492, 565, 682, 690
 - and hide actions 527
 - in Linearized PDF 737
 - M** entry **490**, 509, 798
 - NM** entry **490**
 - P** entry **490**
 - Page** entry (FDF files) **568**
 - Popup** entry **492**, 508
 - Rect** entry **490**, 494, 504, 505, 523, 534, 548
 - StructParent** entry **492**, 601
 - Subtype** entry **490**
 - T** entry **492**, 509, 553
 - Type** entry **490**
- annotation dictionaries (*continued*)
 - See also*
 - circle annotation dictionaries
 - FDF annotation dictionaries
 - file attachment annotation dictionaries
 - free text annotation dictionaries
 - ink annotation dictionaries
 - line annotation dictionaries
 - link annotation dictionaries
 - markup annotation dictionaries
 - movie annotation dictionaries
 - pop-up annotation dictionaries
 - printer's mark annotation dictionaries
 - rubber stamp annotation dictionaries
 - sound annotation dictionaries
 - square annotation dictionaries
 - text annotation dictionaries
 - trap network annotation dictionaries
 - widget annotation dictionaries
- annotation flags 490, **492–494**, 565, 799
 - Hidden **493**, 526, 615, 799, 801
 - Invisible **493**
 - NoRotate **493**, 494, 496, 500
 - NoView **493**
 - NoZoom **493**, 494, 496, 500
 - Print **493**, 682, 690, 799
 - ReadOnly **493**, 682, 690
- annotation handlers **489**, 493, 723
- annotation icons 488, 799
 - Approved 507
 - AsIs 507
 - background color 491
 - for button fields 565, 566
 - Comment 500
 - Confidential 507
 - Departmental 507
 - Draft 507
 - Experimental 507
 - Expired 507
 - for file attachment annotations 509
 - Final 507
 - ForComment 507
 - ForPublicRelease 507
 - Graph 509
 - Help 500
 - Insert 500
 - Key 500
 - Microphone 510
 - NewParagraph 500
 - NotApproved 507
 - Note 500
 - NotForPublicRelease 507
 - Paperclip 509
 - Paragraph 500

- annotation icons (*continued*)
 - PushPin 509
 - for rubber stamp annotations 507
 - scaling 566
 - Sold 507
 - for sound annotations 510
 - Speaker 510
 - Tag 509
 - for text annotations 499, 500
 - TopSecret 507
 - for widget annotations 537, 565, 566
- annotation rectangle **490**
 - and appearance streams 496, 534, 535, 537
 - border style 495
 - for circle annotations 504
 - clipping to 496
 - coordinate system 496
 - highlighting 501, 512
 - for movie annotations 525, 572
 - for printer's mark annotations 680
 - rotation 493, 494
 - scaling 493, 494
 - for signature fields 548
 - for square annotations 504
 - and submit-form actions 552
 - for text fields 543
 - and URI actions 523
 - for widget annotations 565, 566
- annotation types 488, 490, **498–513**, 799–800
 - Circle** 499, **505**
 - dictionary entries for 489
 - FileAttachment** 499, **509**
 - FreeText** 499, **502**
 - handlers for 489
 - and Hidden flag 493
 - Highlight** 499, **506**
 - Ink** 499, **508**
 - Line** 499, **503**
 - Link** 491, 492, 499, **501**, 562
 - Movie** 492, 499, **511**, 562, 799
 - plug-in extensions for 489, 498
 - Popup** 499, **509**
 - PrinterMark** 492, 499, 562, **682**
 - Sound** 499, **510**
 - Square** 499, **505**
 - Squiggly** 499, **506**
 - Stamp** 499, **507**
 - StrikeOut** 499, **506**
 - Text** 499, **500**
 - TrapNet** 492, 499, 513, 562, 690, **691**
 - Underline** 499, **506**
 - unknown 493, 498
 - Widget** 492, 499, **512**, 562
- annotations 1, 9, 23, **488–513**
 - actions for 492, 501, 513
 - activating **488**, 492, 501, 508, 510, 513, 565, 568, 737, 754, 799
 - active area 497, 501, 512, 515, 517, 526, 537
 - additional-actions dictionary 492
 - alternate description 490, 657, 691
 - annotation rectangle. *See* annotation rectangle
 - appearance dictionary 491, 493, **497–498**
 - appearance state 491, **497–498**, 513, 539, 541
 - appearances. *See* appearance streams
 - blend mode 491, 496
 - border style. *See* border styles
 - border width 490, 495, 503, 505, 508
 - color 491
 - corner radii 490
 - dash pattern 490, 491, 495, 503, 505, 508
 - destinations for 474, 476, 501, 799
 - in FDF 557, 558, 562
 - flags. *See* annotation flags
 - handlers **489**, 493, 723
 - hiding and showing 493, 526, 527
 - highlighting 496, 501, 512
 - icon. *See* annotation icons
 - label 492, 553
 - in Linearized PDF 737, 745, 754
 - as logical structure content items 618
 - in logical structure elements 598
 - marked-content sequences, association with 618
 - modification date 490
 - name 4, 490
 - in page content order, sequencing of 618–619
 - and page objects 81
 - plug-in extensions for 489, 723, 784
 - pop-up window 491, 492, 553
 - PostScript conversion to 22
 - printing 493, 497, 799
 - as real content 615
 - and reference XObjects 289
 - rotating 493, 494, 496, 499
 - scaling 493, 494, 496, 499, 566
 - in structural parent tree 590
 - and submit-form actions 552, 553
 - and trap networks 691
 - trigger events for **515**
 - type. *See* annotation types
 - in updating example 774, 775, 777, 778, 779, 780, 782
 - URI actions for 523
 - user interaction 493, 496, 497, 511
 - version compatibility 787
 - See also*
 - annotation dictionaries
 - circle annotations
 - file attachment annotations
 - free text annotations

annotations (*continued*)

See also

- ink annotations
- line annotations
- link annotations
- markup annotations
- movie annotations
- pop-up annotations
- printer's mark annotations
- rubber stamp annotations
- sound annotations
- square annotations
- text annotations
- trap network annotations
- widget annotations

Annots entry

- FDF dictionary **562**
- page object **90**, 489, 513, 618, 689, 691, 737, 775, 778, 779, 780, 807

AnnotStates entry (trap network annotation dictionary)
690, **691**

AntiAlias entry (shading dictionary) **234**

anti-aliasing

- shading patterns **234**
- transparency 411, 421

AP entry

- annotation dictionary **491**, 497, 500, 502, 503, 505, 506, 507, 508, 509, 510, 534, 548, 549, 554, 682, 690, 798
- FDF field dictionary **565**, 803
- name dictionary **93**, 498

API. *See* application programming interface

apostrophe (') character

- in dates 100
- as text-showing operator 134, 302, 305, **311**, 702

appearance characteristics dictionaries

- AC** entry **537**
- BC** entry **536**
- BG** entry **536**
- CA** entry **536**
- I** entry **537**
- IF** entry **537**
- IX** entry **537**
- R** entry **536**
- RC** entry **537**
- RI** entry **537**
- TP** entry **537**

- for widget annotations 512, **536–537**

appearance dictionaries 491, **497–498**

- D** entry **497**, 565, 682, 690
- N** entry **497**, 498, 534, 565, 615, 682, 689, 690, 692
- for pushbutton fields 565

appearance dictionaries (*continued*)

- R** entry **497**, 565, 682, 690
- subdictionaries 491, **497–498**, 682, 690
- and unknown annotation types 493
- for widget annotations 529, 539

appearance states 491, **497–498**, 513

- for checkbox fields 539, 541
- for printer's marks 682
- in trap networks 690, 691

appearance streams **496–498**

- appearance states, selected by 491
- backdrop 496
- for button fields 557
- as content streams 93, 614
- coordinate system 496
- down appearance **497**, 501, 512
- dynamic 512, 529, **533–537**
- and form fields 529, 533
- and Form standard structure type 637
- marked-content sequences in 594
- named 93, 498
- normal appearance **497**, 498, 615
- opacity 491, 492
- pop-up annotations, inapplicable to 508
- and pop-up help systems 493, 526
- for printer's mark annotations 682
- and reference XObjects 289
- resources 96, 801
- rollover appearance **497**
- soft mask 496
- and transparency 496
- transparency groups as 451
- for trap network annotations. *See* trap network appearances
- and unknown annotation types 493
- for widget annotations 529, 534, 539, 637

appearances, annotation

- See* appearance streams

AppendOnly signature flag **530**

Apple Computer, Inc.

- Mac[®] OS operating system 19, 328, 714
- TrueType[®] font format 321
- TrueType Reference Manual* 321, **813**

application data dictionaries 284, **581–582**

- LastModified** entry **582**
- Private** entry **582**

- application/pdf content type (MIME) 553
- application programming interface (API) 19, 20
- application-specific data 19
- application/vnd.fdf content type (MIME) 558
- application/x-www-form-urlencoded content type (MIME) 672

- applications
 - launching 513, 520, 521
 - See also*
 - consumer applications, PDF
 - consumer applications, Tagged PDF
 - producer applications, PDF
 - producer applications, Tagged PDF
 - viewer applications, PDF
- Approved annotation icon 507
- APRef** entry (FDF field dictionary) 565
- Arabic writing systems 619, 642
- architecture, PDF 2
- Arial standard font name 795
- Arial–Bold standard font name 795
- Arial,Bold standard font name 795
- Arial–BoldItalic standard font name 795
- Arial,BoldItalic standard font name 795
- Arial–BoldItalicMT standard font name 795
- Arial–BoldMT standard font name 795
- Arial–Italic standard font name 795
- Arial,Italic standard font name 795
- Arial–ItalicMT standard font name 795
- ArialMT standard font name 795
- array objects 26, 27, 34
 - capacity limit 34, 706
 - as dictionary values 99
 - null elements 28
 - syntax 34
- arrays
 - color space. *See* color space arrays
 - explicit destinations, defining 474, 476
 - related files 123, 125–127, 128
 - See also*
 - array objects
- art box 677
 - and bounding box 288
 - clipping to 679
 - display of 681
 - imposition of pages, ignored in 679
 - in page object 89
 - page placement in another document 679
 - printer’s marks excluded from 680
 - printing, ignored in 679
- Art standard structure type 627, 628, 631
- ArtBox** entry
 - box color information dictionary 681
 - page object 89, 677, 806
- articles 481, 483–485
 - in document catalog 83, 84
 - in Linearized PDF 753, 755
 - and page content order 618
- articles (*continued*)
 - in page objects 89, 791
 - as structure elements (Tagged PDF) 627
 - and text discontinuities 617
 - See also*
 - beads
 - threads
- Artifact marked-content tag 616
- artifact types 616
 - Layout** 616
 - Page** 616
 - Pagination** 616
- artifacts (Tagged PDF) 613, 614, 615–617
 - attached 616
 - bounding box 616, 617
 - incidental 617
 - layout 615
 - logical structure order, excluded from 618
 - page 615
 - page content order, included in 618
 - pagination 615
 - property list 616, 617
 - specification 616–617
 - See also*
 - artifact types
- AS** entry (annotation dictionary) 491, 498, 682, 690, 691, 798
- Ascent** entry (font descriptor) 356, 647
- ASCII (American Standard Code for Information Interchange) 24–25
 - base-85 encoding 15, 41, 42, 43, 44–45
 - character set 25, 26, 31
 - compression 41
 - in file specifications 119, 121, 122
 - filters 41
 - hexadecimal encoding 41, 42, 44, 57, 73
 - LZW encoding 47
 - nonprinting characters 30
 - for PDF representation 15
 - for portability 41
 - strings and streams 25
 - text files 573, 659
 - TIFF tags 697
 - in uniform resource locators (URLs) 127, 664
 - in unique names (Web Capture) 666, 667
 - UTF-8 character encoding 34
- ASCII85Decode** filter 42, 44–45
 - A85** abbreviation 280, 789
 - in inline images 279
- ASCIIHexDecode** filter 41, 42, 44
 - AHx** abbreviation 280, 789
 - in inline images 279
- Asian writing systems 299, 335

- AsIs annotation icon 507
 - ASN. *See* Adobe Solutions Network
 - ASN.1 (Abstract Syntax Notation One) 100
 - Aspect** entry (movie dictionary) 571, 572
 - atan** operator (PostScript) 116, 703
 - AToB* transformation (ICC color profile) 192, 415, 686, 807
 - Attached** entry (property list, Tagged PDF artifact) 616
 - attribute classes 605–606
 - class map 590
 - name 590, 592, 605–606, 607, 638
 - revision number 592, 606, 607
 - and standard attribute owners 638
 - structure elements belonging to 592, 606
 - attribute objects 605
 - for attribute classes 605, 606
 - O** entry 605, 639, 640, 650, 651
 - owner 605, 639
 - revision number 606–607
 - role map 590
 - standard attribute owners 638–639, 640
 - structure elements, associated with 591
 - attribute owners, standard
 - See* standard attribute owners
 - attribute revision numbers 591, 592, 605, 606–607
 - generation numbers, distinguished from 606
 - AU** entry (source information dictionary) 670
 - authenticity of documents, certifying xx, 18
 - Author** entry (document information dictionary) 576
 - Auto height attribute 645, 649
 - Auto line height 646, 647
 - Auto width attribute 645, 649
 - automatic stroke adjustment
 - See* stroke adjustment, automatic
 - Average predictor function (LZW and Flate encoding) 50, 51
 - AvgWidth** entry (font descriptor) 357
 - axial shadings
 - See* type 2 shadings
- B**
- B border style (beveled) 495
 - B** entry
 - page object 89, 484, 557, 737, 791
 - sound object 569, 570
 - thread action dictionary 522
 - B** operator 134, 167, 699
 - and transparent overprinting 462
 - b** operator 134, 162, 167, 699
 - and transparent overprinting 462
 - B*** operator 134, 167, 699
 - and transparent overprinting 462
 - b*** operator 134, 167, 699
 - and transparent overprinting 462
 - B5pc–H predefined CMap 344, 346
 - B5pc–V predefined CMap 344, 346
 - backdrop 410
 - for annotation appearances 496
 - compositing with 133, 171, 412
 - and fully opaque objects 467
 - group. *See* group backdrop
 - immediate (transparency group element). *See* immediate backdrop
 - initial (transparency group). *See* initial backdrop
 - for page group 412, 433, 436, 437
 - for patterns 453, 454
 - and transparent overprinting 462, 463
 - See also*
 - backdrop alpha
 - backdrop color
 - backdrop opacity
 - backdrop shape
 - backdrop alpha
 - in compositing 419, 420
 - notation 414, 429
 - backdrop color 413
 - backdrop fraction 432
 - blending color space, conversion to 415
 - and CompatibleOverprint blend mode 461
 - in compositing 411, 419, 420
 - and nonseparable blend modes 419
 - notation 414, 429, 437
 - and overprinting 459
 - for page group 436
 - removal from compositing computations 432
 - and separable blend modes 416, 417, 418
 - and soft masks 439, 440, 441, 446
 - specifying 441
 - spot color components 457
 - in transparency groups 431
 - backdrop fraction 432
 - backdrop opacity 411, 424
 - notation 423
 - and overprinting 459
 - backdrop shape 424
 - notation 423
 - Background** entry (shading dictionary) 232, 234, 238, 239, 453
 - backslash (\) character 29
 - as DOS (Windows) file name delimiter 120, 521
 - as escape character 30–31, 119, 122, 312
 - escape sequence for 30, 312
 - in unique names (Web Capture) 666–667

- backspace (BS) character
 - escape sequence for 30
- balanced trees 86, 813
- <BASE> body element (universal resource identifier) 524
- base color space (**Indexed** color space) 195, **199–200**, 236, 456, 461
- base encoding 330, 358, 710
- Base** entry (URI dictionary) **524**
- base images 267, **273**, 274, 277, 793
- base URI (URI action) **524**
- BaseEncoding** entry (encoding dictionary) **330**, 332, 333, 796
- BaseFont** entry
 - CIDFont dictionary **338**, 353, 356
 - font dictionary 356
 - font subset **322**, 796
 - multiple master font dictionary **320**
 - TrueType font dictionary **321–322**
 - Type 0 font dictionary **353**
 - Type 1 font dictionary 34, 290, **317**
- baseline shift (ILSEs) 647
- BaselineShift** standard structure attribute 633, 637, 640, **647**, 648, 649
- basic compositing formula
 - See compositing computations
- BBox** entry
 - property list (Tagged PDF artifact) **616**
 - shading dictionary **234**, 237, 241
 - type 1 form dictionary 283, **284**, 288, 451, 534, 648
 - type 1 pattern dictionary **222**, 223
- BBox** standard structure attribute 637, 640, **644**
- BC** entry
 - appearance characteristics dictionary **536**
 - soft-mask dictionary 445, **446**
- BDC** operator 134, 583, **584**, 593, 699
 - property list 583, 585
- bead dictionaries **483–484**
 - in Linearized PDF 734, 737, 740, 755
 - N** entry **484**, 755
 - P** entry **484**, 755
 - R** entry **484**
 - T** entry **484**, 737
 - thread actions, target of 522
 - Type** entry **484**
 - V** entry **484**
- Bead** object type **484**
- beads, article **483**, 484
 - in Linearized PDF 737, 753, 755
 - in page objects 89, 791
 - and thread actions 522
- beads, article (*continued*)
 - See also
 - articles
 - bead dictionaries
 - threads
- Before block alignment **645**
- before edge **625**
 - of allocation rectangle 649
 - in layout 625, 641, 643, 645, 647
- Before** entry (JavaScript dictionary) **563**
- Before placement attribute **641**
- beginbfchar** operator (PostScript) **352**, 354, 369, 371
- beginbfrange** operator (PostScript) **352**, 369, 372
- begincidchar** operator (PostScript) **352**, 354
- begincidrange** operator (PostScript) **352**
- begincmap** operator (PostScript) **351**
- begincodespacerange** operator (PostScript) **351**, 354, 369, 371
- beginnotdefchar** operator (PostScript) **352**, 355
- beginnotdefrange** operator (PostScript) **352**, 355
- beginrearrangedfont** operator (PostScript) 352
- beginusematrix** operator (PostScript) 352
- Bernstein polynomials 258
- bevel line join style 153, **154**, 168
- “Bézier Curve-Based Root-Finder, A” (Schneider) **814**
- Bézier curves, cubic
 - See cubic Bézier curves
- BG** entry
 - appearance characteristics dictionary **536**
 - graphics state parameter dictionary **158**, 218, 379
- BG2** entry (graphics state parameter dictionary) **158**, 218, 379
- BI** operator 134, **278**, 279, 699
- BibEntry standard structure type **634**
- bibliographies 628, 634
- bicubic tensor-product patches 256, **257–259**
- Bidirectional Algorithm, The* (Unicode Standard Annex #9) 642, **816**
- Big Five character encoding 344, 788, 803
- Big Five character set 344
- bilevel output devices **13**, 14
 - halftone screens 383, 390
- bilinear interpolation 250, 251
- binary data 25
- binary files 15, 25
- biometric authentication 538, 547
- bitshift** operator (PostScript) **116**, **704**

- BitsPerComponent** entry
FlateDDecode filter parameter dictionary 49
 image dictionary 268, 271, 276, 277, 448, 480
 inline image object 279
LZWDecode filter parameter dictionary 49
 type 4 shading dictionary 244, 247
 type 5 shading dictionary 249
 type 6 shading dictionary 253
- BitsPerCoordinate** entry
 type 4 shading dictionary 244, 247
 type 5 shading dictionary 249
 type 6 shading dictionary 253
- BitsPerFlag** entry
 type 4 shading dictionary 244, 247
 type 6 shading dictionary 253
- BitsPerSample** entry (type 0 function dictionary) 110, 111
- BI** entry (additional-actions dictionary) 515
- BI** trigger event (annotation) 515
- black color component
 black-generation function 150, 378, 379
DeviceCMYK color space 178, 180
DeviceN color spaces 206
 gray, complement of 377
 grayscale conversion 377
 halftones for 401
 initialization 180
 in multitoness 205
 overprinting 464
 RGB conversion 378, 380
 transfer function 380, 381
 transparent overprinting 464
- black colorant
 overprinting 464
 PANTONE Hexachrome system 205
 printing ink 201
 process colorant 178, 180
 transparent overprinting 464
- black-generation function 150, 378, 379, 380
BG entry (graphics state parameter dictionary) 158
BG2 entry (graphics state parameter dictionary) 158
 and transparency 466, 468, 469
- black point, diffuse 183, 185, 188
- BlackIs1** entry (**CCITTFaxDecode** filter parameter dictionary) 54
- BlackPoint** entry
CalGray color space dictionary 183
CalRGB color space dictionary 185
Lab color space dictionary 188
- bleed box 677
 clipping to 679
 display of 5, 681
 in page object 88
 printing of finished pages, ignored in 679
 in printing of intermediate pages 679
- BleedBox** entry
 box color information dictionary 681
 page object 88, 677, 806
- blend circles (type 3 shading) 240–242
- blend functions 415, 416
 in basic compositing formula 413
 blending color space, assumptions about 415
 in compositing 419, 420
 linear 455
 notation 414, 427, 431
 and overprinting 459, 461
 and subtractive color components 465
 white-preserving blend modes 460
See also
 blend modes
- blend modes 410, 416–419
 additive color representation 416, 465
 for annotations 491, 496
 in basic compositing formula 414
 blending color space, assumptions about 415
Color 419
ColorBurn 418
ColorDodge 418
Compatible 419, 461, 467
CompatibleOverprint 460–462, 465, 466, 467–468
 in compositing 420
 current. *See* current blend mode
Darken 417, 459, 460, 466
Difference 418, 460
Exclusion 418, 460
HardLight 412, 418
Hue 419
 in isolated groups 433
 in knockout groups 434
Lighten 417
Luminosity 419
Multiply 417, 433
 nonseparable 418–419, 460
 non-white-preserving 460
Normal 412, 417, 419, 432, 433, 437, 442, 452, 454, 459, 460, 461, 462, 463, 466, 467, 491, 496
Overlay 417
 and overprinting 459–460, 462, 463
Saturation 419
Screen 417, 462
 separable 416–418, 460
SoftLight 418
 specifying 442
 standard 416–419, 442
 in subtractive color spaces 415
 in transparency groups 411, 425, 431
 white-preserving 460
See also
 blend functions

- blending color space 192, 196, **414–416**
 - for isolated groups 425, 433, 450
 - for nonseparable blend modes 419
 - for page group 442
 - process colors 415
 - specifying **442**, 797
 - spot colors 415
 - for transparency groups 442, 450, 454
- Blinds transition style **486**, 487
- block alignment **645**
 - After **645**
 - Before **645**
 - Justify **645**
 - Middle **645**
- block-level structure, Tagged PDF 632
 - strong 632
 - weak 632
- block-level structure elements (BLSEs) **624**, **629–632**
 - bounding box 644
 - content items in 627
 - direct content items in 625, 632
 - general layout attributes 640
 - illustrations as 624
 - ILSEs contained in 624, 632
 - ILSEs, nested within 632, 647
 - list elements **630**
 - L 628, 629, **630**, 649
 - Lbl 629, **630**, 633, 634, 644, 649, 650
 - LBody 629, **630**, 644
 - LI 628, 629, **630**, 649
 - list elements, nested within 630
 - nesting of 625
 - packing of ILSEs within **625**, 641
 - paragraphlike elements **629–630**, 644
 - H 629, **630**, 631, 632
 - H1–H6 629, **630**, 632
 - P 629, **630**, 631, 632
 - stacking **625**, 632, 641, 642
 - standard layout attributes for 640, **642–646**
 - BBox** 640, **644**
 - BlockAlign** 640, **645**
 - EndIndent** 640, **644**
 - Height** 640, **645**
 - InlineAlign** 640, **646**
 - SpaceAfter** 640, **643**
 - SpaceBefore** 640, **643**
 - StartIndent** 640, **643**
 - TextAlign** 640, **644**
 - TextIndent** 640, **644**
 - Width** 640, **645**
 - table element **631**
 - Table 629, **631**, 640, 644, 645
 - usage guidelines 631–632
- Block placement attribute **641**, 643, 644, 649
- block-progression direction **624**
 - illustrations, height of 649
 - in layout 625, 629, 632, 641, 643, 645, 646
 - shift direction, opposite to 625, 647
 - table expansion 651
 - writing mode 642
- block quotations 627
- BlockAlign** standard structure attribute 640, **645**
- BlockQuote standard structure type **627**
 - Quote, distinguished from 633
- BLSEs. *See* block-level structure elements
- blue color component
 - CMYK conversion 377, 380
 - DeviceRGB** color space 178, 179
 - grayscale conversion 377
 - halftones for 401
 - in **Indexed** color table 199
 - initialization 180
 - and threshold arrays 390
 - transfer function 381
 - yellow, complement of 378
- blue colorant
 - additive primary 178, 179, 180
 - display phosphor 201
- BM** entry (graphics state parameter dictionary) **159**, 442
- BMC** operator 134, 535, 583, **584**, 699
- body, file
 - See* file body
- Bold outline item flag **479**
- bookmarks
 - access permission 75, 77
 - PostScript conversion 22
 - See also*
 - outline items
- boolean objects 27, **28**
 - as dictionary values 99
- boolean operators 28
- Border** entry (annotation dictionary) **490**, 495, 496, 798
- Border** object type **495**
- border style dictionaries 490, **495–496**, 503, 505, 508
 - D** entry **495**
 - S** entry **495**
 - Type** entry **495**
 - W** entry **495**
- border styles 490–491, **495–496**, 503, 505, 508, 799
 - B (beveled) **495**
 - D (dashed) **495**
 - I (inset) **495**
 - S (solid) **495**, 799
 - U (underline) **495**

- bounding box
 - artifact 616, 617
 - BLSE 644
 - font 323, 325, 326, 356, 700
 - form XObject 284, 534
 - glyph **299**
 - illustration 624, 648
 - imported page 288
 - movie 571
 - non-isolated group 451
 - page **474–475**
 - pattern cell 222
 - reference XObject 288
 - shading object 234, 237, 238, 241
 - soft mask 445
 - table 624, 648
 - table cell 648
 - transparency group 452
 - Bounds** entry (type 3 function dictionary) **114**
 - box color information dictionaries 89, **679–680, 681**
 - ArtBox** entry **681**
 - BleedBox** entry **681**
 - CropBox** entry **681**
 - TrimBox** entry **681**
 - box style dictionaries 680, **681**
 - C** entry **681**
 - D** entry **681**
 - S** entry **681**
 - W** entry **681**
 - Box transition style **486, 487**
 - BoxColorInfo** entry (page object) **89, 679**
 - BPC** entry (inline image object) **279**
 - braces ({}) 26
 - as delimiters in PostScript calculator functions 116
 - brackets ([]) 26
 - as array delimiters 34
 - BS** entry
 - annotation dictionary **490, 491, 495, 798**
 - circle annotation dictionary **505**
 - ink annotation dictionary **508**
 - line annotation dictionary **503**
 - square annotation dictionary **505**
 - BT** operator 134, 294, 305, **308, 535, 584, 637, 699**
 - Btn** field type **531, 539, 540**
 - BTtoA** transformation (ICC color profile) 192, 415, 686
 - built-in character encodings 328
 - embedded fonts 330
 - expert fonts 709
 - overriding 329
 - simple fonts 316
 - Symbol font 709, **718–720**
 - symbolic fonts 329, 330, 709
 - TrueType fonts 332
 - built-in character encodings (*continued*)
 - Type 1 fonts 318, 332, 710
 - ZapfDingbats font 709, **721–722**
 - bullet character 714, 796
 - butt line cap style **153, 154, 168**
 - button field flags **538**
 - NoToggleToOff **538, 540**
 - Pushbutton **538, 539, 540**
 - Radio **538, 539, 540**
 - button fields 531, 537, **538–543**
 - alternate (down) caption 537
 - alternate (down) icon 537
 - appearances 557
 - flags. *See* button field flags
 - icon 565, 566
 - icon fit dictionary 537
 - normal caption 536
 - normal icon 537
 - rollover caption 537
 - rollover icon 537
 - scaling 566
 - trigger events inapplicable to 517
 - See also*
 - checkbox fields
 - pushbutton fields
 - radio button fields
 - BX** operator **95, 134, 699**
 - byte order marker (Unicode) 98
 - ByteRange** entry (signature dictionary) **549**
- ## C
- C** entry
 - additional-actions dictionary
 - form field **516, 529**
 - page **515, 800**
 - annotation dictionary **491, 509**
 - box style dictionary **681**
 - hint stream dictionary **736**
 - outline item dictionary **479**
 - sound object **569, 570**
 - source information dictionary **670**
 - structure element dictionary **592, 606, 607, 638**
 - URL alias dictionary **671**
 - Web Capture command settings dictionary **674, 708**
 - Web Capture information dictionary **660, 708**
 - c** operator 134, **163, 165, 699**
 - C programming language 814, 815
 - C++ programming language 606
 - C** trigger event
 - form field **516, 517**
 - page **515**

- C0** entry (type 2 function dictionary) 113
- C1** entry (type 2 function dictionary) 113
- CA** entry
 - annotation dictionary 491, 496
 - appearance characteristics dictionary 536
 - graphics state parameter dictionary 159, 444
- ca** entry (graphics state parameter dictionary) 159, 444
- CalCMYK** color spaces 181
- calculator functions, PostScript
 - See type 4 functions
- CalGray** color space dictionaries 183
 - BlackPoint** entry 183
 - Gamma** entry 183, 184
 - WhitePoint** entry 183, 184
- CalGray** color spaces 176, 181, 182–184, 272
 - as blending color space 415
 - color values 183
 - gamma correction 183
 - and **ICCBased** color spaces, compared 189, 192
 - initial color value 216
 - rendering 374
 - setting color values in 217
 - See also
 - CalGray** color space dictionaries
- calibrated color 181
 - CalGray** color spaces 183, 184
 - CalRGB** color spaces 184
 - CMYK*, as blending color space 415
 - device profiles 375
 - implicit conversion 195–197
 - multitones 210, 213
- CalRGB** color space dictionaries 185, 186
 - BlackPoint** entry 185
 - Gamma** entry 184, 185, 186, 440
 - Matrix** entry 185, 186, 440
 - WhitePoint** entry 185, 186
- CalRGB** color spaces 176, 181, 184–187, 272
 - as base color space 199
 - as blending color space 415
 - color values 184
 - gamma correction 185
 - and **ICCBased** color spaces, compared 189, 192
 - initial color value 216
 - process colors, conversion to 456
 - rendering 374
 - setting color values in 217
 - for soft masks 440
 - sRGB* color space, approximating 193
 - transfer functions 380
 - and transparent overprinting 461, 465
 - See also
 - CalRGB** color space dictionaries
- CanonicalFormat field flag (submit-form field) 553
- CapHeight** entry (font descriptor) 356
- Caption standard structure type 627, 630, 631
- captions 627, 630, 631
- carriage return (CR) character 26
 - in annotations 500, 503, 505, 506, 509
 - in cross-reference tables 65
 - as end-of-line marker 26, 31, 37, 62, 65
 - escape sequence for 30
 - in HTTP requests 674
 - in stream objects 36, 37
 - as white space 24, 32
- catalog
 - document. See document catalog
 - FDf (Forms Data Format). See FDF catalog
- Catalog** object type 83, 758, 760
- CCF** filter abbreviation 280, 789
- CCITT (Comité Consultatif International Téléphonique et Télégraphique)
 - encoding standard 52, 55
 - facsimile compression 15, 42, 52–55, 815
- CCITTFaxDecode** filter 42, 52–55
 - CCF** abbreviation 280, 789
 - end-of-facsimile-block (EOFB) pattern 54
 - parameters. See **CCITTFaxDecode** filter parameter dictionaries
 - return-to-control (RTC) pattern 54
 - in sampled images 268
- CCITTFaxDecode** filter parameter dictionaries 53–54
 - BlackIs1** entry 54
 - Columns** entry 54
 - DamagedRowsBeforeError** entry 53, 54
 - EncodedByteAlign** entry 53, 54
 - EndOfBlock** entry 54
 - EndOfLine** entry 54
 - K** entry 53, 54
 - Rows** entry 54
- CD-ROM, *PDF Reference* 3
- ceiling** operator (PostScript) 116, 703
- cells, halftone 383
 - coordinate system 384
 - frequency 384, 394
 - predefined spot functions 385
 - and spot function 384, 393
 - and threshold array 389–390
 - type 10 halftones 391, 395, 396, 397
 - type 16 halftones 391
- Center inline alignment 646
- Center text alignment 644
- CenterWindow** entry (viewer preferences dictionary) 472
- CFF (Compact Font Format) 5, 365, 367
- CGI (Common Gateway Interface) file format 664
- Ch** field type 531

- character classes
 - Alphabetic** 362, 363
 - AlphaNum** 363
 - Dingbats** 362, 363
 - DingbatsRot** 363
 - Generic** 362, 363
 - GenericRot** 363
 - Hangul** 363
 - Hanja** 363
 - Hanzi** 362
 - HKana** 363
 - HKanaRot** 363
 - HojoKanji** 363
 - HRoman** 362, 363
 - HRomanRot** 362, 363
 - Kana** 362, 363
 - Kanji** 363
 - Proportional** 362, 363
 - ProportionalRot** 362, 363
 - Ruby** 363
- character classes (CIDFonts) **361–364**
- character codes
 - 7-bit ASCII 15
 - character encodings, mapped by 323, 328
 - in CIDFonts 338
 - and CMaps 335, 336, 342, 343, 348, 349, 352, **354**
 - codespace ranges **351, 354**, 355, 369
 - Differences** array 330, 331
 - in font dictionaries 317, 324
 - hexadecimal, in name objects **33**, 124
 - mapping operators 351–352
 - multiple-byte 122, 303, 312, 322, 335, 336
 - notdef mappings 352, 354, 355
 - octal, in literal strings 30, **31**
 - and predefined CMaps 347
 - Shift-JIS encoding 348
 - showing text 292
 - single-byte 303, 312, 316
 - and standard encodings 709
 - in text objects 291
 - and text-showing operators 312
 - and **Tj** operator 294
 - in TrueType fonts 332, 333, 334, 367
 - in Type 0 fonts 353
 - in Type 1 fonts 331, 332
 - in Type 3 fonts 324, 332
 - undefined characters 355
 - undefined widths 357
 - Unicode, mapping to 318, 324, 353, 368–372, 620–621
 - and word spacing 303
- character collections **335**
 - Adobe-CNS1 344, 346, 362, 369, 621, 796
 - Adobe-GB1 344, 346, 362, 369, 621, 796
 - Adobe-Japan1 345, 346–347, 362, 363, 369, 621, 796
- character collections (*continued*)
 - Adobe-Japan2 362, 363
 - Adobe-Korea1 345, 347, 362, 363, 369, 621, 796
 - Chinese (simplified) 346
 - Chinese (traditional) 346
 - for CIDFonts 338, 339, 361, 362
 - CIDSystemInfo** dictionaries 336–337, 349
 - CJK (Chinese, Japanese, and Korean) 346–347
 - for CMaps 336, 343
 - and Identity–H predefined CMap 345
 - and Identity–V predefined CMap 345
 - Japanese 346–347
 - Korean 347
 - ordering **337**, 347, 621
 - for predefined CMaps 346–347, 621
 - registry 621
 - registry identifier **337**, 347
 - supplement number **337**, 347, 369
- character encodings 3, 314, 316, **328–334**, 816
 - base 330, 358, 710
 - Big Five 344, 788, 803
 - built-in. *See* built-in character encodings
 - CJK (Chinese, Japanese, and Korean) 322
 - and CMaps, compared 335–336, 342
 - for composite fonts. *See* CMaps
 - content extraction 16
 - EUC-CN 343
 - EUC-JP 344
 - EUC-KR 345
 - EUC-TW 344
 - for FDF fields 562
 - GBK 343, 803
 - glyph selection 312
 - Hong Kong SCS 344
 - ISO-2022-JP 345
 - ISO Latin 1 98
 - Microsoft Unicode 333
 - for name objects 788
 - named 368
 - natural language 361
 - predefined. *See* predefined character encodings
 - Shift-JIS 336, 344, 348, 788, 803
 - for simple fonts **328–334**
 - and **ToUnicode** CMaps 369
 - for TrueType fonts 332–334, 367, 710
 - for Type 1 fonts 318, 331–332, 710
 - for Type 3 fonts 323, 332
 - UCS-2 343, 344, 345
 - UHC (Unified Hangul Code) 345, 803
 - Unicode, mapping to 621
 - Unicode. *See* Unicode[®] character encoding
 - UTF-8 34, 788
 - UTF-16 98

- character identifiers (CIDs) **335**
 - and character collections 335, 337, 347
 - CID 0 340, 355
 - and CIDFonts 336, 337, 339, 340, 341, 342, 361, 369
 - and CMaps 335, 336, 342, 343, 349, 352, 354
 - and Identity–H predefined CMap 345
 - and Identity–V predefined CMap 345
 - mapping to glyph indices 339
 - maximum value 335, **706**
 - notdef mappings 352, 354, 355
 - and Type 0 fonts 353
 - undefined characters 355
 - Unicode conversion 368, 621
- character names
 - CharProcs** dictionary 324, 325
 - CID-keyed fonts, unused in 335
 - and CMaps 336, 342, 349, 352
 - Differences** array 330
 - in font subsets 357, 796
 - glyph descriptions, TrueType 332, 333, 334
 - glyph descriptions, Type 1 331, 332
 - glyph descriptions, Type 3 323, 332
 - .notdef **331**, 332, 340, 355, 357
 - and standard encodings 709
 - in Type 1 fonts 368
 - in Type 3 fonts 324
 - Unicode conversion 368, 620
- character selectors **336**
 - in CIDFonts 337
 - and CMaps 342, 348, 349, **354**
 - undefined characters 355
- character sequences
 - double left angle bracket (<<) **35**, 67, 560
 - double period (..) 120, 664
 - double right angle bracket (>>) **35**, 67, 560
 - tilde, right angle bracket (~>) **44**, 45
- character sets 3, 16, 292, 816
 - ASCII 25, 26, 31
 - Big Five 344
 - and character collections 335
 - CJK (Chinese, Japanese, and Korean) 334, 335
 - CNS 11643-1992 344
 - encodings for 329
 - ETen 344
 - for font subsets 357
 - Fujitsu FMR 344
 - GB 2312-80 343
 - GB 18030-2000 344
 - GBK 343
 - Hong Kong SCS 344
 - JIS C 6226 344
 - JIS X 0208 344, 345
 - JIS78 344
 - KanjiTalk6 344
 - character sets (*continued*)
 - KanjiTalk7 344
 - KS X 1001:1992 345
 - Latin. *See* Latin character set, standard
 - Mac OS KH 345
 - non-Latin 329, 540, 542
 - PDF **25–26**
 - Unicode, conversion to 368
- character spacing (T_c) parameter 298, 301, **302–303**
 - and horizontal scaling 304
 - and quotation mark (") operator 311
 - Tc** operator 302, 701, 702
 - text matrix, updating of 313–314
- characters **292**
 - accented 329, 356
 - apostrophe (') **100**, 134, 302, 305, **311**, 702
 - backslash (\) 29, **30–31**, 119, 120, 122, 312, 521, 666–667
 - backspace (BS) 30
 - bullet 714, 796
 - and bytes 25
 - carriage return (CR) 24, **26**, 30, 31, 32, 36, 37, 62, 65, 500, 503, 505, 506, 509, 674
 - codes. *See* character codes
 - colon (:) 120, 675
 - currency 714
 - Cyrillic 362, 363
 - delimiter 25, **26**, 27, 32, 33
 - dollar sign (\$) 343
 - ellipsis (...) 782
 - em dash 623
 - encodings. *See* character encodings
 - escape 30
 - euro 714
 - exclamation point (!) 44, **45**
 - form feed (FF) **26**, 30, 32
 - glyphs. *See* glyphs, character
 - Greek 362, 363
 - hangul 363
 - hanzi (kanji, hanja) 361, 362, 363
 - horizontal tab (HT) 24, **26**, 27, 30, 32
 - hyphen (-) 617, 623, 714
 - illuminated 651, 658
 - jamo 363
 - kana (katakana, hiragana) 362, 363
 - Latin 361, 362, 363
 - left angle bracket (<) 26, 29, **32**, **35**, 67, 122, 560
 - left brace ({}), **26**, **116**
 - left bracket ([]) 26, **34**, 808
 - left parenthesis (()) 26, **29**, 30, 312
 - ligatures 329, 371, 651, 658, 709
 - line feed (LF) 24, **26**, 30, 31, 32, 36, 37, 62, 65, 674
 - line-drawing 362, 363
 - minus sign (–) 100

- characters (*continued*)
- names. *See* character names
 - newline 26, 30
 - nonbreaking space 714
 - nonprinting 30, 31
 - null (NUL) 26, 666
 - number sign (#) 33, 34, 322, 664, 787–788
 - numeric 363
 - percent sign (%) 26, 27, 664
 - period (.) 120, 121, 533, 548, 664, 666, 801
 - plus sign (+) 100, 323
 - quotation mark (") 134, 302, 305, 311, 702
 - regular 25, 26, 27, 32
 - right angle bracket (>) 26, 29, 32, 35, 44, 67, 122, 560
 - right brace (}) 26, 116
 - right bracket (]) 26, 34
 - right parenthesis (]) 26, 29, 30, 312
 - ruby 363
 - selectors. *See* character selectors
 - sets. *See* character sets
 - slash (/) 26, 32, 34, 83, 94, 118, 121, 122, 357, 561, 673
 - space (SP) 24, 26, 27, 32, 64, 65, 303, 306, 320, 321, 552, 619–620, 623, 758, 787
 - special symbols 362, 363
 - tab. *See* horizontal tab (HT)
 - underscore (_) 121, 320
 - white-space 24, 25–26, 32, 33, 44, 279, 808
 - yuan symbol (¥) 343
- CharProcs** entry (Type 3 font dictionary) 324, 325, 326, 332
- CharSet** entry (font descriptor) 357, 368
- checkbox field dictionaries
- Opt** entry 540
- checkbox fields 537, 538, 539–540
- normal caption 536
 - Off appearance state 539
 - as radio buttons 541
 - value 539, 540
 - Yes appearance state 539
 - See also*
 - checkbox field dictionaries
- CheckSum** entry (embedded file parameter dictionary) 125
- Chinese
- character collections (simplified) 346
 - character collections (traditional) 346
 - character sets 334, 335
 - CMaps (simplified) 343–344
 - CMaps (traditional) 344
 - fonts 322
 - glyph widths 340
 - hanzi (kanji, hanja) characters 361, 362, 363
 - R2L reading order 472
 - writing systems 642
- choice field dictionaries 546
- I** entry 546
 - Opt** entry 546, 547, 802
 - TI** entry 546
- choice field flags 546
- Combo 546
 - DoNotSpellCheck 546
 - Edit 546
 - MultiSelect 546, 547
 - Sort 546
- choice fields 531, 538, 545–547, 802
- default appearance string 546, 547
 - flags. *See* choice field flags
 - multiple selection 546, 547
 - value 545, 546, 547
 - See also*
 - choice field dictionaries
 - combo box fields
 - list box fields
- chroma-key 277
- chromaticity 186, 415, 450
- chrominance 60
- CICI.SignIt** signature handler 549
- CID-Keyed Font Technology Overview* (Adobe Technical Note #5092) 335, 812
- CID-keyed fonts 5, 335–336, 709
- character collections 335
 - DescendantFonts** array 336
 - embedded 16
 - Encoding** entry 336
 - glyph descriptions 335, 336
 - as Type 0 fonts 336
 - See also*
 - CIDFonts
 - CMaps
- CIDFont dictionaries 314, 338–339
- BaseFont** entry 338, 353, 356
 - in CID-keyed fonts 336
 - CIDSystemInfo** entry 336, 337, 338, 339, 348
 - CIDToGIDMap** entry 339, 345
 - DW** entry 338, 340
 - DW2** entry 338, 341, 368
 - FontDescriptor** entry 338
 - metadata inapplicable to 579
 - Subtype** entry 338
 - and TrueType fonts 367
 - in Type 0 fonts 353
 - Type** entry 338
 - W** entry 338, 340–341
 - W2** entry 338, 341–342, 368
- CIDFont **FD** dictionaries 361–364
- CIDFont files 336

- CIDFont font descriptors 338, **360–364**
 - CIDSet** entry 361, 368
 - FD** entry 361
 - Lang** entry 361
 - Style** entry 361
 - See also*
 - CIDFont **FD** dictionaries
 - CIDFont **Style** dictionaries
- CIDFont **Style** dictionaries **360–361**
 - Panose** entry 360
- CIDFont subtypes
 - See* CIDFont types
- CIDFont types 338
 - CIDFontType0** 315, **338**, 365
 - CIDFontType2** 315, **338**, 365
- CIDFontName** entry (CIDFont program) 338
- CIDFonts 314, 315, 334, **337–342**
 - base font 338
 - character classes **361–364**
 - character collection 338, 339, 361, 362
 - CIDFont files **336**
 - and CMaps 338, 342, 343, 349, 352, 354
 - embedded 336, 339
 - font descriptors. *See* CIDFont font descriptors
 - and fonts, compared 338
 - glyph descriptions 337, 339
 - glyph indices, mapping from CIDs to 339
 - glyph metrics 338, 340–342, 362, 795
 - glyph selection 339–340
 - glyph widths 338
 - and Identity–H predefined CMap 345
 - and Identity–V predefined CMap 345
 - PostScript name 338
 - subsets 361
 - Tf** operator inapplicable to 338
 - Type 0 337, 338, 339, 353
 - Type 0 fonts, descendants of 338, 353, 355, 369
 - Type 2 337, 338, 339–340, 345, 353
 - Unicode mapping 621
 - vertical writing 338, 341–342, 362, 363
 - writing mode 343
 - See also*
 - CIDFont dictionaries
 - CIDFont **FD** dictionaries
 - CIDFont **Style** dictionaries
 - CIDFont types
- CIDFontType0** CIDFont type 315, **338**, 365
- CIDFontType0C** compact font subtype 365, 366
- CIDFontType2** CIDFont type 315, **338**, 365
- CIDs. *See* character identifiers
- CIDSet** entry (CIDFont font descriptor) 361, 368
- CIDSystemInfo** dictionaries 336–337, 339, 349, 361, 362, 369, 621, 796
 - Ordering** entry 337, 345, 362
 - Registry** entry 337, 345, 362
 - Supplement** entry 337, 345, 362
- CIDSystemInfo** entry
 - CIDFont dictionary 336, 337, **338**, 339, 348
 - CMap dictionary 336, 337, **349**, 796
- CIDToGIDMap** entry (CIDFont dictionary) 339, 345
- CIE (Commission Internationale de l'Éclairage) 176
- CIE 1931 XYZ color space
 - and **CalGray** color spaces 182, 183, 184
 - and **CalRGB** color spaces 185
 - CIE-based color spaces, semantics of 181
 - implicit conversion, bypassed in 196
 - and **Lab** color spaces 188
 - soft masks, derivation of 440
- CIE 1976 $L^*a^*b^*$ color space 60, 187, 189
- CIE-based A color spaces **182–183**
 - color values 182
 - decoding functions 183, 184
- CIE-based ABC color spaces **181–182**, 184, 187, 199
 - color values 181
 - decoding functions 181, 184, 186, 188
- CIE-based color spaces **176**, **181–198**
 - as alternate color space 204
 - as base color space 199
 - blending in 416, 455
 - CalCMYK** 181
 - CIE 1931 XYZ. *See* CIE 1931 XYZ color space
 - CIE 1976 $L^*a^*b^*$ 60, 187, 189
 - CIE-based A **182–183**
 - CIE-based ABC **181–182**, 184, 187, 199
 - color conversion, control of 376
 - color mapping mapping function 375
 - and color specification 172
 - decoding functions 181, 183, 184, 186, 188
 - default 178, **194–195**, 456, 686
 - device spaces, conversion to 236, 373, **374–375**
 - diffuse achromatic highlight 185
 - diffuse achromatic shadow 185
 - diffuse white point 183, 184, 185, 188
 - and flattening of transparent content 470
 - gamut mapping function 185, **375**, 380
 - as group color space 455, 456
 - implicit conversion to device colors 195–197
 - initial color value 182
 - inline images, prohibited in 280
 - and overprinting 214
 - for page group 437
 - parameters 182
 - process colors, rendered as 376
 - rendering intents. *See* rendering intents

- CIE-based color spaces (*continued*)
 - setting color values in 217
 - for shadings 234, 236
 - for soft masks 440
 - specification 182
 - specular highlight 185
 - sRGB* (standard *RGB*) 192–193, 456
 - for transparency groups 178, 450
 - See also*
 - CalGray** color spaces
 - CalRGB** color spaces
 - ICCBased** color spaces
 - Lab** color spaces
- CIE colorimetric system 181, 814
- CIP4 (International Cooperation for the Integration of Processes in Prepress, Press and Postpress)
 - JDF Specification* 374, 689, **814**
- circle annotation dictionaries **505**
 - BS** entry **505**
 - Contents** entry **505**
 - IC** entry **505**
 - Subtype** entry **505**
- Circle annotation type 499, **505**
- circle annotations 499, **504–505**
 - border style 490, 495
 - border width 505
 - contents 505
 - dash pattern 505
 - interior color 505
 - See also*
 - circle annotation dictionaries
- Circle line ending style **504**
- Circle list numbering style **650**
- CJK (Chinese, Japanese, and Korean)
 - character collections 346–347
 - character sets 334, 335
 - CMaps **343–345**
 - fonts 322
 - glyph widths 340
 - See also*
 - Chinese
 - Japanese
 - Korean
- CJKV Information Processing* (Lunde) 347, **816**
- class map 590, **606**, 638
 - metadata inapplicable to 579
- ClassMap** entry (structure tree root) **590**, 606, 638
- clear-table marker (LZW compression) 46, 48
- cleartomark** operator (PostScript) 366
- Client-Side JavaScript Reference* (Netscape Communications Corporation) 556, **816**
- Clip marked-content tag 637
- clip** operator (PostScript) 702
- clipping 11, **12**, **171–172**, **305–306**
 - to annotation rectangle 496
 - to art box 679
 - to bleed box 88, 677, 679
 - to crop box 88, 677, 679
 - even-odd rule 172, 702
 - to form bounding box 283, 284
 - to function domain 108, 110, 271
 - to function range 108, 111
 - to function sample table 110
 - to glyph outlines 296, 305
 - in illustration elements (Tagged PDF) 637
 - and marked content **585–588**
 - nonzero winding number rule 172, 306, 702
 - to page boundaries 88, 473, 677, 679
 - paths 131, 132, **171–172**, 305, 306
 - pattern cells 222
 - to reference XObject bounding box 288
 - scan conversion 406
 - shadings 234
 - soft 159, **412**, 421, **439**, 444
 - text rendering mode 305–306
 - to transparency group bounding box 452
 - See also*
 - clipping path operators
 - current clipping path
- clipping objects **585–587**
- clipping path, current
 - See* current clipping path
- clipping path operators **134**, 162, **171–172**
 - W** 134, 169, 171, **172**, 585, 586, 702
 - W*** 134, 169, 171, **172**, 585, 586, 702
- ClosedArrow line ending style **504**
- closepath** operator (PostScript) 699, 700, 701
- ClrF** entry (FDF field dictionary) **565**
- ClrFf** entry (FDF field dictionary) **564**
- cm** operator 134, 139, 148, **156**, 266, 699
- CMap dictionaries 348, **349**
 - in CID-keyed fonts 336
 - CIDSystemInfo** entry 336, 337, **349**, 796
 - CMapName** entry **349**, 353
 - for **ToUnicode** CMaps 369
 - Type** entry **349**
 - UseCMap** entry **349**, **351**, 369
 - WMode** entry 341, **349**
- CMap files **335–336**, 343
 - CIDSystemInfo** dictionary 349
 - example 348, 350–351
 - name 349
 - ToUnicode** CMaps 318, 324, 353, 369
 - writing mode 349
- CMap object type **349**
- “cmap” table (TrueType font) 332, 333–334, 339, 340, 367
- CMapName** entry (CMap dictionary) **349**, 353

- CMaps **312**, 314, 315, 334, **342–352**, 796, 816
 base CMap 349
 and character collections 336, 343
 and character encodings, compared 335–336, 342
 Chinese (simplified) **343–344**
 Chinese (traditional) **344**
 and CIDFonts 338, 342, 343, 349, 352, 354
 CJK (Chinese, Japanese, and Korean) **343–345**
 embedded 336, 347, **348**, 796
 example 348, 350–351
 files. *See* CMap files
 font numbers 336, **342**, 348, 349, 352, 353, 354
 and fonts 342, 349
 Identity–H 345, 368, 621
 Identity–V 345, 368, 621
 Japanese **344–345**
 Korean **345**
 mapping operators 351–352
 notdef mappings 352, 354, 355
 PostScript name 349, 353
 predefined. *See* predefined CMaps
ToUnicode. *See* **ToUnicode** CMaps
 for Type 0 fonts 353, 796
 undefined characters 355
 Unicode mapping 621
 writing mode 343, 349
See also
 CMap dictionaries
- CMS. *See* color management system
- CMYK color representation
 calibrated, as blending color space 415
DCTDecode filter, transformation by 60
DeviceCMYK color space 176, 180
 and grayscale, conversion between 377, 382
 in halftones 383
 and high-fidelity color, compared 205
 in output devices 172, 376
 and output intents 686
RGB, conversion from 150, 377–379, 469
RGB, conversion to 380
 for subtractive color 178
 in transfer functions 380
- CMYK** color space abbreviation (inline image object) **280**
- CNS 11643-1992 character set 344
- CNS–EUC–H predefined CMap **344**, 346
- CNS–EUC–V predefined CMap **344**, 346
- CO** entry
 interactive form dictionary 516, **529**
 sound object **569**
- code, computer program 634
- Code standard structure type **634**
- Codes for the Representation of Names of Countries and Their Subdivisions* (ISO 3166) 99, 100, 653, **814**
- Codes for the Representation of Names of Languages* (ISO 639) 99, 100, 653, **814**
- codespace ranges **351**, **354**, 355
 for **ToUnicode** CMaps 369
- collaborative editing xx, 553
- colon (:) character
 in conversion engine names 675
 as DOS file name delimiter 120
 as Mac OS file name delimiter 120
- color
 annotations 491
 backdrop. *See* backdrop color
 background, dynamic appearance stream 536
 border, dynamic appearance stream 536
 calibrated. *See* calibrated color
 conversion between spaces. *See* color conversion
 current. *See* current color
 duotone **205**, 209–211
 glyph descriptions 326
 gradient fills 151
 group. *See* group color
 group backdrop 429
 high-fidelity 172, 176, **205**
 ICC profiles 181, **189–192**
 interior
 annotations 503, 505
 line endings 504
 inversion 272, 381
 mapping 172, 176, 199
 masking. *See* color key masking
 multitone 172, 176, **205**, 209–213, 792–793
 object (transparent imaging model) 430
 outline items 479
 overprint control **213–216**
 page boundaries 681
 process. *See* process colors
 quadtone **205**, 212–213
 remapping 178, **194–195**, 376, 450, 456, 686
 rendering 173, 373–374
 result (transparent imaging model). *See* result color
 separations. *See* separations, color
 smoothness tolerance 405
 source (transparent imaging model). *See* source color
 specification 172
 tints. *See* tints
YUV 60
YUVK 60
See also
 color components
 color operators
 color representation
 color spaces
 color values
 colorants

Color Appearance Models (Fairchild) 813

color bars 5, 676, 677, 680, 683

as page artifacts 615

as printer's mark annotations 512

Color blend mode 419

color components 173

additive 179, 195, 202, 381, 383

alternate color space 204, 207

DeviceCMYK 202

DeviceGray 202

DeviceN (tints) 205, 206, 207, 706

DeviceRGB 202

halftones for 382, 383, 391, 400, 401, 402

linear 455

nonprimary 381, 393, 395, 398, 400, 401, 402

and nonseparable blend modes 418

nonstandard 393, 395, 398, 400, 401

and output intents 686

and overprinting **464–465**

primary 179, 401

range 415

and separable blend modes **416**

Separation (tints) 202

smoothness tolerance 405

in soft-mask images 447, 448, 449

spot 205, 381, 447, 457, 458

subtractive 180, 195, 201, 202, 381, 383

transfer functions 373, 381

and transparent overprinting 458, 459, 460, 461, **464–465**, 466

See also

black color component

blue color component

cyan color component

gray color component

green color component

magenta color component

red color component

yellow color component

color conversion 374–380

to alternate color space 236, 458

to base color space 236

CIE-based to device 236, 456

CMYK to RGB 380

device color spaces, among 236, 373, **376–380**

device to CIE-based not generally possible 455, 456

grayscale and *CMYK*, between 377, 382

grayscale and *RGB*, between 377

to group color space 450, 452, 454, 455, 456

to page color space 455

in rendering 173, 466

RGB to CMYK 150, 377–379, 469

in shading patterns 236, 458

soft-mask images, preblending of 449

color conversion (continued)

See also

black-generation function

undercolor-removal function

Color entry (version 1.3 OPI dictionary) 696**color functions 235**

type 1 (function-based) shadings 237

type 2 (axial) shadings 238, 239

type 3 (radial) shadings 240, 241

type 4 shadings (free-form Gouraud-shaded triangle meshes) 244, 247

type 5 shadings (lattice-form Gouraud-shaded triangle meshes) 249

type 6 shadings (Coons patch meshes) 253, 255

color key masking 268, 275, 277–278, 793

and object shape 443

and soft masks 444

color management system (CMS) 375

color mapping (**Indexed** color spaces) 172, 176, 199

color mapping functions, CIE-based 375

color operators 134, 156, 216–218

CS 134, 173, 177, 179, 180, **216**, 218, 220, 699

cs 134, 173, 177, 179, 180, **217**, 218, 220, 699

G 134, 173, 177, 178, 179, **217**, 218, 700

g 134, 173, 177, 178, 179, **217**, 218, 264, 295, 700

in glyph descriptions 326

K 134, 173, 177, 178, 180, **218**, 700

k 134, 173, 177, 178, 180, **218**, 700

restrictions on 218

RG 134, 173, 177, 178, 180, **217**, 218, 701

rg 134, 173, 177, 178, 180, **217**, 218, 456, 461, 701

SC 134, 173, 177, 179, 180, 201, **217**, 218, 701

sc 134, 173, 177, 179, 180, 201, **217**, 218, 247, 264, 701

SCN 134, 177, 201, 202, 206, **217**, 218, 220, 224, 227, 701

scn 134, 177, 201, 202, 206, **217**, 218, 220, 224, 227, 228, 701

text, showing 295

in text objects **309**

color patches

bicubic tensor-product 256, **257–259**

Coons **250–252**, 256, 257, 259

color plates 3

Plate 1, *Additive and subtractive color* 178

Plate 2, *Uncalibrated color* 181

Plate 3, *Lab color space* 187

Plate 4, *Color gamuts* 187

Plate 5, *Rendering intents* 197

Plate 6, *Duotone image* 205

Plate 7, *Quadtone image* 205, 212

Plate 8, *Colored tiling pattern* 224

Plate 9, *Uncolored tiling pattern* 228

Plate 10, *Axial shading* 239

- color plates (*continued*)
 - Plate 11, *Radial shadings depicting a cone* 241, 242
 - Plate 12, *Radial shadings depicting a sphere* 241
 - Plate 13, *Radial shadings with extension* 242
 - Plate 14, *Radial shading effect* 242
 - Plate 15, *Coons patch mesh* 250
 - Plate 16, *Transparency groups* 411
 - Plate 17, *Isolated and knockout groups* 433, 434
 - Plate 18, *RGB blend modes* 416
 - Plate 19, *CMYK blend modes* 416
 - Plate 20, *Blending and overprinting* 462
- color profiles, ICC
 - See ICC color profiles
- color representation
 - ICC profiles 181, **189–192**
 - YUV 60
 - YUVK 60
 - See also
 - additive color representation
 - CMYK color representation
 - grayscale color representation
 - RGB color representation
 - subtractive color representation
- Color Separation Conventions for PostScript Language Programs* (Adobe Technical Note #5044) 792, **812**
- color separations
 - See separations, color
- color space arrays 177
 - for CIE-based color spaces 182
 - as **ColorSpace** resources 177, 216
 - content streams, prohibited in 177
 - for **DeviceN** color spaces 206
 - for **ICCBased** color spaces 189
 - for **Indexed** color spaces 199
 - for **Pattern** color spaces 227
 - in PDF objects 177
 - for **Separation** color spaces 202, 207
- color spaces 172–218
 - abbreviations for, in inline images 279–280
 - additive. See additive color representation
 - alternate. See alternate color space
 - arrays. See color space arrays
 - blending. See blending color space
 - CalCMYK** 181
 - CIE 1931 XYZ. See CIE 1931 XYZ color space
 - CIE 1976 $L^*a^*b^*$ 60, 187, 189
 - CIE-based A **182–183**
 - CIE-based ABC **181–182**, 184, 187, 199
 - conversion between. See color conversion
 - current. See current color space
 - Decode** arrays, default 272
 - diffuse achromatic highlight 185
 - diffuse achromatic shadow 185
 - diffuse black point 183, 185, 188
 - diffuse white point 183, 184, 185, 188
 - families **176–177**, 792
 - gamma correction 183, 185
 - gamut 187, 205, **375**, 468, 469
 - group. See group color space
 - for image XObjects 177, 194, 480
 - implicit conversion 195–197
 - for inline images 194, 280
 - in Linearized PDF 748
 - as named resources 97
 - native (output device). See native color space
 - and overprinting **464–465**
 - for page group 437, 452, 455, 466
 - rendering intents. See rendering intents
 - for sampled images 262, 263, 264, 268, 269, 271, 277
 - for separable blend modes 416
 - for shadings 194, 234, 235–236, 237, 238, 240, 244, 247, 249, 253
 - for soft masks 440, **446**, 448, 449
 - specification 177
 - specular highlight 185
 - sRGB (standard RGB) 192–193, 456
 - subtractive. See subtractive color representation
 - for thumbnail images 480
 - for transparency groups. See group color space
 - and transparent overprinting 458, **464–465**
 - for variable-text fields 534
 - See also
 - CalGray** color spaces
 - CalRGB** color spaces
 - CIE-based color spaces
 - default color spaces
 - device color spaces
 - DeviceCMYK** color space
 - DeviceGray** color space
 - DeviceN** color spaces
 - DeviceRGB** color space
 - ICCBased** color spaces
 - Indexed** color spaces
 - Lab** color spaces
 - Pattern** color spaces
 - Separation** color spaces
 - special color spaces
- color table (**Indexed** color space) 199, **200**, 449
- color values 148, **173**
 - background (shadings) 234, 238, 239, 242
 - CalGray** 183
 - CalRGB** 184
 - CIE-based A 182
 - CIE-based ABC 181
 - CIE-based color mapping 375
 - components 173, 373
 - DeviceCMYK** 180

color values (*continued*)**DeviceGray** 179**DeviceN** 206**DeviceRGB** 179**Indexed** 199

interpolation (shadings) 235–236

Lab 187**Pattern** 220

remapping 195

Separation 202

transfer functions, produced by 383

in transparent imaging model 413

for uncolored tiling patterns 227

colorants

additive 201

device 178, 205, 214, 458, 459, 465, 684

DeviceN 206, 207, 706

halftones for 382, 391, 392, 400, 401, 402

misregistration 513, 676, **688**

for OPI proxies 698

primary **201**, 204, 383, 391

for printer's marks 683

process. *See* process colorants

and separations 202, 203, 204, 683

spot. *See* spot colorants

subtractive 201, 203

and transparent overprinting 458

for trap networks 692

See also

black colorant

blue colorant

cyan colorant

green colorant

magenta colorant

orange colorant

red colorant

yellow colorant

Colorants entry**DeviceN** color space attributes dictionary **207**printer's mark form dictionary **683****ColorBurn** blend mode **418****ColorDodge** blend mode **418**colored tiling patterns **221**, **223–227**

in transparent imaging model 454

Colors entry**FlateDecode** filter parameter dictionary **49****LZWDecode** filter parameter dictionary **49****ColorSpace** entryimage dictionary 177, **268**, 269, 271, 276, 448, 449, 461, 480inline image object **279**, **280**resource dictionary **97**, 177, 194, 216, 280, 450separation dictionary **684**shading dictionary 232, **234**, 235, 236**ColorSpace** resource type **97**, 177, 194, 216, 280, 450**ColorTransform** entry (**DCTDecode** filter parameter dictionary) **60****ColorType** entry (version 1.3 OPI dictionary) **696***Colour Measurement and Management in Multimedia Systems and Equipment* (International Electrotechnical Commission) 192, **814****ColSpan** standard structure attribute **651****Columns** entry**CCITTFaxDecode** filter parameter dictionary **54****FlateDecode** filter parameter dictionary **49****LZWDecode** filter parameter dictionary **49**

combo box fields 538, 545, 546

trigger events for 516

Combo field flag (choice field) **546**

command dictionaries, Web Capture

See Web Capture command dictionaries

command settings dictionaries, Web Capture

See Web Capture command settings dictionaries

Comment annotation icon 500

comments 25, 26, **27**, 787OPI. *See* OPI comments**Comments** entry (version 1.3 OPI dictionary) **695**

Commission Internationale de l'Éclairage (International Commission on Illumination) 176

Compact Font Format (CFF) 5, 365, 367

Compact Font Format Specification, The (Adobe Technical Note #5176) 5, 316, **813**

compact font programs

embedded 365, 367

subtypes 365

compact font subtypes 366

CIDFontType0C **365**, **366****Type1C** **365**, **366**

compatibility

blend modes for 419, 461

of file names 121

with other applications 62

of PDF versions. *See* version compatibility, PDFsections **95**, 699, 700

transparency 469–470

compatibility operators **95**, **134****BX** **95**, 134, 699**EX** **95**, 134, 700compatibility sections **95**, 699, 700**Compatible** blend mode **419**, 461

and fully opaque objects 467

CompatibleOverprint blend mode **460–462**, 465

and halftones 467–468

overprint mode ignored by 466

overprint parameter ignored by 466

and transfer functions 467–468

and transparency groups 466

- Composite (group compositing) function **428**
 - backdrop, compositing with 429
 - backdrop removal 432
 - in group compositing computations 430
 - for page group 436, 437
 - recursive application 430
 - soft masks, derivation of 439, 440
 - summary 438
- composite fonts **314, 334–355**
 - in CID-keyed fonts 336
 - descendants 353
 - encoding. *See* CMaps
 - glyph selection 312
 - PostScript and PDF, compared 334
 - Tj operator 294
 - Unicode mapping 621
 - word spacing 303
 - writing mode 299
 - See also*
 - CID-keyed fonts
 - Type 0 fonts
- composite pages **201**
 - separations, generation of 683
 - spot colorants in 457
- compositing 11, 133, 409, **410, 411**
 - of annotation appearances 496
 - blending color space **414–416, 442, 797**
 - computations. *See* compositing computations
 - in isolated groups 433, 451
 - of isolated groups 433
 - in knockout groups 434–435, 451
 - in non-isolated groups 451
 - in non-knockout groups 434
 - and overprinting 459
 - in page group 412, 437
 - of page group 436, 437
 - pattern cells 453
 - shading patterns 453
 - of spot color components 457–458
 - text knockout parameter 307–308
 - tiling patterns 453
 - in transparency groups 411, 412, 425, 427, 428, 440, 450, 452, 454
 - of transparency groups 411, 412, 425, 427, 428, 429, 445, 450, 451, 452, 454, 462, 463
 - See also*
 - alpha
 - blend modes
 - opacity
 - shape
- compositing computations
 - basic **412–424**
 - formula **413–414**
 - notation **413**
 - summary **424**
 - compositing computations (*continued*)
 - group **426–427, 428–432**
 - general groups **426–427**
 - isolated groups **433**
 - knockout groups **434–435**
 - non-isolated groups 436
 - non-isolated, non-knockout groups **430**
 - notation **426**
 - page group **436–437**
 - summary **438–439**
 - linear 455
 - for patterns 454
 - and preblending of soft-mask image data 449
 - simplification of 439
- “Compositing Digital Images” (Porter and Duff) 414, **816**
- compression, data 15, 41
 - CCITT facsimile 15, 42, **52–55, 815**
 - DCT (discrete cosine transform) 42, **59–61**
 - filters 15, 22, **45–61, 569**
 - Flate (zlib/deflate) 15, 42, **45–52**
 - JBIG2 (Joint Bi-Level Image Experts Group) 15, 42, **55–59**
 - JPEG (Joint Photographic Experts Group) 15, 42, 59
 - lossless **42, 55, 193**
 - lossy **42, 55, 59**
 - LZW (Lempel-Ziv-Welch) 15, 41, 42, 43, **45–52**
 - run-length encoding 42, **52**
 - sounds 569
- computation order 516, **529**
- Computer Graphics: Principles and Practice* (Foley et al.) **813**
- concat** operator (PostScript) 699
- Confidential annotation icon 507
- constant opacity 149, 159, **422, 442**
 - for annotations 491–492
 - notation **422, 427, 431**
 - specifying **444–445, 797**
 - in transparency groups 425
- constant shape 149, 159, **421, 442**
 - notation **422, 427, 431**
 - specifying **444–445, 797**
 - in transparency groups 425
- consumer applications, PDF **1**
 - decoding of data 41
 - embedded file streams, processing of 124
 - embedded fonts, copyright restrictions on 365
 - glyph widths, use of 794
 - logical structure, navigation of 589
 - logical structure, usage of 631
 - masked images, treatment of 793
 - output intents, use of 686
 - page tree, handling of 86
 - role map, processing of 593

- consumer applications, Tagged PDF
 - artifacts, treatment of 615–616
 - fragmented BLSEs, recognition of 629
 - hidden page elements, recognition of 617
 - hyphenation 617
 - ILSEs, line height for 647
 - layout 623
 - nonstandard structure types 626
 - page content order 618
 - placement attributes, treatment of negative values 643, 644
 - reverse-order show strings 619, 620
 - standard structure elements, processing of 613
 - text discontinuities, recognition of 617
 - Unicode mapping 621
 - word breaks, recognition of 623
- containing document (reference XObject) **287**
- content
 - extraction. *See* content extraction
 - importing 4, 287–289
 - interchange 9, 592, 605
 - merging *xx*
 - reflow. *See* reflow of content
- content database, Web Capture
 - See* Web Capture content database
- content extraction *xx*, 10, 573
 - access permission for 74, 75, 77
 - for accessibility to disabled users 651
 - from annotations 490, 501, 509, 510, 511, 512, 691
 - character encodings and 16
 - of character properties 620–622
 - of fonts 365
 - from form fields 531
 - of graphics 74, 75, 77, 136
 - lists, autonumbering of 650
 - in PostScript conversion 22
 - from structure elements 592
 - in Tagged PDF 612, 613, 628, 638
 - of text 16, 74, 75, 77, 313, 368, 638, 652
- content items (logical structure) **593–604**
 - annotations as 618
 - direct **625**, 626, 632
 - finding structure elements from 590, 600–604, 655
 - link annotations, association with 634
 - marked-content sequences as 285, 589, 590, 591, **593–598**, 599, 600, 601, 602, 603, 618
 - PDF objects as 589, 590, 591, 593, **598–599**, 601
 - structural parent tree 269, 285
 - structure elements as 589, 591, 593, **599–600**
 - structure elements, associated with 589, 591, 593
 - in Tagged PDF 627, 631
- content rectangle **626**, **648**
 - and allocation rectangle 649
 - in layout 645, 646
- content set subtypes (Web Capture) 668, 669
 - SIS 668**, **669**
 - SPS 668**
- content sets, Web Capture
 - See* Web Capture content sets
- content streams 9, **92–97**, 791
 - annotation appearances, defining 93, 614
 - application-specific data in 19
 - and basic layout model (Tagged PDF) 623
 - color space arrays prohibited in 177
 - color space, selection of 173
 - common programming language features, lack of 21
 - compatibility sections **95**, 699, 700
 - as component of PDF syntax 23–24
 - data syntax 281
 - external objects (XObjects) 261
 - filters, decoding with 94, 786
 - font characteristics 622
 - fonts 93, 293
 - form XObjects 93, 133, 281, 282, 283, 614, 615
 - glyph descriptions 324, 326
 - glyphs, painting 292
 - images, painting 262
 - and **Indexed** color spaces 199
 - indirect object references prohibited in 40
 - inline images 263, 278
 - in Linearized PDF 738, 743, 744, 745, 754, 808
 - and marked content. *See* marked content
 - marked-content sequences confined within 583
 - named resources in **95–96**
 - natural language specification 653, 654, 656
 - operands in 11, **94**, 132
 - operators in 11, **94**, 132, 699
 - at page description level 134
 - pages, describing contents of 89, 93, 614, 615, 757, 758, 760, 762, 786, 807
 - parent, of patterns **220**, 221, 223
 - and **Pattern** color spaces 199
 - patterns, defining 93, 221, 222, 223, 227, 231
 - PostScript, conversion to 22
 - PostScript language fragments 289, 290
 - printer's marks in 680
 - procedure sets 574
 - q** and **Q** operators in 152
 - and resources 92, 95–96, 285
 - as self-describing graphics objects 131, 261
 - shading patterns in 232
 - in structural parent tree 590, 601
 - text operators 309
 - text state parameters 301
 - transparency group XObjects 447, 451–452
 - trap network annotations 689
 - and trap networks 691
 - type 2 (shading) patterns, absent in 157

- content streams (*continued*)
 - uncolored tiling patterns 218
 - unrecognized filters in 789
- content types (Web Capture) 668, 672, 673, 675
- Contents** entry
 - annotation dictionary 490, 509, 657
 - circle annotation dictionary 505
 - free text annotation dictionary 502
 - ink annotation dictionary 508
 - line annotation dictionary 503
 - link annotation dictionary 501
 - markup annotation dictionary 506
 - movie annotation dictionary 511
 - page object 89, 96, 152, 583, 691, 737, 738, 757, 791
 - rubber stamp annotation dictionary 507
 - signature dictionary 62, 549
 - sound annotation dictionary 510
 - square annotation dictionary 505
 - text annotation dictionary 500
 - trap network annotation dictionary 691
 - widget annotation dictionary 512
- continuous-tone reproduction 50, 59, 373, 376, 382
- controller bars (movies) 572
- conversion engines (Web Capture) 674, 675
- Coons patch meshes
 - See type 6 shadings
- Coons patches 250–252, 256, 257, 259
- coordinate spaces
 - See coordinate systems
- coordinate systems 136–147
 - for appearance streams 496
 - coordinate spaces 137–141
 - relationships among 141
 - for soft masks 446
 - for type 1 (function-based) shadings 237
 - See also
 - device space
 - form space
 - glyph space
 - image space
 - pattern space
 - target coordinate space
 - text space
 - transformation matrices
 - user space
- coordinate transformations 139, 141–144
 - cm** operator 139
 - combining 143–144, 146–147
 - on glyphs 291
 - inverting 147
 - reflection 137, 266, 695
 - rotation. See rotation
 - on sampled images 265
- coordinate transformations (*continued*)
 - scaling. See scaling
 - skewing. See skewing
 - translation. See translation
- Coords** entry
 - type 2 shading dictionary 238
 - type 3 shading dictionary 240
- copy** operator (PostScript) 116, 704
- copyright 6
 - permission 6–7
- cos** operator (PostScript) 116, 703
- CosineDot** predefined spot function 386
- Count** entry
 - outline dictionary 478
 - outline item dictionary 478, 740
 - page tree node 86
- country codes (ISO 3166) 99, 100, 653
- Courier standard font 319, 795
- Courier typeface 16, 709
- Courier–Bold standard font 319, 795
- Courier,Bold standard font name 795
- Courier,BoldItalic standard font name 795
- Courier–BoldOblique standard font 319, 795
- Courier,Italic standard font name 795
- Courier–Oblique standard font 319, 795
- CourierNew standard font name 795
- CourierNew–Bold standard font name 795
- CourierNew,Bold standard font name 795
- CourierNew–BoldItalic standard font name 795
- CourierNew,BoldItalic standard font name 795
- CourierNew–Italic standard font name 795
- CourierNew,Italic standard font name 795
- CourierNewPS–BoldItalicMT standard font name 795
- CourierNewPS–BoldMT standard font name 795
- CourierNewPS–ItalicMT standard font name 795
- CourierNewPSMT standard font name 795
- CP** entry (sound object) 569
- Create Thumbnails command (Acrobat) 789, 807
- creation date
 - document 573, 575, 576
 - Web Capture content set 668, 671
- CreationDate** entry
 - document information dictionary 576
 - embedded file parameter dictionary 125
- Creator** entry
 - document information dictionary 576
 - Mac OS file information dictionary 125
- creator signature (Mac OS) 125

- crop box 677, 678
 - and attached artifacts 616
 - and bounding box 288, 474–475
 - clipping to 679
 - display of 5, 681
 - in page imposition 806
 - in page object 88
 - printer's marks excluded from 680
 - in printing 679
- CropBox** entry
 - box color information dictionary 681
 - page object 88, 89, 138, 139, 473, 677
- CropFixed** entry (version 1.3 OPI dictionary) 695
- CropRect** entry
 - version 1.3 OPI dictionary 695
 - version 2.0 OPI dictionary 698
- Cross** predefined spot function 388
- cross-reference sections 64
 - byte offset 67, 68, 69, 780
 - example 757, 774, 775, 777, 779, 780
 - and incremental updates 64, 66, 69
 - length count 68
 - in Linearized PDF 729
 - object numbers in 66
- cross-reference table 17, 61, 64–67, 779, 780
 - entries 64–66, 68, 69, 758, 774, 779, 780
 - in FDF files 558
 - and file trailer 67, 68
 - free entries 65–66, 740, 778, 779
 - in-use entries 65, 740
 - incremental updates 18
 - in Linearized PDF 729, 732, 734, 735, 740–741, 742
 - reconstruction 707
 - sections. *See* cross-reference sections
 - subsections 64, 740, 777
- CS** entry
 - inline image object 279, 280
 - transparency group attributes dictionary 446, 450, 452
- CS** operator 134, 177, 216, 218, 699
 - in content streams 173, 177
 - for **DeviceCMYK** color space 180
 - for **DeviceGray** color space 179
 - for **DeviceRGB** color space 180
 - for **Pattern** color space 220
- cs** operator 134, 177, 217, 218, 699
 - in content streams 173, 177
 - for **DeviceCMYK** color space 180
 - for **DeviceGray** color space 179
 - for **DeviceRGB** color space 180
 - for **Pattern** color space 220
- CSS (Cascading Style Sheets) file format 622, 624
 - standard attribute owner 639
- CSS-1.00 standard attribute owner 639
- CSS-2.00 standard attribute owner 639
- CT** entry
 - Web Capture command dictionary 672, 673–674
 - Web Capture content set 668
- CTM. *See* current transformation matrix
- cubic Bézier curves 164–166, 813, 814, 815
 - control points 163, 164, 165, 233, 253–254, 256–259
 - example 762
 - path construction 163, 699, 701, 702
 - in path objects 132
 - type 6 shadings (Coons patch meshes) 233, 250, 254
 - type 7 shadings (tensor-product patch meshes) 257–259
- cubic spline interpolation 110, 113
- currency character 714
- current alpha constant 149, 444–445
 - and alpha source parameter 159
 - for annotations 491, 496
 - CA** entry (graphics state parameter dictionary) 159
 - ca** entry (graphics state parameter dictionary) 159
 - current color, analogous to 444
 - and fully opaque objects 467
 - ignored by older viewer applications 797
 - initialization 452, 453
 - multiple objects, applied to 444, 445
 - nonstroking. *See* nonstroking alpha constant
 - and overprinting 462, 463
 - setting 444
 - soft-mask images, unaffected by 268
 - stroking. *See* stroking alpha constant
 - and transparency groups 444
- current blend mode 149, 442
 - BM** entry (graphics state parameter dictionary) 159, 442
 - and **CompatibleOverprint** 462
 - and fully opaque objects 467
 - ignored by older viewer applications 797
 - initialization 452, 453
 - and overprinting 460
 - and process colorants 442
 - soft-mask images, unaffected by 268
 - and spot colorants 442
- current clipping path 12, 131, 148, 161
 - clipping path operators 171
 - even-odd rule 169
 - explicit masks, simulating 276
 - glyph outlines 296
 - as “hard clip” 439
 - initialization 161
 - and marked content 585
 - n** operator 167
 - nonzero winding number rule 169

- current clipping path (*continued*)
 - object shape 171, 439
 - sh** operator 232
 - shading patterns 234
 - and soft clipping, compared 159, 412, 439
 - text rendering mode 306
 - transparency groups 171
 - W** operator 169, 172, 702
 - W*** operator 169, 172, 702
 - See also*
 - clipping path operators
- current color 12, 131, **148**, 151
 - B** operator 167
 - for colored tiling patterns 221
 - current alpha constant, analogous to 444
 - “current opacity,” analogous to 422
 - f** operator 147, 173
 - forced into valid range 151
 - initializing 179, 180, 182, 199
 - nonzero overprint mode 215
 - path objects, used by 136
 - Pattern** color spaces 220
 - S** operator 173
 - Separation** color spaces 201
 - setting 151, 177, 179, 180, 202, 206, **217–218**, 700, 701
 - sh** operator, ignored by 232
 - shading patterns 232
 - as source color (transparent imaging model) 441
 - stencil masking 276
 - stroking and nonstroking 151, 167
 - text, showing 295
 - text objects, nesting of 309
 - tiling patterns as 223, 224
 - tints 201
 - and transparent overprinting 458
 - Type 3 glyph descriptions 325, 326
 - for uncolored tiling patterns 221
 - See also*
 - nonstroking color, current
 - stroking color, current
- current color space **148**
 - All** colorant name 203
 - color values interpreted in 173
 - and current color 148
 - DeviceN** color spaces 206
 - and image dictionaries 267
 - nonzero overprint mode 215
 - and overprinting 214
 - Pattern** color spaces 224
 - remapping 194
 - Separation** color spaces 201
 - setting 177, 179, 180, 182, 199, **216–218**, 699, 700, 701
 - stroking and nonstroking 151
 - and transparent overprinting 458, 461–462
- current color space (*continued*)
 - See also*
 - nonstroking color space, current
 - stroking color space, current
- current font 12, 294
 - composite fonts 334
 - setting 293
 - Type 0 fonts 354
 - See also*
 - text font parameter
 - text font size parameter
- current halftone **150**, **381**, 391, 402
 - HT** entry (graphics state parameter dictionary) 159
 - setting 151
 - and transparency 466, 467
- current line (text) 310
- current line width 12, **148**, **152–153**
 - forced into valid range 151
 - LW** entry (graphics state parameter dictionary) 157
 - and miter length 154
 - and projecting square line cap style 153
 - and round line cap style 153
 - and round line join style 154
 - and **S** operator 147, 168
 - setting 151
 - stroke adjustment 149, 152, **407–408**
 - and text rendering mode 305
 - and Type 3 glyph descriptions 325
 - w** operator 156, 702
- current page 11, 138, 139, 412, 425
- current path **162–163**, 166, 172
- current point **163**, 164, 165
- current rendering intent **149**
 - Intent** entry (image dictionary) 268
 - shading patterns, compositing of 453
 - and transparency 466, 468
- current resource dictionary **96**
 - ColorSpace** subdictionary 177, 194, 216, 280, 450
 - ExtGState** subdictionary 156, 157
 - Font** subdictionary 290, 293, 302, 317
 - Pattern** subdictionary 217, 224
 - Properties** subdictionary 584, 585
 - Shading** subdictionary 232
 - XObject** subdictionary 261, 266, 269, 282, 284
- current soft mask **149**, **444**, 445
 - alpha source parameter 159
 - and fully opaque objects 467
 - ignored by older viewer applications 797
 - initialization 452, 453
 - SMask** entry (graphics state parameter dictionary) 159
 - soft-mask images, overridden by 268

current text position 294, 297

current transfer function 150, 381, 393, 395, 398, 400

- TR operator 158
- TR2 operator 159
- and transparency 466, 467

current transformation matrix (CTM) 131, 139, 148

- cm operator 156, 699
- form XObjects, positioning 283
- halftones unaffected by 383
- sampled images, positioning 141, 266
- shading patterns, compositing of 453
- and soft masks 446
- stroking, effect on 152
- and text rendering matrix 308, 313
- text size 294
- and tiling patterns 223
- in Type 3 glyph descriptions 325

current trap network 690

curves, cubic Bézier

- See cubic Bézier curves

curveto operator (PostScript) 699, 701, 702

Custom production condition 685

cut marks 5, 512, 676, 677, 679, 680

- as page artifacts 615

cvi operator (PostScript) 116, 703

cvr operator (PostScript) 116, 703

“cvt_” table (TrueType font) 367

cyan color component

- DeviceCMYK color space 178, 180
- DeviceN color spaces 206
- grayscale conversion 377, 382
- halftones for 401
- initialization 180
- overprinting 464
- red, complement of 378
- RGB conversion 377, 378
- transfer function 380, 381
- transparent overprinting 464
- undercolor removal 150, 378, 379

cyan colorant

- overprinting 464
- PANTONE Hexachrome system 205
- printing ink 201
- process colorant 178, 180
- subtractive primary 178, 180
- transparent overprinting 464

Cyrillic characters 362, 363

D

D border style (dashed) 495

D entry

- additional-actions dictionary 515
- appearance dictionary 497, 565, 682, 690
- border style dictionary 495
- box style dictionary 681
- go-to action dictionary 519
- graphics state parameter dictionary 157
- inline image object 279
- named destination dictionary 476
- remote go-to action dictionary 520
- thread action dictionary 522
- transition dictionary 486, 487
- Windows launch parameter dictionary 521

D guideline style (page boundaries) 681

d operator 134, 156, 700

D trigger event (annotation) 515, 517

d0 operator 134, 324, 325, 326, 700

d1 operator 134, 218, 324, 325, 326, 700

DA entry

- field dictionary 529, 534, 535, 544
- free text annotation dictionary 502
- interactive form dictionary 529

DamagedRowsBeforeError entry (CCITTFaxDecode filter parameter dictionary) 53, 54

Darken blend mode 417

- and overprinting 459, 460, 466

dash array 155, 157

- annotation borders 491, 495, 798
- page boundaries 681

dash phase 155, 157

- annotation borders, unspecified for 491, 495
- page boundaries, unspecified for 681

data

- binary 25
- types for dictionary entries 99

data structures 98–106

- See also
 - dates
 - name trees
 - number trees
 - rectangles
 - text strings

Data Structures and Algorithms (Aho, Hopcroft, and Ullman) 86, 813

dates 100

- creation
 - document 573, 575, 576
 - Web Capture content set 668, 671

- dates (*continued*)
 - as dictionary values 99
 - expiration (Web Capture content set) 670, 671
 - modification
 - annotation 490
 - document 573, 575, 576, 805
 - form XObject 284, 582
 - page 88, 582
 - trap network 691
 - Web Capture content set 670, 671
 - in submit-form actions 553
- DC** entry (additional-actions dictionary) 516
- DC** trigger event (document) 516
- DCS (Desktop Color Separation) images 125, 126
- DCT (discrete cosine transform) compression 42, 59–61
- DCT filter abbreviation 280, 789
- DCTDecode** filter 42, 59–61, 790
 - color key masking, not recommended with 278
 - DCT** abbreviation 280, 789
 - parameters. *See* **DCTDecode** filter parameter dictionaries
 - in sampled images 268
- DCTDecode** filter parameter dictionaries 59, 60
 - ColorTransform** entry 60
- Decimal list numbering style 650
- Decode** arrays
 - color inversion with 272
 - image masks 276
 - sampled images 264, 269, 271–273, 277
 - shadings 244, 247, 249, 253
- Decode** entry
 - image dictionary 264, 269, 276, 448, 480
 - inline image object 279
 - type 0 function dictionary 109, 110, 112
 - type 4 shading dictionary 244, 247
 - type 5 shading dictionary 249
 - type 6 shading dictionary 253
- DecodeParms** entry
 - inline image object 279
 - stream dictionary 38, 41
 - DP** abbreviation 788
- decoding filters 41–61, 73, 736, 788–790
 - See also*
 - ASCII85Decode** filter
 - ASCIIHexDecode** filter
 - CCITTFaxDecode** filter
 - DCTDecode** filter
 - FlateDecode** filter
 - JBIG2Decode** filter
 - LZWDecode** filter
 - RunLengthDecode** filter
- decoding functions 181, 183, 184, 186, 188
- default appearance strings
 - FDf field 565
 - form field 534, 535, 546, 547
 - free text annotation 502
- default color spaces 178, 194–195, 456, 686
 - DefaultCMYK** 194, 197, 218, 375, 450
 - DefaultGray** 194, 217, 375, 450
 - DefaultRGB** 194, 195, 217, 375, 450
- Default** entry
 - type 5 halftone dictionary 401
- Default graphics state parameter value
 - black-generation function 158
 - halftone parameter 159
 - transfer function 159
 - undercolor-removal function 158
- default user space 138, 139
 - for annotations 490, 494, 496, 503, 506
 - BLSEs, layout of 643, 644, 645, 646, 647
 - current transformation matrix (CTM) 148
 - in destinations 474
 - glyph space, mapping from 647
 - glyphs, scaling of 294
 - halftone angles 393
 - for page boundaries 88, 89, 677
 - page size limits 707, 808
 - pattern matrix 220
 - for Web Capture pages 676
- DefaultCMYK** default color space 194, 197, 218, 375, 450
- DefaultForPrinting** entry (alternate image dictionary) 274
- DefaultGray** default color space 194, 217, 375, 450
- DefaultRGB** default color space 194, 195, 217, 375, 450
- DEFLATE Compressed Data Format Specification* (Internet RFC 1951) 46, 815
- delimiter characters 25, 26, 27, 32, 33
- Department of Commerce, U.S. 72
- Departmental annotation icon 507
- descendant fonts (Type 0 font) 334, 353
- DescendantFonts** entry
 - CID-keyed font dictionary 336
 - Type 0 font dictionary 336, 353, 354
- Descent** entry (font descriptor) 356, 647
- Dest** entry
 - link annotation dictionary 492, 501, 519
 - outline item dictionary 478, 479, 519
- destination handlers 723
- destination profile (PDF/X output intent dictionary) 686, 807
- destinations 84, 474–477, 513
 - explicit 474–476, 520
 - for go-to actions 474, 519
 - handlers 723
 - for link annotations 474, 476, 501, 799

- destinations (*continued*)
 - magnification (zoom) factor 474, 475
 - metadata inapplicable to 579
 - named. *See* named destinations
 - for outline items 474, 476, 477, 478, 479, 798
 - plug-in extensions for 723
 - for remote go-to actions 474, 476, 520
- DestOutputProfile** entry (PDF/X output intent dictionary) 685, 686
- Dests** entry
 - document catalog 84, 476, 740, 753
 - name dictionary 93, 476
- device color spaces 176, 178–180
 - as alternate color space 204
 - as base color space 199
 - blending in 416, 455
 - CIE-based spaces, conversion from 236, 373, 374–375
 - and color specification 172
 - conversion among 236, 373, 376–380
 - and **DeviceN** spaces 205
 - flattening of transparent content to 470
 - implicit conversion of CIE-based colors to 195–197
 - in inline images 280
 - and overprinting 214
 - for page group 437, 455
 - process colors, rendered as 376
 - and rendering intents 149
 - and separations 201
 - setting color values in 217
 - for shadings 234, 236
 - for soft masks 441
 - in transparency groups 178, 450, 456
 - See also*
 - DeviceCMYK** color space
 - DeviceGray** color space
 - DeviceRGB** color space
- device colorants 178, 205
 - and overprinting 214, 465
 - for pre-separated pages 684
 - and transparent overprinting 458, 459
- device-dependent graphics state parameters 147, 150–151, 223, 374
- device gamut 197, 198, 375
- device-independent graphics state parameters 147, 148–149, 152
- device profiles 375
- device space 137–138
 - current transformation matrix (CTM) 131, 141, 148
 - form space, mapping from 283
 - halftone cells, orientation relative to 383, 384, 393, 394, 395, 399
 - halftones defined in 383
 - resolution 383
- device space (*continued*)
 - scan conversion in 405–406
 - stroke adjustment in 407–408
 - text space, relationship with 313
 - threshold arrays defined in 390, 394, 398, 399
 - type 6 shadings (Coons patch meshes) 251, 252
 - URI actions, mouse position for 523
- DeviceCMY** process color model 692
- DeviceCMYK** color space 176, 178, 180, 272
 - as alternate color space 190
 - as blending color space 415
 - CMYK** abbreviation 280
 - color values 180
 - and **DeviceGray**, conversion between 377, 382
 - and **DeviceN** color spaces, compared 205
 - DeviceRGB**, conversion from 377–379, 469
 - DeviceRGB**, conversion to 380
 - in dynamic appearance streams 536
 - halftones for 401
 - implicit conversion from CIE-based 196
 - initial color value 180, 216
 - in inline image objects 280
 - as native color space 374, 376
 - overprint mode 150
 - and overprinting 214–216, 464–465
 - for page group 437
 - process colors, specification of 456
 - remapping to alternate color space 194
 - in sampled images 262
 - and **Separation** color spaces, compared 202
 - setting 177, 216
 - setting color values in 217, 218
 - for soft masks 441
 - specification 180
 - spot color components, effect on in transparency groups 458
 - substituted for **CalCMYK** 181
 - tint transformation function 115
 - transfer functions 382
 - in transparency groups 456
 - and transparent overprinting 458, 461, 462, 464–465
- DeviceCMYK** process color model 692
- DeviceColorant** entry (separation dictionary) 684
- DeviceGray** color space 176, 178, 179, 272
 - as alternate color space 190
 - and **DeviceRGB**, conversion between 377
 - as blending color space 415
 - color values 179
 - and **DeviceCMYK**, conversion between 377, 382
 - in dynamic appearance streams 536
 - G** abbreviation 280
 - halftones for 401
 - initial color value 179, 216
 - in inline image objects 280

- DeviceGray** color space (*continued*)
- as native color space 374, 376
 - and overprinting 465
 - for pre-separated pages 683
 - remapping to alternate color space 194
 - in sampled images 262
 - and **Separation** color spaces, compared 202
 - setting 177, 216
 - setting color values in 217
 - for soft masks 441
 - specification 179
 - for thumbnail images 480
 - transfer functions 382
 - in transparency groups 456
- DeviceGray** process color model 692
- DeviceN** color space attributes dictionaries 207
- Colorants** entry 207
 - metadata inapplicable to 579
- DeviceN** color spaces 176, 199, 205–209, 272
- All** colorant name prohibited in 206
 - alternate color space for 115, 206–207, 216, 236
 - alternate color space, prohibited as 204
 - attributes. *See* **DeviceN** color space attributes dictionaries
 - as base color space 199, 200
 - blending color space, prohibited as 450
 - color values 206
 - colorant names 206
 - halftones for 401
 - initial color value 206, 217
 - None** colorant name 207
 - nonzero overprint mode 215
 - number of components 206, 706
 - and overprinting 214, 464–465
 - parameters 206–207
 - for pre-separated pages 684
 - remapping of alternate color space 195
 - in sampled images 262
 - and **Separation** color spaces, compared 206–207
 - setting color values in 217
 - for shadings 236
 - in soft masks 447
 - specification 206
 - spot color components in 457
 - for spot colorants 376, 415
 - tint transformation function 115, 207, 216, 236
 - tints 206, 207, 217
 - in transparency groups 196
 - and transparent overprinting 462, 465
- DeviceN** process color model 692
- DeviceRGB** color space 176, 178, 179–180, 272
- as alternate color space 190
 - and **DeviceGray**, conversion between 377
 - for annotations 491, 503, 505
- DeviceRGB** color space (*continued*)
- as base color space 199, 200
 - as blending color space 415
 - color values 179
 - DeviceCMYK**, conversion from 380
 - DeviceCMYK**, conversion to 377–379, 469
 - in dynamic appearance streams 536
 - halftones for 401
 - initial color value 180, 216
 - in inline image objects 280
 - as native color space 374, 376
 - for outline items 479
 - and overprinting 465
 - for page boundaries 681
 - for page group 437
 - remapping to alternate color space 194
 - RGB** abbreviation 280
 - in sampled images 262
 - and **Separation** color spaces, compared 202
 - setting 177, 216
 - setting color values in 217
 - for soft masks 441
 - specification 179
 - for thumbnail images 480
 - transfer functions 380
 - in transparency groups 456
- DeviceRGB** process color model 692
- DeviceRGBK** process color model 692
- devices, output
- See* output devices
- Di** entry (transition dictionary) 486, 487
- Diamond line ending style 504
- Diamond** predefined spot function 389
- dictionaries
- See* dictionary objects
 - See also*
 - action dictionaries
 - additional-actions dictionaries
 - alternate image dictionaries
 - annotation dictionaries
 - appearance characteristics dictionaries
 - appearance dictionaries
 - application data dictionaries
 - attribute objects
 - bead dictionaries
 - border style dictionaries
 - box color information dictionaries
 - box style dictionaries
 - button field dictionaries
 - CalGray** color space dictionaries
 - CalRGB** color space dictionaries
 - CCITTFaxDecode** filter parameter dictionaries
 - checkbox field dictionaries

dictionaries (*continued*)*See also*

- choice field dictionaries
- CIDFont dictionaries
- CIDFont **FD** dictionaries
- CIDFont **Style** dictionaries
- CIDSystemInfo** dictionaries
- circle annotation dictionaries
- class map
- CMap dictionaries
- DCTDecode** filter parameter dictionaries
- DeviceN** color space attributes dictionaries
- document catalog
- document information dictionary
- embedded file parameter dictionaries
- embedded file stream dictionaries
- embedded font stream dictionaries
- encoding dictionaries
- encryption dictionaries
- FDF annotation dictionaries
- FDF catalog
- FDF dictionary
- FDF field dictionaries
- FDF named page reference dictionaries
- FDF page dictionaries
- FDF page information dictionaries
- FDF template dictionaries
- FDF trailer dictionary
- field dictionaries
- file attachment annotation dictionaries
- file specification dictionaries
- file trailer dictionary
- filter parameter dictionaries
- FlateDecode** filter parameter dictionaries
- font descriptors
- font dictionaries
- form dictionaries
- free text annotation dictionaries
- function dictionaries
- go-to action dictionaries
- graphics state parameter dictionaries
- group attributes dictionaries
- halftone dictionaries
- hide action dictionaries
- hint stream dictionaries
- ICC profile stream dictionaries
- icon fit dictionaries
- image dictionaries
- import-data action dictionaries
- ink annotation dictionaries
- interactive form dictionary
- JavaScript action dictionaries
- JavaScript dictionary
- JBIG2Decode** filter parameter dictionaries
- Lab** color space dictionaries

dictionaries (*continued*)*See also*

- launch action dictionaries
- line annotation dictionaries
- linearization parameter dictionary
- link annotation dictionaries
- LZWDecode** filter parameter dictionaries
- Mac OS file information dictionaries
- mark information dictionary
- marked-content reference dictionaries
- markup annotation dictionaries
- metadata stream dictionaries
- movie action dictionaries
- movie activation dictionaries
- movie annotation dictionaries
- movie dictionaries
- multiple master font dictionaries
- name dictionary
- name tree nodes
- named-action dictionaries
- named destination dictionaries
- number tree nodes
- object reference dictionaries
- OPI dictionaries
- OPI version dictionaries
- outline dictionary
- outline item dictionaries
- output intent dictionaries
- page label dictionaries
- page objects
- page-piece dictionaries
- page tree nodes
- pattern dictionaries
- PDF/X output intent dictionaries
- pop-up annotation dictionaries
- PostScriptXObject dictionaries
- printer's mark annotation dictionaries
- printer's mark form dictionaries
- property lists
- radio button field dictionaries
- reference dictionaries
- remote go-to action dictionaries
- reset-form action dictionaries
- resource dictionaries
- role map
- rubber stamp annotation dictionaries
- separation dictionaries
- shading dictionaries
- signature dictionaries
- soft-mask dictionaries
- soft-mask image dictionaries
- sound action dictionaries
- sound annotation dictionaries
- sound objects
- source information dictionaries

dictionaries (*continued*)*See also*

- square annotation dictionaries
- stream dictionaries
- structure element dictionaries
- structure tree root
- submit-form action dictionaries
- text annotation dictionaries
- text field dictionaries
- thread action dictionaries
- thread dictionaries
- thread information dictionaries
- transition dictionaries
- transparency group attributes dictionaries
- trap network annotation dictionaries
- trap network appearance stream dictionaries
- TrueType font dictionaries
- Type 0 font dictionaries
- Type 0 function dictionaries (sampled)
- Type 1 font dictionaries
- type 1 form dictionaries
- type 1 halftone dictionaries
- type 1 pattern dictionaries (tiling)
- type 1 shading dictionaries (function-based)
- type 2 function dictionaries (exponential interpolation)
- type 2 pattern dictionaries (shading)
- type 2 shading dictionaries (axial)
- Type 3 font dictionaries
- type 3 function dictionaries (stitching)
- type 3 shading dictionaries (radial)
- type 4 shading dictionaries (free-form Gouraud-shaded triangle mesh)
- type 5 halftone dictionaries
- type 5 shading dictionaries (lattice-form Gouraud-shaded triangle mesh)
- type 6 halftone dictionaries
- type 6 shading dictionaries (Coons patch mesh)
- type 7 shading dictionaries (tensor-product patch mesh)
- type 10 halftone dictionaries
- type 16 halftone dictionaries
- URI action dictionaries
- URI dictionaries
- URL alias dictionaries
- viewer preferences dictionary
- Web Capture command dictionaries
- Web Capture command settings dictionaries
- Web Capture content sets
- Web Capture image sets
- Web Capture information dictionary
- Web Capture page sets
- widget annotation dictionaries
- Windows launch parameter dictionaries

dictionary objects 35–36

- adding new entries to 723, 786
- as attribute objects 605
- capacity limit 35, 101, **706**
- as dictionary values 99
- duplicate keys 35
- entries 35
- keys 25, **35**, 786
- metadata associated with 578, 580
- null entries 28
- as operands 94
- syntax **35–36**
- values 35
- version compatibility 787

Difference blend mode 418

- not white-preserving 460

Differences entry

- encoding dictionary 324, **330**, 620
- FDf dictionary 552, **562**

differencing (image compression) 49

- diffuse achromatic highlight 185
- diffuse achromatic shadow 185
- diffuse black point 183, 185, 188
- diffuse white point 183, 184, 185, 188

Digital Compression and Coding of Continuous-Tone Still Images (ISO/IEC 10918-1) **815**

- digital identifiers (Web Capture) 660, **664–665**, 805
 - in content database 661, 668, 670
 - for image 269, 675
 - in name dictionary 93, 661
 - for page 90, 675
 - in unique name generation 666

digital signatures

- See* signatures, digital

Dingbats character class 362, 363

Dingbats typeface

- See* ITC Zapf Dingbats® typeface

DingbatsRot character class 363

- direct content items **625**, 632
 - allocation rectangle 626
 - content rectangle 626

direct objects

- in FDF files 560
- stream dictionaries 36

Direction entry (viewer preferences dictionary) 472Disc list numbering style **650**displacement vector (glyph) **299**

- DW2** entry (CIDFont) 341
- horizontal scaling 304
- W2** entry (CIDFont) 341–342
- See also*
 - glyph displacement

- display duration 90, **485**, 487
- DisplayDocTitle** entry (viewer preferences dictionary) **472**
- displays, raster-scan 12
 - and halftones 382, 383, 390
 - primary colorants 201
 - resolution 13
 - scan conversion for 403
 - and **Separation** color spaces 203
 - stroke adjustment for 407
- Dissolve transition style **486**
- Distiller[®], Acrobat
 - See Acrobat Distiller[®]
- div** operator (PostScript) **116**, **703**
- Div standard structure type **627**, 628, 631
- Dm** entry (transition dictionary) 486, **487**
- Do** operator 134, **261**, 451, 452, 700
 - base images 267
 - and black-generation functions 469
 - colored tiling patterns 224
 - form XObjects 220, 282, 283
 - and fully opaque objects 467
 - image XObjects 262
 - and logical structure elements 595, 596, 599
 - and marked content 586, 595
 - PostScript XObjects 290
 - and rendering intents 468
 - shading patterns 232
 - uncolored tiling patterns 228
 - and undercolor-removal functions 469
- Doc** entry (JavaScript dictionary) **563**
- document catalog 81, **83–85**, 791
 - AA** entry **85**, 515, 791
 - AA** entry (obsolete) 791
 - AcroForm** entry **85**, 529, 734
 - Dests** entry **84**, 476, 740, 753
 - example 758, 760, 762
 - in file trailer 68
 - Lang** entry **85**, 652, 653, 654
 - in Linearized PDF 728, 730, 732, **734**, 740, 753
 - MarkInfo** entry **85**, 613
 - Metadata** entry **85**, 578
 - metadata inapplicable to 579
 - Names** entry **83**, 92, 740, 753
 - OpenAction** entry **84**, 474, 513, 515, 523, 734, 737, 748, 751, 808
 - Outlines** entry **84**, 477, 757
 - OutputIntents** entry **85**, 684
 - PageLabels** entry **83**, 482, 798
 - PageLayout** entry **84**, 791, 798
 - PageMode** entry **84**, 472, 730, 734, 737
 - Pages** entry **83**
 - private data in 575, 784
 - SpiderInfo** entry **85**, 660
- document catalog (*continued*)
 - StructTreeRoot** entry **85**, 589, 628
 - Threads** entry **84**, 483, 522, 734, 753
 - Type** entry **83**
 - URI** entry **85**, 524
 - Version** entry 63, 70, **83**, 784, 785, 791
 - ViewerPreferences** entry **84**, 471, 734
- document information dictionary 68, 484, **575–577**
 - Author** entry **576**
 - CreationDate** entry **576**
 - Creator** entry **576**
 - and file identifiers 581
 - keys in 575
 - Keywords** entry **576**
 - in Linearized PDF 734, 740
 - metadata inapplicable to 579
 - metadata streams, compared with 577
 - ModDate** entry **576**
 - Producer** entry **576**
 - registered names not required in 724
 - Subject** entry **576**
 - Title** entry **576**
 - Trapped** entry **576**
 - version compatibility, use for 804–805
- document interchange 1, 3, 22, 132, **573–698**
 - pdfmark** language extension (PostScript) 21
 - version compatibility 783
 - See also
 - accessibility to disabled users
 - file identifiers
 - logical structure
 - marked content
 - metadata
 - page-piece dictionaries
 - prepress production
 - procedure sets
 - Tagged PDF
 - Web Capture plug-in extension
- document outline 4, 83, 84, 474, **477–480**, 798
 - hiding and showing 84, 472
 - hierarchy. See outline hierarchy
 - items. See outline items
 - outline dictionary 84, 757, 758, 760, 762
- Document Properties dialog box (Acrobat) 804
- Document standard structure type **627**, 628, 631
- document structuring conventions, PostScript (DSC) 27
- document windows
 - centering on screen 472
 - and destinations 474, 475
 - fitting to document 472
 - and remote go-to actions 520
 - title bar 472

- documents 9, 24
 - additional-actions dictionary 85, 514–516
 - application-specific data 19
 - article threads 83, 84
 - authenticity, certification of xx, 18
 - author 573, 575, 576
 - catalog. *See* document catalog
 - closing 516
 - collaborative editing xx, 553
 - creation date 573, 575, 576
 - creator application 576
 - encryption 4, 18, 31, 71–81, 790
 - extensibility 18–19
 - extraction of content. *See* content extraction
 - fidelity, preservation of xx
 - incremental updates. *See* incremental updates
 - information dictionary 68, 484, 575–577, 579
 - interactive form dictionary 85, 734, 750
 - interchange. *See* document interchange
 - keywords 576
 - language identifier 85, 592
 - logical structure. *See* logical structure
 - mark information dictionary 85, 613–614
 - metadata 85, 573, 575–580, 804–805
 - modification date 573, 575, 576, 805
 - name dictionary 83, 92–93
 - named halftones 392
 - natural language specification 652, 653, 654
 - open action 84, 474, 513
 - opening 513, 520, 521
 - outline. *See* document outline
 - output intent dictionaries 85
 - page labels 83
 - page layout 84, 800
 - page mode 84, 472
 - page objects. *See* page objects
 - page tree 81, 83, 86–92, 557, 757, 765
 - page tree root 83
 - printing 75, 365, 374, 516, 520, 521
 - producer application 576
 - reading order 472
 - saving 516
 - scan conversion 403
 - security xx, 18
 - structure 23, 81–93
 - structure tree root 85
 - subject 576
 - title 573, 575, 576
 - trapping status 576
 - trigger events for 516
 - undoing changes 18
 - URI dictionary 85, 524
 - viewer preferences 84
 - viewing. *See* viewing of documents
 - Web Capture information dictionary 85
- dollar sign (\$) character 343
- domain
 - function 107, 108, 109, 110, 114, 237, 238, 240, 244, 249, 253
 - shading 115, 237, 238, 239, 240, 241
- Domain** entry
 - function dictionary 108, 109, 113, 114, 117
 - type 1 shading dictionary 237
 - type 2 shading dictionary 238, 239
 - type 3 shading dictionary 240, 241
- DoNotScroll field flag (text field) 543
- DoNotSpellCheck field flag
 - choice field 546
 - text field 543
- DOS (Disk Operating System) 20, 120
 - file names 121, 580
 - file system 122
- DOS** entry (file specification dictionary) 122, 123
- dot-matrix printers 12
 - resolution 13
- double angle brackets (<< >>)
 - as dictionary delimiters 35, 67, 560
- double left angle bracket (<<) character sequence
 - as dictionary delimiter 35, 67, 560
- double period (..) character sequence
 - in relative file specifications 120
 - in uniform resource locators (URLs) 664
- Double** predefined spot function 386
- double right angle bracket (>>) character sequence
 - as dictionary delimiter 35, 67, 560
- DoubleDot** predefined spot function 385
- down appearance (annotation) 497, 501, 512
- DP** entry
 - additional-actions dictionary 516
 - inline image object 279
 - stream dictionary (abbreviation for **DecodeParms**) 788
- DP** operator 134, 583, 584, 700
 - property list 583, 585
- DP** trigger event (document) 516
- DR** entry
 - field dictionary 529, 534, 535, 801
 - interactive form dictionary 529
- Draft annotation icon 507
- drag-and-drop 680
- driver, printer 19, 20
- dropped capitals 649, 658
- DS** entry (additional-actions dictionary) 516
- DS** trigger event (document) 516
- DSC. *See* document structuring conventions, PostScript
- duotone color 205
 - examples 209–211

- dup** operator (PostScript) **116, 704**
 - Dur** entry (page object) **90, 485–486, 487**
 - Duration** entry (movie activation dictionary) **571, 572**
 - DV** entry (field dictionary) **532, 533, 539, 548, 554**
 - DVI (Device Independent) file format **19**
 - DW** entry (CIDFont dictionary) **338, 340**
 - DW2** entry (CIDFont dictionary) **338, 341, 368**
 - dynamic appearance streams **512, 529, 533–537**
 - alternate (down) caption **537**
 - alternate (down) icon **537**
 - background color **536**
 - border color **536**
 - for choice fields **546, 547, 565**
 - icon fit dictionary **537**
 - normal caption **536**
 - normal icon **537**
 - rollover caption **537**
 - rollover icon **537**
 - for text fields **544**
- E**
- E** entry
 - additional-actions dictionary **515**
 - hint stream dictionary **736**
 - linearization parameter dictionary **733, 736, 808**
 - property list **616, 621, 622, 633, 659**
 - sound object **569, 570**
 - source information dictionary **670, 671**
 - E** trigger event (annotation) **515, 517, 526**
 - EarlyChange** entry (LZWDecode filter parameter dictionary) **49**
 - edge flags **244, 245–247, 253, 254–256, 259–260**
 - Edit field flag (choice field) **546**
 - editing, collaborative **xx, 553**
 - EF** entry (file specification dictionary) **123, 124, 128, 129**
 - EI** operator **134, 278, 279, 280, 700**
 - element identifiers (logical structure) **590, 591**
 - Ellipse** predefined spot function **387**
 - EllipseA** predefined spot function **387**
 - EllipseB** predefined spot function **388**
 - EllipseC** predefined spot function **388**
 - ellipsis (...) character **782**
 - em dash character
 - as word separator **623**
 - embedded CIDFonts **336, 339**
 - embedded CMaps **336, 347, 348, 796**
 - embedded file parameter dictionaries **124, 125**
 - Checksum** entry **125**
 - CreationDate** entry **125**
 - Mac** entry **125**
 - embedded file parameter dictionaries (*continued*)
 - ModDate** entry **125**
 - Size** entry **125**
 - embedded file stream dictionaries **124, 562–563**
 - EncryptionRevision** entry **563**
 - Params** entry **124**
 - Subtype** entry **124**
 - Type** entry **124**
 - embedded file streams **4, 123–127**
 - checksum **125**
 - encryption **71, 562–563**
 - FDF **4, 562**
 - and file attachment annotations **508**
 - maintenance of file specifications **128**
 - named **93**
 - platform-specific **123**
 - and reference XObjects **287**
 - related files arrays **123, 125–127, 128**
 - See also*
 - embedded file stream dictionaries
 - embedded font programs **16, 291, 292, 314, 315, 357, 364–368, 796–797**
 - base encoding **330**
 - built-in encoding **330**
 - compact **365, 367**
 - filters in **366**
 - font descriptors and **356**
 - in Linearized PDF **738, 754**
 - organization, by font type **365**
 - overriding standard fonts **319**
 - PaintType** entry ignored **367**
 - for portability **797**
 - snapshots (multiple master) **321**
 - TrueType **321, 333, 339, 345, 357, 365, 366, 367–368**
 - Type 0 compact CIDFonts **365, 366, 367**
 - Type 1 **357, 365, 366, 794**
 - Type 1 compact fonts **365, 366, 367**
 - unrecognized filters in **789**
 - embedded font stream dictionaries **357, 366**
 - Length1** entry **366, 367**
 - Length2** entry **366, 367**
 - Length3** entry **366, 367**
 - Metadata** entry **366**
 - Subtype** entry **357, 365, 366**
 - embedded font subsets **16, 368**
 - EmbeddedFDFs** entry (FDF dictionary) **562**
 - EmbeddedFile** object type **124**
 - EmbeddedFiles** entry (name dictionary) **93, 124**
 - EMC** operator **134, 535, 583, 584, 593, 700**
 - Encapsulated PostScript (EPS) **412**
 - Encode** entry
 - type 0 function dictionary **109, 110**
 - type 3 function dictionary **114**

- EncodedByteAlign** entry (**CCITTFaxDecode** filter parameter dictionary) 53, 54
- Encoding** array (Type 1 font program) 332
- encoding dictionaries 318, 324, 330, 710, 714
- BaseEncoding** entry 330, 332, 333, 796
- Differences** entry 324, 330, 620
- metadata inapplicable to 579
- Type** entry 330
- Encoding** entry 329, 342
- CID-keyed font dictionary 336
- CIDFonts, absent in 338
- FDF dictionary 562, 803
- field resource dictionary 801
- TrueType font dictionary 321, 332, 333, 334
- Type 0 font dictionary 336, 345, 349, 352, 353, 354
- Type 1 font dictionary 318, 332
- Type 3 font dictionary 324, 332, 796
- encoding filters 73
- encoding formats, sound 569
- ALaw** 569
- muLaw** 569, 570
- Raw** 569, 570
- Signed** 569, 570
- Encoding** object type 330
- Encoding** resource type 738
- encodings
- ASCII base-85 15, 41, 42, 43, 44–45
- ASCII hexadecimal 42, 44, 57, 73
- character. *See* character encodings
- Encrypt** entry (file trailer dictionary) 68, 71, 734
- encryption 4, 18, 31, 71–81, 790
- algorithm revision number (FDF) 563
- algorithms 72
- ASCII filters not useful with 41
- dictionary. *See* encryption dictionary
- in FDF files 562–563
- general algorithm 72–73
- key algorithm 78
- keys. *See* encryption keys
- metadata streams, not recommended for 577
- password algorithms 79–81, 563, 790
- passwords 79–81, 563, 790
- security handlers 71, 72, 74–81, 723
- signature fields 538, 547
- encryption dictionary 68, 71–72, 73, 74
- Filter** entry 71, 72, 79
- Length** entry 71, 72, 73, 78, 79
- in Linearized PDF 734
- metadata inapplicable to 579
- O** entry 76, 78, 79, 81
- P** entry 76, 78, 79
- R** entry 76, 79
- for standard security handler 76, 79
- encryption dictionary (*continued*)
- standard 75–77, 790
- U** entry 76, 79–80
- V** entry 71, 72, 73, 76, 79, 790
- encryption keys 73
- computing 76, 78
- and encryption revision number 72
- for FDF files 563
- length 71, 72
- owner password, authenticating 81
- owner password, computing 79
- user password, computing 79, 80
- using 73
- EncryptionRevision** entry (embedded file stream dictionary) 563
- end edge 625
- of allocation rectangle 649
- and content rectangle 648
- in layout 625, 641, 644, 646
- End inline alignment 646
- end-of-data (EOD) marker 37, 46, 47, 53
- for **ASCII85Decode** (~>) 44, 45
- for **ASCIIHexDecode** (>) 44
- for **RunLengthDecode** 52
- end-of-facsimile-block (EOFB) pattern (**CCITTFaxDecode** filter) 54
- end-of-file (EOF) marker (%%EOF) 67, 69, 560, 732, 790
- end-of-line (EOL)
- conventions 15, 25
- markers 15, 26, 30, 31, 37, 38, 62, 64, 65
- End placement attribute 641, 644, 649
- End text alignment 644
- endbfchar** operator (PostScript) 352, 354, 369
- endbfrange** operator (PostScript) 352, 369
- endcidchar** operator (PostScript) 352, 354
- endcidrange** operator (PostScript) 352
- endcmap** operator (PostScript) 351
- endcodespacerange** operator (PostScript) 351, 354, 369, 371
- EndIndent** standard structure attribute 624, 640, 644
- endnotdefchar** operator (PostScript) 352, 355
- endnotdefrange** operator (PostScript) 352, 355
- endobj** keyword 39, 707
- EndOfBlock** entry (**CCITTFaxDecode** filter parameter dictionary) 54
- EndOfLine** entry (**CCITTFaxDecode** filter parameter dictionary) 54
- endrearrangedfont** operator (PostScript) 352
- endstream** keyword 36–37, 38, 707, 788
- endusematrix** operator (PostScript) 352

- entries, dictionary 35
- Entrust.PPKEF** signature handler 549
- eoclip** operator (PostScript) 702
- EOD. *See* end-of-data
- EOF. *See* end-of-file
- eofill** operator (PostScript) 699, 700
- EOL. *See* end-of-line
- ePaper[®] Solutions network publishing software xix–xx, 7
- EPS. *See* Encapsulated PostScript
- eq** operator (PostScript) 116, 704
- error reporting (Acrobat) 784–785, 786, 789, 790, 791, 792, 793, 796, 798, 799, 800, 801, 804
- escape character 30
- escape sequences 30–31, 73
 - backslash (\) 30, 312
 - backspace (BS) 30
 - carriage return (CR) 30
 - form feed (FF) 30
 - horizontal tab (HT) 30
 - left parenthesis (() 30, 312
 - line feed (LF) 30
 - octal character code 30, 31
 - right parenthesis ()) 30, 312
 - Unicode natural language escape 99–100, 652, 657, 658, 659
- ET** operator 134, 305, 308, 535, 584, 637, 700
- ETen character set 344
- ETen–B5–H predefined CMap 344, 346
- ETen–B5–V predefined CMap 344, 346
- ETenms–B5–H predefined CMap 344, 346
- ETenms–B5–V predefined CMap 344, 346
- EUC-CN character encoding 343
- EUC–H predefined CMap 344, 346
- EUC-JP character encoding 344
- EUC-KR character encoding 345
- EUC–V predefined CMap 344, 346
- EUC-TW character encoding 344
- euro character 714
- even-odd rule 170–171
 - clipping 172, 702
 - filling 167, 169, 699, 700
- events, trigger
 - See* trigger events
- EX** operator 95, 134, 700
- examples
 - abbreviation expansion 659
 - alternate image 274–275
 - appearance dictionary 498
 - article 484–485
 - CalGray** color spaces 184
 - CalRGB** color space 186
 - examples (*continued*)
 - character encoding 330–331
 - checkbox field 539–540
 - choice field 547
 - CIDFont **FD** dictionary 363–364
 - CIDFont, **W** entry 341
 - CIDFont, **W2** entry 342
 - CMap 348, 350–351
 - conversion engine, internal name 675
 - cross-reference sections 66–67
 - DeviceCMYK** color space, color specification 180
 - DeviceGray** color space, color specification 179
 - DeviceN** color space 208
 - DeviceRGB** color space, color specification 179
 - dictionary syntax 35
 - document catalog 85
 - document information dictionary 577
 - document outline 479–480
 - duotone color spaces (**DeviceN**) 209–211
 - embedded font program 366–367
 - file specification string 119
 - file trailer 67–68
 - filter pipeline 41
 - font definition 293–294
 - font descriptor, **Style** entry 361
 - form XObject 286
 - glyph positioning 298–299
 - go-to action 519
 - graphics state parameter dictionaries 160–161
 - ICCBased** color space 193–194
 - image coordinate system, transformation of 266–267
 - incremental updates 774–782
 - Indexed** color space 200
 - indirect object reference 39–40
 - indirect object specification 39
 - inline image 281
 - JBIG2 encoding 57–59
 - Lab** color space 189
 - language specification hierarchy 653–654, 654–655, 656
 - Linearized PDF 729–731
 - link annotation 502
 - link element (Tagged PDF) 635–636
 - logical structure 607–612
 - LZW (Lempel-Ziv-Welch) encoding 47–48
 - marked content and clipping 585, 586, 587–588
 - marked-content reference 594–595
 - multiple master font 320
 - name tree 103–105
 - nonzero overprint mode 215
 - outline hierarchy 770–774
 - output intent dictionary 687–688
 - page labels 482
 - page object 90–91
 - page tree 87, 765–770
 - partial field name 533

- examples (*continued*)
 - PDF files
 - graphics 762–765
 - minimal 757–759
 - text 760–762
 - presentation parameters 488
 - quadtone color space (**DeviceN**) 212–213
 - radio button field 541–542
 - related files array 126–127
 - relative file specifications 120
 - resource dictionary 96–97
 - reverse-order show string 619
 - sampled image 269–270
 - Separation** color space 204–205
 - showing of text 293
 - stream, encoded 43
 - stream, indirect length specification 40–41
 - stream, unencoded 43–44
 - structure elements
 - as content items 599–600
 - finding from content items 602–604
 - form XObjects in 595–598
 - text annotation 500
 - text field 544–545
 - text, special graphical effects for 295–297
 - thumbnail image 480–481
 - tiling pattern, colored 224–227
 - tiling pattern, uncolored 228–230
 - tint transformation function 211
 - ToUnicode** CMap 370–371
 - transfer function 381
 - TrueType font 321–322
 - Type 0 font 352–353
 - type 0 (sampled) function 111–112
 - Type 1 font 318–319
 - type 1 halftone 394
 - Type 3 font 327–328
 - type 3 (radial) shading 242–243
 - type 4 (PostScript calculator) function 116–117
 - type 5 halftone 402–403
 - unique name specification (Web Capture) 666–667
 - URL specification 127
 - URLs, same path 673
- exch** operator (PostScript) **116, 704**
- exclamation point (!) character
 - in ASCII base-85 encoding 44, **45**
- ExclFKey field flag (submit-form field) **553**
- ExclNonUserAnnots field flag (submit-form field) **553**
- Exclusion** blend mode **418**
 - not white-preserving 460
- exp** operator (PostScript) **116, 703**
- Experimental annotation icon 507
- expert character set, standard 709, 710, **715–717**
- expert fonts 329, 710
- expiration date (Web Capture content set) 670, 671
- Expired annotation icon 507
- explicit destinations **474–476**
 - for remote go-to actions 520
 - syntax **475**
- explicit masking 275, **277, 793**
 - and object shape 421, 443
 - simulation of 276
 - and soft masks 444
 - See also*
 - image masks
- exponential interpolation functions
 - See* type 2 functions
- ExportFormat field flag (submit-form field) **552, 553, 554**
- exporting
 - PDF fields 557, 558, 561, 803
 - interactive form fields 9, 528, 531, 532, 542, 543, 546, 548
 - Tagged PDF 638, 639, 640, 647
 - text 368, 620, 622, 628
- Ext–RKSJ–H predefined CMap **344, 347**
- Ext–RKSJ–V predefined CMap **344, 347**
- Extend** entry
 - type 2 shading dictionary **238, 239**
 - type 3 shading dictionary **240, 241**
- extensibility of documents 18–19
- Extensible Markup Language (XML) 1.0* (World Wide Web Consortium) 554, 653, **816**
- Extensible Stylesheet Language (XSL) 1.0* (World Wide Web Consortium) 622, **816**
- extensions, plug-in
 - See* plug-in extensions
- extent, stream 37
- external objects (XObjects) **132–133, 261, 793**
 - in logical structure elements 598
 - and marked content 586
 - as named resources 97, 132
 - painting 261, 700
 - in structural parent tree 590
 - subtypes. *See* XObject subtypes
 - in Web Capture content database 661
 - See also*
 - form XObjects
 - group XObjects
 - image XObjects
 - PostScript XObjects
 - reference XObjects
 - transparency group XObjects
 - XObject operator
- external streams 37, 38, 41, 788
 - digital identifiers, not used in 805

- ExtGState** entry
 - resource dictionary 97, 156, 157
 - type 2 pattern dictionary 231, 453
- ExtGState** object type 157
- ExtGState** resource type 97, 156, 157
- extraction of document content
 - See content extraction
- Extreme™ printing systems 804

- F**
- F** entry
 - additional-actions dictionary 516
 - annotation dictionary 490, 492, 565, 682, 690
 - FDf dictionary 553, 561, 802, 803
 - FDf field dictionary 565
 - FDf named page reference dictionary 568
 - file specification dictionary 122, 127
 - import-data action dictionary 556, 802
 - inline image object 279
 - launch action dictionary 520, 521
 - movie dictionary 571
 - outline item dictionary 479
 - reference dictionary 288
 - remote go-to action dictionary 476, 520
 - stream dictionary 38, 568, 569
 - submit-form action dictionary 551
 - thread action dictionary 476, 522
 - thread dictionary 483, 484
 - version 1.3 OPI dictionary 694, 807
 - version 2.0 OPI dictionary 697, 807
 - Web Capture command dictionary 672
 - Windows launch parameter dictionary 521
- f** keyword 65, 779
- F** operator 134, 167, 700
- f** operator 12, 134, 166, 167, 168–169, 700
 - and current color 173
 - ending path 162
 - and graphics state parameters 147
 - and patterns 219, 221, 224, 228, 232
 - and subpaths 169
- F** trigger event (form field) 516, 517
- f*** operator 134, 167, 169, 700
- facsimile compression, CCITT
 - See compression, CCITT facsimile
- false** (boolean object) 28
- false** operator (PostScript) 116, 704
- “fauxing” of fonts 691, 794
- fax compression, CCITT
 - See compression, CCITT facsimile
- FD** dictionaries
 - See CIDFont **FD** dictionaries
- FD** entry (CIDFont font descriptor) 361
- FDcodeParms** entry (stream dictionary) 38, 41
- FDf. See Forms Data Format
- FDf annotation dictionaries 562, 568
- FDf catalog 559, 560–563, 803
 - FDf** entry 560, 561
 - Version** entry 559, 561, 802
- FDf dictionary 560–562, 803
 - Annots** entry 562
 - Differences** entry 552, 562
 - EmbeddedFDfs** entry 562
 - Encoding** entry 562, 803
 - F** entry 553, 561, 802, 803
 - Fields** entry 561, 562
 - ID** entry 561, 802
 - JavaScript** entry 562
 - Pages** entry 561, 566
 - Status** entry 561, 803
 - Target** entry 562
- FDf entry (FDf catalog) 560, 561
- FDf field dictionaries 561, 564–565
 - A** entry 565
 - AA** entry 565
 - AP** entry 565, 803
 - APRef** entry 565
 - ClrF** entry 565
 - ClrFf** entry 564
 - F** entry 565
 - Ff** entry 564
 - field dictionaries, compared with 564
 - IF** entry 565, 566
 - Kids** entry 564
 - Opt** entry 562, 565
 - SetF** entry 565
 - SetFf** entry 564
 - T** entry 564, 803
 - in template pages 567
 - V** entry 544, 562, 564, 803
 - widget annotation dictionaries, compared with 564
- FDf fields
 - See fields, FDf
- FDf files 4
 - body 558, 559–560
 - cross-reference table 558
 - document structure 558
 - file name extension (Windows and Unix) 558
 - file type (Mac OS) 558
 - generation numbers in 558
 - header 558, 559, 561, 802
 - in import-data actions 556, 802, 803
 - incremental updates not permitted in 558
 - object numbers in 558
 - source file 561

- FDF files (*continued*)
 - structure **558–560**
 - in submit-form actions 803
 - target file **561**
 - trailer 558, **560**
 - version specification 558, 559, 561, 802
- FDF named page reference dictionaries 565, 567, **568**
- F** entry **568**
- Name** entry **568**
- FDF page dictionaries 561, **566–567**
- Info** entry **567**
- Templates** entry **567**
- FDF page information dictionaries **567**
- FDF template dictionaries **567**
- Fields** entry **567**
- Rename** entry **567**, 804
- TRef** entry **567**, 568
- FDF trailer dictionary **560**
- Root** entry **560**
- Ff** entry
 - FDF field dictionary **564**
 - field dictionary **531**, 532, 538, 543, 546, 554, 564
- FFilter** entry (stream dictionary) **38**, 41
- fidelity, preservation of xx
- field dictionaries 528, **530–537**
- AA** entry 514, **532**
- DA** entry 529, **534**, 535, 544
- DR** entry 529, **534**, 535, 801
- DV** entry **532**, 533, 539, 548, **554**
- FDF field dictionaries, compared with 564
- Ff** entry **531**, 532, 538, 543, 546, 554, 564
- FT** entry **531**, 533, 548
- Kids** entry **531**, 541, 542, 543
- Parent** entry **531**
- Q** entry 529, **534**, 535
- in reset-form actions 555
- in submit-form actions 551, 553
- T** entry **531**, 532, 533, 548
- TM** entry **531**, 552
- TU** entry **531**, 657
- V** entry **531**, 532, 533, 539, 540, 541, 542, 543, 544, 546, 547, 548, 551, 554
- for variable-text fields **534**
- widget annotation dictionaries, merged with 511, 528, 548
- field flags 531, **532**, 564
- NoExport **532**, 551, 554
- ReadOnly 493, **532**
- Required **532**
- See also*
 - button field flags
 - choice field flags
 - reset-form field flag
- field flags (*continued*)
 - See also*
 - signature flags
 - submit-form field flags
 - text field flags
- field hierarchy **530**, 531
- FDF 561, 567
- inheritance of attributes 528, 530
- in Linearized PDF 740
- and reset-form actions 555
- and submit-form actions 551, 553
- field names **532–533**, 801
- alternate 531, 657
- fully qualified. *See* fully qualified field names
- mapping name **531**, 552
- partial 531, **532–533**, 564, 801
- renaming 567, 804
- in submit-form actions 551, 552, 554
- field types **537–550**
- Btn** **531**, 539, 540
- Ch** **531**
- Sig** **531**, 548
- Tx** **531**
- field values 531, 564
- for checkbox fields 539, 540
- for choice fields 545, 546, 547
- and default appearance strings 535
- and dynamic appearance streams 535
- for FDF fields 562
- for radio button fields 541, 542, 543
- for signature fields 548
- in submit-form actions 551, 552, 553, 554
- for text fields 543, 544
- fields, FDF **564–566**, 803–804
- actions for 565
- additional-actions dictionary 565
- default appearance strings 565
- exporting 557, 558, 561, 803
- hierarchy. *See* field hierarchy
- importing 557, 558, 561, 563, 564, 565, 567, 803
- name. *See* field names
- renaming 567, 804
- root 561
- value. *See* field values
- fields, interactive form 9, 528–529
- access permissions for 75, 77
- additional-actions dictionary 532
- computation order 516, **529**
- default appearance string 534, **535**, 546, 547
- default value **532**, 554
- dynamic appearance streams 512, 529, **533–537**
- exporting 9, 528, 531, 532, 542, 543, 546, 548
- file-select controls 4, 543, **544**
- flags. *See* field flags

- fields, interactive form (*continued*)
 - Form standard structure type 637
 - formatting 516, 517
 - and hide actions 527
 - hierarchy. *See* field hierarchy
 - importing 9, 528
 - inheritance of attributes 528, 530
 - and JavaScript actions 556
 - mapping name 531, 552
 - name. *See* field names
 - nonterminal 531
 - quadding 534, 535
 - read-only 532
 - recalculation 516, 517, 529
 - root 529, 567
 - terminal 528, 531
 - trigger events for 516, 517, 532
 - type. *See* field types
 - Unicode specification of export values 4
 - unique name (Web Capture) 665, 666, 667
 - validation 516, 517
 - value. *See* field values
 - variable text. *See* variable text
- See also*
 - button fields
 - checkbox fields
 - choice fields
 - combo box fields
 - field dictionaries
 - list box fields
 - pushbutton fields
 - radio button fields
 - signature fields
 - text fields
 - widget annotations
- Fields** entry
 - FDf dictionary 561, 562
 - FDf template dictionary 567
 - interactive form dictionary 529
 - reset-form action dictionary 555
 - submit-form action dictionary 551, 553, 554
- Figure standard structure type 637, 638
 - standard layout attributes for 640, 644, 645, 649
- file attachment annotation dictionaries 509
 - Subtype** entry 509
- file attachment annotations 9, 499, 508–509
 - contents 509
 - See also*
 - file attachment annotation dictionaries
- file body
 - FDf 558, 559–560
 - PDF 61, 64, 757, 774
- File Format for Color Profiles* (International Color Consortium) 814
- file formats
 - ACFM (Adobe Composite Font Metrics) 300
 - Acrobat[®] products, native 787, 792
 - Adobe Type 1. *See* Type 1 fonts
 - AFM (Adobe Font Metrics) 300, 319, 812
 - AIFF (Audio Interchange File Format) 569
 - AIFF-C (Audio Interchange File Format, Compressed) 569
 - ASCII (American Standard Code for Information Interchange). *See* ASCII
 - CFF (Compact Font Format) 5, 365, 367
 - CGI (Common Gateway Interface) 664
 - CSS (Cascading Style Sheets) 622, 624, 639
 - DVI (Device Independent) 19
 - FDf (Forms Data Format). *See* Forms Data Format
 - GIF (Graphics Interchange Format) 573, 659, 660, 661, 662, 665
 - HPGL (Hewlett-Packard Graphics Language) 19
 - HTML (Hypertext Markup Language)
 - See*
 - HTML
 - HTML Form format
 - JDF (Job Definition Format) 374, 677, 689, 692, 806
 - JPEG (Joint Photographic Experts Group) 573, 659
 - OEB (Open eBook[™]) 639
 - PCL (Printer Command Language) 19
 - PDF/X (Portable Document Format, Exchange) 684, 685, 686
 - PJTF (Portable Job Ticket Format) 24, 374, 677, 689, 692, 806
 - PostScript. *See* PostScript[®] page description language
 - RIFF (Resource Interchange File Format) 569
 - RTF (Rich Text Format) 613, 624, 627, 639
 - SGML (Standard Generalized Markup Language) 589
 - snd 569
 - TIFF (Tag Image File Format) 46, 50, 697
 - TrueType. *See* TrueType[®] fonts
 - XML (Extensible Markup Language). *See* XML
 - XSL (Extensible Stylesheet Language) 622, 624
- file header
 - FDf 558, 559, 561, 802
 - PDF 61, 63, 70, 83, 784, 785, 786, 790, 791
- file identifiers 573, 580–581, 805
 - encryption 78, 80
 - FDf 561
 - file specifications 123
 - file trailer 68
 - reference XObjects 288
- file names
 - compatibility 121
 - DOS/Windows 121
 - drive identifier 120
 - empty 119
 - extension 121, 125, 558

- file names (*continued*)
 - and file identifiers 580
 - in file-select controls 543, 544
 - in file specifications 119
 - Mac OS 121
 - network resource name 120
 - platform-dependent 119, 120–121
 - server name 120
 - volume name 120
- file-select controls 4, 543, **544**
- file specification dictionaries **122–123**, 128
 - DOS** entry 122, **123**
 - EF** entry **123**, 124, 128, 129
 - F** entry **122**, 127
 - FS** entry **122**, 127
 - ID** entry **123**
 - Mac** entry 122, **123**
 - metadata inapplicable to 579
 - RF** entry **123**, 126, 128
 - Type** entry **122**, 123, 128
 - Unix** entry 122, **123**
 - V** entry **123**
- file specification strings **118–122**
 - DOS 123
 - Mac OS 123
 - UNIX 123
- file specifications **118–129**
 - absolute **119–120**
 - conversion to platform-dependent file names 120–121
 - dictionaries. *See* file specification dictionaries
 - as dictionary values 99
 - embedded file streams **123–127**, 562
 - in file-select controls 544
 - full **118**
 - maintenance 128–129
 - for movie files 571
 - multiple-byte strings in 122
 - related files arrays 123, **125–127**, 128
 - relative **119–120**, 802, 803
 - simple **118**
 - for sound files 568, 800
 - strings. *See* file specification strings
 - URL 119, **127**, 551
 - volatile files 123
- file streams, embedded
 - See* embedded file streams
- file structure
 - PDF** **558–560**
 - PDF** **61–71**, 790
- file systems 122
 - CD-ROM 121
 - handlers 723
- file systems (*continued*)
 - hierarchy 120
 - local 118
 - naming conventions 118, 121
 - plug-in extensions for 723
 - URL** 122
 - URL-based 119
- file trailer
 - PDF** 558, **560**
 - Linearized PDF 729, **732**, 734, 741
 - PDF **61**, **67–68**, 69, 790
 - example 757, 774, 775, 777, 780, 782
 - See also*
 - file trailer dictionary
- file trailer dictionary 67, **68**
 - custom data prohibited in 723
 - Encrypt** entry **68**, 71, 734
 - ID** entry **68**, 78, 80, 561, 581, 805
 - Info** entry **68**, 575
 - metadata inapplicable to 579
 - Prev** entry **68**, 69, 732, 733, 740, 741, 775, 777, 780, 782
 - Root** entry 63, **68**, 83, 732
 - Size** entry **68**, 732, 741
- file type (Mac OS) 125, 558
- FileAttachment** annotation type 499, **509**
- files
 - attached. *See* file attachment annotations
 - binary 15, 25
 - CIDFont **336**
 - CMap. *See* CMap files
 - creator signature (Mac OS) 125
 - PDF**. *See* PDF files
 - file type (Mac OS) 125, 558
 - file-select controls. *See* file-select controls
 - font. *See* font files
 - formats. *See* file formats
 - movie 571
 - PDF. *See* PDF files
 - related **125–127**, 128
 - remote go-to action, target of 520
 - resource fork (Mac OS) 125
 - sampled image 694, 697
 - sound **568–569**
 - specifications 99, **118–129**, 544
 - text 15, 25
 - thread action, target of 522
 - URL specifications 119, **127**, 551
 - volatile 123
- FileSelect field flag (text field) **543**, **544**
- Filespec** object type **122**, 128
- fill** operator (PostScript) 699, 700

filling

- color. *See* nonstroking color, current
- color space. *See* nonstroking color space, current
- even-odd rule 167, 169, **170–171**, 699, 700
- glyphs 305, 367, 793
- nonzero winding number rule 167, 169, **169–170**, 699, 700
- paths 12, 131, 132, 151, 167, **168–171**, 699, 700, 762
- scan conversion 406
- text 12, 132, 305
- text rendering mode 305, 367, 793
- and transparent overprinting 461, 462–463, 466

Filter entry

- encryption dictionary 71, **72**, 79
- inline image object **279**
- signature dictionary **549**
- stream dictionary **38**, 41, 366, 448, 569

filter parameter dictionaries 38, 41

See also

- CCITTFaxDecode** filter parameter dictionaries
- DCTDecode** filter parameter dictionaries
- FlateDecode** filter parameter dictionaries
- JBIG2Decode** filter parameter dictionaries
- LZWDecode** filter parameter dictionaries

filters 37, 38, **41–61**, 788–790

- abbreviations for **279–280**, 788–789
- ASCII **41**, **44–45**
- compression 15, 22, 569
- in content streams 94, 786, 789
- decoding **41–61**, 73, 736, 788–790
- decompression **41**, **45–61**
- in embedded font programs 366, 789
- encoding 73
- in form XObjects 789
- in image streams 268, 789
- in inline images 280, 789
- metadata streams, not recommended for 577
- parameters. *See* filter parameter dictionaries
- pipeline 41
- standard **42**
- in thumbnail images 789
- in Type 3 glyph descriptions 789
- unrecognized 789

See also

- ASCII85Decode** filter
- ASCIIHexDecode** filter
- CCITTFaxDecode** filter
- DCTDecode** filter
- FlateDecode** filter
- JBIG2Decode** filter
- LZWDecode** filter
- RunLengthDecode** filter

Final annotation icon 507

Find command (Acrobat) 789

findfont operator (PostScript) 290

fingerprints (user authentication) 538, 547

first-class names 124, **723–724****First** entry

- outline dictionary 477, **478**
- outline item dictionary 477, **478**

first-page cross-reference table (Linearized PDF) 729, **732**, 734, 744

- document-level objects indexed by 734
- hint streams in 735
- linearization parameter dictionary in 732
- and main cross-reference table 740
- startxref** line and 740, 751

first-page section (Linearized PDF) 730, **736–738**, 745, 808

- and primary hint stream, ordering of 735, 736, 751, 752

first-page trailer (Linearized PDF) 729, **732**, 734**FirstChar** entry

- Type 1 font dictionary **317**, 318, 319, 795
- Type 3 font dictionary **324**

FirstPage named action 527

See also named-action dictionaries

FitWindow entry (viewer preferences dictionary) 472fixed-pitch fonts **297**, 358FixedPitch font flag **358**, 622**FL** entry (graphics state parameter dictionary) **159**, 404**FI** filter abbreviation **280**, 789

flags

See

- access flags
- annotation flags
- button field flags
- choice field flags
- edge flags
- field flags
- font flags
- outline item flags
- reset-form field flag
- signature flags
- submit-form field flags
- text field flags
- Web Capture command flags

Flags entry

- font descriptor **356**, 357, 622
- reset-form action dictionary 554, **555**
- submit-form action dictionary 550, **551**, 553

Flate (zlib/deflate) compression 15, 42, **45–52**

- predictor functions 46, 49, **50–52**, 268

- FlateDecode** filter 42, 45–52, 193
 - Fl** abbreviation 280, 789
 - parameters. *See* **FlateDecode** filter parameter dictionaries
 - predictor functions 50–52
 - in sampled images 268
- FlateDecode** filter parameter dictionaries 48–49
 - BitsPerComponent** entry 49
 - Colors** entry 49
 - Columns** entry 49
 - Predictor** entry 49, 50–51
- flatness tolerance 150, 404
 - FL** entry (graphics state parameter dictionary) 159, 404
 - i* operator 156, 404, 700
 - and smoothness tolerance, compared 405
- floating elements 626, 641–642, 643, 644
 - bounding box 624
- floating windows
 - for movies 525, 572
- floor** operator (PostScript) 116, 703
- Fo** entry (additional-actions dictionary) 515
- Fo** trigger event (annotation) 515
- focus, input
 - See* input focus
- fold marks 679
- folios 615
- FOND resource (Mac OS) 321
- font characteristics 614, 622
 - Italic 622
 - Proportional 622
 - Serifed 622
 - Smallcap 622
- font descriptors 16, 316, 355–364, 760, 796
 - Ascent** entry 356, 647
 - AvgWidth** entry 357
 - CapHeight** entry 356
 - CharSet** entry 357, 368
 - for CIDFonts. *See* CIDFont font descriptors
 - Descent** entry 356, 647
 - embedded font programs 357, 365
 - FD** dictionary. *See* CIDFont **FD** dictionaries
 - Flags** entry 356, 357, 622
 - flags. *See* font flags
 - for font subsets 322, 796
 - FontBBox** entry 356
 - FontFile** entry 357, 365, 366, 738
 - FontFile2** entry 357, 365
 - FontFile3** entry 357, 365, 366, 796
 - FontName** entry 322, 356, 357, 361
 - ItalicAngle** entry 356
 - Leading** entry 356
 - in Linearized PDF 738
- font descriptors (*continued*)
 - MaxWidth** entry 357
 - metadata inapplicable to 579
 - MissingWidth** entry 317, 357
 - StemH** entry 357
 - StemV** entry 357
 - Style** dictionary. *See* CIDFont **Style** dictionaries
 - for TrueType fonts 357
 - Type 0 fonts, lacking in 356
 - for Type 1 fonts 317–318, 357
 - Type 3 fonts, lacking in 356
 - Type** entry 356
 - XHeight** entry 357
- font dictionaries 36, 291, 292, 314–315, 316
 - BaseFont** entry 356
 - in CID-keyed fonts 336
 - compact fonts 367
 - encoding 328, 329
 - font matrix 140
 - glyph metrics in 297, 298
 - metadata inapplicable to 579, 580
 - as named resources 97, 293
 - in PostScript 315
 - standard fonts 319
 - Subtype** entry 36, 292, 314
 - text font parameter 158
 - ToUnicode** entry 369, 620
 - in trap networks 691
 - Type** entry 36
 - Widths** entry 357
 - See also*
 - CIDFont dictionaries
 - multiple master font dictionaries
 - TrueType font dictionaries
 - Type 0 font dictionaries
 - Type 1 font dictionaries
 - Type 3 font dictionaries
- Font** entry
 - graphics state parameter dictionary 158
 - resource dictionary 97, 290, 293, 302, 317, 338, 534, 535
- font files 292, 364
 - embedded 292
 - external 291, 292
 - in font descriptors 316, 738
 - metadata 580
- font flags 329, 356, 357–360
 - AllCap 358
 - FixedPitch 358, 622
 - ForceBold 358, 359
 - Italic 358, 622
 - Nonsymbolic 358
 - Script 358
 - Serif 358, 622

- font flags (*continued*)
 - SmallCap 358, 622
 - Symbolic 358
- font management 316
- font matrix 140, 298, 323–324
 - rotation 324
 - and Type 3 glyph descriptions 325
- font names 293
 - conventions 316
 - font subsets 322–323
 - multiple master 320
 - PostScript 314, 317, 320, 321–322, 353, 356
 - Type 0 fonts 353
 - Type 1 fonts 317
 - Type 3 fonts 323
- Font Naming Issues* (Adobe Technical Note #5088) 320, 812
- font numbers 336, 342, 348, 349, 352, 353, 354
- Font** object type 36, 317, 323, 338, 353, 760
- font programs 5, 291, 292, 314, 315–316, 323
 - compact 365, 367
 - copyright permissions 364–365
 - embedded. *See* embedded font programs
 - encoding 316
 - external 314, 316
 - font dictionaries, defined in 293
 - glyph metrics in 297, 298, 317
 - hints 316, 323
 - multiple master 319
 - in PostScript 315
 - TrueType 321
 - Type 1 316, 321, 794
- Font** resource type 97, 290, 293, 302, 317, 338, 534, 535, 738
- font subsets 322–323, 796
 - BaseFont** entry 322, 796
 - character set 357
 - embedded 16, 368
 - font descriptors for 322, 796
 - merging 322
 - name 322–323
 - PostScript name 322–323, 338
 - tag 322–323, 357, 361, 796
- font subtypes
 - See* font types
- font types 292, 314, 315
 - MMType1** 315, 320, 365
 - TrueType** 36, 315, 321, 365
 - Type0** 315, 334, 353
 - Type1** 36, 315, 317, 365
 - Type3** 315, 323
 - See also*
 - CIDFont types
- FontBBox** entry
 - font descriptor 356
 - Type 3 font dictionary 323
- FontDescriptor** entry
 - CIDFont dictionary 338
 - Type 1 font dictionary 317, 318, 319, 795
- FontDescriptor** object type 356
- FontDescriptor** resource type 738
- FontFauxing** entry (trap network annotation dictionary) 691
- FontFile** entry (font descriptor) 357, 365, 366, 738
- FontFile2** entry (font descriptor) 357, 365
- FontFile3** entry (font descriptor) 357, 365, 366, 796
- FontMatrix** entry (Type 3 font dictionary) 298, 323, 324, 325
- FontName** entry
 - font descriptor 322, 356, 357, 361
 - Type 1 font program 317
- fonts 2, 12, 23, 291, 292, 314–368, 760
 - all-cap 358
 - ascent 356
 - availability 315
 - average glyph width 357
 - bounding box 323, 325, 326, 356, 700
 - cap height 356
 - CFF (Compact Font Format) 5, 365, 367
 - character collections 335
 - character sets 292, 329, 335
 - characteristics. *See* font characteristics
 - CJK (Chinese, Japanese, and Korean) 322
 - and CMaps 342, 349
 - content streams 93
 - current 12, 293, 294, 334, 354
 - data structures 291, 314–316
 - descent 356
 - descriptor 16, 316
 - encoding. *See* character encodings
 - expert 329, 710
 - “fauxing” 691
 - files. *See* font files
 - fixed-pitch 297, 358
 - flags. *See* font flags
 - formats 16
 - glyph selection 316
 - glyph space 140, 298, 323, 356
 - glyphs in 25
 - in Linearized PDF 738, 739, 754
 - interpreter 315, 316
 - italic 358
 - italic angle 356
 - kerning information 300
 - leading 356
 - management 15–16

fonts (*continued*)

- matrix **140, 298**, 323–324, 325
- maximum glyph width 357
- metadata for 580
- metrics 16, 22, 297–300, 316, 317, 319, 355
- monospaced **297**
- multiple master **319–321**, 796
- name. *See* font names
- nonsymbolic **329**, 330, 333, 358, 359
- number 336, **342**, 348, 349, 352, 353, 354
- organization and use 292
- PostScript files 22
- proportional **297**, 358
- resources 293, 302
- sans serif 358
- scaling 294, 302
- script 358
- selection 291, 294
- serifed 358
- size 294
- small-cap 358
- standard. *See* standard 14 fonts
- stem height 357
- stem width 357
- style information 16
- subsets 16, **322–323**, 368, 796
- substitution 16, 22, 316, 319, 356, 362, 738, 754, 794, 796–797
- subtype. *See* font types
- symbolic **329**, 330, 333, 358, 359
- Tagged PDF, determination of characteristics in 622
- and text operators 131
- type. *See* font types
- variable-pitch **297**, 358
- for variable-text fields 534
- x height 357

See also

- CID-keyed fonts
- composite fonts
- font descriptors
- font dictionaries
- font programs
- simple fonts
- TrueType[®] fonts
- Type 0 fonts
- Type 1 fonts
- Type 3 fonts
- Type 42 fonts

footnotes

- as BLSEs 629
- marked content 583
- Note standard structure type 633
- and page content order 618
- placement of 633

footnotes (*continued*)

- as structure elements 591, 592
- Tagged PDF 589
- ForceBold font flag **358, 359**
- ForComment annotation icon 507
- form actions
 - See*
 - import-data actions
 - JavaScript actions
 - reset-form actions
 - submit-form actions
- form dictionaries 282, **283–286**, 793
 - for dynamic appearance streams 534
 - See also*
 - printer's mark form dictionaries
 - type 1 form dictionaries
- form feed (FF) character **26, 32**
 - escape sequence for 30
- form fields
 - See* fields, interactive form
- form matrix **141**, 283, 284
- form space **141**, 220, **283**
 - and annotation appearances 496
 - and dynamic appearance streams 534
- Form standard structure type 618, **637**
 - standard layout attributes for 640, 644, 645, 649
- form types **283**, 284
 - type 1 **283–286**
- Form** XObject subtype 261, **284**
- form XObjects **133, 261, 281–289**, 449
 - annotation appearances 282, 496, 534, 682
 - annotation icons 537
 - bounding box 284, 534
 - clipping to bounding box 283, 284
 - content stream 93, 281, 282, 283, 614, 615
 - defining 282
 - dictionaries. *See* form dictionaries
 - form matrix **141**, 283, 284
 - form space **141**, 220, **283**, 496, 534
 - form type **283**, 284
 - for importing content 282
 - and interactive forms, distinguished 281–282, 528
 - in logical structure elements 595–598
 - marked-content sequences in 594
 - metadata 285
 - modification date 284, 582
 - name 284
 - as OPI proxies 693, 807
 - OPI version dictionary 285
 - page-piece dictionary 285
 - as page templates 282
 - painting 282–283, 451, 452
 - patterns and 220

- form XObjects (*continued*)
 - private data associated with 581, 582
 - for repeated graphical elements 282
 - resource dictionary 284–285, 691
 - resources 96, 284
 - soft masks 282
 - transparency groups 282, 786
 - trap networks 689, 692
 - unrecognized filters in 789
 - uses 282
 - See also*
 - group XObjects
 - reference XObjects
 - transparency group XObjects
- forms
 - See*
 - form XObjects
 - interactive forms
- Forms Data Format (FDF) **557–568**
 - annotations in 557, 558, 562
 - catalog. *See* FDF catalog
 - differences stream 562
 - digital signatures 4, 562
 - encryption 562–563
 - exporting 557, 558, 561, 803
 - fields. *See* fields, FDF
 - in file-select controls 544
 - files. *See* FDF files
 - in import-data actions 550, 555, 557
 - importing 557, 558, 561, 563, 564, 565, 567, 803
 - objects 558
 - options 562, 565
 - pages **566–568**, 804
 - PDF, compared with 558
 - PDF 1.4 enhancements 4
 - PDF syntax 24, 558
 - remote collaboration 4
 - in submit-form actions 552, 553, 554, 562, 802
 - template pages **567–568**, 804
 - trailer. *See* FDF trailer dictionary
 - and trigger events 517
 - version specification 558, 559, 561, 802
 - See also*
 - FDF annotation dictionaries
 - FDF dictionary
 - FDF field dictionaries
 - FDF named page reference dictionaries
 - FDF page dictionaries
 - FDF template dictionaries
- FormType** entry (form dictionary) **284**
- Formula standard structure type **637**
 - standard layout attributes for 640, 644, 645, 649
- ForPublicRelease annotation icon 507
- FP** entry (additional-actions dictionary, obsolete) 800
- “fpgm” table (TrueType font) 367
- FrameMaker[®] document publishing software 576
- free-form Gouraud-shaded triangle meshes
 - See* type 4 shadings
- free text annotation dictionaries **502**
 - Contents** entry **502**
 - DA** entry **502**
 - Q** entry **502**
 - Subtype** entry **502**
- free text annotations 499, **502**
 - contents 502
 - default appearance string 502
 - quadding 502
 - and text annotations, compared 502
 - See also*
 - free text annotation dictionaries
- FreeText** annotation type 499, **502**
- frequency (halftone screen) **384**, 391, 392, 393, 394, 395, 397
- Frequency** entry (type 1 halftone dictionary) **393**
- “from CIE” information (ICC color profile) 192, 686
- FS** entry
 - file specification dictionary **122**, 127
- FT** entry (field dictionary) **531**, 533, 548
- FTP (File Transfer Protocol) 664
- Fujitsu FMR character set 344
- full file specifications **118**
- FullScreen page mode **84**, 472
- fully opaque objects **467**
- fully qualified field names **532–533**
 - for FDF fields 564, 803
 - in hide actions 527
 - in reset-form actions 555
 - in submit-form actions 551, 553, 554
- function-based shadings
 - See* type 1 shadings
- function dictionaries **107–108**
 - Domain** entry **108**, 109, 113, 114, 117
 - FunctionType** entry **108**
 - metadata inapplicable to 579
 - Range** entry **108**, 109, 111, 113, 117
 - See also*
 - type 0 function dictionaries (sampled)
 - type 2 function dictionaries (exponential interpolation)
 - type 3 function dictionaries (stitching)
 - type 4 function dictionaries (PostScript calculator)
- Function** entry
 - shading dictionary 235, 236
 - type 1 shading dictionary **237**
 - type 2 shading dictionary **238**, 239
 - type 3 shading dictionary **240**, 241

Function entry (*continued*)

- type 4 shading dictionary 244, 245, 247
- type 5 shading dictionary 249
- type 6 shading dictionary 253, 255

function objects 106, 107, 235, 381

function types 107–108

- type 0 (sampled) 107, 108, 109–113, 792
- type 2 (exponential interpolation) 107, 108, 113, 792
- type 3 (stitching) 108, 113–115, 792
- type 4 (PostScript calculator) 3, 108, 115–118, 792

functions 106–118

- black-generation 379
- clipping to domain 108, 110
- clipping to range 108, 111
- color (shadings). *See* color functions
- color mapping 375
- decoding 181, 183, 184, 186, 188
- dictionaries. *See* function dictionaries
- as dictionary values 99
- dimensionality 108, 109
- domain 107, 108, 109, 110, 114, 237, 238, 240, 244, 249, 253
- function objects 106, 107, 235, 381
- gamma 183, 184, 186, 455
- gamut mapping 185, 375, 380
- interpolation 243
- range 107, 108, 109, 110, 111
- spot. *See* spot functions
- tint transformation. *See* tint transformation functions
- transfer. *See* transfer functions
- type. *See* function types
- undercolor-removal 379
- See also*
 - type 0 functions (sampled)
 - type 2 functions (exponential interpolation)
 - type 3 functions (stitching)
 - type 4 functions (PostScript calculator)

Functions entry (type 3 function dictionary) 114

FunctionType entry (function dictionary) 108

Fundamentals of Interactive Computer Graphics (Foley and van Dam) 813

FWPosition entry (movie activation dictionary) 572

FWScale entry (movie activation dictionary) 572

G

G color space abbreviation (inline image object) 280

G entry

- soft-mask dictionary 445, 446, 450
- Web Capture command settings dictionary 674

G operator 134, 173, 177, 178, 179, 217, 218, 700

g operator 134, 173, 177, 178, 179, 217, 218, 264, 295, 700

gamma correction 173, 373, 380, 381, 382

CalGray color spaces 183

CalRGB color spaces 185

gamut mapping functions, distinguished from 380

Gamma entry

CalGray color space dictionary 183, 184

CalRGB color space dictionary 184, 185, 186, 440

gamma functions 183, 184, 186, 455

gamut

color space 187, 205, 375, 468, 469

device 197, 198, 375

source (page) 375

gamut mapping functions 185, 375

gamma correction, distinguished from 380

garbage collection 805

GB 2312-80 character set 343

GB 18030-2000 character set 344

GB-EUC-H predefined CMap 343, 346

GB-EUC-V predefined CMap 343, 346

GBK character encoding 343, 803

GBK character set 343

GBK-EUC-H predefined CMap 343, 346

GBK-EUC-V predefined CMap 343, 346

GBKp-EUC-H predefined CMap 343, 346

GBKp-EUC-V predefined CMap 343, 346

GBK2K-H predefined CMap 344, 346

GBK2K-V predefined CMap 344, 346

GBpc-EUC-H predefined CMap 343, 346

GBpc-EUC-V predefined CMap 343, 346

GDI (Graphics Device Interface) imaging model 19, 20

ge operator (PostScript) 116, 704

general graphics state operators 134

d 134, 156, 700

gs 134, 156, 157, 159, 218, 301, 307, 374, 446, 700

i 134, 156, 404, 700

J 134, 156, 700

j 134, 156, 700

M 134, 156, 700

ri 134, 156, 197, 216, 218, 701

in text objects 309

w 134, 151, 156, 296, 702

general layout attributes 640–642

Placement 626, 629, 632, 637, 638, 640, 641, 643, 644, 649

WritingMode 624, 640, 642, 647, 651

generation numbers 39

attribute revision numbers, distinguished from 606

in cross-reference table 65, 66, 707

and encryption 73

in FDF files 558

- generation numbers (*continued*)
 - and incremental updates 69
 - in Linearized PDF 729
 - in updating example 774, 775, 779, 780
- Generic** character class 362, 363
- generic hint tables (Linearized PDF) 741, **749–750**, 751
- GenericRot** character class 363
- “Geometrically Continuous Cubic Bézier Curves” (Seidel) 813
- Geschke, Chuck xx
- GET request (HTTP) 552, 670, 673
- GetMethod field flag (submit-form field) **552**, 553
- GIF (Graphics Interchange Format) file format 573, 659, 660, 661, 662, 665
- Glitter transition style **486**, 487
- “glyph” table (TrueType font) 367
- glyph coordinate system
 - See glyph space
- glyph descriptions 291, **292**
 - for character collections 336, 337
 - and character encodings 328
 - in CID-keyed fonts 335, 336
 - in CIDFonts 337, 339
 - color 326
 - in font subsets 16
 - and graphics state 325
 - restrictions on 218
 - text objects in 309
 - and Tj operator 294
 - in TrueType fonts 332, 333, 334, 339, 345, 367
 - in Type 1 fonts 316
 - in Type 3 fonts 323, 324, 325, 789
- glyph displacement **297**, **299**, 340
 - character spacing 302
 - in CIDFonts 341
 - displacement vector **299**, 304
 - DW2** entry (CIDFont) 341
 - horizontal scaling 304
 - in right-to-left writing systems 619
 - simple fonts 316
 - text matrix, updating of 313–314
 - and text-positioning operators 310
 - text space 313
 - in Type 3 fonts 326
 - W2** entry (CIDFont) 341–342
 - word spacing 303
- glyph indices **339**, 345
- glyph origin **298**
 - positioning 309, 341
 - in right-to-left writing systems 619
 - and writing mode 299, 340
- glyph space **140**, **298**, 323, 356
 - displacement vector 299
 - font descriptors expressed in 356
 - glyph widths 324
 - origin **298**, 299, 309, 340, 341, 619
 - text space, relationship with 298, 313, 323
 - and Type 3 glyph descriptions 325, 326
 - units 294, 298, 323
 - user space, mapping to 325, 647
- glyph widths 297, **299**
 - in CIDFonts 338, 340–341
 - CJK (Chinese, Japanese, and Korean) 340
 - d0** operator 326, 700
 - d1** operator 326, 700
 - default 338, 340
 - and horizontal scaling 304
 - in printing 794
 - in reflow of content 794
 - in right-to-left writing systems 619
 - in searching of text 794
 - in simple fonts 316
 - in Type 1 fonts 317, 794–795
 - in Type 3 fonts 324, 325, 326, 700
 - undefined 357
 - in viewing of documents 794
- glyphs, character 2, 291, **292**
 - Adobe imaging model 11
 - bold 359
 - bounding box **299**
 - and characters, contrasted 25, 292
 - as clipping path 296, 305
 - descriptions. See glyph descriptions
 - displacement. See glyph displacement
 - emulating tiling patterns with 219
 - filling 305, 367, 793
 - fixed-pitch 358
 - font management 16
 - in font subsets 796
 - font substitution 16
 - in fonts 292, 314
 - indices **339**, 345
 - italic 358
 - metrics 297–300, 316, 343
 - object shape 443
 - origin **298**, 299, 309, 340, 341, 619
 - painting 293, 294, 305–306, 368, 462, 794
 - position vector **299**, 341–342
 - positioning 297–300, 309–311
 - in right-to-left writing systems 619
 - scaling 291, 302, 701
 - scan conversion 13, 14, 407
 - script 358
 - selection 316
 - serifs 358

- glyphs, character (*continued*)
 - shading patterns, painting with 232
 - size 294
 - small capitals 358
 - stencil masking 276–277
 - stroking 296, 305, 367, 794
 - text knockout flag 159
 - text knockout parameter 308
 - in text objects 12
 - text operators 131
 - width. *See* glyph widths
- go-to action dictionaries **519**
 - D** entry **519**
 - S** entry **519**
- go-to actions 518, **519**
 - destination 474, 519
 - named destinations, targets of 476
 - and URI actions 501
 - See also*
 - go-to action dictionaries
 - remote go-to actions
- GoTo** action type 518, **519**
- GoToR** action type 518, **520**
- Gouraud interpolation **243**, 247, 255, 813
- Gouraud-shaded triangle meshes
 - free-form. *See* type 4 shadings
 - lattice-form. *See* type 5 shadings
- gradient fills **219**
 - color conversion in 236
 - geometry independent of object painted 232
 - interpolation algorithms 235
 - sh** operator 232
 - shading dictionaries 233
 - shading objects, defined by 231
 - smoothness tolerance 151
 - streams, defined by 235
 - in tiling patterns 232
- Graph annotation icon 509
- graphics 2, **131–290**, 757
 - and rendering, distinguished 132, 373
 - special text effects 295–297
 - See also*
 - clipping
 - color spaces
 - color values
 - coordinate systems
 - coordinate transformations
 - cubic Bézier curves
 - device space
 - external objects (XObjects)
 - graphics objects
 - graphics operators
 - graphics state
- graphics (*continued*)
 - See also*
 - images, sampled
 - paths
 - patterns
 - rendering
 - text
 - transformation matrices
 - transparent imaging model
 - user space
 - Graphics Gems* (Glassner, ed.) **814**
 - Graphics Gems II* (Arvo, ed.) **813**
 - graphics objects 9, 11–12, 19, **132–136**
 - artifacts 615
 - clipping of 162, 172
 - color spaces for 194
 - colors of 172, 192
 - in composite pages 683
 - compositing of 441, 452
 - coordinate spaces for 137, 138
 - coordinate transformations, unaffected by subsequent 145
 - in form XObjects 261, 281, 282, 283
 - in glyph descriptions 309, 324
 - in illustration elements (Tagged PDF) 637
 - in ILSEs (Tagged PDF) 648
 - invisible **586**, 587
 - logical structure, independent of 589
 - marked-content operators prohibited within 583
 - in marked-content sequences 583, 594
 - page content order 618
 - in pattern coordinate space 220, 223
 - rendering of 147, 374
 - shape (transparent imaging model) 171
 - in table cells (Tagged PDF) 648
 - in tiling patterns 219, 223, 453
 - in transparency groups 451
 - in transparent imaging model 459
 - and transparent overprinting 461
 - in trap networks 689, 692
 - types 132–133
 - visible **586**, 587
 - graphics operators 2, 131–132
 - in glyph descriptions 323
 - in logical structure content 593
 - See also*
 - clipping path operators
 - color operators
 - graphics state operators
 - inline image operators
 - path construction operators
 - path-painting operators
 - shading operator
 - text object operators

graphics operators (*continued*)*See also*

- text-positioning operators
- text-showing operators
- text state operators
- Type 3 font operators
- XObject operator

graphics state 12, **131**, **147–161**

- and annotation appearances 496
- compatibility operators 95
- and form XObjects 281, 283, 452
- initialization 147, 452, 453
- and marked-content sequences 583
- and OPI proxies 693
- page description level 134
- parameter dictionaries 97
- saving and restoring 152, 156, 162, 172, 223, 266, 283, 325, 452, 701
- and shading patterns 231
- stack 152, 156
- and transparent patterns 453, 454
- and Type 3 glyph descriptions 325

See also

- graphics state operators
- graphics state parameter dictionaries
- graphics state parameters
- text state

graphics state operators 21, 131, 151, **156**, 157

- cm** 134, 139, 148, **156**, 266, 699
- d** 134, **156**, 700
- in default appearance strings 534, 535
- general **134**, 309
- gs** 134, **156**, 157, 159, 218, 301, 307, 374, 446, 700
- i** 134, **156**, 404, 700
- J** 134, **156**, 700
- j** 134, **156**, 700
- M** 134, **156**, 700
- and marked-content operators 583
- Q** 134, 152, **156**, 172, 223, 266, 283, 296, 306, 535, 587, 701, 706, 807
- q** 134, 152, **156**, 172, 223, 266, 283, 535, 587, 701, 706, 807
- ri** 134, **156**, 197, 216, 218, 701
- special **134**
- w** 134, 151, **156**, 296, 702

graphics state parameter dictionaries 151, 156, **157–161**, 231, 700, 786

- AIS** entry 159, 443
- BG** entry 158, 218, 379
- BG2** entry 158, 218, 379
- BM** entry 159, 442
- CA** entry 159, 444
- ca** entry 159, 444
- D** entry 157
- FL** entry 159, 404
- Font** entry 158
- HT** entry 159, 218, 391
- LC** entry 157
- LJ** entry 157
- LW** entry 151, 157
- metadata inapplicable to 579
- ML** entry 157
- as named resources 97
- OP** entry 158, 213
- op** entry 158, 213
- OPM** entry 158, 214
- RI** entry 158
- SA** entry 159, 408
- SM** entry 159, 405
- SMask** entry 159, 444
- TK** entry 159, 301, 307
- TR** entry 158, 218, 381
- TR2** entry 159, 218, 381
- Type** entry 157
- UCR** entry 158, 218, 379
- UCR2** entry 158, 218, 379

graphics state parameter dictionaries (*continued*)

- device-dependent 147, **150–151**, 223, 374
- device-independent 147, **148–149**, 152
- dictionaries. *See* graphics state parameter dictionaries
- for filling 167
- initialization 147, 452, 453
- and **S** operator 166
- and sampled images 264
- setting 156, 157, 699–702
- for shading patterns 231
- for stroking 166, 167
- transparency-related 453, 797

See also

- alpha source parameter
- black-generation function
- character spacing (T_c) parameter
- current alpha constant
- current blend mode
- current clipping path
- current color
- current color space
- current halftone
- current line width
- current rendering intent
- current soft mask
- current transfer function
- current transformation matrix (CTM)
- flatness tolerance
- horizontal scaling (T_h) parameter

graphics state parameters (*continued*)

See also

- leading (T_l) parameter
- line cap style
- line dash pattern
- line join style
- miter limit
- overprint mode
- overprint parameter
- smoothness tolerance
- stroke adjustment parameter
- text font (T_f) parameter
- text font size (T_{fs}) parameter
- text knockout (T_k) parameter
- text line matrix (T_{lm})
- text matrix (T_m)
- text rendering matrix (T_{rm})
- text rendering mode (T_{mode})
- text rise (T_{rise}) parameter
- text state parameters
- undercolor-removal function
- word spacing (T_w) parameter

graphics state stack 152, 156

- depth limit 706, 807

gray color component 179

- black, complement of 377
- CMYK conversion 377
- halftones for 401, 402
- RGB conversion 377
- transfer function 381, 402

gray levels

- CMYK conversion 377
- color values 173
- DeviceGray** color space 179
- G** operator 217, 700
- g** operator 217, 700
- halftones, approximation with 382–385, 390, 393
- pixel depth 13
- RGB conversion 377

gray ramps 5, 676, 680

GrayMap entry (version 1.3 OPI dictionary) 696

grayscale color representation

- and CMYK, conversion between 377, 382
- DeviceGray** color space 176, 179
- in halftones 383
- multitone components, specifying 205
- in output devices 172, 376
- rendering 14
- and RGB, conversion between 377

Greek characters 362, 363

green color component

- CMYK conversion 377, 380
- DeviceRGB** color space 178, 179

green color component (*continued*)

- grayscale conversion 377
- halftones for 401
- in **Indexed** color table 199
- initialization 180
- magenta, complement of 378
- and threshold arrays 390
- transfer function 381

green colorant

- additive primary 178, 179, 180
- display phosphor 201
- PANTONE Hexachrome system 205

grestore operator (PostScript) 701, 706, 807

group alpha

- group backdrop, removal of 432
- in isolated groups 433
- notation 429, 431, 437
- and overprinting 463
- in page group 437
- soft masks, deriving from 439–440, 445, 446

group attributes dictionaries 286–287

- form XObject 285
- metadata inapplicable to 579
- page 89, 449, 470
- S** entry 286, 287, 449, 452
- transparency group XObject 449, 452
- Type** entry 287

See also

- transparency group attributes dictionaries

group backdrop 411, 428

- blending color space 425
- compositing with 425, 427, 429, 444, 450, 452, 454
- isolated groups, unused in 433
- in knockout groups 435
- in non-isolated groups 451
- removal from compositing computations 432, 440
- for soft masks 440, 445

group color

- in compositing 425, 427, 451, 452
- group backdrop, removal of 432
- in isolated groups 433
- in knockout groups 434
- notation 429, 437
- in page group 437
- for soft masks 439, 440

group color space 178, 192, 196, 449, 450–451, 452, 454–456

- CIE-based 455, 456
- and CompatibleOverprint blend mode 461
- in isolated groups 455
- in non-isolated groups 455
- and overprinting 459, 460, 465
- in page group 455

group color space (*continued*)
 process colors, conversion to and from 456
 rendering intents, target of 468
 spot colors not converted to 457

group compositing function (Composite)
 See Composite function

Group entry
 page object 89, 449, 470
 type 1 form dictionary 286, 449, 452, 285, 288

group hierarchy 411, 425

group luminosity
 soft masks, deriving from 412, 440–441, 445, 446

Group object type 287

group opacity 421
 in compositing 411, 425, 427, 451, 452
 in knockout groups 434
 for soft masks 439
 and spot color components 457

group shape 171, 421
 in compositing 411, 425, 427, 451, 452
 group backdrop, removal of 432
 in isolated groups 433
 notation 429, 431, 437
 for soft masks 439, 440, 441
 and spot color components 457

group stack 425, 427–428
 in isolated groups 433
 in knockout groups 434

group subtypes 286, 287, 449, 450
Transparency 286, 287, 449, 450, 452

group XObjects 133, 261, 285, 286–287, 449
 subtype 286, 287, 449, 450
 See also
 transparency group XObjects

grouping elements, standard
 See standard grouping elements

groups, transparency
 See transparency groups

gs operator 134, 156, 157, 159, 218, 301, 307, 374, 446, 700

gsave operator (PostScript) 701, 706, 807

gt operator (PostScript) 116, 704

GTS_PDFX output intent subtype 685

guideline styles (page boundaries)
 D 681
 S 681

guillemotleft character name, misspelled 714
 guillemotright character name, misspelled 714

H

H entry
 hide action dictionary 527
 inline image object 279
 linearization parameter dictionary 733, 741, 808
 link annotation dictionary 501
 Web Capture command dictionary 672, 674
 widget annotation dictionary 512

h operator 134, 162, 163, 168, 700

H predefined CMap 345, 347

H standard structure type 629, 630, 631, 632

H1–H6 standard structure types 629, 630, 632

halftone dictionaries 150, 159, 381, 391–403, 797
HalftoneName entry 392
HalftoneType entry 391, 392
Type entry 392
 See also
 type 1 halftone dictionaries
 type 5 halftone dictionaries
 type 6 halftone dictionaries
 type 10 halftone dictionaries
 type 16 halftone dictionaries

Halftone object type 392, 393, 395, 398, 400, 401

halftone screens 14, 383
 accurate screens algorithm 393–394
 angle 384, 391, 392, 393, 394, 395, 397, 399
 cells. See cells, halftone
 current transformation matrix (CTM), unaffected by 383
 device space, defined in 383
 frequency 384, 391, 392, 393, 394, 395, 397
 spot function 384–389, 391, 392, 393, 797
 See also predefined spot functions
 threshold array 389–390, 391, 392, 394, 395, 399, 400
 transfer functions for 393, 395, 398, 400, 401
 in type 5 halftones 391, 401

halftone streams 150, 159

halftone types
 threshold arrays for 390
 type 1 391, 392–394
 type 5 391, 393, 395, 398, 400–403
 type 6 391, 394–395, 398
 type 10 391, 394–398, 399
 type 16 391, 399–400

HalftoneName entry 392
 type 1 halftone dictionary 393
 type 5 halftone dictionary 401
 type 6 halftone dictionary 395
 type 10 halftone dictionary 398
 type 16 halftone dictionary 400

- halftones 14, 173, 373, 376, **382–403**
 - accurate screens algorithm 393–394
 - cells. *See* cells, halftone
 - current transformation matrix (CTM), unaffected by 383
 - device space, defined in 383
 - name 392, 393, 395, 398, 400, 401
 - proprietary 392
 - spot function 115, **384–389**, 391, 392, 393, 797
 - See also* predefined spot functions
 - threshold array **389–390**, 391, 392, 394, 395, 399, 400
 - transfer functions, applied after 380, 381, 382
 - and transparency 467–468
 - See also*
 - current halftone
 - halftone dictionaries
 - halftone screens
 - halftone types
 - type 1 halftones
 - type 5 halftones
 - type 6 halftones
 - type 10 halftones
 - type 16 halftones
- HalftoneType** entry 391, 392
 - type 1 halftone dictionary **393**
 - type 5 halftone dictionary **401**
 - type 6 halftone dictionary **395**
 - type 10 halftone dictionary **398**
 - type 16 halftone dictionary **400**
- handlers
 - action 723
 - annotation **489**, 493, 723
 - destination 723
 - file system 723
 - security **71**, 72, 74–81, 723
 - signature **547–548**, 549, 550
- hanging indent 644
- Hangul** character class 363
- hangul characters 363
- Hanja** character class 363
- hanja (hanzi, kanji) characters 361, 362, 363
- Hanzi** character class 362
- hanzi (kanji, hanja) characters 361, 362, 363
- hard hyphen character (Unicode) 617
- HardLight** blend mode 412, **418**
- “head” table (TrueType font) 367
- header, file
 - See* file header
- headings 630
- Hebrew writing systems 619, 642
- Height** entry
 - image dictionary **267**, 277, 448, 480
 - inline image object **279**
 - type 6 halftone dictionary 394, **395**, 398
 - type 16 halftone dictionary 399, **400**
- Height** standard structure attribute 637, 640, **645**, 648, 649
- Height2** entry (type 16 halftone dictionary) 399, **400**
- Help annotation icon 500
- help systems 493, 526
- Helvetica* typeface 16, 292, 293, 294, 709, 760
- Helvetica standard font **319**, **795**
- Helvetica–Bold standard font **319**, **795**
- Helvetica,Bold standard font name **795**
- Helvetica–BoldItalic standard font name **795**
- Helvetica,BoldItalic standard font name **795**
- Helvetica–BoldOblique standard font **319**, **795**
- Helvetica–Italic standard font name **795**
- Helvetica,Italic standard font name **795**
- Helvetica–Oblique standard font **319**, **795**
- HeWlett-Packard Company
 - PANOSE Classification Metrics Guide* 360, **814**
- Hexachrome™ color system, PANTONE® 205
- hexadecimal strings 29, **32**
- “hhea” table (TrueType font) 367
- Hid** entry
 - page object (obsolete) 791
- Hidden annotation flag **493**, 526, 799, 801
 - and real content 615
- hidden page elements 617
- hide action dictionaries **527**
 - H** entry **527**
 - S** entry **527**
 - T** entry **527**
- Hide** action type 518, **527**, 801
- hide actions 518, **526–527**, 801
 - and pop-up help systems 526
 - See also*
 - hide action dictionaries
- HideMenuBar** entry (viewer preferences dictionary) **472**
- HideToolBar** entry (viewer preferences dictionary) **472**
- HideWindowUI** entry (viewer preferences dictionary) **472**
- hiding and showing
 - annotations 493, 526, 527
 - document outline 84, 472
 - menu bar 472
 - navigation controls 472
 - scroll bars 472
 - thumbnail images 84, 472
 - tool bars 472
- high-fidelity color 172, 176, **205**

- highlight
 - diffuse achromatic 185
 - specular 185
- highlight annotation dictionaries
 - See markup annotation dictionaries
- Highlight** annotation type 499, **506**
- highlight annotations
 - See markup annotations
- highlighting mode (annotation) **501, 512**
 - I (invert) **501, 512**
 - N (none) **501, 512**
 - O (outline) **501, 512**
 - P (push) **501, 512**
 - T (toggle) **512**
- hint stream dictionaries **735–736**
 - A** entry **736**
 - C** entry **736**
 - E** entry **736**
 - I** entry **736**
 - L** entry **736**
 - O** entry **736**
 - S** entry **736**
 - T** entry **736**
 - V** entry **736**
- hint streams (Linearized PDF) **728, 734–736**
 - length **733, 741, 742**
 - offset **733, 741, 742**
 - See also
 - hint stream dictionaries
 - overflow hint stream
 - primary hint stream
- hint tables (Linearized PDF) **726, 741–751**
 - and document retrieval **751, 752, 753**
 - generic **741, 749–750, 751**
 - in hint streams **728, 734–735, 741**
 - information dictionary **736, 751**
 - interactive form **736, 745, 750**
 - logical structure **736, 745, 750**
 - named destination **736, 751**
 - and one-pass file generation **753**
 - outline **736, 751**
 - page label **736, 751**
 - page offset **736, 739, 741, 742–745, 747, 753, 754, 808**
 - pages, locating from **736, 738**
 - shared object **736, 739, 744, 745–747, 750, 753, 808, 809**
 - standard **736**
 - thread information **736, 751**
 - thumbnail **736, 747–749**
- hints
 - in font programs **316, 323**
 - in Linearized PDF
 - See
 - hint streams
 - hint tables
 - scan conversion **407**
- hiragana characters **362, 363**
- HKana** character class **363**
- HKanaRot** character class **363**
- HKscs–B5–H** predefined CMap **344, 346**
- HKscs–B5–V** predefined CMap **344, 346**
- “hmtx” table (TrueType font) **367**
- HoyoKanji** character class **363**
- Hong Kong SCS character encoding **344**
- Hong Kong SCS character set **344**
- horizontal scaling (T_h) parameter **301, 304**
 - text matrix, updating of **313–314**
 - text space **309**
 - TJ** operator **794**
 - Tz** operator **302, 701**
- horizontal tab (HT) character **26**
 - in comments **27**
 - escape sequence for **30**
 - as white space **24, 32**
- HPGL (Hewlett-Packard Graphics Language) file format **19**
- <href> tag (HTML) **562**
- HRoman** character class **362, 363**
- HRomanRot** character class **362, 363**
- HSL** (hue-saturation-luminance) color representation
 - for nonseparable blend modes **418**
- HSV** (hue-saturation-value) color representation
 - blending color space, prohibited for **415**
- HT** entry (graphics state parameter dictionary) **159, 218, 391**
- HTML (Hypertext Markup Language)
 - digital identifiers **665**
 - <href> tag **562**
 - hypertext links **634, 635**
 - importation of **573**
 - layout model **624**
 - and Linearized PDF **727**
 - PDF logical structure compared with **589**
 - standard attribute owners **639**
 - Tagged PDF, conversion from **613, 627**
 - target attribute **562**
 - “unsafe” characters in **664**
 - weakly structured document organization **632**
 - Web Capture **659, 660, 661, 667, 676**

- HTML-3.20** standard attribute owner **639**
 - HTML-4.01** standard attribute owner **639**
 - HTML Form format
 - in file-select controls 544
 - interactive form fields, converted to (Web Capture) 665
 - in submit-form actions 552, 554, 803–804
 - HTTP (Hypertext Transfer Protocol) 69, 664, 671
 - GET request 552, 670, 673
 - and Linearized PDF 726, 727–728
 - POST request 552, 670, 673
 - request headers 672, 674
 - Hue** blend mode **419**
 - Huffman coding 46
 - hypertext links 1, 22
 - link annotations 9, 501
 - link elements (Tagged PDF) 634, 635
 - named destinations, converted to (Web Capture) 665
 - pdfmark** language extension (PostScript) 21
 - plug-in extensions for 784
 - PostScript conversion 22
 - uniform resource identifiers (URIs) 523
 - Hypertext Markup Language 2.0 Proposed Standard* (Internet RFC 1866) 524, 554, **815**
 - Hypertext Transfer Protocol—HTTP/1.1* (Internet RFC 2068) 674, **815**
 - hyphen character (-)
 - hard 617
 - soft 617, 714
 - as word separator 623
 - hyphenation 617
 - and packing of ILSEs 625
- I**
- l border style (inset) **495**
 - l color space abbreviation (inline image object) **280**
 - l entry
 - appearance characteristics dictionary 537
 - choice field dictionary 546
 - hint stream dictionary 736
 - inline image object 279
 - thread dictionary 484, 724, 740
 - transparency group attributes dictionary 450, 451
 - l highlighting mode (invert) **501, 512**
 - i operator 134, **156**, 404, 700
 - IANA (Internet Assigned Numbers Authority) 653
 - IC** entry
 - circle annotation dictionary 505
 - line annotation dictionary 503
 - ICC. *See* International Color Consortium
 - ICC Characterization Data Registry* (International Color Consortium) 685, 687, **814**
 - ICC color profiles **189–192**
 - A*To*B* transformation 192, 415, 686, 807
 - for blending color spaces 192
 - B*To*A* transformation 192, 415, 686
 - color spaces 191, 192
 - device classes 191
 - “from CIE” information 192, 686
 - for **ICCBased** color spaces 181, 189–192
 - metadata 190, 580
 - for output devices 375
 - for output intents 192, 685, 686, 807
 - profile types 191
 - rendering intents 192
 - “to CIE” information 192, 686, 807
 - versions 190–191
 - ICC profile stream dictionaries **190**
 - Alternate** entry **190, 191**
 - Metadata** entry **190**
 - N** entry **190**
 - Range** entry **190, 216, 272**
 - ICCBased** color spaces 176, 181, **189–194**, 272, 686
 - alternate color space for 190, 191
 - bidirectional 415
 - as blending color space 415, 450
 - CIE-based *A* color spaces, representing 182
 - CIE-based *ABC* color spaces, representing 182
 - color profile **189–192**, 375
 - as default color space 195
 - as group color space 458
 - implicit conversion 196, 197
 - initial color value 216
 - metadata for 580
 - process colors, conversion to 456
 - rendering 374
 - setting color values in 217, 701
 - specification 189–190
 - spot color components, effect on in transparency groups 458
 - sRGB* (standard *RGB*) 192–193, 456
 - and transparent overprinting 465
 - icon fit dictionaries 565, **566**
 - A** entry **566**
 - for dynamic appearance streams 537
 - S** entry **566**
 - SW** entry **566**
 - icons, annotation
 - See* annotation icons
 - icSigCmykData ('CMYK') ICC profile color space 191
 - icSigColorSpaceClass ('spac') ICC profile device class 191
 - icSigDisplayClass ('mnr') ICC profile device class 191
 - icSigGrayData ('GRAY') ICC profile color space 191

- icSigInputClass ('scnr') ICC profile device class 191
 icSigLabData ('Lab ') ICC profile color space 191
 icSigOutputClass ('prtr') ICC profile device class 191
 icSigRgbData ('RGB ') ICC profile color space 191
- ID** entry
 FDF dictionary 561, 802
 file specification dictionary 123
 file trailer dictionary 68, 78, 80, 561, 581, 805
 image dictionary 269, 448, 675
 page object 90, 675
 reference dictionary 288
 structure element dictionary 591
 version 1.3 OPI dictionary 695
 Web Capture content set 668
- ID** operator 134, 278, 279, 280, 700
- Identity** mapping (**CIDToGIDMap**) 339, 345
- Identity** transfer function 158, 393, 395, 398, 400, 446
- Identity–H** predefined CMap 345, 368, 621
- Identity–V** predefined CMap 345, 368, 621
- idiv** operator (PostScript) 116, 703
- IDS** entry (name dictionary) 93, 661, 662, 663, 664, 668, 675
- IDTree** entry (structure tree root) 590, 591
- IEC. *See* International Electrotechnical Commission
- IF** entry
 appearance characteristics dictionary 537
 FDF field dictionary 565, 566
- if** operator (PostScript) 28, 116, 704
- ifelse** operator (PostScript) 28, 116, 704
- illuminated characters 651, 658
- illustration elements, standard
 See standard illustration elements
- illustrations
 bounding box 624, 648
 and page content order 618
- Illustrator[®] graphics software 411, 412, 436
- ILSEs. *See* inline-level structure elements
- IM** entry (inline image object) 279
- image coordinate system 265–267
- image dictionaries 263, 264, 267–275, 279, 793
 Alternates entry 269, 448
 BitsPerComponent entry 268, 271, 276, 277, 448, 480
 ColorSpace entry 177, 268, 269, 271, 276, 448, 449, 461, 480
 Decode entry 269, 276, 448, 480
 decoding of sample data 244, 247, 249, 253
 Height entry 267, 277, 448, 480
 ID entry 269, 448, 675
 ImageMask entry 264, 268, 269, 275, 276, 448
 Intent entry 197, 268, 448
 Interpolate entry 269, 273, 448
- image dictionaries (*continued*)
 Mask entry 268, 275, 276, 277, 444, 448, 793
 Metadata entry 269
 Name entry 269, 448, 793
 OPI entry 269, 448, 693
 SMask entry 149, 268, 276, 444, 447, 448, 467, 797
 StructParent entry 269, 448, 601
 Subtype entry 267, 279, 448, 480
 for thumbnail images 480
 Type entry 267, 279, 448
 Width entry 267, 277, 448, 480
- Image** entry (alternate image dictionary) 274
- image masks
 color operators, exception to limitations on 218, 222
 with colored tiling patterns 224
 Decode array 276
 explicit masking 277, 421
 image dictionaries for 267
 image XObjects as 262
 ImageMask entry (image dictionary) 268
 Mask entry (image dictionary) 268
 object shape 443
 and sampled images, compared 276
 with shading patterns 232
 stencil masking 276–277
 in Type 3 glyph descriptions 326
 with uncolored tiling patterns 227, 228
 See also
 explicit masking
 soft-mask images
 stencil masking
- image objects
 See image XObjects
- image sets, Web Capture
 See Web Capture image sets
- image space 141, 263, 265, 266, 394, 398, 399
- image streams 36, 262, 263, 268
 filters in 268, 789
- Image XObject** subtype 261, 267, 448, 480
- image XObjects 12, 132, 261, 262–263
 as alternate images 262, 267, 273–275
 alternate images for 269
 color space 177, 194, 480
 fully opaque 467
 in glyph descriptions 324
 as image masks 262
 and **JBIG2Decode** filter 56–57, 58
 in Linearized PDF 738, 739, 745, 754
 name 269
 as OPI proxies 693, 807
 painting 262
 parameters 263
 parent content set 269, 675

- image XObjects (*continued*)
 - as poster images (movies) 571
 - reference counts (Web Capture) 669, 805
 - as soft-mask images **447–449**, 797
 - in Tagged PDF 627
 - as thumbnail images 263, 267, 480
 - in Web Capture content database 661, 662, 669, 670, 805
 - See also*
 - image dictionaries
 - images, sampled
- ImageB** procedure set 574
- ImageC** procedure set 574
- %%ImageCropRect OPI comment (PostScript) **698**
- %%ImageDimensions OPI comment (PostScript) **697**
- %%ImageFilename OPI comment (PostScript) **697**
- ImageI** procedure set 574
- %%ImageInks OPI comment (PostScript) **698**
- ImageMask** entry
 - image dictionary 264, **268**, 269, 275, 276, 448
 - inline image object **279**
- %%ImageOverprint OPI comment (PostScript) **698**
- images, sampled 11, 25, 131, **262–281**
 - alternate images 262, 267, 269, **273–275**
 - base images 267, **273**, 274, 277
 - color inversion 272
 - color space 262, 263, 264, 268, 269, 271, 277
 - color specification 173, 199
 - compression 41, 42, 49, 52, 55, 56, 59
 - coordinate system 265–267
 - data format 263, 264
 - DCS 125, 126
 - Decode** array 264, 269, **271–273**, 277
 - dictionaries. *See* image dictionaries
 - embedded file streams 124
 - encoding 41
 - fully opaque 467
 - height 267
 - image masks. *See* image masks
 - image space **141**, 263, 265, 266, 394, 398, 399
 - image streams 36, 262, 263, 268
 - inline 56, **133**, **263**, **278–281**, 699, 700
 - interpolation 269, **273**, 277
 - in Linearized PDF 738, 739
 - masking **275–278**, 421, 793
 - metadata 269
 - and multitone color 205
 - nonzero overprint mode, unaffected by 215
 - object shape 443
 - objects. *See* image XObjects
 - OPI proxies 693
 - OPI version dictionary 269
 - painting 263
- images, sampled (*continued*)
 - parameters 263
 - in pattern cells 221, 223
 - poster (movie) 571
 - rendering intents 268
 - restrictions on painting **218**
 - sample values. *See* sample values
 - scan conversion 406
 - soft-mask 268, **444**, 445, **447–449**, 797
 - specification 262–263
 - threshold arrays compared to 390
 - thumbnail. *See* thumbnail images
 - tint values, source of 202, 206
 - and transparent overprinting 461, 466
 - Type 3 glyph descriptions, prohibited in 326
 - Web Capture content set 269
 - width 267
 - See also*
 - image XObjects
- imagesetters 12, 137, 201, 203
- ImageType** entry (version 1.3 OPI dictionary) **696**
- imaging model **10–14**
 - See also*
 - Adobe imaging model
 - opaque imaging model
 - transparent imaging model
- immediate backdrop (transparency group element) 410, 426, **428**
 - in knockout groups 434, 435
 - in non-isolated groups 436
- implementation limits 3, **705–708**
 - architectural **705**, **706–707**
 - array capacity 34, **706**
 - character identifier (CID) value 335, **706**
 - clipping paths, complexity of 276
 - DeviceN** tint components 206, **706**
 - dictionary capacity 35, 101, **706**
 - file size **705**
 - graphics state nesting depth **706**, 807
 - indirect objects, number of **706**
 - magnification (zoom) factor **707**
 - memory **705**, 706, **707**
 - name length 33, 34, **706**
 - numeric range and precision 28, 117, 354, 571, **706**
 - page size **707**, 808
 - sampled functions, dimensionality of 109
 - string length 29, 36, **706**
 - thumbnail image samples **707**
 - Web Capture 660, **708**
- implicit color conversion 195–197
- import-data action dictionaries **556**
 - F** entry 556, 802
 - S** entry 556

- import-data actions 518, 528, 550, **555–556**, 802
 - and named pages 557
 - See also*
 - import-data action dictionaries
- ImportData** action type 518, **556**
- importing
 - content 4, 287–289
 - FDF fields 557, 558, 561, 563, 564, 565, 567, 803
 - interactive form fields 9, 528
 - pages 285, 287, 288
- incidental artifacts **617**
 - hidden page elements 617
 - hyphenation 617
 - text discontinuities 617
- Include/Exclude field flag
 - reset-form field **555**
 - submit-form field **551**, 553, 554
- IncludeAnnotations field flag (submit-form field) **552**, 553
- IncludeAppendSaves field flag (submit-form field) **552**, 562
- IncludedImageDimensions** entry (version 2.0 OPI dictionary) **698**
- %%IncludedImageDimensions OPI comment (PostScript) **698**
- IncludedImageQuality** entry (version 2.0 OPI dictionary) **698**
- %%IncludedImageQuality OPI comment (PostScript) **698**
- IncludeNoValueFields field flag (submit-form field) **551**, 553, 554
- incremental updates **18**, **68–71**, 757
 - cross-reference sections 64
 - cross-reference subsections 64
 - and digital signatures 69, 562, 785
 - FDF files, not permitted in 558
 - and file identifiers 581
 - file structure for 61
 - generation numbers 39, 69
 - and Linearized PDF 725–726, 735, 751, **755**
 - and submit-form actions 552
 - and version numbers 63, 83, 785–786
- independent software vendors (ISVs) 48, 73
- InDesign™ page layout software 436, 793
- index** operator (PostScript) **116**, **704**
- Index standard structure type **628**, 632
- Indexed** color spaces 176, **199–201**, 272
 - alternate color space, prohibited as 204
 - base color space of 195, **199–200**, 236, 456, 461
 - base color space, prohibited as 199
 - blending color space, prohibited as 450
 - color table 199, **200**, 449
- Indexed** color spaces (*continued*)
 - color values 199
 - default color space, prohibited as 195
 - I abbreviation 280
 - initial color value 199, 216
 - in inline image objects 280
 - maximum index value 200
 - parameters 199–200
 - remapping of base color space 195
 - setting color values in 217
 - for shadings 236
 - for soft-mask images 449
 - specification 199–201
 - for thumbnail images 480
 - type 1, 2, and 3 shadings, prohibited in 237, 238, 240
 - unrecognized filters and 789
- indexes 628, 632
- indexing of text 368, 613
- indices, page
 - See* page indices
- generation numbers
 - in indirect object references 40
- indirect object references **39–40**
 - hint stream dictionaries (Linearized PDF), prohibited in 735
 - in name trees 101
 - to nonexistent object 39
 - operands, prohibited as 94, 95
 - and progressive document retrieval (Linearized PDF) 754
 - version compatibility 787
- indirect objects 27, **39–41**, 64
 - cross-reference table 61
 - definition 39
 - in FDF files 559, 560
 - generation number. *See* generation numbers
 - number, limit on **706**
 - object identifier. *See* object identifiers
 - object number. *See* object numbers
 - operands, not permitted as 94, 95
 - random access 64
 - references. *See* indirect object references
 - for stream lengths, not permitted in FDF files 558
 - streams 36
- Info** entry
 - FDF page dictionary **567**
 - file trailer dictionary **68**, 575
 - PDF/X output intent dictionary 685, **686**
- information dictionary, document
 - See* document information dictionary
- information dictionary hint table (Linearized PDF) 736, **751**

- inheritance
 - field attributes 528, 530
 - page attributes 87, 88, 91–92, 96
- initial backdrop (transparency group) 428
 - compositing with 425, 427, 431
 - in isolated groups 428, 433, 451
 - in knockout groups 434, 435, 451
 - in non-isolated groups 428, 436
 - notation 426
- ink annotation dictionaries 508
 - BS** entry 508
 - Contents** entry 508
 - InkList** entry 508
 - Subtype** entry 508
- Ink** annotation type 499, 507–508
- ink annotations 499, 507–508, 799
 - border style 490, 495
 - border width 508
 - contents 508
 - dash pattern 508
 - ink list 508
 - See also*
 - ink annotation dictionaries
- ink-jet printers 12
 - resolution 13
- InkList** entry (ink annotation dictionary) 508
- Inks** entry (version 2.0 OPI dictionary) 698
- inline alignment 646
 - Center 646
 - End 646
 - Start 646
- inline image objects 279–281
 - BitsPerComponent** entry 279
 - color spaces, abbreviations for 279–280
 - CMYK (DeviceCMYK)** 280
 - G (DeviceGray)** 280
 - I (Indexed)** 280
 - RGB (DeviceRGB)** 280
 - ColorSpace** entry 279, 280
 - Decode** entry 279
 - DecodeParms** entry 279
 - dictionary entries, abbreviations for 279
 - BPC (BitsPerComponent)** 279
 - CS (ColorSpace)** 279, 280
 - D (Decode)** 279
 - DP (DecodeParms)** 279
 - F (Filter)** 279
 - H (Height)** 279
 - I (Interpolate)** 279
 - IM (ImageMask)** 279
 - W (Width)** 279
 - inline image objects (*continued*)
 - Filter** entry 279
 - filters, abbreviations for 279–280, 788–789
 - A85 (ASCII85Decode)** 280
 - AHx (ASCIIHexDecode)** 280
 - CCF (CCITTFaxDecode)** 280
 - DCT (DCTDecode)** 280
 - FI (FlateDecode)** 280
 - LZW (LZWDecode)** 280
 - RL (RunLengthDecode)** 280
 - Height** entry 279
 - image data 278, 279, 280–281
 - ImageMask** entry 279
 - Intent** entry 279
 - Interpolate** entry 279
 - Width** entry 279
 - See also*
 - inline images
 - inline image operators 134, 278–279
 - BI** 134, 278, 279, 699
 - EI** 134, 278, 279, 280, 700
 - ID** 134, 278, 279, 280, 700
 - inline images 56, 133, 263, 278–281, 699, 700
 - color space 194, 280
 - filters in 280, 789
 - in Linearized PDF 727
 - parameters 263
 - See also*
 - inline image objects
- inline-level structure elements (ILSEs) 624, 629, 632–636
 - baseline shift 647
 - BibEntry** 634
 - BLSEs, contained in 624, 632
 - BLSEs nested within 632, 647
 - BLSEs, treated as 632, 641, 643
 - Code** 634
 - content items in 627
 - general layout attributes for 640
 - illustrations as 624
 - Link** 618, 634
 - link annotations, association with 634
 - link elements 634–636
 - Note** 633
 - packing 625, 641
 - Quote** 633
 - Reference** 628, 633
 - Span** 633, 652
 - standard layout attributes for 640, 646–647
 - BaselineShift** 640, 647
 - LineHeight** 640, 646
 - TextDecorationType** 640, 647
 - usage guidelines 631
- Inline placement attribute 629, 640, 641, 643, 649

- inline-progression direction **624**
- illustrations, width of 649
- in layout 625, 632, 641, 643, 644, 645, 646
- local override 642
- table expansion 651
- writing mode 642
- InlineAlign** standard structure attribute 640, **646**
- input focus 515, **517**
- Insert annotation icon 500
- insideness 169–171
 - even-odd rule **170–171**
 - nonzero winding number rule **169–170**
 - and object shape 411, 435
 - and scan conversion 405
- integer objects 27
 - as dictionary values 99
 - as number tree keys 105
 - range limits 28, **706**
 - syntax **28**
- intellectual property 6–7
- Intent** entry
 - image dictionary 197, **268**, 448
 - inline image object **279**
- interactive features 2, 9, 22, 23, **471–572**
 - pdfmark** language extension (PostScript) 21
 - See*
 - actions
 - annotations
 - articles
 - destinations
 - document outline
 - interactive forms
 - movies
 - page labels
 - presentations
 - sounds
 - thumbnail images
 - viewer preferences
- interactive form dictionary 85, **529–530**
 - CO** entry 516, **529**
 - DA** entry **529**
 - DR** entry **529**
 - Fields** entry **529**
 - in Linearized PDF 734, 750
 - NeedAppearances** entry 517, **529**
 - Q** entry **529**
 - SigFlags** entry **529**, 530
 - signature flags. *See* signature flags
- interactive form fields
 - See* fields, interactive form
- interactive form hint table (Linearized PDF) 736, 745, **750**
- interactive forms 4, 9, **528–568**
 - computation order 516, **529**
 - FDF (Forms Data Format). *See* Forms Data Format fields. *See* fields, interactive form
 - and form XObjects, distinguished 281–282, 528
 - and import-data actions 802
 - interactive form dictionary 85, 734, 750
 - in Linearized PDF 750
 - named pages 93, **557**
 - template pages 93, **567–568**, 804
 - XML submission 4
- interchange
 - of content 9, 592, 605
 - of documents. *See* document interchange
- interior color
 - annotations 503, 505
 - line endings 504
- International Color Consortium (ICC) 189, 197, 814
- International Commission on Illumination 176
- International Electrotechnical Commission (IEC) 192, 814
- International Organization for Standardization (ISO) 55, 56, 59, 100, 361, 653
 - ISO 639 (*Codes for the Representation of Names of Languages*) 99, 100, 361, 653, **814**
 - ISO 3166 (*Codes for the Representation of Names of Countries and Their Subdivisions*) 99, 100, 653, **814**
 - ISO/IEC 8824-1 (*Abstract Syntax Notation One (ASN.1): Specification of Basic Notation*) 100, **815**
 - ISO/IEC 10918-1 (*Digital Compression and Coding of Continuous-Tone Still Images*) **815**
- International Telecommunications Union (ITU) 52, 815
- Internet xix
 - uniform resource identifiers (URIs) 523
 - Web Capture plug-in extension 573, 659
- Internet Assigned Numbers Authority (IANA) 653
- Internet RFCs (Requests for Comment)
 - 1321 (*The MD5 Message-Digest Algorithm*) 72, 125, 581, 665, **815**
 - 1738 (*Uniform Resource Locators*) 119, 127, 664, **815**
 - 1766 (*Tags for the Identification of Languages*) 653, **815**
 - 1808 (*Relative Uniform Resource Locators*) 119, 664, 673, **815**
 - 1866 (*Hypertext Markup Language 2.0 Proposed Standard*) 524, 554, **815**
 - 1950 (*ZLIB Compressed Data Format Specification*) 46, **815**
 - 1951 (*DEFLATE Compressed Data Format Specification*) 46, **815**
 - 2045 (*Multipurpose Internet Mail Extensions (MIME), Part One: Format of Internet Message Bodies*) 544, 553, 668, 673, **815**

- Internet RFCs (Requests for Comment) (*continued*)
- 2046 (*Multipurpose Internet Mail Extensions (MIME), Part Two: Media Types*) 124, **815**
 - 2068 (*Hypertext Transfer Protocol—HTTP/1.1*) 674, **815**
 - 2083 (*PNG (Portable Network Graphics) Specification*) 50, **815**
- Interpolate** entry
- image dictionary **269**, 273, 448
 - inline image object **279**
- Interpolate function **109**, 110, 114, **271**
- interpolation
- bilinear 250, 251
 - cubic spline 110, 113
 - Gouraud **243**, 247, 255, 813
 - linear 110, 251, 792
 - in sampled images 269, **273**, 277
- interpolation functions 243
- “Interpolation Using Bézier Curves” (Elber) **815**
- InvertedDouble** predefined spot function **386**
- InvertedDoubleDot** predefined spot function **385**
- InvertedEllipseA** predefined spot function **387**
- InvertedEllipseC** predefined spot function **388**
- InvertedSimpleDot** predefined spot function **385**
- Invisible annotation flag **493**
- IsMap** entry (URI action dictionary) **523**
- ISO. *See* International Organization for Standardization
- ISO-2022-JP character encoding 345
- ISO Latin 1 character encoding 98
- isolated groups 425, 431, **433**, 434, 451
- blend mode 433
 - blending color space 425, 433, 450
 - compositing in 433, 451
 - compositing of 433
 - group alpha 433
 - group backdrop unused in 433
 - group color 433
 - group color space, explicit 455
 - group compositing formulas **433**, 438
 - group shape 433
 - group stack 433
 - initial backdrop 428, 433, 451
 - knockout 435
 - object alpha 433
 - object shape 433
 - page group as 436, 437, 450
 - pattern cells as 454
 - for soft masks 440
 - source color 433
 - and white-preserving blend modes 460
- Italic font characteristic **622**
- Italic font flag **358**, 622
- italic fonts 358
- Italic outline item flag **479**
- ItalicAngle** entry (font descriptor) **356**
- ITC Zapf Dingbats® typeface 16, 224, 709
- items, outline. *See* outline items
- ITU (International Telecommunications Union)
- Recommendations T.4 and T.6 52, **815**
- IX** entry (appearance characteristics dictionary) **537**
- ## J
- J** operator 134, **156**, 700
- j** operator 134, **156**, 700
- jamo characters 363
- Japanese
- character collections 346–347
 - character sets 334, 335
 - CMaps **344–345**
 - fonts 322
 - glyph widths 340
 - kana (katakana, hiragana) characters 362, 363
 - kanji (Chinese) characters 361, 362, 363
 - R2L reading order 472
 - ruby characters 363
 - Shift-JIS character encoding 336, 348
 - writing systems 642
- Java™ programming language 606
- JavaScript action dictionaries **556**
- JS** entry **556**
 - S** entry **556**
- JavaScript** action type 518, **556**
- JavaScript actions 518, 528, 550, **556–557**, 802
- document-level 93, 124, 556
 - and named pages 557
 - trigger events for 516, 517
- See also*
- JavaScript action dictionaries
- JavaScript dictionary 562, **563**
- After** entry **563**
 - Before** entry **563**
 - Doc** entry **563**
- JavaScript** entry
- FDf dictionary **562**
 - name dictionary **93**, 556, 563
- JavaScript® scripting language
- functions 556, 563
 - interpreter 550, 556
 - scripts 550, 553, 556, 557, 562, 563
 - Unicode, incompatible with 802
- See also*
- JavaScript actions
 - JavaScript dictionary

- JBIG (Joint Bi-Level Image Experts Group) 15, 55
 - JBIG2 compression 15, 42, 55–59
 - JBIG2Decode** filter 4, 42, 55–59
 - inline images, prohibited in 279
 - parameters. *See* **JBIG2Decode** filter parameter dictionaries
 - in sampled images 268
 - JBIG2Decode** filter parameter dictionaries 57
 - JBIG2Globals** entry 57, 58, 59
 - JBIG2Globals** entry (**JBIG2Decode** filter parameter dictionary) 57, 58, 59
 - JDF (Job Definition Format) 374, 677, 689, 692, 806
 - JDF Specification* (CIP4) 374, 689, 814
 - JIS C 6226 character set 344
 - JIS X 0208 character set 344, 345
 - JIS78 character set 344
 - job tickets 374, 689, 692, 806
 - JDF 374, 677, 689, 692, 806
 - PJTF 24, 374, 677, 689, 692, 806
 - JPEG (Joint Photographic Experts Group) 59
 - baseline format 59, 61
 - compression 15, 42, 59
 - file format 573, 659
 - progressive extension 61, 790
 - JPEG: Still Image Data Compression Standard* (Pennebaker and Mitchell) 59, 816
 - JS** entry
 - JavaScript action dictionary 556
 - Justify block alignment 645
 - Justify text alignment 644
- ## K
- K** entry
 - additional-actions dictionary 516
 - CCITTFaxDecode** filter parameter dictionary 53, 54
 - structure element dictionary 591, 593, 598, 599
 - structure tree root 590, 628
 - transparency group attributes dictionary 451, 452
 - K** operator 134, 173, 177, 178, 180, 218, 700
 - k** operator 134, 173, 177, 178, 180, 218, 700
 - K** trigger event (form field) 516, 517
 - Kana** character class 362, 363
 - kana (katakana, hiragana) characters 362, 363
 - Kanji** character class 363
 - kanji (hanzi, hanja) characters 361, 362, 363
 - KanjiTalk6 character set 344
 - KanjiTalk7 character set 344
 - katakana characters 362, 363
 - kerning 300
 - Key annotation icon 500
 - keyboard 2, 9
 - annotations 488
 - checkbox fields 539
 - text fields 533, 538, 543
 - trigger events 516
 - keys
 - dictionary 25, 35, 786
 - encryption. *See* encryption keys
 - name tree 101, 102
 - number tree 105, 106
 - keywords
 - endobj** 39, 707
 - endstream** 36–37, 38, 707, 788
 - f** 65, 779
 - false** 28, 34
 - n** 65
 - null** 39
 - obj** 34, 39
 - operators 94
 - R** 40
 - startxref** 67, 732, 740, 751, 775, 777, 780, 782
 - stream** 36–37, 38, 788
 - trailer** 67, 560
 - true** 28, 34
 - xref** 64, 67, 733
 - Keywords** entry (document information dictionary) 576
 - Kids** entry
 - FDF field dictionary 564
 - field dictionary 531, 541, 542, 543
 - name tree node 101, 102
 - number tree node 106
 - page tree node 86
 - knockout groups 425, 434–436, 451
 - and backdrop color removal 432
 - blend mode 434
 - compositing in 434–435, 451
 - group backdrop 428, 435
 - group color 434
 - group compositing formulas 434–435, 438
 - group opacity 434
 - group stack 434
 - immediate backdrop (group elements) 434, 435
 - initial backdrop 434, 435, 451
 - isolated 435
 - non-isolated 435
 - non-isolated group, parent of 451
 - non-isolated groups nested within 436
 - object opacity 434
 - object shape 434
 - and overprinting 463
 - result alpha 434, 435
 - result color 434, 435
 - result opacity 435

- knockout groups (*continued*)
 - shading patterns implicitly enclosed in 453
 - soft clipping 444
 - for soft masks 440
 - source alpha 435
 - source opacity 435
 - source shape 434, 435
 - text knockout parameter equivalent to 308
 - Korean
 - character collections 347
 - character sets 334, 335
 - CMaps 345
 - fonts 322
 - glyph widths 340
 - hangul characters 363
 - jamo characters 363
 - R2L reading order 472
 - KS X 1001:1992 character set 345
 - KSC–EUC–H predefined CMap 345, 347
 - KSC–EUC–V predefined CMap 345, 347
 - KSCms–UHC–H predefined CMap 345, 347
 - KSCms–UHC–HW–H predefined CMap 345, 347
 - KSCms–UHC–HW–V predefined CMap 345, 347
 - KSCms–UHC–V predefined CMap 345, 347
 - KSCpc–EUC–H predefined CMap 345, 347
- L**
- L entry
 - hint stream dictionary 736
 - line annotation dictionary 503
 - linearization parameter dictionary 733
 - Web Capture command dictionary 672
 - I operator 134, 163, 168, 700
 - L standard structure type 629, 630, 649
 - in tables of contents 628
 - L2R reading order 472
 - $L^*a^*b^*$ color representation 191
 - blending color space, prohibited for 415
 - Lab** color space dictionaries 188
 - BlackPoint** entry 188
 - Range** entry 187, 188, 216, 272
 - WhitePoint** entry 188
 - Lab** color spaces 176, 181, 187–189, 272
 - as base color space 199
 - blending color space, prohibited as 415, 450
 - color values 187
 - default color space, prohibited as 195
 - and **ICCBased** color spaces, compared 189, 192
 - initial color value 216
 - rendering 374
 - setting color values in 217
 - Lab** color spaces (*continued*)
 - See also*
 - Lab** color space dictionaries
 - labeling ranges, page 83, 482, 483
 - labels, page
 - See* page labels
 - Lang** entry 616, 650, 652, 658
 - CIDFont font descriptor 361
 - document catalog 85, 652, 653, 654
 - property list 633, 652, 653, 654, 656
 - structure element dictionary 592, 638, 652, 654
 - language, natural
 - See* natural language specification
 - language codes
 - IANA 653
 - ISO 639 99, 100, 361, 653
 - language identifiers 85, 592, 652–653
 - laser printers 12
 - resolution 13
 - Last** entry
 - outline dictionary 477, 478
 - outline item dictionary 477, 478
 - LastChar** entry
 - Type 1 font dictionary 317, 318, 319, 795
 - Type 3 font dictionary 324
 - LastModified** entry
 - application data dictionary 582
 - page object 88, 690
 - trap network annotation dictionary 690, 691
 - type 1 form dictionary 284
 - LastPage** named action 527
 - See also*
 - named-action dictionaries
 - Latin character set, standard 709, 710, 711–714, 796
 - character names 331–332, 368, 620
 - encodings 329
 - glyph displacements 299
 - name objects limited to 540
 - nonsymbolic fonts 329, 358, 359, 738
 - standard fonts 16, 292
 - Tj operator 294
 - Latin characters 361, 362, 363
 - metrics 362
 - Latin writing systems 329, 359
 - lattice-form Gouraud-shaded triangle meshes
 - See* type 5 shadings
 - lattices, pseudorectangular 248
 - launch action dictionaries 521
 - F** entry 520, 521
 - Mac** entry 520, 521
 - NewWindow** entry 521
 - S** entry 521

- launch action dictionaries (*continued*)
 - Unix** entry 520, 521
 - Win** entry 520, 521, 800
- Launch** action type 518, 521
- launch actions 518, 520–521, 800
 - See also*
 - launch action dictionaries
 - Windows launch parameter dictionaries
- layout, page 84, 800
- Layout** artifact type 616
- layout artifacts 615
- layout attributes, standard
 - See* standard layout attributes
- Layout** standard attribute owner 639, 640
- Lbl standard structure type 629, 630, 633, 634, 649, 650
 - standard layout attributes for 644
- LBody standard structure type 629, 630
 - standard layout attributes for 644
- LC** entry (graphics state parameter dictionary) 157
- LE** entry (line annotation dictionary) 503
- le** operator (PostScript) 116, 704
- Leading** entry (font descriptor) 356
- leading (T_l) parameter 301, 304–305
 - T*** operator 310
 - TD** operator 310, 701
 - TL** operator 302, 701
- left angle bracket (<) character 26
 - double, as dictionary delimiter 35, 67, 560
 - as hexadecimal string delimiter 29, 32, 35, 122
- left brace ({} character 26
 - as delimiter in PostScript calculator functions 116
- left bracket ([]) character 26
 - as array delimiter 34, 808
- left parenthesis (()) character 26
 - escape sequence for 30, 312
 - as literal string delimiter 29
- Length** entry
 - encryption dictionary 71, 72, 73, 78, 79
 - stream dictionary 37, 38, 40, 73, 235, 279, 366, 757
- Length1** entry (embedded font stream dictionary) 366, 367
- Length2** entry (embedded font stream dictionary) 366, 367
- Length3** entry (embedded font stream dictionary) 366, 367
- Level1** entry (PostScript XObject dictionary) 290
- lexical conventions 24–27
- Ll standard structure type 629, 630, 649
 - in tables of contents 628
- ligatures 329, 371, 651, 658, 709
- Lighten** blend mode 417
- lightness 415, 450
- limits, implementation
 - See* implementation limits
- Limits** entry
 - name tree node 102
 - number tree node 106
- line annotation dictionaries 503
 - BS** entry 503
 - Contents** entry 503
 - IC** entry 503
 - L** entry 503
 - LE** entry 503
 - Subtype** entry 503
- Line** annotation type 499, 503
- line annotations 499, 503–504
 - border style 490, 495
 - contents 503
 - dash pattern 503
 - interior color (line endings) 503, 504
 - line ending style. *See* line ending styles
 - line width 503
 - See also*
 - line annotation dictionaries
- Line Breaking Properties* (Unicode Standard Annex #14) 623, 816
- line cap style 148, 153
 - butt 153, 154, 168
 - and dash pattern 155
 - J** operator 156, 700
 - LC** entry (graphics state parameter dictionary) 157
 - projecting square 153, 168
 - round 153, 168
 - and **S** operator 168
 - and Type 3 glyph descriptions 325
- line dash pattern 149, 155
 - for annotation borders 490, 491, 495
 - for circle annotations 505
 - D** entry (graphics state parameter dictionary) 157
 - d** operator 156, 700
 - dash array 155, 157, 491, 495, 681, 798
 - dash phase 155, 157, 491, 495, 681
 - for ink annotations 508
 - for line annotations 503
 - for page boundaries 681
 - and **S** operator 168
 - for square annotations 505
 - and Type 3 glyph descriptions 325
- line ending styles 503, 504
 - Circle 504
 - ClosedArrow 504
 - Diamond 504
 - None 504
 - OpenArrow 504
 - Square 504

- line feed (LF) character **26**
 - in cross-reference tables **65**
 - as end-of-line marker **26, 31, 37, 62, 65**
 - escape sequence for **30**
 - in HTTP requests **674**
 - in stream objects **36, 37**
 - as white space **24, 32**
- line height **646–647**
 - Auto **646, 647**
 - Normal **646, 647**
- line join style **148, 153–154**
 - bevel **153, 154, 168**
 - and dash pattern **155**
 - j operator **156, 700**
 - L** entry (graphics state parameter dictionary) **157**
 - miter **153, 154, 168**
 - round **154, 168**
 - and **S** operator **168**
 - and Type 3 glyph descriptions **325**
- Line** predefined spot function **384, 386**
- line width, current
 - See* current line width
- linear interpolation **110, 251, 792**
- linearization parameter dictionary **729, 732–733, 808**
 - E** entry **733, 736, 808**
 - file length in **751, 755**
 - first-page object number in **733, 737**
 - H** entry **733, 741, 808**
 - hint stream offsets in **735, 742**
 - L** entry **733**
 - Linearized** entry **733**
 - N** entry **733**
 - O** entry **733, 736, 744, 747**
 - P** entry **733, 737**
 - T** entry **733, 753**
- Linearized** entry (linearization parameter dictionary) **733**
- Linearized PDF **3, 17, 62, 64, 91, 725–755**
 - access strategies **751–755**
 - background and assumptions **726–728**
 - cross-reference tables **729, 742**
 - first-page **729, 732, 734, 735, 740, 744, 751**
 - main **731, 732, 740–741**
 - document catalog **728, 730, 732, 734, 740, 753**
 - document structure **728–741**
 - first-page section **730, 735, 736–738, 745, 751, 752, 808**
 - generation numbers **729**
 - header **731**
 - hint streams. *See* hint streams
 - hint tables. *See* hint tables
 - HTML (Hypertext Markup Language) **727**
 - HTTP (Hypertext Transfer Protocol) **726, 727–728**
 - incremental updates and **725–726, 735, 751, 755**
 - Linearized PDF (*continued*)
 - indirect objects, numbering of **728–729, 735**
 - inline images, retrieval of **727**
 - linearization parameter dictionary **729, 732–733, 735, 737, 742, 751, 755, 808**
 - MIME (Multipurpose Internet Mail Extensions) **726, 727**
 - shared object signatures **747, 809**
 - shared objects section **738, 739, 745, 747, 808**
 - thumbnail shared objects section **739, 740, 748**
 - trailer
 - first-page **729, 732, 734**
 - main **741**
 - URLs (uniform resource locators) **726, 727**
 - version identification **733**
 - and World Wide Web **726**
- LineHeight** standard structure attribute **633, 640, 643, 646, 648**
- lines (text) **625**
 - stacking within parent BLSE **625, 632**
- LineThrough text decoration type **647**
- lineto** operator (PostScript) **22, 700**
- LineX** predefined spot function **386**
- LineY** predefined spot function **386**
- link annotation dictionaries **501**
 - Contents** entry **501**
 - Dest** entry **492, 501, 519**
 - H** entry **501**
 - PA** entry **501**
 - Subtype** entry **501**
- Link** annotation type **491, 499, 501, 562**
 - opacity inapplicable to **492**
- link annotations **9, 499, 501–502, 514, 799**
 - actions for **492**
 - border color **491**
 - border style **491, 496**
 - contents **501**
 - destination **474, 476, 501, 799**
 - and go-to actions **519**
 - highlighting mode **501**
 - and link elements (Tagged PDF) **634–636**
 - movie actions associated with **525**
 - and trigger events **515**
 - and URI actions **501, 800**
 - and Web Capture **501**
 - See also*
 - link annotation dictionaries
- link elements (Tagged PDF) **634–636**
- Link standard structure type **618, 634**
- links, hypertext
 - See* hypertext links
- list attribute, standard
 - See* standard list attribute

- list box fields 4, 538, 545, 546
 - trigger events for 516
 - variable text in 533
 - list elements, standard
 - See standard list elements
 - list numbering style 650
 - Circle 650
 - Decimal 650
 - Disc 650
 - LowerAlpha 650
 - LowerRoman 650
 - None 650
 - Square 650
 - UpperAlpha 650
 - UpperRoman 650
 - List standard attribute owner 639, 650
 - ListNumbering standard structure attribute 649, 650
 - literal strings 29–31
 - continuation lines 30–31
 - escape sequences 30–31
 - octal character codes in 30, 31
 - LJ entry (graphics state parameter dictionary) 157
 - In operator (PostScript) 116, 703
 - “loca” table (TrueType font) 367
 - Location entry (signature dictionary) 549
 - log operator (PostScript) 116, 703
 - LOGFONT structure (Windows) 321
 - logical structure 2, 9, 22, 23, 573, 588–612
 - annotations, sequencing of 618
 - content 593–604
 - example 607–612
 - fragmented BLSEs, recognition of 629
 - incidental artifacts not included in 617
 - link elements (Tagged PDF) 634
 - and page content order 651
 - page tree, distinguished from 86
 - pdfmark language extension (PostScript) 21
 - and real content 614
 - and reference XObjects 289
 - structural parent tree 90, 269, 285, 492
 - Tagged PDF and 573, 612, 613, 615
 - text discontinuities, recognition of 617
 - visible content, separation from 589
 - See also
 - content items
 - structure attributes
 - structure elements
 - structure hierarchy
 - structure tree root
 - structure types
 - logical structure hint table (Linearized PDF) 736, 745, 750
 - logical structure order 618
 - annotations, sequencing of 618
 - artifacts 618
 - lossless filters 42, 55, 193
 - lossy filters 42, 55, 59, 278
 - LowerAlpha list numbering style 650
 - LowerRoman list numbering style 650
 - LP entry (additional-actions dictionary, obsolete) 800
 - LrTb writing mode 642
 - It operator (PostScript) 116, 704
 - luminance 60
 - Luminosity blend mode 419
 - Luminosity soft-mask subtype 445, 446, 450
 - $L^*u^*v^*$ color representation
 - blending color space, prohibited for 415
 - LW entry (graphics state parameter dictionary) 151, 157
 - LZW (Lempel-Ziv-Welch) compression 15, 41, 42, 43, 45–52
 - clear-table marker 46, 48
 - predictor functions 46, 49, 50–52, 268
 - LZW filter abbreviation 280, 789
 - LZWDecode filter 42, 45–52
 - LZW abbreviation 280, 789
 - parameters. See LZWDecode filter parameter dictionaries
 - predictor functions 50–52
 - in sampled images 268
 - LZWDecode filter parameter dictionaries 48–49
 - BitsPerComponent entry 49
 - Colors entry 49
 - Columns entry 49
 - EarlyChange entry 49
 - Predictor entry 49, 50–51
- ## M
- M entry
 - annotation dictionary 490, 509, 798
 - signature dictionary 549
 - transition dictionary 486, 487
 - M operator 134, 156, 700
 - m operator 134, 162, 163, 168, 169, 700
 - Mac entry
 - embedded file parameter dictionary 125
 - file specification dictionary 122, 123
 - launch action dictionary 520, 521
 - Mac OS file information dictionaries 125
 - Creator entry 125
 - ResFork entry 125
 - Subtype entry 125

- Mac OS KH character set 345
- Mac[®] OS operating system
 - application launch parameters 520, 521
 - character encoding 328, 329, 710, 714
 - file information 125
 - file names 120
 - file system 122
 - FOND resource 321
 - font names 321
 - PDF Writer printer driver 20
 - Preferences folder 802
 - QuickDraw imaging model 19
 - Script Manager 343, 344, 345
 - TrueType[®] font format 332
- MacExpertEncoding** predefined character encoding **329**, 709, **710**
 - as base encoding 330
 - for TrueType fonts 333
 - for Type 1 fonts 318
 - and Unicode mapping 368, 620
- MacRomanEncoding** predefined character encoding **329**, 709, **710**, 714
 - as base encoding 330
 - for TrueType fonts 333
 - for Type 1 fonts 318
 - and Unicode mapping 368, 620
- magenta color component
 - DeviceCMYK** color space 178, 180
 - DeviceN** color spaces 206
 - grayscale conversion 377, 382
 - green, complement of 378
 - halftones for 401
 - initialization 180
 - in multitone 205
 - overprinting 464
 - RGB conversion 377, 378
 - transfer function 380, 381
 - transparent overprinting 464
 - undercolor removal 150, 378, 379
- magenta colorant
 - overprinting 464
 - PANTONE Hexachrome system 205
 - printing ink 201
 - process colorant 178, 180
 - subtractive primary 178, 180
 - transparent overprinting 464
- magnification (zoom) factor 84, 90, 493, 494, 676
 - in destinations 474, 475
 - implementation limits **707**
 - for movies 572
- main cross-reference table (Linearized PDF) 731, 732, **740–741**
 - and page retrieval 753
- main trailer (Linearized PDF) 741
- MainImage** entry (version 2.0 OPI dictionary) **697**
- %%MainImage OPI comment (PostScript) **697**
- mapping name (form field) **531**, 552
- mark information dictionary 85, **613–614**
 - Marked** entry 613, **614**
- marked clipping sequences **585–587**, 588
 - in illustration elements (Tagged PDF) 637
- marked content 19, 136, 573, **583–588**
 - and clipping **585–588**
 - in dynamic appearance streams 536
 - elements. *See* marked content elements
 - language identifiers 85, 592
 - metadata for 580
 - operators. *See* marked-content operators
 - property lists 583, 584, **585**, 699, 700
 - and Tagged PDF 573
- marked-content elements **583**
 - empty **586**, 587, 588
 - tags 583
 - See also*
 - marked-content points
 - marked-content sequences
- marked-content identifiers **593–594**
 - and natural language specification 654
 - parent structure element, finding from 601, 603
 - small values recommended for 601
 - in structure elements 591, 598
- marked-content operators 132, **134**, 573, **583–584**
 - BDC** 134, 583, **584**, 585, 593, 699
 - BMC** 134, 535, 583, **584**, 699
 - DP** 134, 583, **584**, 585, 700
 - EMC** 134, 535, 583, **584**, 593, 700
 - MP** 134, 583, **584**, 700
 - tags 583
 - text object operators, combined with 584
 - in text objects 309
- marked-content points **583**, 584, 700
 - and clipping 588
 - empty 586
- marked-content reference dictionaries 591, **594**
 - MCID** entry **594**
 - metadata inapplicable to 579
 - Pg** entry **594**, 598
 - Stm** entry **594**
 - StmOwn** entry **594**
 - Type** entry **594**
- marked-content sequences 285, **583**, 584, 699, 700
 - annotations, association with 618
 - annotations, sequencing of 618–619
 - in appearance streams 594
 - for artifact specification 616
 - and clipping 585–587

- marked-content sequences (*continued*)
 - empty 586
 - in form XObjects 594
 - identifiers. *See* marked content identifiers
 - as logical structure content items 285, 589, 590, 591, 593–598, 599, 600, 601, 602, 603, 618
 - marked clipping sequences 585–587, 588, 637
 - natural language specification 652, 653–655, 656, 657, 659
 - nesting of 583
 - reference dictionaries 591
 - for reverse-order show strings 619
 - Span tag 633
- marked-content tags 583, 584, 723
 - Artifact 616
 - Clip 637
 - ReversedChars 619
 - Span 633, 652, 653, 654, 656, 657, 659
 - and structure types 594
- Marked** entry (mark information dictionary) 613, 614
- MarkInfo** entry (document catalog) 85, 613
- MarkStyle** entry (printer’s mark form dictionary) 683
- markup annotation dictionaries 506
 - Contents** entry 506
 - QuadPoints** entry 506, 799
 - Subtype** entry 506
- markup annotations 9, 505–506, 799
 - contents 506
 - See also*
 - markup annotation dictionaries
- Mask** entry (image dictionary) 268, 275, 276, 277, 444, 448, 793
- Mask** object type 446
- mask opacity 149, 159, 268, 421, 442
 - notation 422, 427, 431
 - soft masks 439
 - specifying 443–444, 797
 - in transparency groups 425
- mask shape 149, 159, 268, 421, 442
 - notation 422, 427, 431
 - soft masks 439
 - specifying 443–444, 797
 - in transparency groups 425
- masked images 275–278, 793
 - shape (transparent imaging model) 421
 - See also*
 - color key masking
 - explicit masking
 - soft masks
 - stencil masking
- matrices, transformation
 - See* transformation matrices
- Matrix** entry
 - CalRGB** color space dictionary 185, 186, 440
 - type 1 form dictionary 283, 284, 288, 446, 496
 - type 1 pattern dictionary 222, 223, 283
 - type 1 shading dictionary 237
 - type 2 pattern dictionary 231, 283
- matte color (soft-mask image) 447, 449
- Matte** entry (soft-mask image dictionary) 447, 448, 449
- MaxLen** entry (text field dictionary) 544
- “maxp” table (TrueType font) 367
- MaxWidth** entry (font descriptor) 357
- MCID** entry
 - marked-content reference dictionary 594
 - property list 593, 633, 654, 656
- MCR** object type 594
- MD5 message-digest algorithm 72
 - checksum, embedded files 125
 - for digital identifiers (Web Capture) 664–665
 - for file identifiers 581, 805
 - hash function 73, 78, 79, 80, 563, 660, 747
 - for shared object signatures (Linearized PDF) 747, 809
- MD5 Message-Digest Algorithm, The* (Internet RFC 1321) 72, 125, 581, 665, 815
- media box 677, 678
 - inheritance of 91
 - for media selection 679
 - in page imposition 806
 - in page object 88
 - page placement, ignored in 679
 - in printing 679
 - in rendering 374
- MediaBox** entry (page object) 88, 91, 677, 737
- medium, output
 - See* output medium
- menu bar, hiding and showing 472
- menu items 530
 - as named actions 801
- merging of content xx
- meshes
 - Coons patch. *See* type 6 shadings
 - free-form Gouraud-shaded triangle. *See* type 4 shadings
 - lattice-form Gouraud-shaded triangle. *See* type 5 shadings
 - tensor-product patch. *See* type 7 shadings
- metadata 573, 575–580, 804–805
 - constructs that do not take 579
 - date stamp 804–805
 - document information dictionary 575–577
 - for documents 85
 - for form XObjects 285
 - for **ICCBased** color spaces 190

- metadata (*continued*)
 - for marked content 580
 - for pages 90
 - for sampled images 269
 - version compatibility 804
 - See also*
 - document information dictionary
 - metadata streams
- Metadata** entry 578, 580
 - document catalog 85, 578
 - embedded font stream dictionary 366
 - ICC profile stream dictionary 190
 - image dictionary 269
 - page object 90
 - property list 580
 - type 1 form dictionary 285
- Metadata** object type 578
- metadata stream dictionaries 578
 - in property lists 580
 - Subtype** entry 578
 - Type** entry 578
- metadata streams 4, 575, 577–580, 804–805
 - document information dictionary, compared with 577
 - for documents 85
 - for embedded font programs 366
 - encryption not recommended for 577
 - filters not recommended for 577
 - for form XObjects 285
 - for **ICCBased** color spaces 190
 - for pages 90
 - for sampled images 269
 - See also*
 - metadata stream dictionaries
- metadata subtypes
 - XML** 578
- Microphone annotation icon 510
- Microsoft Corporation
 - TrueType 1.0 Font Files Technical Specification* 321, 360, 816
 - Windows[®] operating system 19, 321, 328
- Microsoft Unicode character encoding 333
- Middle block alignment 645
- MIME (Multipurpose Internet Mail Extensions)
 - application/pdf content type 553
 - application/vnd.fdf content type 558
 - application/x-www-form-urlencoded content type 672
 - and Linearized PDF 726, 727
 - media type name 124
 - multipart/form-data content type 544
 - text/html content type 668
- Minion[®] typeface 943
- minus sign (–) character
 - in dates 100
- misregistration of colorants 513, 676, 688
- MissingWidth** entry (font descriptor) 317, 357
- miter limit 148, 153–154
 - forced into valid range 151
 - M** operator 156, 700
 - ML** entry (graphics state parameter dictionary) 157
 - and **S** operator 168
- miter line join style 153, 154, 168
- Mix** entry (sound action dictionary) 525
- MK** entry (widget annotation dictionary) 512, 536
- ML** entry (graphics state parameter dictionary) 157
- MMType1** font type 315, 320, 365
- MN** entry (printer's mark annotation dictionary) 682
- mod** operator (PostScript) 116, 703
- ModDate** entry
 - document information dictionary 576
 - embedded file parameter dictionary 125
- Mode** entry
 - movie action dictionary 526
 - movie activation dictionary 572
- modification date
 - annotation 490
 - document 573, 575, 576, 805
 - form XObject 284, 582
 - page 88, 582
 - page-piece dictionaries and 88, 284, 582
 - trap network 691
 - Web Capture content set 670, 671
- monospaced fonts 297
- mouse 2, 9, 497, 517
 - annotations 488, 493, 497, 501, 512, 514
 - button fields 537, 538
 - checkbox fields 539
 - document-level navigation 474
 - outline items 477
 - pop-up help labels 526
 - read-only form fields unresponsive to 532
 - submit-form actions, tracking in 552
 - thumbnail images 480
 - trigger events related to 515, 517
 - URI actions, tracking in 523–524
 - widget annotations 537
- moveto** operator (PostScript) 22, 700
- movie action dictionaries 525–526
 - Annot** entry 526
 - Mode** entry 526
 - and movie activation dictionaries, compared 525
 - Operation** entry 526
 - S** entry 526
 - Start** entry 526
 - T** entry 526
- Movie** action type 518, 526, 801

- movie actions 518, **525–526**, 801
 - and movie annotations 525, 526
 - operations. *See* movie operations
 - See also*
 - movie action dictionaries
 - movie annotations
 - movies
- movie activation dictionaries 570, 571–572
 - Duration** entry 571, 572
 - FWPosition** entry 572
 - FWScale** entry 572
 - Mode** entry 572
 - and movie action dictionaries, compared 525
 - in movie annotations 511
 - Rate** entry 572
 - ShowControls** entry 572
 - Start** entry 571, 572
 - Synchronous** entry 572
 - Volume** entry 572
- movie annotation dictionaries 511
 - A** entry 492, **511**, 570
 - Contents** entry 511
 - Movie** entry 511, 570
 - Subtype** entry 511
- Movie** annotation type 499, **511**, 562, 799
 - opacity inapplicable to 492
- movie annotations 9, 499, **510–511**, 570, 799–800
 - annotation rectangle 525, 572
 - contents 511
 - and file specifications 123
 - and movie actions 525, 526
 - plug-in extensions 498
 - title 526
 - See also*
 - movie actions
 - movie annotation dictionaries
 - movies
- movie dictionaries 570, 571
 - Aspect** entry 571, 572
 - F** entry 571
 - in movie annotations 511
 - Poster** entry 571
 - Rotate** entry 571
- Movie** entry (movie annotation dictionary) **511**, 570
- movie files 571
- movie operations **526**
 - Pause **526**
 - Play **526**
 - Resume **526**
 - Stop **526**
- movies 9, 488, **570–572**
 - asynchronous 572
 - bounding box 571
 - controller bar 572
 - and file specifications 123
 - magnification (zoom) factor 572
 - and movie actions 514, 525
 - and movie annotations 510
 - operations. *See* movie operations
 - play mode **572**
 - poster image 571
 - rotation 571
 - synchronous 572
 - time scale **571**
 - See also*
 - movie actions
 - movie activation dictionaries
 - movie annotations
 - movie dictionaries
- MP** operator 134, 583, **584**, 700
- mul** operator (PostScript) **116**, **703**
- mulLaw sound encoding format **569**, 570
- Multiline field flag (text field) **543**
- multipart/form-data content type (MIME) 544
- multiple-byte character codes
 - in CID-keyed fonts 335, 336
 - in file specifications 122
 - in font names 322
 - and text-showing operators 312
 - and word spacing 303
- multiple master font dictionaries **320**
 - BaseFont** entry **320**
 - Subtype** entry **320**
- multiple master fonts **319–321**
 - instances **320**, 321
 - naming conventions 320
 - PostScript name 320
 - snapshots 321
 - substitution 796
 - See also*
 - multiple master font dictionaries
- Multiply** blend mode **417**, 433
- Multipurpose Internet Mail Extensions (MIME), Part One: Format of Internet Message Bodies* (Internet RFC 2045) 544, 553, 668, 673, **815**
- Multipurpose Internet Mail Extensions (MIME), Part Two: Media Types* (Internet RFC 2046) 124, **815**
- MultiSelect field flag (choice field) **546**, 547
- multitone color 172, 176, **205**
 - duotone **205**, 209–211
 - examples 209–213
 - quadtone **205**, 212–213
- Myriad® typeface 943

N

- N entry
 - appearance dictionary 497, 498, 534, 565, 615, 682, 689, 690, 692
 - bead dictionary 484, 755
 - ICC profile stream dictionary 190
 - linearization parameter dictionary 733
 - named-action dictionary 528
 - type 2 function dictionary 113
- N highlighting mode (none) 501, 512
- n keyword 65
- n operator 134, 167, 172, 585, 586, 700
- name dictionary 83, 92–93
 - AP entry 93, 498
 - Dests entry 93, 476
 - EmbeddedFiles entry 93, 124
 - IDS entry 93, 661, 662, 663, 664, 668, 675
 - JavaScript entry 93, 556, 563
 - in Linearized PDF 753
 - metadata inapplicable to 579
 - Pages entry 93, 557
 - Templates entry 93, 557
 - URLS entry 93, 661, 662, 663, 664, 668
- Name entry
 - FDF named page reference dictionary 568
 - image dictionary 269, 448, 793
 - rubber stamp annotation dictionary 507
 - signature dictionary 549
 - sound annotation dictionary 510
 - text annotation dictionary 500
 - Type 1 font dictionary 317, 794
 - type 1 form dictionary 284, 793
 - Type 3 font dictionary 323
- name objects 26, 27, 32–34, 787–788
 - character encodings for 788
 - for destinations 476, 477
 - as dictionary keys 35, 392, 401
 - as dictionary values 99
 - hexadecimal character codes in 33, 124
 - length limit 33, 34, 706
 - syntax 32–34, 787–788
 - UTF-8 encoding 34
 - for version specifications 561
- name registry, PDF 3, 723–724
- “name” table (TrueType font) 321
- name tree nodes 101–102
 - intermediate 102
 - Kids entry 101, 102
 - leaf 102
 - Limits entry 102
 - metadata inapplicable to 579
 - Names entry 101, 102
 - root 101, 102
- name trees 101–105
 - for appearance streams 93, 498
 - for destinations 93, 476
 - dictionaries, compared with 101
 - as dictionary values 99
 - for element identifiers 590
 - for embedded file streams 93, 124
 - for JavaScript actions 93, 124, 556
 - keys in 101, 102
 - in name dictionary 92, 93
 - for named pages 93, 557
 - nodes. *See* name tree nodes
 - number trees, compared with 106
 - for template pages 93, 557
 - values in 101, 102
 - for Web Capture content sets 93, 661, 664, 668, 675
- named-action dictionaries 528
 - N entry 528
 - S entry 528
- Named action type 518, 528, 801
- named actions 518, 527–528, 801
 - FirstPage 527
 - LastPage 527
 - NextPage 527
 - PrevPage 527
 - See also*
 - named-action dictionaries
- named destination dictionaries 476
 - D entry 476
- named destination hint table (Linearized PDF) 736, 751
- named destinations 474, 476–477
 - in document catalog 83, 84
 - go-to actions as targets of 476
 - in Linearized PDF 734, 740, 749, 752, 753
 - in name dictionary 93
 - unique name (Web Capture) 665, 666
 - See also*
 - named destination dictionaries
- named page reference dictionaries
 - See* FDF named page reference dictionaries
- named pages 93, 557
 - in import-data actions 557
 - invisible. *See* template pages
 - in JavaScript actions 557
- named resources 95–96
 - color spaces as 97
 - external objects (XObject) as 97, 132
 - font dictionaries as 97, 293
 - in form XObjects 284–285
 - graphics state parameter dictionaries as 97
 - in Linearized PDF 739
 - patterns as 97
 - procedure sets as 97, 574
 - property lists as 97, 580, 584, 585

- named resources (*continued*)
 - shading dictionaries as 97
 - in type 3 fonts 324
- names
 - annotations 4
 - appearance states 539, 541, 542, 691
 - attribute classes 590, 592, 605–606, 607, 638
 - attribute owners 605
 - blend modes 416, 442
 - character classes 362
 - character encodings 318, 329, 330, 333, 368
 - characters. *See* character names
 - CMaps, predefined 343–345, 349, 352, 353
 - color space families 177, 182, 199, 202, 206
 - color spaces 177, 199, 204, 216, 280
 - colorants 202, 203, 204, 206, 207, 391, 392, 401, 698
 - conversion engines (Web Capture) 674
 - destinations **476–477**
 - dictionary 83, 92–93
 - embedded spaces in 787
 - first-class 124, **723–724**
 - fonts. *See* font names
 - form XObjects 282, 284
 - graphics state parameter dictionaries 156, 157
 - halftones 392, 393, 395, 398, 400, 401
 - images 269
 - JavaScript scripts 563
 - marked-content tags 584
 - object subtypes 36
 - object types 36
 - patterns 217, 224, 227, 228
 - registered 36, 72, 547, 583, 675, 685, 786
 - See also* name registry, PDF
 - rendering intents 158, 197, 268
 - resources 284, 324, 796
 - second-class **724**
 - shadings 232
 - spot functions, predefined 384, 797
 - structure types 34, 591, 592, 593
 - third-class **724**
 - XX** prefix **724**
 - See also*
 - name objects
- Names** entry
 - document catalog **83**, 92, 740, 753
 - name tree node 101, **102**
- native color space (output device) 236, 373, 374, **376**
 - CIE-based color mapping 375
 - and flattening of transparent content 470
 - and halftones 382, 401
 - and overprinting 465, 466
 - page group, inherited by 437, 442, 452
 - process colors, specification of 456, 458
- native color space (output device) (*continued*)
 - rendering intents, target of 468
 - transfer functions 380, 381, 382
 - and transparent overprinting 460
- natural language specification 573, 651, **652–657**
 - for CIDFont character encodings 361
 - for documents 652, 653, 654
 - hierarchy **653–657**
 - language identifiers 85, 592, **652–653**
 - for marked-content sequences 652, 653–655, 656, 657, 659
 - for structure elements 652, 653, 654, 656–657
 - in Tagged PDF 616
 - in Unicode 99–100, 652, 657, 658, 659
- navigation **474–488**
 - document-level **474–481**
 - See also*
 - destinations
 - document outline
 - thumbnail images
 - page-level **481–488**
 - See also*
 - articles
 - page labels
 - presentations
- navigation controls, hiding and showing 472
- ne** operator (PostScript) **116, 704**
- NeedAppearances** entry (interactive form dictionary) 517, **529**
- neg** operator (PostScript) **116, 703**
- Netscape Communications Corporation
 - Client-Side JavaScript Reference* 556, **816**
- network access 62, 725, 726, 727, **751–755**
 - See also* Forms Data Format (FDF)
- Network Publishing xix
- New York typeface 321
- newline characters **26**, 30
- NewParagraph annotation icon 500
- NewWindow** entry
 - launch action dictionary **521**
 - remote go-to action dictionary **520**
- Next** entry
 - action dictionary **514**
 - outline item dictionary 477, **478**
- NextPage** named action **527**
 - See also*
 - named-action dictionaries
- NM** entry (annotation dictionary) **490**
- no-op actions (obsolete) 518
- NoExport field flag **532**, 551, 554
- nonbreaking space character 714

- None** colorant name
 - DeviceN** color spaces 207
 - in **Separation** color spaces 203
- None line ending style 504
- None list numbering style 650
- None predictor function (LZW and Flate encoding) 50, 51
- None text decoration type 647
- NonFullScreenPageMode** entry (viewer preferences dictionary) 472
- non-isolated groups 425, 431, 434
 - and backdrop color removal 432
 - bounding box 451
 - and **CompatibleOverprint** blend mode 462
 - compositing in 451
 - group backdrop 451
 - group color space, inherited from parent group 455
 - group compositing formulas 436
 - immediate backdrop (group elements) 436
 - initial backdrop 428, 436
 - knockout 435
 - knockout groups, nested within 436
 - and overprinting 462, 463
 - page group as 436, 450
 - painting 451
 - patterns implicitly enclosed in 453
- non-knockout groups 425
 - and **CompatibleOverprint** blend mode 462
 - compositing in 434
 - and overprinting 462
 - tiling patterns implicitly enclosed in 453
- non-Latin character sets 329
 - in checkbox fields 540, 542
- non-Latin writing systems 329
 - in checkbox fields 540
- nonprinting characters 30, 31
- nonseparable blend modes 418–419
 - spot colors, inapplicable to 460
- nonstroking alpha constant, current 149, 159, 444
 - for annotations 491
 - and fully opaque objects 467
 - initialization 452
 - and overprinting 462, 463
 - setting 444
 - and transparency groups 444
- nonstroking color, current 148, 151
 - DeviceCMYK** color space 180, 218, 700
 - DeviceGray** color space 179, 217, 700
 - DeviceN** color spaces 206
 - DeviceRGB** color space 180, 217, 701
 - f operator 168
 - Pattern** color spaces 224, 228
 - sampled images 268
 - Separation** color spaces 202
 - nonstroking color, current (*continued*)
 - setting 173, 177, 217–218, 700, 701
 - stencil masking 264
 - text, showing 295
- nonstroking color space, current 148
 - CIE-based color spaces 182
 - DeviceCMYK** color space 180, 218, 700
 - DeviceGray** color space 179, 217, 700
 - DeviceRGB** color space 180, 217, 701
 - Indexed** color spaces 199
 - Pattern** color spaces 228
 - setting 173, 177, 217–218, 699
- NonStruct** standard structure type 628
- Nonsymbolic font flag 358
- nonsymbolic fonts 329, 333, 358, 359
 - base encoding 330
- nonterminal fields 531
- non-white-preserving blend modes
 - spot colors, inapplicable to 460
- nonzero overprint mode 196, 215–216
 - and transparency 461
- nonzero winding number rule 169–170
 - clipping 172, 306, 702
 - filling 167, 169, 699, 700
- normal appearance (annotation) 497
 - and real content 615
 - for unknown annotation types 498
- Normal** blend mode 412, 417
 - for annotations 491, 496
 - and backdrop color removal 432
 - blend function 419
 - Compatible** blend mode equivalent to 419, 461
 - and **CompatibleOverprint** 462, 466
 - current blend mode initialized to 452
 - as default blend mode 442
 - and fully opaque objects 467
 - in isolated groups 433
 - and overprinting 459, 460, 463
 - in page group 437
 - patterns, painting of 454
 - for spot colors 460
- Normal line height 646, 647
- NoRotate annotation flag 493, 494, 496, 500
- not** operator (PostScript) 116, 704
- NotApproved annotation icon 507
- .notdef character name 331, 332, 340, 355, 357
- notdef mappings 352, 354, 355
- Note annotation icon 500
- Note standard structure type 633
- NotForPublicRelease annotation icon 507
- NoToggleToOff field flag (button field) 538, 540

- NoView annotation flag **493**
 - NoZoom annotation flag **493**, 494, 496, 500
 - NP** entry (additional-actions dictionary, obsolete) 800
 - NTSC (National Television Standards Committee) video standard 377
 - null (NUL) character **26**
 - in unique names (Web Capture) 666
 - null object (**null**) 27, 38, **39**
 - in **AnnotStates** arrays (trap networks) 691
 - as choice field value 547
 - in **CIDSystemInfo** arrays 349
 - as dictionary value 35, 39, 99
 - as indirect reference to nonexistent object 39, 40
 - number sign (#) character
 - as hexadecimal escape character in names **33**, 34, 322, 787–788
 - in uniform resource locators (URLs) 664
 - number tree nodes **106**
 - intermediate 106
 - Kids** entry **106**
 - leaf 105, 106
 - Limits** entry **106**
 - metadata inapplicable to 579
 - Nums** entry **106**
 - root 105, 106
 - number trees **105–106**
 - as dictionary values 99
 - keys 105, 106
 - name trees, compared with 106
 - nodes. *See* number tree nodes
 - for page labeling ranges 83, 482
 - structural parent tree 590, 601
 - values 106
 - numbers **29**
 - See also* numeric objects
 - numeric characters 363
 - numeric objects **28–29**
 - as dictionary values 99
 - integer **28**
 - range and precision 28
 - real **28**
 - Nums** entry (number tree node) **106**
- O**
- O** entry
 - additional-actions dictionary 515, 800
 - attribute object **605**, 639, 640, 650, 651
 - encryption dictionary **76**, 78, 79, 81
 - hint stream dictionary **736**
 - linearization parameter dictionary **733**, 736, 744, 747
 - Web Capture content set 667, **668**, 669, 675, 805
 - Windows launch parameter dictionary **521**
 - O** highlighting mode (outline) **501**, **512**
 - O** trigger event (page) **515**
 - Obj** entry (object reference dictionary) **599**
 - obj** keyword 34, **39**
 - object alpha 429, 430
 - in isolated groups 433
 - notation **431**
 - object color 430
 - and rendering intents 468
 - and soft masks 439
 - object hierarchy
 - PDF 560
 - PDF 81
 - object identifiers **39**
 - cross-reference table, reconstruction of 707
 - and encryption 73
 - and incremental updates 69
 - shared (Linearized PDF) 743, **744**, 750
 - in updating example 774, 775, 777, 779, 780, 782
 - object numbers **39**
 - in cross-reference table 64–66, 68, 69
 - and encryption 73
 - in FDF files 558
 - in indirect object references 40
 - in updating example 775, 779, 780
 - object opacity **421**, 442
 - in knockout groups 434
 - notation **422**, **427**
 - and overprinting 459
 - patterns 453
 - specifying **443**
 - and tiling patterns 422
 - object reference dictionaries 591, **598–599**
 - metadata inapplicable to 579
 - Obj** entry **599**
 - Pg** entry **599**
 - Type** entry **599**
 - object references (logical structure) 601, 602
 - and Form standard structure type 637
 - and link elements (Tagged PDF) 634
 - See also*
 - object reference dictionaries
 - object shape **421**, 430, 442
 - current clipping path 439
 - glyphs 443
 - image masks 443
 - in isolated groups 433
 - in knockout groups 434
 - notation **422**, **427**, **431**
 - path objects 443
 - patterns 443, 453
 - sample images 443
 - sh** operator 443

- object shape (*continued*)
 - shading patterns 443
 - specifying **443**
 - tiling patterns 422, 443
 - and topmost object 467
 - in transparency groups 429
- object subtypes **35–36**
 - embedded files 124
 - external objects (XObjects) 261, 290
- object types **35–36**
 - Action** 514
 - Annot** 490, 775, 780
 - Bead** 484
 - Border** 495
 - Catalog** 83, 758, 760
 - CMap** 349
 - EmbeddedFile** 124
 - Encoding** 330
 - ExtGState** 157
 - Filespec** 122, 128
 - Font** 36, 317, 323, 338, 353, 760
 - FontDescriptor** 356
 - Group** 287
 - Halftone** 392, 393, 395, 398, 400, 401
 - Mask** 446
 - MCR** 594
 - Metadata** 578
 - OBJR** 599
 - OPI** 694, 697
 - Outlines** 478, 758, 760
 - OutputIntent** 685
 - Page** 88, 557, 758, 760, 775
 - PageLabel** 483
 - Pages** 86, 758, 760
 - Pattern** 221, 231
 - Sig** 549
 - Sound** 569
 - SpiderContentSet** 668
 - StructElem** 591
 - StructTreeRoot** 590
 - Template** 557
 - Thread** 484
 - Trans** 486
 - XObject** 261, 267, 284, 290, 448
- objects **9, 23, 24**
 - in FDF 558
 - fully opaque **467**
 - generation number. *See* generation numbers
 - hierarchy 81, 560
 - identifier. *See* object identifiers
 - indirect references. *See* indirect object references
 - length 17
 - as logical structure content items 589, 590, 591, 593, **598–599**, 601
- objects (*continued*)
 - number. *See* object numbers
 - processing 17
 - subtype. *See* object subtypes
 - syntax **27–41**
 - topmost **467**
 - type. *See* object types
 - See also*
 - array objects
 - boolean objects
 - dictionary objects
 - direct objects
 - external objects (XObjects)
 - form XObjects
 - function objects
 - graphics objects
 - group XObjects
 - image XObjects
 - indirect objects
 - inline image objects
 - integer objects
 - null object (**null**)
 - numeric objects
 - page objects
 - path objects
 - PostScript XObjects
 - real objects
 - reference XObjects
 - shading objects
 - stream objects
 - string objects
 - text objects
 - transparency group XObjects
 - OBJR** object type **599**
 - OEB (Open eBook™) file format
 - standard attribute owner 639
 - OEB-1.00** standard attribute owner **639**
 - Off appearance state
 - checkbox field 539
 - radio button field 541
 - offset printing presses 688
 - OLE (Object Linking and Embedding) 69
 - Once play mode (movie) **572**
 - OneColumn page layout **84**
 - OP** entry (graphics state parameter dictionary) **158, 213**
 - op** entry (graphics state parameter dictionary) **158, 213**
 - opacity 11, 133, **410–411**
 - alpha source parameter 149, 159, 268
 - anti-aliasing 421
 - backdrop 423, 424
 - in basic compositing formula 414
 - computation **420–424**
 - current alpha constant 149, 159, 268

- opacity (*continued*)
 - fully opaque objects **467**
 - notation for 413
 - soft masks 412, 421, 443
 - specifying **442–445**
 - See also*
 - constant opacity
 - group opacity
 - mask opacity
 - object opacity
 - result opacity
 - source opacity
- opacity constant **422**
- opaque imaging model 4, **11**, 410, 797
 - clipping 412, 439
 - graphics objects, painting of 133
 - graphics state, initialization for patterns 453
 - knockout groups compared to 434
 - masked images 275
 - overprinting 214, 458, 459, **463–466**
 - page group, flattening of 437, 469
 - shading patterns 234
 - spot colors 457
- Open** entry
 - text annotation dictionary **500**
- open paths **162**
- Open play mode (movie) **572**
- Open Prepress Interface (OPI) 676, **693–698**, 807
 - proxies 288, 289, 574, 676, **693–698**
 - server 693, 695
 - versions 693, 694, 697
 - 1.3 693, **694–697**
 - 2.0 693, **697–698**, 807
 - See also*
 - OPI comments
 - OPI dictionaries
 - OPI version dictionaries
- Open Prepress Interface (OPI) Specification* (Adobe Technical Note #5660) 694, 812, **813**
- OpenAction** entry (document catalog) **84**, 474, 513, 515, 734, 737, 748, 751, 808
 - URI actions ignored for 523
- OpenArrow line ending style **504**
- operands 11, **94**, 132
- Operation** entry (movie action dictionary) **526**
- operators, PDF 11, 132, **699–702**
 - ' (apostrophe) 134, 302, 305, **311**, 702
 - " (quotation mark) 134, 302, 305, **311**, 702
 - B** 134, **167**, 462, 699
 - b** 134, 162, **167**, 462, 699
 - B*** 134, **167**, 462, 699
 - b*** 134, **167**, 462, 699
 - BDC** 134, 583, **584**, 585, 593, 699
- operators, PDF (*continued*)
 - BI** 134, **278**, 279, 699
 - BMC** 134, 535, 583, **584**, 699
 - boolean 28
 - BT** 134, 294, 305, **308**, 535, 584, 637, 699
 - BX** **95**, 134, 699
 - c** 134, **163**, 165, 699
 - categories **134**
 - cm** 134, 139, 148, **156**, 266, 699
 - CS** 134, 173, 177, 179, 180, **216**, 218, 220, 699
 - cs** 134, 173, 177, 179, 180, **217**, 218, 220, 699
 - d** 134, **156**, 700
 - d0** 134, 324, 325, **326**, 700
 - d1** 134, 218, 324, 325, **326**, 700
 - defined **94**
 - Do** 134, 220, 224, 228, 232, **261**, 262, 267, 282, 283, 290, 451, 452, 467, 468, 469, 586, 595, 596, 599, 700
 - DP** 134, 583, **584**, 585, 700
 - EI** 134, **278**, 279, 280, 700
 - EMC** 134, 535, 583, **584**, 593, 700
 - ET** 134, 305, **308**, 535, 584, 637, 700
 - EX** **95**, 134, 700
 - F** 134, **167**, 700
 - f** 12, 134, 147, 162, 166, **167**, **168–169**, 173, 219, 221, 224, 228, 232, 700
 - f*** 134, **167**, 169, 700
 - G** 134, 173, 177, 178, 179, **217**, 218, 700
 - g** 134, 173, 177, 178, 179, **217**, 218, 264, 295, 700
 - gs** 134, **156**, 157, 159, 218, 301, 307, 374, 446, 700
 - h** 134, 162, **163**, 168, 700
 - i** 134, **156**, 404, 700
 - ID** 134, **278**, 279, 280, 700
 - implementations of 574
 - J** 134, **156**, 700
 - j** 134, **156**, 700
 - K** 134, 173, 177, 178, 180, **218**, 700
 - k** 134, 173, 177, 178, 180, **218**, 700
 - l** 134, **163**, 168, 700
 - M** 134, **156**, 700
 - m** 134, 162, **163**, 168, 169, 700
 - MP** 134, 583, **584**, 700
 - n** 134, **167**, 172, 585, 586, 700
 - ordering rules 133–136
 - painting 11, 12, 21, 151, 171, 203, 204, 207, 214, 215, 221, 222, 223, 232
 - postfix notation 94, 132
 - procedure sets 573, 574, 804
 - Q** 134, 152, **156**, 172, 223, 266, 283, 296, 306, 535, 587, 701, 706, 807
 - q** 134, 152, **156**, 172, 223, 266, 283, 535, 587, 701, 706, 807
 - re** 134, 162, 163, **164**, 701
 - relational 28
 - RG** 134, 173, 177, 178, 180, **217**, 218, 701
 - rg** 134, 173, 177, 178, 180, **217**, 218, 456, 461, 701

operators, PDF (*continued*)

ri 134, **156**, 197, 216, 218, 701
S 12, 134, 147, 162, 166, **167**, 168, 173, 219, 221, 232, 325, 701
s 134, **167**, 701
SC 134, 173, 177, 179, 180, 201, **217**, 218, 701
sc 134, 173, 177, 179, 180, 201, **217**, 218, 247, 264, 701
SCN 134, 177, 201, 202, 206, **217**, 218, 220, 224, 227, 701
scn 134, 177, 201, 202, 206, **217**, 218, 220, 224, 227, 228, 701
sh 134, 218, **232**, 233, 234, 405, 443, 453, 701
summary 3
T* 134, 302, 305, **310**, 701
Tc 134, **302**, 701
TD 134, 305, **310**, 701
Td 134, 294, **310**, 701
Tf 36, 134, 158, 293, 294, **302**, 338, 535, 701
TJ 134, 298, 302, 304, **311**, 312, 314, 701, 794
Tj 134, 219, 224, 228, 232, 294, 295, 296, 297, 298, 302, **311**, 312, 354, 701, 793
TL 134, **302**, 701
Tm 134, **310**, 535, 701
Tr 134, 296, **302**, 701
Ts 134, **302**, 701
Tw 134, **302**, 701
Tz 134, **302**, 701
v 134, **163**, 165, 166, 701
W 134, 162, 169, 171, **172**, 585, 586, 702
w 134, 151, **156**, 296, 702
W* 134, 162, 169, 171, **172**, 585, 586, 702
y 134, **163**, 165, 166, 702

See also

- clipping path operators
- color operators
- compatibility operators
- graphics operators
- graphics state operators
- inline image operators
- marked-content operators
- path construction operators
- path-painting operators
- shading operator
- text object operators
- text-positioning operators
- text-showing operators
- text state operators
- Type 3 font operators
- XObject operator

operators, PostScript 292, **699–702**

abs 116, **703**
add 116, **703**
and 116, **704**
atan 116, **703**

operators, PostScript (*continued*)

beginbfchar 352, 354, 369, 371
beginbfrange 352, 369, 372
begincidchar 352, 354
begincidrange 352
begincmap 351
begincodespacerange 351, 354, 369, 371
beginnotdefchar 352, 355
beginnotdefrange 352, 355
beginrearrangedfont 352
beginusematrix 352
bitshift 116, **704**
ceiling 116, **703**
cleartomark 366
clip 702
closepath 699, 700, 701
concat 699
copy 116, **704**
cos 116, **703**
curveto 699, 701, 702
cvi 116, **703**
cvr 116, **703**
div 116, **703**
dup 116, **704**
endbfchar 352, 354, 369
endbfrange 352, 369
endcidchar 352, 354
endcidrange 352
endcmap 351
endcodespacerange 351, 354, 369, 371
endnotdefchar 352, 355
endnotdefrange 352, 355
endrearrangedfont 352
endusematrix 352
eoclip 702
eofill 699, 700
eq 116, **704**
exch 116, **704**
exp 116, **703**
false 116, **704**
fill 699, 700
findfont 290
floor 116, **703**
ge 116, **704**
grestore 701, 706, 807
gsave 701, 706, 807
gt 116, **704**
idiv 116, **703**
if 28, 116, **704**
ifelse 28, 116, **704**
index 116, **704**
le 116, **704**
lineto 22, 700
ln 116, **703**
log 116, **703**

operators, PostScript (*continued*)

lt 116, 704
mod 116, 703
moveto 22, 700
mul 116, 703
ne 116, 704
neg 116, 703
not 116, 704
or 116, 704
pop 116, 704
restore 807
roll 116, 704
round 116, 703
save 807
selectfont 701
setcachedevice 325, 700
setcharwidth 325, 700
setcmykcolor 700
setcolor 701
setcolorspace 699
setdash 700
setflat 700
setgray 700
sethalftone 797
setlinecap 700
setlinejoin 700
setlinewidth 702
setmiterlimit 700
setrgbcolor 701
setscreen 797
shfill 701
show 701
sin 116, 703
sqrt 116, 703
stroke 699, 701
sub 116, 703
true 116, 704
truncate 116, 703
 in type 4 (PostScript calculator) functions 116, 703–704
usecmap 351
usefont 352
xor 116, 704

OPI. *See* Open Prepress Interface

OPI color types 696

Process 696
 Separation 696
 Spot 696

OPI comments (PostScript) 289, 693, 694–698, 807

%ALDImageAsciiTag 697
%ALDImageColor 696
%ALDImageColorType 696
%ALDImageCropFixed 695
%ALDImageCropRect 695

OPI comments (PostScript) (*continued*)

%ALDImageDimensions 695
%ALDImageFilename 694
%ALDImageGrayMap 696
%ALDImageID 695
%ALDImageOverprint 696
%ALDImagePosition 695
%ALDImageResolution 696
%ALDImageTint 696
%ALDImageTransparency 696
%ALDImageType 696
%ALDObjectComments 695
%%ImageCropRect 698
%%ImageDimensions 697
%%ImageFilename 697
%%ImageInks 698
%%ImageOverprint 698
%%IncludedImageDimensions 698
%%IncludedImageQuality 698
%%MainImage 697
%%TIFFASCIIIDTag 697

OPI dictionaries 123, 288, 693, 694–698, 807

metadata inapplicable to 579

and trap networks 691

version 1.3 694–697

Color entry 696

ColorType entry 696

Comments entry 695

CropFixed entry 695

CropRect entry 695

F entry 694, 807

GrayMap entry 696

ID entry 695

ImageType entry 696

Overprint entry 696

Position entry 695

Resolution entry 696

Size entry 695

Tags entry 697

Tint entry 696

Transparency entry 696

Type entry 694

Version entry 694

version 2.0 697–698

CropRect entry 698

F entry 697, 807

IncludedImageDimensions entry 698

IncludedImageQuality entry 698

Inks entry 698

MainImage entry 697

Overprint entry 698

Size entry 697, 698

Tags entry 697

Type entry 697

Version entry 697

- OPI entry
 - image dictionary 269, 448, 693
 - type 1 form dictionary 285, 693
- OPI object type 694, 697
- OPI: *Open Prepress Interface Specification* (Adobe Systems Incorporated) 694, 812
- OPI version dictionaries 269, 285, 693–694, 697
 - 1.3 entry 693, 694
 - 2.0 entry 693, 694
- OPM entry (graphics state parameter dictionary) 158, 214
- Opt entry
 - checkbox field dictionary 540
 - choice field dictionary 546, 547, 802
 - FDF field dictionary 562, 565
 - radio button field dictionary 542, 543
- optimization of PDF files 17
- or operator (PostScript) 116, 704
- orange colorant
 - PANTONE Hexachrome system 205
- Order entry (type 0 function dictionary) 110, 113, 792
- Ordering entry (CIDSystemInfo dictionary) 337, 345, 362
- “OS/2” table (TrueType font) 360
- outline, document
 - See document outline
- outline dictionary 84, 477–478, 757, 758, 760, 762
 - Count entry 478
 - First entry 477, 478
 - Last entry 477, 478
 - Type entry 478
- outline hierarchy 84, 477, 478
 - example 757, 770
 - in Linearized PDF 730, 734, 737, 740
 - root 477
- outline hint table (Linearized PDF) 736, 751
- outline item dictionaries 477–479
 - A entry 478, 513, 519
 - AA entry (obsolete) 798
 - C entry 479
 - Count entry 478, 740
 - Dest entry 478, 479, 519
 - F entry 479
 - First entry 477, 478
 - Last entry 477, 478
 - Next entry 477, 478
 - Parent entry 478
 - Prev entry 477, 478
 - SE entry 479
 - Title entry 478
- outline item flags 479
 - Bold 479
 - Italic 479
- outline items 477–479, 770, 772
 - actions for 477, 478, 513
 - activating 477, 478, 513, 798
 - closed 477, 478, 772
 - color 479
 - destinations for 474, 476, 477, 478, 479, 798
 - flags 479
 - and go-to actions 519
 - in Linearized PDF 749
 - movie actions associated with 525
 - open 477, 478, 770
 - structure elements associated with 479
 - and trigger events 515
 - URI actions ignored for 523
- Outlines entry (document catalog) 84, 477, 757
- Outlines object type 478, 758, 760
- output devices
 - additive 201, 203, 383
 - bilevel 13, 14, 383, 390
 - black-generation function 379
 - CMYK 196, 381
 - color 14, 178, 383, 390
 - continuous-tone 376, 382
 - default halftone 382
 - device space 137–138, 141
 - displays. See displays, raster-scan
 - gamut 197, 198, 375
 - halftones 382, 383, 392
 - high-resolution 393
 - ICC profile 375
 - monochrome 376, 381, 382, 390
 - native color space. See native color space
 - output intents 684–688, 807
 - overprinting 214
 - paper-based 180
 - physical limitations 677
 - PostScript 22, 289, 317, 469, 574, 793, 797, 807
 - printers. See printers
 - process color model 214, 373, 376
 - raster 12–13, 137, 151, 263, 373, 374
 - rendering intents 197
 - resolution 13, 137, 149, 151, 234, 273, 394, 395
 - RGB 381
 - smoothness tolerance, limits on 405
 - subtractive 201, 203, 383
 - transfer functions 380
 - undercolor-removal function 379
- output intent dictionaries 85, 196, 684–688
 - See also PDF/X output intent dictionaries
- output intent subtypes 684, 685
 - GTS_PDFX 685
- output intents 5, 375, 574, 676, 684–688, 807
 - ICC color profiles for 192
 - output intent dictionaries 85, 196
 - subtype 684, 685

- output medium 202, 676, 679
 - backdrop color for compositing 436, 437
 - crop box 88, 677, 806
 - film 202, 676
 - imposition of pages on 436–437, 450, 451, 679, 806
 - media box 88, 677
 - paper 676
 - plates 676
 - properties of 374
 - OutputCondition** entry (PDF/X output intent dictionary) 685
 - OutputConditionIdentifier** entry (PDF/X output intent dictionary) 685, 686
 - OutputIntent** object type 685
 - OutputIntents** entry (document catalog) 85, 684
 - over** operator (Porter and Duff) 414
 - overflow hint stream (Linearized PDF) 731, 735
 - hint table offsets in 736
 - in linearization parameter dictionary 733, 808
 - object offsets unaffected by 742
 - and one-pass file generation 753
 - primary hint stream, concatenated with 735, 741
 - use discouraged 753
 - Overlay** blend mode 417
 - Overline text decoration type 647
 - overprint control 213–216
 - See also*
 - overprint mode
 - overprint parameter
 - Overprint** entry
 - version 1.3 OPI dictionary 696
 - version 2.0 OPI dictionary 698
 - overprint mode 150, 214–216
 - CompatibleOverprint blend mode, ignored by 466
 - K** operator 218
 - k** operator 218
 - nonzero 196, 215–216, 461
 - OPM** entry (graphics state parameter dictionary) 158
 - summary 463, 464–465
 - and transparency 458, 461, 462
 - overprint parameter 150, 151, 213–214
 - CompatibleOverprint blend mode, ignored by 466
 - and halftones 467–468
 - nonstroking 150, 158, 467
 - OP** entry (graphics state parameter dictionary) 158
 - op** entry (graphics state parameter dictionary) 158
 - stroking 150, 158, 467
 - summary 463, 464–465
 - and transfer functions 467–468
 - and transparency 458, 459, 461, 462, 463
 - overprinting 11
 - and alternate color space 204
 - and blend modes 459–460, 462, 463
 - overprinting (*continued*)
 - and **DeviceN** color spaces 206
 - in EPS files 792
 - opaque and transparent, compatibility between 460–462, 463
 - OPI proxies 696, 698
 - summary 463–466
 - and transparency 458–466
 - owner password 74, 76, 790
 - authenticating (algorithm) 81
 - computing (algorithm) 79
- ## P
- P** entry
 - annotation dictionary 490
 - bead dictionary 484, 755
 - encryption dictionary 76, 78, 79
 - linearization parameter dictionary 733, 737
 - page label dictionary 483
 - structure element dictionary 591
 - Web Capture command dictionary 672, 673
 - Windows launch parameter dictionary 521
 - P** highlighting mode (push) 501, 512
 - P** standard structure type 629, 630, 631, 632
 - PA** entry (link annotation dictionary) 501
 - packing of ILSEs 625, 641
 - floating elements exempt from 626
 - Paeth predictor function (LZW and Flate encoding) 50, 51
 - Page** artifact type 616
 - page artifacts 615
 - page boundaries 574, 676–680, 806
 - and bounding box 288
 - box colors 89
 - clipping to 473
 - color 681
 - dash pattern 681
 - display of 5, 679–680
 - in printing 139, 473, 679
 - in viewing of documents 139, 473
 - See also*
 - art box
 - bleed box
 - crop box
 - media box
 - trim box
 - page content order 613, 614, 618–620, 651
 - annotations, sequencing of 618–619
 - artifacts 618
 - illustration elements 624
 - and nested BLSEs 629
 - for reverse-order show strings 619
 - and text discontinuities 617

- page description languages **10**
 - See also* PostScript® page description language
- page description level **134, 136, 216**
- page descriptions **2, 14, 23, 36**
- page device parameters (PostScript) **692**
 - ProcessColorModel** **692**
- Page** entry
 - annotation dictionary (FDF files) **568**
 - reference dictionary **288**
- page group **89, 412, 436–437, 449–450**
 - backdrop **412, 433, 436, 437**
 - backdrop color **436**
 - and black-generation functions **469**
 - color space **437, 442, 452, 455, 466**
 - compositing in **412, 437**
 - compositing of **436, 437**
 - group alpha **437**
 - group color **437**
 - group color space, explicit **455**
 - group compositing formula **436–437**
 - isolated **436, 437, 450**
 - non-isolated **436, 450**
 - notation **437**
 - and overprinting **466**
 - and rendering intents **468**
 - transparency stack **412, 467**
 - and undercolor-removal functions **469**
- page indices **83, 481, 482**
 - in reference XObjects **288**
 - See also*
 - page labels
- page label dictionaries **83, 482–483**
 - P** entry **483**
 - S** entry **483**
 - St** entry **483**
 - Type** entry **483**
- page label hint table (Linearized PDF) **736, 751**
- page labels **83, 481–483, 798**
 - label prefix **482, 483**
 - labeling ranges **83, 482, 483**
 - numbering style **483**
 - in reference XObjects **288**
 - See also*
 - page indices
 - page label dictionaries
- page layout **84, 800**
- page mode **84, 472**
- Page** object type **88, 557, 758, 760, 775**
- page objects **81, 86, 87–91, 758, 760, 762, 775, 791**
 - AA** entry **90, 514**
 - and annotations **490**
 - Annots** entry **90, 489, 513, 618, 689, 691, 737, 775, 778, 779, 780, 807**
- page objects (*continued*)
 - ArtBox** entry **89, 677, 806**
 - article beads **484**
 - B** entry **89, 484, 557, 737, 791**
 - BleedBox** entry **88, 677, 806**
 - BoxColorInfo** entry **89, 679**
 - Contents** entry **89, 96, 152, 583, 691, 737, 738, 757, 791**
 - CropBox** entry **88, 89, 138, 139, 473, 677**
 - in destinations **474, 476**
 - Dur** entry **90, 485–486, 487**
 - Group** entry **89, 449, 470**
 - Hid** entry (obsolete) **791**
 - ID** entry **90, 675**
 - inheritance of attributes **87, 88, 91–92, 96**
 - LastModified** entry **88, 690**
 - in Linearized PDF **737, 738, 742–745**
 - MediaBox** entry **88, 91, 677, 737**
 - Metadata** entry **90**
 - parent content set **90, 675**
 - Parent** entry **88, 557**
 - PieceInfo** entry **88, 90, 581**
 - for pre-separated pages **684**
 - PZ** entry **90, 676**
 - Resources** entry **88, 96, 737, 739, 757**
 - Rotate** entry **89, 138, 487, 494, 806**
 - SeparationInfo** entry **90, 683, 684**
 - StructParents** entry **90, 601, 603**
 - in structural parent tree **590**
 - Thumb** entry **89, 480, 737**
 - Trans** entry **90, 485**
 - TrimBox** entry **89, 677, 806**
 - Type** entry **88**
 - in Web Capture content database **661, 662, 667, 670**
- page offset hint table (Linearized PDF) **742–745, 808**
 - header **742, 743, 747, 808**
 - and page retrieval **753, 754**
 - per-page entries **742, 744–745, 808**
 - primary hint stream, first table in **736, 741**
 - shared object hint table, references to **739**
- page-piece dictionaries **573, 581–582**
 - for form Xobjects **285**
 - and modification dates **88, 284, 582**
 - for page objects **90**
 - See also*
 - application data dictionaries
- page sets, Web Capture
 - See* Web Capture page sets
- page tree **81, 86–92, 757, 765**
 - in Linearized PDF **737, 739**
 - named pages in **557**
 - nodes. *See* page tree nodes
 - page objects. *See* page objects
 - root **83**

- page tree nodes **86**, 87, 88, 91, 758, 760, 762, 765
- Count** entry **86**
 - Kids** entry **86**
 - in Linearized PDF 734, 737
 - metadata inapplicable to 579
 - Parent** entry **86**
 - root 83
 - Type** entry **86**
- PageLabel** object type **483**
- PageLabels** entry (document catalog) **83**, 482, 798
- PageLayout** entry (document catalog) **84**, 791, 798
- PageMaker® page layout software 793
- PageMode** entry (document catalog) **84**, 472, 730, 734, 737
- pages **9**, 23
- additional-actions dictionary 90, 791
 - annotation dictionaries 90
 - art box. *See* art box
 - article beads 89, 484, 791
 - bleed box. *See* bleed box
 - boundaries. *See* page boundaries
 - bounding box **474–475**
 - box colors 89
 - composite **201**, 457, 683
 - content streams in 89, 93, 614, 615, 757, 758, 760, 762, 786, 807
 - crop box. *See* crop box
 - current **11**, 138, 139, 412, 425
 - in destinations 474, 475
 - display duration 90, **485**, 487
 - FDF (Forms Data Format) **566–568**, 804
 - fully opaque objects 467
 - gamut **375**
 - importing 285, 287, 288
 - imposition on output medium 436–437, 450, 451, 679, 806
 - indices 83, 288, **481**, 482
 - labels 83, 288, **481–483**, 798
 - logical structure elements on 591, 594, 598, 599
 - magnification factor 84, 90, 493, 494, 676
 - media box. *See* media box
 - metadata 90
 - modification date 88, 582
 - movie actions associated with 525
 - named 93, **557**
 - See also* template pages
 - output medium 374
 - page-piece dictionary 90
 - placement in another document 436, 450, 451, 679
 - positioning on output medium 677, 679
 - private data associated with 581, 582
 - resource dictionary 88, 284, 324, 796
 - rotation 89, 493, 494
 - separation dictionary 90
 - size limits **707**, 808
- pages (*continued*)
- in structural parent tree 90, 590
 - in Tagged PDF 614
 - template 93, **567–568**, 804
 - thumbnail image 89
 - transition dictionary 90, **486–487**
 - transparency group 89, 412, **436–437**, 449–450
 - trap network annotation 689
 - trigger events for **515**
 - trim box. *See* trim box
 - Web Capture content set 90
 - See also*
 - page objects
- Pages** entry
- document catalog **83**
 - FDF dictionary **561**, 566
 - name dictionary **93**, 557
 - separation dictionary **684**
- Pages** object type **86**, 758, 760
- Pagination** artifact type **616**
- pagination artifacts **615**
- paint types (tiling patterns) 221–222
- type 1 (colored) **221**, 223
 - type 2 (uncolored) **221–222**, 227
- painting
- external objects (XObject) 261, 700
 - filling. *See* filling
 - form XObjects 282–283, 451, 452
 - glyphs 293, 294, 305–306, 368, 462, 794
 - images 262, 263
 - non-isolated groups 451
 - nonzero overprint mode 215
 - opaque imaging model 410, 411, 412
 - overprint parameter 214
 - paths 12, 131, 132, **166–171**, 232, 462–463
 - scan conversion 405–407
 - stroking. *See* stroking
 - transparency groups 451–452
 - transparent imaging model 410, 412
- painting operators
- All** colorant name 203
 - clipping of 12, 171
 - current page 11
 - DeviceN** color spaces 207
 - filling 167, **168–171**, 699, 700
 - and graphics state 12, 151
 - None** colorant name 203
 - parameters 12
 - pattern cells 221, 222, 223
 - PostScript and PDF compared 21
 - Separation** color spaces 203
 - shading patterns 232
 - stroking 12, 131, 132, **166–168**, 699, 700, 701
 - tint transformation functions 204

- PaintType** entry
 - Type 1 font program 367
 - type 1 pattern dictionary 221, 454
- Palindrome play mode (movie) 572
- PANOSE Classification Metrics Guide* (Hewlett-Packard Company) 360, 814
- PANOSE™ classification system 360
- Panose** entry (CIDFont **Style** dictionary) 360
- PANTONE® Hexachrome™ color system 205
- Paperclip annotation icon 509
- Paragraph annotation icon 500
- paragraphlike elements, standard
 - See standard paragraphlike elements
- parameters, graphics state
 - See graphics state parameters
- Params** entry (embedded file stream dictionary) 124
- parent content set (Web Capture)
 - image 269, 675
 - page 90, 675
- Parent** entry
 - field dictionary 531
 - outline item dictionary 478
 - page object 88, 557
 - page tree node 86
- parent fonts (Type 0 font) 334
- parent tree, structural
 - See structural parent tree
- parentheses (()) 26, 312
 - as literal string delimiters 29
 - unbalanced 29, 30, 31
- ParentTree** entry (structure tree root) 590, 601, 602
- ParentTreeNextKey** entry (structure tree root) 590, 601
- Part standard structure type 627, 628, 631
- partial field names 531, 532–533, 564, 801
- Password field flag (text field) 543
- passwords 74, 76, 78
 - computing (algorithms) 79–81, 563, 790
 - for FDF encryption 563
 - owner 74, 76, 79, 81, 790
 - in text fields 543
 - user 74, 76, 78, 79–80
- patches, color
 - bicubic tensor-product 256, 257–259
 - Coons 250–252, 256, 257, 259
- Patent Clarification Notice, Adobe 7, 811
- patents, Adobe 7
- path construction operators 131, 134, 162–164
 - b** 162
 - c** 134, 163, 165, 699
 - in clipping paths 171
 - f** 162, 166, 173
- path construction operators (*continued*)
 - h** 134, 162, 163, 168, 700
 - l** 134, 163, 168, 700
 - m** 134, 162, 163, 168, 169, 700
 - in path objects 12
 - re** 134, 162, 163, 164, 701
 - S** 166
 - in Type 3 glyph descriptions 325
 - v** 134, 163, 165, 166, 701
 - W** 162
 - W*** 162
 - y** 134, 163, 165, 166, 702
- PATH environment variable (Windows) 802
- path objects 12, 132
 - as clipping paths 171, 585, 586
 - in glyph descriptions 324
 - graphics state, dependence on 136
 - m** operator 134
 - object shape 443
 - and path operators 162
 - re** operator 134
 - in Tagged PDF 627
- path-painting operators 131, 134, 162, 166–171
 - B** 134, 167, 462, 699
 - b** 134, 167, 462, 699
 - B*** 134, 167, 462, 699
 - b*** 134, 167, 462, 699
 - in clipping paths 171
 - ending path 132
 - F** 134, 167, 700
 - f** 12, 134, 147, 167, 168–169, 219, 221, 224, 228, 232, 700
 - f*** 134, 167, 169, 700
 - filling 167, 168–171, 699, 700
 - n** 134, 167, 172, 585, 586, 700
 - object shape 443
 - in path objects 12
 - S** 12, 134, 147, 162, 166, 167, 168, 173, 219, 221, 232, 325, 701
 - s** 134, 167, 701
 - shading patterns, compositing of 453
 - stroking 166–168, 699, 700, 701
 - in Type 3 glyph descriptions 325
 - and transparent overprinting 462–463
- paths 131, 151, 161–172, 232
 - clipping 131, 132, 171–172, 305, 306
 - construction 12, 131, 162–166, 699, 700, 701, 702
 - current 162–163, 166, 172
 - current point 163, 164, 165
 - filling 12, 131, 132, 167, 168–171, 699, 700, 762
 - for ink annotations 508
 - open 162
 - painting 12, 131, 132, 166–171, 232, 462–463
 - scan conversion 405–406

- paths (*continued*)
 - stroking 12, 131, 132, 148, 149, 152, 153, 155, 166–168, 508, 699, 700, 701
 - subpaths 162, 163, 164, 171, 306, 700
 - See also
 - path objects
- pattern cells 219, 221, 223, 227
 - bounding box 222
 - clipping 222
 - colors 221, 222
 - compositing in 453
 - compositing of 453
 - and fully opaque objects 467
 - as isolated groups 454
 - key 223
 - spacing 222
 - text objects in 309
 - transparent objects in 453
- Pattern** color spaces 176, 199, 219, 272
 - alternate color space, prohibited as 190, 204
 - base color space, prohibited as 199
 - blending color space, prohibited as 450
 - colored tiling patterns 224
 - default color space, prohibited as 195
 - initial color value 217
 - remapping of underlying color space 195
 - sampled images, prohibited in 268
 - setting 177, 216
 - setting color values in 217, 220
 - shadings, prohibited in 234
 - uncolored tiling patterns 227
 - underlying color space for 195, 217, 227, 456
 - underlying color space, prohibited as 227
- pattern dictionaries 219, 220, 232
 - PatternType** entry 220
 - See also
 - type 1 pattern dictionaries (tiling)
 - type 2 pattern dictionaries (shading)
- Pattern** entry (resource dictionary) 97, 217, 220, 224
- pattern matrix 141, 220, 221, 222, 223, 231
- Pattern** object type 221, 231
- pattern objects 219
- Pattern** resource type 97, 217, 220, 224
- pattern space 141, 220, 222, 223, 232, 233
- pattern types 221, 231
 - type 1 (tiling) 199, 220, 221–230
 - type 2 (shading) 157, 199, 220, 231–260
- patterns 11, 172, 176, 219–260, 793
 - as color values 220
 - content streams 93, 221, 222, 223, 227, 231
 - dictionaries. See pattern dictionaries
 - explicit masks, simulating 276
 - and form XObjects 220
- patterns (*continued*)
 - general properties 220
 - as named resources 97
 - object shape 443
 - page coordinate system, alignment with 220
 - parent content stream 220, 221, 223
 - Pattern** color spaces 199
 - pattern matrix 141, 220, 221, 222, 223, 231
 - pattern objects 219
 - pattern space. See pattern space
 - resources for 96
 - and transparency 453–454
 - for variable-text fields 534
 - See also
 - shading patterns (type 2)
 - tiling patterns (type 1)
- PatternType** entry 220
 - type 1 pattern dictionary 221
 - type 2 pattern dictionary 231
- Pause movie operation 526
- PCL (Printer Command Language) file format 19
- PCM** entry (trap network appearance stream dictionary) 692
- PDF files 9
 - body 61, 64, 757, 774
 - conversion from PostScript 20–21
 - cross-reference table. See cross-reference table
 - embedded file streams. See embedded file streams
 - encryption 4, 18, 31, 71–81, 790
 - header 61, 63, 70, 83, 784, 785, 786, 790, 791
 - incremental updates. See incremental updates
 - indirect generation 19–21
 - job tickets 374
 - network access 62, 725, 726, 727, 751–755
 - See also Forms Data Format (FDF)
 - optimization 17
 - portability 14–15, 25, 382, 570, 571, 797
 - pre-separated 683, 693
 - random access 17, 21, 61, 64, 71
 - single-pass generation 17
 - structure 21, 23, 61–71, 790
 - trailer 61, 67–68, 69, 790
 - example 757, 774, 775, 777, 780, 782
 - translation from other file formats 19
 - See also
 - file identifiers
- PDF name registry
 - See name registry, PDF
- PDF** procedure set 574
- PDF Public-Key Digital Signature and Encryption Specification* (Adobe Systems Incorporated) 548, 812
- PDF versions
 - See versions, PDF

- PDF Writer printer driver 20, 86
- PDF/X (Portable Document Format, Exchange) file
 - format 684, 685, 686
- PDF/X output intent dictionaries 685–686
 - DestOutputProfile** entry 685, 686
 - Info** entry 685, 686
 - OutputCondition** entry 685
 - OutputConditionIdentifier** entry 685
 - RegistryName** entry 686
 - S** entry 684, 685
 - Type** entry 685
- PDFDocEncoding** predefined character encoding 98, 127, 709, 710
 - for alternate descriptions 657
 - euro character 714
 - for FDF fields 562, 803
 - for JavaScript scripts 556, 802
 - for status strings 561
 - for text annotations 782
- pdfmark** language extension (PostScript) 21
- percent sign (%) character 26
 - as comment delimiter 27
 - in uniform resource locators, “unsafe” 664
- Perceptual** rendering intent 198
- period (.) character
 - double, in relative file specifications 120
 - double, in uniform resource locators (URLs) 664
 - in field names 533, 801
 - in file names 121
 - in handler names 548
 - in uniform resource locators (URLs) 664
 - in unique names (Web Capture) 666
- permissions, access
 - See* access permissions
- Pg** entry
 - marked-content reference dictionary 594, 598
 - object reference dictionary 599
 - structure element dictionary 591, 594, 598, 599
- photographs 12, 185, 197, 261
 - halftoning 14
- Photoshop® image editing software 581, 792
- PieceInfo** entry
 - page object 88, 90, 581
 - type 1 form dictionary 284, 285, 581
- pixels 13
 - in halftone screens 383–385, 389–390, 393
 - representation in memory 13–14
 - scan conversion 405–406
- PJTF (Portable Job Ticket Format) 24, 374, 677, 689, 692, 806
- placement attributes 641–642
 - Before 641
 - Block 641, 643, 644, 649
 - End 641, 644, 649
 - Inline 629, 640, 641, 643, 649
 - Start 641, 643, 649
- Placement** standard structure attribute 626, 629, 632, 637, 638, 640, 641, 643, 644, 649
- plates, color 3
 - Plate 1, *Additive and subtractive color* 178
 - Plate 2, *Uncalibrated color* 181
 - Plate 3, *Lab color space* 187
 - Plate 4, *Color gamuts* 187
 - Plate 5, *Rendering intents* 197
 - Plate 6, *Duotone image* 205
 - Plate 7, *Quadtone image* 205, 212
 - Plate 8, *Colored tiling pattern* 224
 - Plate 9, *Uncolored tiling pattern* 228
 - Plate 10, *Axial shading* 239
 - Plate 11, *Radial shadings depicting a cone* 241, 242
 - Plate 12, *Radial shadings depicting a sphere* 241
 - Plate 13, *Radial shadings with extension* 242
 - Plate 14, *Radial shading effect* 242
 - Plate 15, *Coons patch mesh* 250
 - Plate 16, *Transparency groups* 411
 - Plate 17, *Isolated and knockout groups* 433, 434
 - Plate 18, *RGB blend modes* 416
 - Plate 19, *CMYK blend modes* 416
 - Plate 20, *Blending and overprinting* 462
- play mode (movie) 572
 - Once 572
 - Open 572
 - Palindrome 572
 - Repeat 572
- Play movie operation 526
- plug-in extensions 3
 - action types 518
 - for actions 800, 801
 - annotation handlers 489, 498
 - for annotations 799
 - file systems 122
 - and Linearized PDF 734, 736, 741, 749
 - and logical structure 605
 - and marked content 583, 585
 - and metadata 575
 - modification dates, maintenance of 690
 - names, registering 723
 - output intents 807
 - for RGB output 684, 723
 - signature handlers 547
 - sound formats 570
 - Web Capture. *See* Web Capture plug-in extension
 - WebLink 800
 - and version compatibility 783, 784

- plus sign (+) character
 - in dates 100
 - in font subset names 323
- PNG (Portable Network Graphics) predictor functions
 - 50–51**
 - algorithm tags 51
 - Average **50, 51**
 - None **50, 51**
 - Paeth **50, 51**
 - Sub **50, 51**
 - Up **50, 51**
- PNG (Portable Network Graphics) Specification* (Internet RFC 2083) **50, 815**
- Poetica[®] typeface 323
- points (printers' unit) 139
- pop** operator (PostScript) **116, 704**
- pop-up annotation dictionaries **509**
 - Subtype** entry **509**
- pop-up annotations 492, 499, **508–509**
 - contents 509
 - parent annotation **508, 509**
 - See also*
 - pop-up annotation dictionaries
- pop-up help systems 493, 526
- pop-up windows 488, 491, 492, 502, 553
 - for circle annotations 504, 505
 - for file attachment annotations 509
 - for ink annotations 507, 508
 - for line annotations 503
 - for markup annotations 505, 506
 - for pop-up annotations 508
 - for rubber stamp annotations 507
 - for sound annotations 510, 657
 - for square annotations 504, 505
 - for text annotations 499, 500
- Popup** annotation type 499, **509**
- Popup** entry (annotation dictionary) **492, 508**
- portability of PDF files 14–15, 25, 382, 570, 571, 797
- Portable Job Ticket Format (PJTF) 24, 374, 677, 689, 692, 806
- Portable Job Ticket Format* (Adobe Technical Note #5620) 24, 374, 689, 692, **813**
- Position** entry (version 1.3 OPI dictionary) **695**
- position vector (glyph) **299**
 - in CIDFonts 341–342
 - DW2** entry (CIDFont) 341
 - W2** entry (CIDFont) 341–342
- POST request (HTTP) 552, 670, 673
- “post” table (TrueType font) 332, 334
- Poster** entry (movie dictionary) **571**
- poster images (movies) 571
- postfix notation 94, 132
- PostScript calculator functions
 - See* type 4 functions
- PostScript Language Document Structuring Conventions Specification* (Adobe Technical Note #5001) **812**
- PostScript Language Reference* (Adobe Systems Incorporated) 5, 116, 257, 317, 392, 689, 692, 703, 705, **812**
- PostScript[®] page description language xix
 - CMap files 351
 - CMap names 349
 - color rendering dictionary 375
 - composite fonts 334
 - conversion to PDF 19
 - current path 162
 - default user space 139
 - dictionary keys 35
 - document structuring conventions (DSC) 27
 - Encapsulated PostScript (EPS) 412
 - files 20, 807
 - font dictionaries 315
 - font names 314, 317, 320, 321–322, 322–323, 338, 353, 356
 - forms 282
 - halftone dictionaries 392, 401
 - image space 141
 - imaging model 1, 2, 5, 10
 - implementation limits 705
 - Indexed** color spaces 200
 - interpreter 118, 316, 351, 705
 - LanguageLevel 1 290, 793
 - LanguageLevel 2 793
 - LanguageLevel 3 793
 - names, compatibility with 787–788
 - null object 35
 - number syntax 28
 - OPI comments 289, 693
 - output devices 22, 289, 317, 469, 574, 793, 797, 807
 - page descriptions 20
 - page group, flattening of 437
 - patterns 220
 - and PDF, compared 21–22
 - postfix notation 94
 - predefined spot functions, definitions of 385–389
 - procedure sets 574
 - ProcessColorModel** page device parameter 692
 - scanner 117
 - sequential execution model 132
 - spot functions 392
 - transfer functions 392
 - transparent imaging model, compatibility with **469–470**
 - trapping instructions 689
 - Type 1 font programs 316, 366

- PostScript[®] page description language (*continued*)
- Type 3 fonts 325
 - type 7 shadings, data format 257
 - Type 42 font format 321
 - See also
 - operators, PostScript
 - PostScript XObjects
 - type 4 functions (PostScript calculator)
- PostScript XObject dictionaries **290**
- Level1** entry **290**
 - Subtype** entry **290**
 - Type** entry **290**
- PostScript XObjects **133, 261, 289–290**
- See also PostScript XObject dictionaries
- PP** entry (additional-actions dictionary, obsolete) 800
- PPK**. See public/private-key authentication
- preblending of soft-mask image data 447–449
- predefined character encodings **709–722, 796**
- for Symbol font, built-in 709, **718–720**
 - for ZapfDingbats font, built-in 709, **721–722**
- See also
- MacExpertEncoding** predefined character encoding
 - MacRomanEncoding** predefined character encoding
 - PDFDocEncoding** predefined character encoding
 - StandardEncoding** predefined character encoding
 - WinAnsiEncoding** predefined character encoding
- predefined CMaps 4, 340, **343–348, 709**
- as base CMap 349
 - character collections for 346–347, 621
 - character identification 368
 - Identity–H 345, 368, 621
 - Identity–V 345, 368, 621
 - with Type 0 fonts 353
 - Unicode mapping 621
- predefined spot functions **384–389, 393, 797**
- CosineDot** **386**
 - Cross** **388**
 - Diamond** **389**
 - Double** **386**
 - DoubleDot** **385**
 - Ellipse** **387**
 - EllipseA** **387**
 - EllipseB** **388**
 - EllipseC** **388**
 - InvertedDouble** **386**
 - InvertedDoubleDot** **385**
 - InvertedEllipseA** **387**
 - InvertedEllipseC** **388**
 - InvertedSimpleDot** **385**
 - Line** **384, 386**
 - LineX** **386**
 - LineY** **386**
 - Rhomboid** **388**
- predefined spot functions (*continued*)
- Round** **387**
 - SimpleDot** **384, 385**
 - Square** **388**
- Predictor** entry
- FlateDecode** filter parameter dictionary **49, 50–51**
 - LZWDecode** filter parameter dictionary **49, 50–51**
- predictor functions (LZW and Flate encoding) 46, 49, **50–52, 268**
- Average **50, 51**
 - None **50, 51**
 - Paeth **50, 51**
 - PNG (Portable Network Graphics) **50–51**
 - Sub **50, 51**
 - TIFF (Tag Image File Format) Predictor 2 **50, 51**
 - Up **50, 51**
- preferences, viewer
- See viewer preferences
- Preferences folder (Mac OS) 802
- premultiplied alpha
- See preblending of soft-mask image data
- “prep” table (TrueType font) 367
- prepress production 10, 574, **676–698**
- See also
- Open Prepress Interface (OPI)
 - output intents
 - page boundaries
 - printer’s marks
 - separation dictionaries
 - trapping
- presentations 481, **485–488**
- display duration 90, **485, 487**
 - transition dictionaries 90, **486–487**
 - transition style **486**
- pre-separated files 683
- and trapping 693
- Prev** entry
- file trailer dictionary **68, 69, 732, 733, 740, 741, 775, 777, 780, 782**
 - outline item dictionary 477, **478**
- PrevPage** named action **527**
- See also named-action dictionaries
- primary colorants **201, 204**
- and halftones 383, 391
- primary hint stream (Linearized PDF) **730, 735**
- in first-page cross-reference table 732
 - and first-page section, ordering of 735, 736, 751, 752
 - hint table offsets in 735, 741, 749
 - in linearization parameter dictionary 733, 735
 - object offsets, ignored by 741, 742, 744
 - and one-pass file generation 753
 - overflow hint stream, concatenated with 735, 741

- Print annotation flag 493, 682, 690, 799
- Print Setup dialog 806
- PrintArea** entry (viewer preferences dictionary) 473
- PrintClip** entry (viewer preferences dictionary) 473
- printer driver 19, 20
- printer's mark annotation dictionaries 682
 - MN** entry 682
 - Subtype** entry 682
- printer's mark annotations 499, 512, 680–682
 - annotation rectangle 680
 - appearance streams for 682
 - See also*
 - printer's mark annotation dictionaries
- printer's mark form dictionaries 682–683
 - Colorants** entry 683
 - MarkStyle** entry 683
- printer's marks 5, 574, 676, 677, 680–683
 - See also* printer's mark annotations
- PrinterMark** annotation type 499, 562, 682
 - opacity inapplicable to 492
- printers 137, 180
 - color 383, 400
 - dot-matrix 12, 13
 - and halftones 383
 - ink-jet 12, 13
 - laser 12, 13
 - process colorants 201
 - separations 203
- printing 75
 - access permission 75, 77
 - alternate images 274
 - annotations 493, 497, 799
 - embedded fonts, copyright restrictions on 365
 - glyph widths in 794
 - by launch actions 520, 521
 - list numbering 650
 - n*-up 472
 - OPI proxies 807
 - output medium, dialog with user on 374
 - page boundaries 473, 679
 - PostScript XObjects 289
 - predefined spot functions 797
 - Print Setup dialog 806
 - procedure sets and 804
 - R2L reading order 472
 - reference XObjects 288–289
 - trigger events associated with 516
- printing presses, offset 688
- Private** entry
 - application data dictionary 582
- Private standard structure type 628
- procedure sets 22, 573, 574, 758, 760, 762, 804
 - ImageB** 574
 - ImageC** 574
 - ImageI** 574
 - as named resources 97, 574
 - PDF** 574
 - and PostScript XObjects 290
 - Text** 309, 574
 - trap networks, excluded from 691
- process color components
 - and overprinting 464–465
 - spot colors, treating as 458
 - and transparent overprinting 461, 464–465, 466
- process color model 214, 373, 376, 692
 - DeviceCMY** 692
 - DeviceCMYK** 692
 - DeviceGray** 692
 - DeviceN** 692
 - DeviceRGB** 692
 - DeviceRGBK** 692
- process colorants
 - additive devices, inapplicable to 203
 - All** colorant name 203
 - and alternate color space 204
 - in composite pages 201, 683
 - and current blend mode 442
 - DeviceCMYK** color space 180
 - halftones for 401
 - and high-fidelity color 205
 - and overprinting 464–465
 - process color model 376
 - Separation** color spaces 201
 - spot colorants, approximation of 457
 - transfer functions 158
 - and transparency 456, 459
 - and transparent overprinting 460, 464–465
 - See also*
 - black colorant
 - cyan colorant
 - magenta colorant
 - yellow colorant
- process colors 376
 - and blending color space 415
 - and flattening of transparent content 470
 - group color space, conversion to and from 456
 - separations, previewing of 684
 - and transparency 456
 - and transparent overprinting 460
- Process OPI color type 696
- ProcessColorModel** page device parameter (PostScript) 692
- ProcSet** entry (resource dictionary) 97, 574, 757, 758
- ProcSet** resource type 97, 574, 757, 758

- producer applications, PDF **1**
 - accessibility to disabled users 651
 - artifacts, generation of 615
 - encoding of data 41
 - glyph widths, specification of 794
 - logical structure, use of 589
 - names, embedded spaces in 787
 - names, registering 723
 - page tree, handling of 86
 - PDF version, updating 63, 83
 - predefined CMaps, support for 347
 - printer's marks, generation of 680
 - procedure sets, specification of 574
 - ToUnicode** CMaps 16
 - producer applications, Tagged PDF
 - annotations, sequencing of 619
 - footnotes, placement of 633
 - hyphenation, specification of 617
 - logical structure, definition of 618
 - page content order, establishment of 618
 - Private grouping element 628
 - standard structure elements, role mapping to 613
 - Unicode mapping 621
 - Producer** entry (document information dictionary) **576**
 - production, prepress
 - See prepress production
 - production conditions 684, 685–686
 - Custom 685
 - registry 685, **686**
 - profiles, ICC color
 - See ICC color profiles
 - projecting square line cap style **153**, 168
 - Properties** entry (resource dictionary) **97**, 584, 585
 - Properties** resource type **97**, 584, 585
 - property lists 583, 584, **585**, 699, 700
 - for artifacts **616**, 617
 - Attached** entry (Tagged PDF artifact) **616**
 - BBox** entry (Tagged PDF artifact) **616**
 - E** entry 616, 621, 622, 633, **659**
 - Lang** entry 633, 652, 653, 654, 656
 - for logical structure content items 593
 - MCID** entry **593**, 633, 654, 656
 - Metadata** entry **580**
 - as named resources 97, 580, 584, 585
 - Type** entry (Tagged PDF artifact) **616**
 - Proportional** character class 362, 363
 - Proportional font characteristic **622**
 - proportional fonts **297**, 358
 - proportional scaling **566**
 - ProportionalRot** character class 362, 363
 - proxies
 - OPI 288, 289, 574, 676, **693–698**
 - reference XObject 285, **287**, 288, 289, 754
 - PS** XObject subtype 261, **290**
 - pseudorectangular lattices 248
 - public/private-key (PPK) authentication 538, 547, 548
 - publications, related 5–6
 - Pushbutton field flag (button field) **538**, 539, 540
 - pushbutton fields 537, 538, **539**
 - appearances for 565
 - and reset-form actions 554
 - and submit-form actions 554
 - pushbuttons 530
 - PushPin annotation icon 509
 - PZ** entry (page object) **90**, 676
- ## Q
- Q** entry
 - field dictionary 529, **534**, 535
 - free text annotation dictionary **502**
 - interactive form dictionary **529**
 - Q** operator 134, **156**, 701
 - and clipping path 172, 296, 306
 - and current transformation matrix (CTM) 266
 - and dynamic appearance streams 535
 - and form XObjects 283
 - and graphics state stack 152
 - implementation limit 706, 807
 - marked clipping sequences, prohibited in 587
 - and tiling patterns 223
 - q** operator 134, **156**, 701
 - and clipping path 172
 - and current transformation matrix (CTM) 266
 - and dynamic appearance streams 535
 - and form XObjects 283
 - and graphics state stack 152
 - implementation limit 706, 807
 - marked clipping sequences, prohibited in 587
 - and tiling patterns 223
 - quadding
 - form field 534, 535
 - free text annotation 502
 - QuadPoints** entry (markup annotation dictionary) **506**, 799
 - quadtone color **205**
 - example 212–213
 - QuarkXPress[®] publishing software 793
 - QuickDraw[®] imaging model 19, 20
 - quotation mark (") character
 - as text-showing operator 134, 302, 305, **311**, 702

- quotations
 - block 627
 - inline 633
- Quote standard structure type 633
 - BlockQuote, distinguished from 633
- R**
- R entry
 - appearance characteristics dictionary 536
 - appearance dictionary 497, 565, 682, 690
 - bead dictionary 484
 - encryption dictionary 76, 79
 - sound object 569, 570
 - structure element dictionary 592, 606
 - Web Capture image set 669
- R keyword 40
- R transition style 486
- R2L reading order 472
- radial shadings
 - See type 3 shadings
- radio button field dictionaries
 - Opt entry 542, 543
- radio button fields 537, 538, 540–543
 - normal caption 536
 - Off appearance state 541
 - value 541, 542, 543
 - See also
 - radio button field dictionaries
- Radio field flag (button field) 538, 539, 540
- raised capitals 649
- random access to PDF files 17, 21, 61, 64, 71
- range, function 107, 108, 109, 110, 111
- Range entry
 - function dictionary 108, 109, 111, 113, 117
 - ICC profile stream dictionary 190, 216, 272
 - Lab color space dictionary 187, 188, 216, 272
- raster 13
- raster image processor (RIP) 689
- raster output devices 12–13
 - device space 137
 - graphics state parameters 151
 - rendering 373
 - scan conversion 374
- Rate entry (movie activation dictionary) 572
- Raw sound encoding format 569, 570
- RC entry (appearance characteristics dictionary) 537
- RC4 encryption algorithm 72, 73, 79, 80, 81
 - copyright 73
 - for FDF 563
- re operator 134, 162, 163, 164, 701
- reading order 472
- ReadOnly annotation flag 493, 682, 690
- ReadOnly field flag 532
 - and widget annotations 493, 532
- real objects 27
 - precision limits 28, 706
 - range limits 28, 706
 - syntax 28
- Reason entry (signature dictionary) 549
- Rect entry (annotation dictionary) 490, 494, 504, 505, 523, 534, 548
- rectangles 101
 - as dictionary values 99
 - path construction 164, 701
- red color component
 - CMYK conversion 377, 380
 - cyan, complement of 378
 - DeviceRGB color space 178, 179
 - grayscale conversion 377
 - halftones for 401
 - in Indexed color table 199
 - initialization 180
 - and threshold arrays 390
 - transfer function 380, 381
- red colorant
 - additive primary 178, 179, 180
 - display phosphor 201
- Ref entry
 - type 1 form dictionary 285, 287
- reference area 624
 - and allocation rectangle 648
 - and floating elements 626
 - layout within 641, 642, 643, 644, 645, 646, 647
 - and nested BLSEs 625
 - stacking of BLSEs 625, 641
- reference counts (Web Capture image set) 669, 805
- reference dictionaries 285, 287–288
 - F entry 288
 - ID entry 288
 - Page entry 288
- Reference standard structure type 628, 633
 - in tables of contents 628
- reference XObjects 4, 133, 261, 285, 287–289
 - and annotations 289
 - bounding box 288
 - clipping to bounding box 288
 - containing document 287
 - in Linearized PDF 754
 - and logical structure 289
 - printing 288–289
 - proxy 285, 287, 288, 289, 754
 - target document 287, 288

- reflection
 - images 266
 - OPI proxies 695
 - transformation matrices 137
- reflow of content 613
 - artifacts 617
 - glyph widths 794
 - hidden page elements 617
 - illustrations 637
 - list numbering 650
 - page content order 614, 618
 - standard structure attributes 638, 640
 - standard structure types 627
 - word breaks 622
- registered names
 - conversion engines (Web Capture) 675
 - dictionary keys 786
 - first-class 124, **723–724**
 - marked-content tags 583
 - object types 36
 - output intent subtypes 685
 - second-class **724**
 - security handlers 72
 - signature handlers 547
 - third-class **724**
- registration targets 5, 203, 676, 680
 - as printer's mark annotations 512
- Registry** entry (CIDSystemInfo dictionary) **337, 345, 362**
- RegistryName** entry (PDF/X output intent dictionary) **686**
- regular characters 25, **26, 27, 32**
- related files arrays 123, **125–127, 128**
- related publications 5–6
- relational operators 28
- relative file specifications **119–120, 802, 803**
- Relative Uniform Resource Locators* (Internet RFC 1808) 119, 664, 673, **815**
- RelativeColorimetric** rendering intent 197, **198**
- remapping of colors 178, **194–195, 376, 450, 456, 686**
- remote go-to action dictionaries **520**
 - D** entry **520**
 - F** entry 476, **520**
 - NewWindow** entry **520**
 - S** entry **520**
- remote go-to actions 518, **520**
 - destination 474, 476, 520
 - and Linearized PDF 752
 - target file 520
 - See also*
 - go-to actions
 - remote go-to action dictionaries
- Rename** entry (FDF template dictionary) **567, 804**
- rendering 2, 10, **373–408**
 - alternate color spaces 204
 - of CIE-based colors 374
 - color 173, 373–374
 - color conversion. *See* color conversion
 - coordinate transformations, inverting 147
 - current page 11, 14
 - curves 164
 - flatness tolerance. *See* flatness tolerance
 - and graphics, distinguished 132, 373
 - graphics state parameters, device-dependent 147
 - halftones 14, 150, 373, **382–403, 467–468**
 - image interpolation 273
 - images 263
 - implicit color conversion 195
 - imported pages 289
 - intents. *See* rendering intents
 - marking 374
 - order of transformations 380
 - overprint control 214
 - scan conversion 13, 374, **403–408**
 - smoothness tolerance. *See* smoothness tolerance
 - transfer functions 373, **380–382, 392, 467–468**
 - and transparency 466–469
- rendering intents 194, **197–198, 375**
 - AbsoluteColorimetric 198**
 - current **149, 268, 453, 466, 468**
 - ICC color profiles 192
 - and nested transparency groups 468–469
 - Perceptual 198**
 - RelativeColorimetric 197, 198**
 - RI** entry (graphics state parameter dictionary) 158
 - ri** operator 156, 701
 - for sampled images 268
 - Saturation 198**
 - and transparency 468–469
- Repeat** entry (sound action dictionary) **525**
- Repeat play mode (movie) **572**
- replacement text 651, **658**
 - font characteristics unavailable for 622
 - for structure elements 592, 658
 - in Tagged PDF 616
 - and Unicode natural language escape 658
 - word breaks 623
- Reproduction of Colour, The* (Hunt) **814**
- Required field flag **532**
- reserved words 25
- reset-form action dictionaries **555**
 - Fields** entry **555**
 - Flags** entry **554, 555**
 - S** entry **555**

- reset-form actions 518, 528, 550, 554–555
 - default value 532
 - flag. *See* reset-form field flag
 - and signature fields 548
 - See also*
 - reset-form action dictionaries
- reset-form field flag 555
 - Include/Exclude 555
- ResetForm** action type 518, 555
- ResFork** entry (Mac OS file information dictionary) 125
- resolution (output devices) 13, 137, 149, 151, 234, 273, 394, 395
- Resolution** entry (version 1.3 OPI dictionary) 696
- resource dictionaries 95–97, 757
 - base images in 273
 - ColorSpace** entry 97, 177, 194, 216, 280, 450
 - and content streams 92, 285
 - current. *See* current resource dictionary
 - ExtGState** entry 97, 156, 157
 - Font** entry 97, 290, 293, 302, 317, 338, 534, 535
 - for form XObjects 284–285, 691
 - in Linearized PDF 737
 - metadata inapplicable to 579
 - for pages 88, 284, 324, 691, 796
 - Pattern** entry 97, 217, 220, 224
 - ProcSet** entry 97, 574, 757, 758
 - Properties** entry 97, 584, 585
 - Shading** entry 97, 232
 - for Type 3 fonts 324, 796
 - for variable-text fields 534, 535, 801
 - XObject** entry 97, 261, 266, 269, 282, 284
- resource fork (Mac OS) 125
- resource types 96, 97
 - ColorSpace** 97, 177, 194, 216, 280, 450
 - Encoding** 738
 - ExtGState** 97, 156, 157
 - Font** 97, 290, 293, 302, 317, 338, 534, 535, 738
 - FontDescriptor** 738
 - Pattern** 97, 217, 220, 224
 - ProcSet** 97, 574, 757, 758
 - Properties** 97, 584, 585
 - Shading** 97, 232
 - XObject** 97, 261, 266, 269, 282, 284
- resources 22, 24, 88, 691
 - in Linearized PDF 737, 738, 739
 - See also*
 - named resources
 - resource dictionaries
 - resource types
- Resources** entry
 - page object 88, 96, 737, 739, 757
 - stream dictionary 96
 - type 1 form dictionary 284, 534, 535
- Resources** entry (*continued*)
 - type 1 pattern dictionary 222
 - Type 3 font dictionary 324
- restore** operator (PostScript) 807
- result alpha
 - in compositing 420
 - in knockout groups 434, 435
 - notation 414, 427
- result color (transparent imaging model) 413
 - in compositing 414, 419, 420, 423, 452
 - in knockout groups 434, 435
 - notation 414, 427, 437
 - and overprinting 459, 460
 - and separable blend modes 416, 417
- result opacity 423–424
 - in knockout groups 435
 - notation 423, 427
 - soft clipping 444
- result shape 423–424
 - notation 423, 427
 - soft clipping 444
- Resume movie operation 526
- retinal scans (user authentication) 538, 547
- return-to-control (RTC) pattern (**CCITTFaxDecode** filter) 54
- reverse-order show strings 619–620
- ReversedChars** marked-content tag 619
- revision numbers
 - FDf encryption algorithm 563
 - security handler 74–75, 76
 - structure attributes 591, 592, 605, 606–607
 - structure elements 592, 606, 607
- RF** entry (file specification dictionary) 123, 126, 128
- RFCs (Requests for Comment), Internet
 - See* Internet RFCs
- RG** operator 134, 173, 177, 178, 180, 217, 218, 701
- rg** operator 134, 173, 177, 178, 180, 217, 218, 456, 461, 701
- RGB** color representation
 - for additive color 178
 - and **CMYK**, compared 180
 - CMYK**, conversion from 380
 - CMYK**, conversion to 150, 377–379, 469
 - DCTDecode** filter, transformation by 60
 - DeviceRGB** color space 176, 179
 - and grayscale, conversion between 377
 - in halftones 383
 - in output devices 172, 376
- RGB** color space abbreviation (inline image object) 280
- Rhomboid** predefined spot function 388
- Rl** entry
 - appearance characteristics dictionary 537
 - graphics state parameter dictionary 158

- ri** operator 134, 156, 197, 216, 218, 701
 - RIFF (Resource Interchange File Format) 569
 - right angle bracket (>) character 26
 - double, as dictionary delimiter 35, 67, 560
 - as EOD marker 44
 - as hexadecimal string delimiter 29, 32, 35, 122
 - right brace (}) character 26
 - as delimiter in PostScript calculator functions 116
 - right bracket (]) character 26
 - as array delimiter 34
 - right parenthesis (]) character 26
 - escape sequence for 30, 312
 - as literal string delimiter 29
 - right-to-left writing systems 619
 - RIP (raster image processor) 689
 - RL** filter abbreviation 280, 789
 - RITb writing mode 642
 - role map 590, 593
 - metadata inapplicable to 579
 - and Tagged PDF 613, 626, 627, 629, 637
 - RoleMap** entry (structure tree root) 590
 - roll** operator (PostScript) 116, 704
 - rollover appearance (annotation) 497
 - Root** entry
 - PDF trailer dictionary 560
 - file trailer dictionary 63, 68, 83, 732
 - root fields (interactive form) 529, 567
 - root font (Type 0 font) 334
 - Rotate** entry
 - movie dictionary 571
 - page object 89, 138, 487, 494, 806
 - rotation
 - annotations 493, 494, 496, 499
 - font matrix 324
 - images 266
 - movies 571
 - OPI proxies 695
 - order of transformations 143
 - pages 89, 493, 494
 - text space 309
 - transformation matrices 137, 141, 142
 - user space 494
 - round line cap style 153, 168
 - round line join style 154, 168
 - round** operator (PostScript) 116, 703
 - Round** predefined spot function 387
 - Rows** entry (CCITTFaxDecode filter parameter dictionary) 54
 - RowSpan** standard structure attribute 651
 - RSA Security, Inc. 73
 - RTF (Rich Text Format)
 - layout model 624
 - standard attribute owner 639
 - Tagged PDF, conversion from 613, 627
 - RTF-1.05** standard attribute owner 639
 - rubber stamp annotation dictionaries 507
 - Contents** entry 507
 - Name** entry 507
 - Subtype** entry 507
 - rubber stamp annotations 499, 507
 - contents 507
 - See also*
 - rubber stamp annotation dictionaries
 - Ruby** character class 363
 - ruby characters 363
 - run-length encoding compression 42, 52
 - RunLengthDecode** filter 42, 52
 - RL** abbreviation 280, 789
 - in sampled images 268
 - running heads 615
- ## S
- S border style (solid) 495, 799
 - S** entry 35
 - action dictionary 514
 - border style dictionary 495
 - box style dictionary 681
 - go-to action dictionary 519
 - group attributes dictionary 286, 287, 449, 452
 - hide action dictionary 527
 - hint stream dictionary 736
 - icon fit dictionary 566
 - import-data action dictionary 556
 - JavaScript action dictionary 556
 - launch action dictionary 521
 - movie action dictionary 526
 - named-action dictionary 528
 - page label dictionary 483
 - PDF/X output intent dictionary 684, 685
 - remote go-to action dictionary 520
 - reset-form action dictionary 555
 - soft-mask dictionary 445, 446
 - sound action dictionary 525
 - source information dictionary 670
 - structure element dictionary 591, 626
 - submit-form action dictionary 551
 - thread action dictionary 522
 - transition dictionary 486
 - transparency group attributes dictionary 450
 - URI action dictionary 523
 - Web Capture command dictionary 672, 674
 - Web Capture content set 668
 - Web Capture image set 669
 - Web Capture page set 668

- S guideline style (page boundaries) **681**
- S operator 12, 134, 166, **167**, 701
 - and current color 173
 - ending path 162
 - in glyph descriptions 325
 - and graphics state parameters 147, 166–168
 - and patterns 219, 221, 232
 - and subpaths 168
- s operator 134, **167**, 701
- SA** entry (graphics state parameter dictionary) **159**, 408
- SamePath Web Capture command flag **673**
- SameSite Web Capture command flag **673**
- sample data
 - sounds 569, 570
 - type 0 functions 109, 110, 111
- sample values (images) **12**, **262**
 - decoding 269, 271–273
 - in image masks 276
 - in image space 265
 - in inline images 278
 - order of specification 394, 398, 399
 - representation 264
 - in soft-mask images 447
 - source colors (transparent imaging model) 441
 - stream data 262, 263
- sampled functions
 - See type 0 functions
- sampled images
 - See images, sampled
- sans serif fonts 358
- Saturation** blend mode **419**
- Saturation** rendering intent **198**
- save** operator (PostScript) 807
- SC** operator 134, 177, **217**, 218, 701
 - in content streams 173
 - in **DeviceCMYK** color space 180
 - in **DeviceGray** color space 179
 - in **DeviceRGB** color space 180
 - in **Indexed** color spaces 201
 - and sampled images 173
- sc** operator 134, 177, **217**, 218, 701
 - in content streams 173
 - and **Decode** arrays 264
 - in **DeviceCMYK** color space 180
 - in **DeviceGray** color space 179
 - in **DeviceRGB** color space 180
 - in **Indexed** color spaces 201
 - and sampled images 173
 - and type 4 shadings 247
- scaling
 - anamorphic **566**
 - annotations 493, 494, 496, 499, 566
 - scaling (*continued*)
 - fonts 294, 302
 - icons 566
 - line width 152
 - OPI proxies 695
 - order of transformations 143
 - of pages 806
 - proportional **566**
 - text space 295, 302, 304, 309
 - transformation matrices 137, 141, **142**
 - user space 494
 - Web Capture pages 676
 - scan conversion **13–14**, 147, 173, 374, **403–408**
 - in Adobe Acrobat products 404
 - clipping 406
 - filling 406
 - flatness tolerance. See flatness tolerance
 - glyphs 14, 292, 407
 - images 406
 - for raster-scan displays 403
 - rules **405–407**
 - smoothness tolerance. See smoothness tolerance
 - stroke adjustment **149**, 152, 159, 168, **407–408**
 - stroking 406
 - SCN** operator 134, 177, **217**, 218, 701
 - in **DeviceN** color spaces 206
 - in **Indexed** color spaces 201
 - in **Pattern** color spaces 220, 224, 227
 - in **Separation** color spaces 202
 - scn** operator 134, 177, **217**, 218, 701
 - in **DeviceN** color spaces 206
 - in **Indexed** color spaces 201
 - in **Pattern** color spaces 220, 224, 227, 228
 - in **Separation** color spaces 202
 - Screen** blend mode **417**, 462
 - screens, halftone
 - See halftone screens
 - Script font flag **358**
 - script fonts 358
 - scroll bars, hiding and showing 472
 - scrolling, text fields 543
 - SE** entry (outline item dictionary) **479**
 - searching of text 10, 368, 613, 620, 794
 - second-class names **724**
 - Sect standard structure type **627**, 628, 631
 - security, document xx, 18
 - See also
 - encryption
 - signatures, digital
 - security handlers **71**, 72, 74–81, 723
 - revision number 74–75, 76
 - standard 71, 72, **74–81**, 790

- selectfont** operator (PostScript) 701
- separable blend modes **416–418**
 - color space 416
 - for spot colors 416, 460
 - white-preserving 460
- Separation** color spaces 176, 199, **201–205**, 272
 - All** colorant name **203**
 - alternate color space for 204, 207, 236
 - alternate color space, prohibited as 204
 - as base color space 199, 200
 - blending color space, prohibited as 450
 - color values 202
 - colorant name 202–203, 204
 - DeviceN** color spaces, colorant information for 207
 - and **DeviceN** color spaces, compared 205, 206–207
 - halftones for 401
 - initial color value 202, 217
 - None** colorant name **203**, 207
 - other color components, effect on in transparency
 - groups 458
 - and overprinting 214, **464–465**
 - parameters 202–204
 - for pre-separated pages 684
 - for printer's marks 683
 - remapping of alternate color space 195
 - setting color values in 217
 - for shadings 236
 - in soft masks 447
 - specification 202
 - spot color components in 457, 458
 - for spot colorants 376, 415
 - tint transformation function 204, 207, 236
 - tints **201**, 202, 203, 217
 - in transparency groups 196
 - and transparent overprinting 461, 462, **465**
- separation dictionaries 90, **683–684**
 - ColorSpace** entry **684**
 - DeviceColorant** entry **684**
 - metadata inapplicable to 579
 - Pages** entry **684**
- Separation OPI color type **696**
- SeparationColorNames** entry (trap network appearance stream dictionary) **692**
- SeparationInfo** entry (page object) **90**, 683, 684
- separations, color 172, 176, **201**, 574, 676, 683
 - accurate screens algorithm 393
 - and colorants 202
 - composite pages, generation from 683
 - halftones for 400
 - and overprinting 213
 - pre-separated files 683, 693
 - for spot color components 457
 - See also*
 - separation dictionaries
- Serif font flag **358**, 622
- Serified font characteristic **622**
- serified fonts 358
- set-state actions (obsolete) 518
- setcachedevice** operator (PostScript) 325, 700
- setcharwidth** operator (PostScript) 325, 700
- setcmymcolor** operator (PostScript) 700
- setcolor** operator (PostScript) 701
- setcolorspace** operator (PostScript) 699
- setdash** operator (PostScript) 700
- SetF** entry (FDF field dictionary) **565**
- SetFf** entry (FDF field dictionary) **564**
- setflat** operator (PostScript) 700
- setgray** operator (PostScript) 700
- sethalftone** operator (PostScript) 797
- setlinecap** operator (PostScript) 700
- setlinejoin** operator (PostScript) 700
- setlinewidth** operator (PostScript) 702
- setmiterlimit** operator (PostScript) 700
- setrgbcolor** operator (PostScript) 701
- setscreen** operator (PostScript) 797
- sFamilyClass field (TrueType font) 360
- SGML (Standard Generalized Markup Language)
 - PDF logical structure compared with 589
- sh** operator 134, 218, **232**, 701
 - background color ignored by 234
 - compositing 453
 - object shape 443
 - smoothness tolerance 405
 - target coordinate space 233
- shading dictionaries 115, **231**, 232, **233–235**, 236
 - AntiAlias** entry **234**
 - Background** entry 232, **234**, 238, 239, 453
 - BBox** entry **234**, 237, 241
 - ColorSpace** entry 232, **234**, 235, 236
 - Function** entry 235, 236
 - metadata 580
 - as named resources 97
 - sh** operator 232
 - ShadingType** entry **234**, 237, 257
 - See also*
 - type 1 shading dictionaries (function-based)
 - type 2 shading dictionaries (axial)
 - type 3 shading dictionaries (radial)
 - type 4 shading dictionaries (free-form Gouraud-shaded triangle meshes)
 - type 5 shading dictionaries (lattice-form Gouraud-shaded triangle meshes)
 - type 6 shading dictionaries (Coons patch meshes)
 - type 7 shading dictionaries (tensor-product patch meshes)

- Shading** entry
 resource dictionary 97, 232
 type 2 pattern dictionary 231
- shading objects 133
 anti-aliasing 234
 background color 234, 238, 239, 242
 bounding box 234, 237, 238, 241
 clipping 234
 color space 194, 234, 235–236, 237, 238, 240, 244, 247, 249, 253
 domain 115, 237, 238, 239, 240, 241
 shading operator 232, 701
 shading type 236–260
 target coordinate space. *See* target coordinate space
See also
 shading dictionaries
 type 1 shadings (function-based)
 type 2 shadings (axial)
 type 3 shadings (radial)
 type 4 shadings (free-form Gouraud-shaded triangle meshes)
 type 5 shadings (lattice-form Gouraud-shaded triangle meshes)
 type 6 shadings (Coons patch meshes)
 type 7 shadings (tensor-product patch meshes)
- shading operator 134, 232, 701
sh 134, 218, 232, 233, 234, 405, 443, 453, 701
- shading patterns (type 2) 131, 199, 219, 231–260
 color values, interpolation of 235–236
 compositing of 453
 extended graphics state 231, 453
 gradient fill. *See* gradient fills
 and graphics state parameter dictionaries 157
 metadata for 580
 nonzero overprint mode, unaffected by 215
 object shape 443
 smoothness tolerance 405
 and transparent overprinting 461, 466
 and type 3 (stitching) functions 115
See also
 shading objects
 type 2 pattern dictionaries
- Shading** resource type 97, 232
- shading types 236–260
 type 1 (function-based) 233, 234, 237–238
 type 2 (axial) 233, 234, 236, 238–239
 type 3 (radial) 233, 234, 236, 239–243
 type 4 (free-form Gouraud-shaded triangle meshes) 233, 234, 243–247, 253
 type 5 (lattice-form Gouraud-shaded triangle meshes) 233, 234, 248–250, 253
 type 6 (Coons patch meshes) 233, 234, 250–256, 256
 type 7 (tensor-product patch meshes) 233, 234, 256–260
- ShadingType** entry (shading dictionary) 234, 237, 257
- shadow, diffuse achromatic 185
- shape (transparent imaging model) 171, 410–411
 alpha source parameter 149, 159, 268
 anti-aliasing 421
 backdrop 423, 424
 in basic compositing formula 414
 computation 420–424
 current alpha constant 149, 159, 268
 notation for 413
 soft masks 412, 421, 443
 specifying 442–445
See also
 constant shape
 group shape
 mask shape
 object shape
 result shape
 source shape
- shape constant 421
- shared object hint table (Linearized PDF) 736, 745–747, 809
 group entries 739, 744, 746, 747, 808, 809
 header 746, 747, 809
 and page retrieval 753
 shared object identifiers 744, 750
- shared object identifiers (Linearized PDF) 743, 744, 750
- shared object references (Linearized PDF) 743, 744, 750
- shared object signatures (Linearized PDF) 747, 809
 signature fields, distinguished from 747
- ShellExecute function (Windows) 800
- shfill** operator (PostScript) 701
- shift direction 625, 647
- Shift-JIS character encoding 336, 344, 348, 788, 803
- show** operator (PostScript) 701
- ShowControls** entry (movie activation dictionary) 572
- showing of text 292, 293–295, 757, 760
 character encodings 328
 CMaps 354, 355
 Identity–H predefined CMap 345
 Identity–V predefined CMap 345
 operators 311–313, 794
 simple fonts 316
 ToUnicode CMaps 368
 transparent overprinting 463
 Type 3 fonts 324–325
- SI** entry (Web Capture content set) 668, 669, 670
- Sig** field type 531, 548
- Sig** object type 549
- SigFlags** entry (interactive form dictionary) 529, 530

- signature dictionaries 548–549
 - ByteRange entry 549
 - Contents entry 62, 549
 - Filter entry 549
 - Location entry 549
 - M entry 549
 - Name entry 549
 - Reason entry 549
 - SubFilter entry 549
 - Type entry 549
- signature field dictionaries 548
 - DV entry 548
 - FT entry 548
 - T entry 548
 - V entry 548, 551
- signature fields 529, 530, 531, 538, 547–550
 - access permissions 75, 77
 - appearance 549–550
 - flags. *See* signature flags
 - invalid appearance 550
 - shared object signatures (Linearized PDF), distinguished from 747
 - unvalidated appearance 549, 550
 - valid appearance 549, 550
 - value 548
 - See also*
 - signature field dictionaries
- signature flags 529, 530
 - AppendOnly 530
 - SignaturesExist 530
- signature handlers 547–548, 549, 550
 - Adobe.PPKLite 549
 - CICI.SignIt 549
 - Entrust.PPKF 549
 - name 548, 549
 - VeriSign.PPKVS 549
- signatures, digital xx, 18, 77, 538, 547
 - in FDF files 4, 562
 - handlers 547–548, 549, 550
 - and incremental updates 69, 562, 785
 - and interactive forms 75, 77, 530
 - public/private-key (PPK) authentication 538, 547, 548
 - See also*
 - signature fields
- SignaturesExist signature flag 530
- Signed sound encoding format 569, 570
- simple file specifications 118
- simple fonts 314, 316–334
 - in CID-keyed fonts 336
 - descriptor 316
 - encodings. *See* character encodings
 - glyph selection 316
 - subsets 322–323, 796
 - substitution 796
- simple fonts (*continued*)
 - Tj operator 294
 - ToUnicode CMaps, codespace ranges for 369
 - Unicode, mapping to 620
 - word spacing 303
 - See also*
 - TrueType® fonts
 - Type 1 fonts
 - Type 3 fonts
- SimpleDot predefined spot function 384, 385
- sin operator (PostScript) 116, 703
- single-byte character codes
 - in simple fonts 316
 - and text-showing operators 312
 - and word spacing 303
- single-pass file generation 17
- SinglePage page layout 84
- SIS content set subtype (Web Capture) 668, 669
- Size entry
 - embedded file parameter dictionary 125
 - file trailer dictionary 68, 732, 741
 - type 0 function dictionary 109, 110, 111, 113
 - version 1.3 OPI dictionary 695
 - version 2.0 OPI dictionary 697, 698
- skewing
 - images 266
 - OPI proxies 695
 - order of transformations 143
 - transformation matrices 137, 142
- slash (/) character 26, 94
 - as file specification delimiter 118, 122
 - as name delimiter 32, 34, 83, 357, 561
 - in uniform resource locators (URLs) 673
 - as UNIX file name delimiter 121
- SM entry (graphics state parameter dictionary) 159, 405
- small-cap fonts 358
- Smallcap font characteristic 622
- SmallCap font flag 358, 622
- SMask entry
 - graphics state parameter dictionary 159, 444
 - image dictionary 149, 268, 276, 444, 447, 448, 467, 797
- smoothness tolerance 151, 405
 - and color conversion 405
 - and flatness tolerance, compared 405
 - shading patterns 235
 - SM entry (graphics state parameter dictionary) 159
- snd file format 569
- soft clipping 159, 412, 421, 439, 444
 - and current clipping path, compared 159, 412, 439
 - and knockout groups 444
 - of multiple objects 444
 - and transparency groups 444

- soft hyphen character (Unicode) 617, 621, 714
- soft-mask dictionaries 149, 444, **445–447**, 450, 797
 - BC** entry 445, **446**
 - G** entry 445, **446**, 450
 - S** entry 445, **446**
 - TR** entry 445, **446**
 - Type** entry **446**
- soft-mask image dictionaries 447, **448–449**
 - Matte** entry 447, 448, **449**
- soft-mask images 268, **444**, 445, **447–449**, 797
 - height 448
 - image data, preblending of 447–449
 - matte color **447**, 449
 - source color 447
 - width 448
 - See also*
 - soft-mask image dictionaries
- soft masks 276, **412**, **439–441**
 - Alpha** subtype 445, **446**
 - alternate color space 447
 - for annotations 496
 - backdrop color 439, 440, 441, 446
 - bounding box 445
 - color space 440, **446**, 448, 449
 - coordinate system 446
 - current. *See* current soft mask
 - group backdrop 440, 445
 - Luminosity** subtype 445, **446**, 450
 - mask values 445, 446
 - object color 439
 - opacity, as source of 412, 421, 443
 - shape, as source of 412, 421, 443
 - soft clipping 159, **412**, 421, **439**, 444
 - source color 439
 - specifying 444, **445–449**
 - spot color components unavailable in 447
 - spot colors unavailable in 457
 - subtype 445, **446**
 - transfer functions for 440, 441, 445, 446
 - transparency groups, deriving from 425, **445**, 446
 - group alpha **439–440**, **445**, 446
 - group luminosity 412, **440–441**, **445**, 446
 - See also*
 - soft-mask dictionaries
 - soft-mask images
- SoftLight** blend mode **418**
- Sold annotation icon 507
- “Solving the Nearest-Point-On-Curve Problem” (Schneider) **814**
- “Some Properties of Bézier Curves” (Goldman) **814**
- Sony Trinitron[®] display 186
- Sort field flag (choice field) **546**
- sound action dictionaries **525**
 - Mix** entry **525**
 - Repeat** entry **525**
 - S** entry **525**
 - Sound** entry **525**, 568, 800
 - Synchronous** entry **525**
 - Volume** entry **525**
- Sound** action type 518, **525**
- sound actions 518, **524–525**, 800
 - volume 525
 - See also*
 - sound action dictionaries
 - sounds
- sound annotation dictionaries **510**
 - Contents** entry **510**
 - Name** entry **510**
 - Sound** entry **510**, 568
 - Subtype** entry **510**
- Sound** annotation type 499, **510**
- sound annotations 9, 499, **510**
 - alternate text description 657
 - contents 510
 - and text annotations, compared 510
 - See also*
 - sound annotation dictionaries
 - sounds
- Sound** entry
 - sound action dictionary **525**, 568, 800
 - sound annotation dictionary **510**, 568
- sound files **568–569**
 - AIFF 569
 - AIFF-C 569
 - RIFF 569
 - snd 569
- Sound** object type **569**
- sound objects **568–570**
 - B** entry **569**, 570
 - C** entry **569**, 570
 - CO** entry **569**
 - CP** entry **569**
 - E** entry **569**, 570
 - R** entry **569**, 570
 - and sound actions 525
 - and sound annotations 510
 - Type** entry **569**
- sounds 9, 488, 510, 513, 514, 524, **568–570**, 804
 - asynchronous 525
 - encoding format. *See* encoding formats, sound mixing 525
 - in movies 510, 570, 571, 572
 - synchronous 525
 - volume 525

- sounds (*continued*)
 - See also*
 - sound actions
 - sound annotations
 - sound objects
- source alpha
 - in compositing 419, 420
 - in knockout groups 435
 - notation **414, 427, 431**
 - and overprinting 462
- source color (transparent imaging model) **413**
 - blending color space, conversion to 415
 - and CompatibleOverprint blend mode 461
 - compositing 419, 420
 - in isolated groups 433
 - and nonseparable blend modes 419
 - notation **414, 427, 431**
 - and overprinting 458, 459, 460
 - in patterns 453
 - and separable blend modes 416, 417, 418
 - and soft-mask images 447
 - and soft masks 439
 - specifying **441**
 - in transparency groups 429, 431, 454
- source gamut (page) **375**
- source information dictionaries (Web Capture) 668, **669–671**
 - AU** entry **670**
 - C** entry **670**
 - E** entry **670, 671**
 - implementation limits 708
 - S** entry **670**
 - TS** entry **670, 671**
- source opacity **421–422, 424**
 - in knockout groups 435
 - notation **422, 423, 427**
- source shape **421–422, 424**
 - in knockout groups 434, 435
 - notation **422, 423, 427, 431**
- space (SP) character 24, **26**
 - clipping 306
 - in comments 27
 - in cross-reference table 64, 65, 758
 - in font names 320, 321
 - in form field mapping names 552
 - in hexadecimal strings 32
 - in names 787
 - nonbreaking 714
 - in reverse-order show strings 619–620
 - as word separator 623
 - word spacing 303
- SpaceAfter** standard structure attribute 640, **643, 649**
- SpaceBefore** standard structure attribute 640, **643, 649**
- Span marked-content tag 633, 652, 653, 654, 656, 657, 659
- Span standard structure type **633, 652**
- Speaker annotation icon 510
- special color spaces **176, 199–213**
 - blending color space, prohibited as 450
 - as default color spaces 195
 - inline images, prohibited in 280
 - for shadings 234
 - in transparency groups 456
 - See also*
 - DeviceN** color spaces
 - Indexed** color spaces
 - Pattern** color spaces
 - Separation** color spaces
- special graphics state operators **134**
 - cm** 134, 139, 148, **156, 266, 699**
 - Q** 134, 152, **156, 172, 223, 266, 283, 296, 306, 535, 587, 701, 706, 807**
 - q** 134, 152, **156, 172, 223, 266, 283, 535, 587, 701, 706, 807**
- specular highlight 185
- spell-checking 543, 546, 613, 622
- SpiderContentSet** object type **668**
- SpiderInfo** entry (document catalog) **85, 660**
- Split transition style **486, 487**
- spot color components 457
 - and alternate color space 457, 458
 - compositing of 457–458
 - in **DeviceN** color spaces 457
 - and group opacity 457
 - and group shape 457
 - in multitones 205
 - and overprinting 464–465, 466
 - in **Separation** color spaces 457, 458
 - separations for 457
 - soft masks, unavailable in 447
 - and tint transformation functions 457
 - tints for 457
 - transfer function 381
 - in transparency groups, painting of 457–458
 - and transparent overprinting 461, 464–465, 466
- spot colorants **376**
 - in composite pages 201, 457, 683
 - and current blend mode 442
 - and flattening of transparent content 470
 - and halftones 391, 401
 - in multitones 205
 - and overprinting 464–465
 - process colorants, approximation with 457
 - in **Separation** color spaces 201
 - and transparent overprinting 459, 460, 464–465

- spot colors
 - blending color space, not converted to 415
 - group color space, not converted to 457
 - in opaque imaging model 457
 - and separable blend modes 416, 460
 - soft masks, unavailable in 457
 - and transparency 456–458
 - and transparent overprinting 457, 460, 465
 - and white-preserving blend modes 460
- spot functions 115, 384–389, 391, 392, 393, 797
 - for color displays 390
 - predefined 115
 - and threshold arrays, compared 390
 - threshold arrays, converted internally to 390
 - See also*
 - predefined spot functions
- Spot OPI color type 696
- SpotFunction** entry (type 1 halftone dictionary) 393
- SPS** content set subtype (Web Capture) 668
- sqrt** operator (PostScript) 116, 703
- square annotation dictionaries 505
 - BS** entry 505
 - Contents** entry 505
 - Subtype** entry 505
- Square** annotation type 499, 505
- square annotations 499, 504–505
 - border style 490, 495
 - border width 505
 - contents 505
 - dash pattern 505
 - interior color 505
 - See also*
 - square annotation dictionaries
- Square line ending style 504
- Square list numbering style 650
- Square** predefined spot function 388
- Squiggly** annotation type 499, 506
- squiggly-underline annotation dictionaries
 - See* markup annotation dictionaries
- squiggly-underline annotations
 - See* markup annotations
- sRGB* (standard *RGB*) color space 192–193
 - group color space, unsuitable as 456
- St** entry (page label dictionary) 483
- stack
 - graphics state 152, 156
 - transparency. *See* transparency stack
- stacking of BLSEs 625, 632, 641, 642
 - floating elements exempt from 626
- Stamp** annotation type 499, 507
- standard 14 fonts 16, 317, 318, 319, 329, 622, 709, 760, 795
 - Courier 319, 795
 - Courier–Bold 319, 795
 - Courier–BoldOblique 319, 795
 - Courier–Oblique 319, 795
 - Helvetica 319, 795
 - Helvetica–Bold 319, 795
 - Helvetica–BoldOblique 319, 795
 - Helvetica–Oblique 319, 795
 - MacExpertEncoding** not used for 710
 - Symbol 319, 329, 795, 797
 - Times–Bold 319, 795
 - Times–BoldItalic 319, 795
 - Times–Italic 319, 795
 - Times–Roman 319, 795
 - ZapfDingbats 319, 329, 795, 797
- standard attribute owners 605, 638–639
 - CSS-1.00** 639
 - CSS-2.00** 639
 - HTML-3.20** 639
 - HTML-4.01** 639
 - Layout** 639, 640
 - List** 639, 650
 - OEB-1.00** 639
 - RTF-1.05** 639
 - Table** 639, 651
 - XML-1.00** 639
- standard blend modes 416–419, 442
- standard character sets 709–722
 - expert 709, 710, 715–717
 - Latin 709, 710, 711–714, 796
 - for Symbol font 709, 718–720
 - for ZapfDingbats font 709, 721–722
- standard grouping elements 624, 627–628
 - Art** 627, 628, 631
 - BlockQuote** 627, 633
 - Caption** 627, 630, 631
 - Div** 627, 628, 631
 - Document** 627, 628, 631
 - Index** 628, 632
 - NonStruct** 628
 - Part** 627, 628, 631
 - Private** 628
 - Sect** 627, 628, 631
 - strong structure 632
 - TOC** 628, 632
 - TOCI** 628
 - usage guidelines 631
 - weak structure 632
- standard illustration elements 624, 637–638
 - clipping in 637
 - Figure** 637, 638, 640, 644, 645, 649

- standard illustration elements (*continued*)
 - Form 618, **637**, 640, 644, 645, 649
 - Formula **637**, 640, 644, 645, 649
 - standard layout attributes for 640, **644–645**, 649
- standard layout attributes **639–649**
 - for BLSEs 640, **642–646**
 - BBox** 637, 640, **644**
 - BlockAlign** 640, **645**
 - EndIndent** 624, 640, **644**
 - Height** 637, 640, **645**, 648, 649
 - InlineAlign** 640, **646**
 - SpaceAfter** 640, **643**, 649
 - SpaceBefore** 640, **643**, 649
 - StartIndent** 624, 640, **643**, 644
 - TextAlign** 640, **644**
 - TextIndent** 640, 643, **644**
 - Width** 637, 640, **645**, 648, 649
 - general **640–642**
 - Placement** 626, 629, 632, 637, 638, 640, **641**, 643, 644, 649
 - WritingMode** 624, 640, **642**, 647, 651
 - for illustrations 640, **644–645**, 649
 - for ILSEs 640, **646–647**
 - BaselineShift** 633, 637, 640, **647**, 648, 649
 - LineHeight** 633, 640, 643, **646**, 648
 - TextDecorationType** 633, 640, **647**
 - for tables **644–646**
- standard list attribute **649–650**
 - ListNumbering** 649, **650**
- standard list elements **630**, 632
 - L 628, 629, **630**, 649
 - Lbl 629, **630**, 633, 634, 644, 649, 650
 - LBody 629, **630**, 644
 - LI 628, 629, **630**, 649
- standard paragraphlike elements **629–630**, 644
 - H 629, **630**, 631, 632
 - H1–H6 629, **630**, 632
 - P 629, **630**, 631, 632
- standard RGB (*sRGB*) color space 192–193
 - group color space, unsuitable as 456
- standard security handler (**Standard**) 71, 72, **74–81**, 790
- standard structure attributes 5, 605, 612, 613, 626, **638–651**
 - and basic layout model 623
 - See also*
 - standard attribute owners
 - standard layout attributes
 - standard list attribute
 - standard table attributes
- standard structure types 5, 592, 612, 613, **626–638**
 - and basic layout model 623
 - role map 590, 593
 - usage guidelines 631
- standard structure types (*continued*)
 - See also*
 - block-level structure elements (BLSEs)
 - inline-level structure elements (ILSEs)
 - standard grouping elements
 - standard illustration elements
- standard table attributes **650–651**
 - ColSpan** **651**
 - RowSpan** **651**
- standard table elements **631**, 632
 - standard layout attributes for **644–646**
 - Table 629, **631**, 640, 644, 645
 - TD 629, **631**, 640, 642, 644, 645, 646, 648, 650
 - TH 629, **631**, 640, 644, 645, 646, 648, 650
 - TR 629, **631**, 642
- StandardEncoding** predefined character encoding **329**, 709, **710**
 - as implicit base encoding 330
- start edge **625**
 - of allocation rectangle 649
 - in layout 625, 641, 643, 644, 646
- Start** entry
 - movie action dictionary 526
 - movie activation dictionary **571**, 572
- Start inline alignment **646**
- Start placement attribute **641**, 643, 649
- Start text alignment **644**
- StartIndent** standard structure attribute 624, 640, **643**, 644
- startxref** keyword **67**, 732, 740, 751, 775, 777, 780, 782
- state, graphics. *See* graphics state
- Status** entry (FDF dictionary) **561**, 803
- StemH** entry (font descriptor) **357**
- StemV** entry (font descriptor) **357**
- stencil, uncolored tiling pattern as 221
- stencil masking 264, 275, **276–277**
 - character glyphs, painting 276–277
 - and image interpolation 277
- See also*
 - image masks
- stitching functions
 - See* type 3 functions
- Stm** entry (marked-content reference dictionary) **594**
- StmOwn** entry (marked-content reference dictionary) **594**
- Stop movie operation **526**
- stream dictionaries 36, 37, **38**, 235, 279
 - DecodeParms** entry **38**, 41
 - DP** abbreviation 788
 - as direct objects 36
- F** entry **38**, 568, 569
- FDecodeParms** entry **38**, 41
- FFilter** entry **38**, 41

- stream dictionaries (*continued*)
 - Filter** entry 38, 41, 366, 448, 569
 - Length** entry 37, 38, 40, 73, 235, 279, 366, 757
 - Resources** entry 96
 - See also*
 - attribute objects
 - embedded file stream dictionaries
 - ICC profile stream dictionaries
 - hint stream dictionaries
 - image dictionaries
 - metadata stream dictionaries
 - PostScriptXObject dictionaries
 - printer's mark form dictionaries
 - sound objects
 - trap network appearance stream dictionaries
 - type 1 form dictionaries
 - type 1 pattern dictionaries (tiling)
 - type 4 shading dictionaries
 - type 5 shading dictionaries
 - type 6 halftone dictionaries
 - type 6 shading dictionaries
 - type 10 halftone dictionaries
- stream** keyword 36–37, 38, 788
- stream objects 25, 27, 36–38, 71, 788
 - as attribute objects 605
 - in cross-reference table reconstruction 707
 - data 36, 38, 280
 - as dictionary values 99
 - extent 37
 - indirect objects 36
 - length 36, 558
 - metadata associated with 578, 580
 - strings, compared with 36
 - syntax 36–38, 788
 - See also*
 - stream dictionaries
- streams
 - appearance. *See* appearance streams
 - CIDFont subsets 361
 - CMap files 318, 324, 343, 349, 353, 369
 - color table 200
 - content. *See* content streams
 - embedded CIDFont programs 336
 - embedded CMaps 336, 348
 - embedded file. *See* embedded file streams
 - embedded font programs 292, 314, 321, 357, 364, 366
 - encryption 71
 - external 37, 38, 41, 788, 805
 - FDf differences 562
 - glyph descriptions 323, 325
 - halftone 391, 394, 398, 399
 - hint (Linearized PDF). *See* hint streams
 - HTTP form submissions 673
 - ICC profile 190, 686
 - See also* ICC profile stream dictionaries
- streams (*continued*)
 - image 36, 262, 263, 268
 - JavaScript scripts 556, 563
 - and JBIG2 encoding 56, 57
 - metadata. *See* metadata streams
 - page descriptions 36
 - pattern 219, 224
 - poster image (movie) 571
 - PostScript LanguageLevel 1 290
 - shading 231, 235, 244
 - sound objects 568–570
 - threshold arrays 392
 - ToUnicode** CMaps 369
 - trap networks 692
 - type 0 (sampled) functions 109, 110
 - type 4 (PostScript calculator) functions 115
 - See also*
 - stream objects
 - stream dictionaries
- strikeout annotation dictionaries
 - See* markup annotation dictionaries
- StrikeOut** annotation type 499, 506
- strikeout annotations
 - See* markup annotations
- string objects 25, 26, 27, 29–32
 - as dictionary values 99
 - length limit 29, 706
 - as name tree keys 101, 124, 476, 498, 556, 557
 - syntax 29–32
- strings
 - color table 200
 - default appearance. *See* default appearance strings
 - destinations, names of 476–477
 - element identifiers (logical structure) 591
 - encryption 71, 72
 - file identifiers 581
 - file specification 118–122, 123
 - hexadecimal 29, 32
 - JavaScript scripts 556, 563
 - literal 29–31
 - production condition names 685
 - production condition registry 686
 - reverse-order show strings 619–620
 - showing. *See* showing of text
 - text. *See* text strings
 - text objects 291
 - Web Capture content types 668
 - See also*
 - string objects
- stroke adjustment, automatic 152, 407–408
 - for raster-scan displays 407
 - See also*
 - stroke adjustment parameter

- stroke adjustment parameter 149
 - S operator 168
 - SA entry (graphics state parameter dictionary) 159, 408
- stroke** operator (PostScript) 699, 701
- stroking
 - color. *See* stroking color, current
 - color space. *See* stroking color space, current
 - glyphs 305, 367, 794
 - ink annotations 508
 - paths 12, 131, 132, 148, 149, 151, 152, 153, 155, 166–168, 508, 699, 700, 701, 762
 - scan conversion 406
 - stroke adjustment 149, 152, 159, 168, 407–408
 - text 12, 132, 296, 305
 - text rendering mode 305, 367, 794
 - and transparent overprinting 461, 462–463, 466
- stroking alpha constant, current 149, 159, 444
 - for annotations 491
 - and fully opaque objects 467
 - initialization 452
 - and overprinting 462, 463
 - setting 444
- stroking color, current 148, 151
 - DeviceCMYK** color space 180, 218, 700
 - DeviceGray** color space 179, 217, 700
 - DeviceN** color spaces 206
 - DeviceRGB** color space 180, 217, 701
 - initial value 216–217
 - and S operator 168
 - Separation** color spaces 202
 - setting 173, 177, 217–218, 700, 701
 - text, showing 295, 305
- stroking color space, current 148
 - CIE-based color spaces 182
 - color components, number of 217
 - DeviceCMYK** color space 180, 218, 700
 - DeviceGray** color space 179, 217, 700
 - DeviceRGB** color space 180, 217, 701
 - Indexed** color spaces 199
 - and S operator 168
 - setting 173, 177, 216–218, 699
- StructElem** object type 591
- StructParent** entry 590, 601, 602
 - annotation dictionary 492, 601
 - image dictionary 269, 448, 601
 - type 1 form dictionary 285, 601
- StructParents** entry 590, 601, 602, 603, 655
 - page object 90, 601, 603
 - type 1 form dictionary 285, 601
- StructTreeRoot** entry
 - document catalog 85, 589, 628
- StructTreeRoot** object type 590
- structural parent tree 590, 600–604
 - annotations 492
 - form XObjects 285
 - image XObjects 269
 - next key 590, 601
 - page objects 90
- structure, logical
 - See* logical structure
- structure attributes 605–607
 - attribute classes 590, 605–606
 - owner 605
 - revision number 591, 605, 606–607
 - standard. *See* standard structure attributes
- structure element dictionaries 591–592
 - A entry 591, 592, 605, 606, 607, 638
 - ActualText** entry 592, 616, 621, 622, 623, 638, 658
 - Alt** entry 592, 616, 621, 622, 638, 657, 658
 - C entry 592, 606, 607, 638
 - ID entry 591
 - K entry 591, 593, 598, 599
 - Lang** entry 592, 638, 652, 654
 - P entry 591
 - Pg** entry 591, 594, 598, 599
 - R entry 592, 606
 - S entry 591, 626
 - T entry 592
 - Type** entry 591
- structure elements 589
 - access, dictionary entries for 602
 - alternate description 592, 615, 657, 658
 - annotations, sequencing of 618–619
 - attribute classes 592, 606
 - attribute objects associated with 591
 - and basic layout model (Tagged PDF) 623
 - block-level (BLSEs). *See* block-level structure elements
 - as content items 589, 591, 593, 599–600
 - content items associated with 589, 591, 593
 - content items, finding from 590, 600–604, 655
 - element identifier 590, 591
 - form XObjects in 595–598
 - inline-level (ILSEs). *See* inline-level structure elements
 - language identifier 85, 592
 - natural language specification 652, 653, 654, 656–657
 - outline items, associated with 479
 - replacement text 592, 658
 - revision number 592, 606, 607
 - structure type 591, 626
 - Tagged PDF layout 627
 - title 592
 - See also*
 - structure attributes
 - structure element dictionaries
 - structure types

- structure hierarchy **589–592**
 - and accessibility to disabled users 651, 652, 653, 654, 656
 - in Linearized PDF 740, 750
 - logical structure order 618
 - table of contents parallel to 628
 - See also*
 - structure elements
 - structure tree root
- structure tree
 - See* structure hierarchy
- structure tree root 85, **589–590**
 - class map 590, 606, 638
 - ClassMap** entry **590**, 606, 638
 - and content extraction 628
 - IDTree** entry **590**, 591
 - K** entry **590**, 628
 - in Linearized PDF 750
 - metadata inapplicable to 579
 - ParentTree** entry **590**, 601, 602
 - ParentTreeNextKey** entry **590**, 601
 - role map 590, 593
 - RoleMap** entry **590**
 - Type** entry **590**
- structure types 34, 591, **592–593**
 - and marked-content tags 594
 - role map 590
 - standard. *See* standard structure types
- Style** dictionaries
 - See* CIDFont **Style** dictionaries
- Style** entry (CIDFont font descriptor) **361**
- sub** operator (PostScript) **116**, **703**
- Sub predictor function (LZW and Flate encoding) **50**, 51
- SubFilter** entry (signature dictionary) **549**
- Subject** entry (document information dictionary) **576**
- submit-form action dictionaries **551**
 - F** entry **551**
 - Fields** entry **551**, 553, 554
 - Flags** entry **550**, **551**, 553
 - S** entry **551**
- submit-form actions 518, 528, **550–554**, 802, 803
 - FDf differences stream 562
 - field flags, effects on 532
 - flags. *See* submit-form field flags
 - status string 561
 - See also*
 - submit-form action dictionaries
- submit-form field flags **550–553**
 - CanonicalFormat **553**
 - IncludeNoValueFields **551**, 553, 554
 - ExclFKey **553**
 - ExclNonUserAnnots **553**
 - ExportFormat **552**, 553, 554
 - submit-form field flags (*continued*)
 - GetMethod **552**, 553
 - Include/Exclude **551**, 553, 554
 - IncludeAnnotations **552**, 553
 - IncludeAppendSaves **552**, 562
 - SubmitCoordinates **552**
 - SubmitPDF **553**
 - XML **552**, 553, 554
 - Submit Web Capture command flag **673**
 - SubmitCoordinates field flag (submit-form field) **552**
 - SubmitForm** action type 518, **551**
 - SubmitPDF field flag (submit-form field) **553**
 - subpaths **162**, 163, 164, 171, 306, 700
 - subscripts 307
 - shift direction 647
 - subsets, font
 - See* font subsets
 - subtractive color components
 - and blend functions 465
 - subtractive color representation **178**
 - in blending color space 415, 450
 - and default color spaces 195
 - DeviceCMYK** color space 180
 - and halftones 383
 - and overprinting 465
 - process color components 178
 - in soft-mask images 448
 - tints 202, 457, 465
 - transfer functions, input to 381
 - subtractive colorants 201, 203
 - See also*
 - black colorant
 - cyan colorant
 - magenta colorant
 - yellow colorant
 - subtractive output devices 201, 203, 383
 - Subtype** entry **35–36**
 - annotation dictionary **490**
 - CIDFont dictionary **338**
 - circle annotation dictionary **505**
 - embedded file stream dictionary **124**
 - embedded font stream dictionary 357, 365, **366**
 - external object (XObject) 261, 290
 - file attachment annotation dictionary **509**
 - font dictionary 36, 292, 314
 - free text annotation dictionary **502**
 - image dictionary **267**, 279, 448, 480
 - ink annotation dictionary **508**
 - line annotation dictionary **503**
 - link annotation dictionary **501**
 - Mac OS file information dictionary **125**
 - markup annotation dictionary **506**
 - metadata stream dictionary **578**

Subtype entry (*continued*)

movie annotation dictionary 511
 multiple master font dictionary 320
 pop-up annotation dictionary 509
 PostScriptXObject dictionary 290
 printer's mark annotation dictionary 682
 rubber stamp annotation dictionary 507
 sound annotation dictionary 510
 square annotation dictionary 505
 text annotation dictionary 500
 trap network annotation dictionary 513, 691
 TrueType font dictionary 321
 Type 0 font dictionary 334, 353
 Type 1 font dictionary 317
 type 1 form dictionary 284
 Type 3 font dictionary 323
 widget annotation dictionary 512

subtypes, object

See object subtypes

superscripts 307

shift direction 625, 647

Supplement entry (**CIDSystemInfo** dictionary) 337, 345, 362

supplement number (character collection) 337, 347, 369

Supporting the DCT Filters in PostScript Level 2 (Adobe Technical Note #5116) 813

SW entry (icon fit dictionary) 566

Symbol standard font 319, 329, 795, 797

character encoding, built-in 709, 718–720

character names 368, 620

character set 709, 718–720

Symbol typeface 16

Symbolic font flag 358

symbolic fonts 329, 333, 358, 359

base encoding 330

built-in encoding 329, 330, 709

TrueType font encodings 334

Synchronous entry

movie activation dictionary 572

sound action dictionary 525

syntax, PDF 2, 23–129

array objects 34

boolean objects 28

character set 25–26

comments 27, 787

data structures 98–106

dates 100

dictionary objects 35–36

document structure 81–93

encryption 71–81, 790

file specifications 118–129

file structure 61–71, 790

filters 41–61, 788–790

syntax, PDF (*continued*)

functions 106–118

indirect objects 39–41

integer objects 28

lexical conventions 24–27

name objects 32–34, 787–788

name trees 101–105

null object 39

number trees 105–106

numeric objects 28–29

objects 27–41

real objects 28

rectangles 101

stream objects 36–38, 788

string objects 29–32

text strings 98–100

T

T entry

annotation dictionary 492, 509, 553

bead dictionary 484, 737

FDF field dictionary 564, 803

field dictionary 531, 532, 533, 548

hide action dictionary 527

hint stream dictionary 736

linearization parameter dictionary 733, 753

movie action dictionary 526

structure element dictionary 592

Web Capture page set 668

T highlighting mode (toggle) 512

T* operator 134, 302, 305, 310, 701

tab character

See horizontal tab (HT) character

table attributes, standard

See standard table attributes

table elements, standard

See standard table elements

Table standard attribute owner 639, 651**Table** standard structure type 629, 631

standard layout attributes for 640, 644, 645

tables of contents 628, 632

Tag annotation icon 509

Tagged PDF 573, 612–651

accessibility to disabled users 613, 614, 616, 627, 638

annotations, sequencing of 618–619

artifacts. See artifacts

basic layout model 623–626

character properties, extraction of 620–622

content reflow. See reflow of content

exporting 638, 639, 640, 647

font characteristics, determination of 622

- Tagged PDF (*continued*)
- hidden page elements 617
 - hyphenation 617, 625
 - and logical structure 573, 612, 613, 615
 - logical structure order **618**
 - mark information dictionary 85, **613–614**
 - page content 613, 614–623
 - page content order 613, 614, 617, **618–620**, 651
 - real content 614, **615**
 - reverse-order show strings **619–620**
 - standard structure attributes. *See* standard structure attributes
 - standard structure types. *See* standard structure types
 - text discontinuities 617
 - Unicode, mapping to 613, 614, 620–621
 - word breaks, identifying 612, 613, 614, 619, 622–623
- tags
- algorithm (PNG predictor functions) 51
 - font subset **322–323**, 357, 361, 796
 - marked-content **583**, 584, 723
 - TIFF 697
- Tags** entry
- version 1.3 OPI dictionary **697**
 - version 2.0 OPI dictionary **697**
- Tags for the Identification of Languages* (Internet RFC 1766) 653, **815**
- target attribute (HTML) 562
- target coordinate space **233**, 234, 237, 238, 240, 245, 249, 253, 259
- target document (reference XObject) **287**, 288
- Target** entry (FDF dictionary) **562**
- TbRl writing mode **642**
- Tc operator 134, **302**, 701
- TD operator 134, 305, **310**, 701
- Td operator 134, 294, **310**, 701
- TD standard structure type 629, **631**
- content rectangle 648
 - stacking direction 642
 - standard layout attributes for 640, 644, 645, 646
 - standard table attributes for 650
- Technical Notes, Adobe
- See* Adobe Technical Notes
- Template** object type 557
- template pages 93, **567–568**, 804
- Templates** entry
- FDF page dictionary **567**
 - name dictionary **93**, 557
- tensor-product patch meshes
- See* type 7 shadings
- tensor-product patches, bicubic 256, **257–259**
- terminal fields **528**, 531
- text **291–372**
- alignment. *See* text alignment
 - exporting 368, 620, 622, 628
 - extraction of 16, 74, 75, 77, 313, 368, 638, 652
 - filling 12, 132, 305
 - indexing 368, 613
 - in pattern cells 221
 - positioning 294
 - searching 10, 313, 368, 613, 620, 794
 - special graphical effects 295–297
 - spell-checking 543, 546, 613, 622
 - stroking 12, 132, 305
 - subscripts 307, 647
 - superscripts 307, 625, 647
 - and transparent overprinting 461, 466
 - See also*
 - fonts
 - glyphs, character
 - showing of text
 - text line matrix
 - text matrix
 - text objects
 - text operators
 - text rendering matrix
 - text rendering mode
 - text space
 - text state
 - writing mode
- text alignment **644**
- Center **644**
 - End **644**
 - Justify **644**
 - Start **644**
- text annotation dictionaries **500**
- Contents** entry **500**
 - Name** entry **500**
 - Open** entry **500**
 - Subtype** entry **500**
- Text** annotation type 499, **500**
- text annotations 69, 77, **499–500**, 799
- access permissions for 75
 - contents 500
 - font 499
 - and free text annotations, compared 502
 - modification date, updating 798
 - rotation, not subject to 499
 - scaling, not subject to 499
 - and sound annotations, compared 510
 - text size 499
 - in updating example 774, 775, 777, 778, 779, 780, 782
 - See also*
 - text annotation dictionaries

- text decoration type **647**
 - LineThrough **647**
 - None **647**
 - Overline **647**
 - Underline **647**
- text discontinuities **617**
- text field dictionaries **544**
 - MaxLen** entry **544**
- text field flags **543**
 - DoNotScroll **543**
 - DoNotSpellCheck **543**
 - FileSelect **543, 544**
 - Multiline **543**
 - Password **543**
- text fields **531, 538, 543–545**
 - flags. *See* text field flags
 - trigger events for **516**
 - value **543, 544**
 - variable text in **533**
 - See also*
 - text field dictionaries
- text files **15, 25**
- text font (T_f) parameter **301**
 - Font** entry (graphics state parameter dictionary) **158**
 - showing of text **293**
 - Tf** operator **302, 701**
- text font size (T_{fs}) parameter **301**
 - Font** entry (graphics state parameter dictionary) **158**
 - showing of text **293**
 - text matrix, updating of **313–314**
 - text space **309**
 - Tf** operator **294, 302, 701**
 - unscaled text space units unaffected by **302**
- text/html content type (MIME) **668**
- text identifiers (Web Capture page sets) **665, 667, 668**
- text knockout (T_k) parameter **301, 307–308**
 - TK** entry (graphics state parameter dictionary) **159**
 - transparent overprinting independent of **463**
- text line matrix (T_{lm}) **301, 308, 310–311**
- text matrix (T_m) **140, 294, 301, 302, 308, 309–310, 313–314**
 - text size **294**
 - translation **311**
- text object operators **134, 308**
 - BT** **134, 294, 305, 308, 535, 584, 637, 699**
 - ET** **134, 305, 308, 535, 584, 637, 700**
 - marked-content operators, combined with **584**
- text objects **12, 132, 291, 293, 308–314, 699, 700**
 - as clipping paths **172, 305–306, 585, 586**
 - and color operators **216**
 - and default appearance strings **535**
 - in glyph descriptions **324**
- text objects (*continued*)
 - and graphics state **136**
 - illustration elements (Tagged PDF) prohibited within **637**
 - operators in **309**
 - showing text **293**
 - text knockout parameter **159, 308, 441**
 - text line matrix **301**
 - text matrix **301**
 - text rendering matrix **301**
 - text state operators **301**
 - text state parameters **301**
 - Type 3 fonts in **309**
 - See also*
 - text object operators
- text operators **131, 760**
 - in text objects **291, 308**
 - Text** procedure set **309**
 - See also*
 - text object operators
 - text-positioning operators
 - text-showing operators
 - text state operators
 - Type 3 font operators
- text position, current
 - See* current text position
- text-positioning operators **134, 309–311**
 - in dynamic appearance streams **535**
 - T*** **134, 302, 305, 701, 310**
 - TD** **134, 305, 310, 701**
 - Td** **134, 294, 310, 701**
 - in text objects **309, 311**
 - Tm** **134, 310, 535, 701**
- Text** procedure set **309, 574**
- text rendering matrix (T_{rm}) **301, 308, 313**
- text rendering mode (T_{mode}) **301, 305–306, 367, 793–794**
 - and marked content **585, 586**
 - special graphical effects **295, 296**
- Tr** operator **302, 701**
 - and transparent overprinting **462–463**
 - Type 3 fonts unaffected by **296, 305**
- text rise (T_{rise}) parameter **301, 307**
 - text space **309, 313**
 - Ts** operator **302, 701**
- text-showing operators **134, 298, 311–313, 313, 701–702, 794**
 - ' (apostrophe) **134, 302, 305, 311, 702**
 - " (quotation mark) **134, 302, 305, 311, 702**
 - and CMaps **336, 354**
 - and composite fonts **334**
 - in dynamic appearance streams **535**
 - glyph positioning **298, 299, 309**
 - glyph selection **316**

- text-showing operators (*continued*)
- object shape 443
 - in text objects 309
 - TJ** 134, 298, 302, 304, **311**, 312, 314, 701, 794
 - Tj** 134, 219, 224, 228, 232, 294, 295, 296, 297, 298, 302, **311**, 312, 354, 701, 793
 - and Type 3 fonts 324–325, 326
- text space **140**, 158, **309**, **313–314**
- device space, relationship with 313
 - glyph positioning 299
 - glyph space, relationship with 298, 313, 323
 - origin 298, 309
 - rotation 309
 - scaling 295, 302, 304, 309
 - text size 294–295
 - and text state parameters 305
 - translation 297, 299, 309, 311
 - and Type 3 glyph descriptions 325
 - units 294–295, 311, 317
 - user space, relationship with 309
- text state 291, **300**
- See also*
 - text state operators
 - text state parameters
- text state operators **134**, 156, **301–302**
- in default appearance strings 534, 535
 - initialization 301
 - Tc** 134, **302**, 701
 - in text objects 309
 - Tf** 36, 134, 158, 293, 294, **302**, 338, 535, 701
 - TL** 134, **302**, 701
 - Tr** 134, 296, **302**, 701
 - Ts** 134, **302**, 701
 - Tw** 134, **302**, 701
 - Tz** 134, **302**, 701
- text state parameters 148, 291, 294, **300–301**, 302–308
- See also*
 - character spacing (T_c) parameter
 - horizontal scaling (T_h) parameter
 - leading (T_l) parameter
 - text font (T_f) parameter
 - text font size (T_{fs}) parameter
 - text knockout (T_k) parameter
 - text line matrix (T_{lm})
 - text matrix (T_m)
 - text rendering matrix (T_{rm})
 - text rendering mode (T_{mode})
 - text rise (T_{rise}) parameter
 - word spacing (T_w) parameter
- text strings **98–100**
- as annotation names 98, 490
 - as article names 98
 - as bookmark names 98
- text strings (*continued*)
- as choice field options 546, 547, 802
 - as dictionary values 99
 - encodings for 98, 710
 - as FDF option names 565
 - as field names 532
 - as field values 540, 542, 543, 544
 - as name tree keys 476
 - as page set titles (Web Capture) 668
 - as production condition names 685, 686
 - as structure element titles 592
 - as trap network descriptions 692
- text-to-speech conversion 651
- abbreviations and acronyms 658, 659
 - alternate descriptions 657
 - artifacts 615
 - hidden page elements 617
 - natural language specification 573
 - Unicode 620
 - word breaks 622
- TextAlign** standard structure attribute 640, **644**
- TextDecorationType** standard structure attribute 633, 640, **647**
- TextIndent** standard structure attribute 640, 643, **644**
- Tf** operator 36, 134, 158, 293, 294, **302**, 701
- CIDFonts, inapplicable to 338
 - in default appearance strings 535
- TH standard structure type 629, **631**
- content rectangle 648
 - standard layout attributes for 640, 644, 645, 646
 - standard table attributes for 650
- third-class names **724**
- third-party applications, development of xx
- thread action dictionaries **522**
- B** entry **522**
 - D** entry **522**
 - F** entry 476, **522**
 - S** entry **522**
- Thread** action type 518, **522**
- thread actions 518, **522**
- and Linearized PDF 752
 - and named destinations 476
 - target bead 522
 - target file 522
 - target thread 522
 - See also*
 - thread action dictionaries
- thread dictionaries 84, **483–484**
- F** entry 483, **484**
 - I** entry **484**, 724, 740
 - in Linearized PDF 734, 737, 740
 - thread actions, target of 522
 - Type** entry **484**

- thread information dictionaries **484**
 - in Linearized PDF 734, 740
 - registered names not required in 724
 - thread actions, target of 522
- thread information hint table (Linearized PDF) 736, **751**
- Thread** object type **484**
- threads, article 481, **483**, 484
 - in document catalog 83, 84
 - in Linearized PDF 737, 749, 752, 753, 754, 755
 - and thread actions 522
 - See also*
 - articles
 - beads
 - thread dictionaries
 - thread information dictionaries
- Threads** entry (document catalog) **84**, 483, 522, 734, 753
- threshold arrays **389–390**, 392
 - device space, defined in 390, 394, 398, 399
 - height 395
 - and spot functions, compared 390
 - spot functions converted internally to 390
 - type 6 halftones 391, 394, 395
 - type 10 halftones 391, 394, 395
 - type 16 halftones 391, 399, 400
 - width 395
- Thumb** entry (page object) **89**, 480, 737
- thumbnail hint table (Linearized PDF) 736, **747–749**
 - header 747, **748**
 - per-page entries 747–748, **749**
- thumbnail images 474, **480–481**
 - access permission 75, 77
 - color space 480
 - display of 83
 - hiding and showing 84, 472
 - image XObjects as 263, 267, 480
 - in Linearized PDF 739, 740, **747–749**
 - in page object 81, 89
 - sample limit **707**
 - unrecognized filters in 789
- Tl** entry (choice field dictionary) **546**
- TID** entry (Web Capture page set) 667, **668**
- TIFF (Tag Image File Format) Predictor 2 predictor function 51
- TIFF (Tag Image File Format) standard 46, 50, 697
- %%TIFFASCII Tag OPI comment (PostScript) **697**
- tilde, right angle bracket (~>) character sequence
 - as EOD marker **44**, 45
- tiling patterns (type 1) **199**, **219**, **221–230**
 - bounding box 222
 - colored **221**, **223–227**, 454
- tiling patterns (type 1) (*continued*)
 - compositing in 453
 - compositing of 453
 - and fully opaque objects 467
 - gradient fills in 232
 - key pattern cell **223**
 - metadata for 580
 - object opacity 422
 - object shape 422, 443
 - painting with 223
 - pattern cell. *See* pattern cells
 - resources 222
 - spacing 222
 - and text objects 309
 - uncolored 217, **221–222**, **227–230**, 454
 - See also*
 - type 1 pattern dictionaries
- tiling types (tiling patterns)
 - type 1 (constant spacing) 222
 - type 2 (no distortion) 222
 - type 3 (constant spacing and faster tiling) 222
- TilingType** entry (type 1 pattern dictionary) **222**
- time scale (movie) **571**
- Times* typeface 16, 292, 709
- Times–Bold standard font **319**, **795**
- Times–BoldItalic standard font **319**, **795**
- Times–Italic standard font **319**, **795**
- Times–Roman standard font **319**, **795**
- TimesNewRoman standard font name **795**
- TimesNewRoman–Bold standard font name **795**
- TimesNewRoman,Bold standard font name **795**
- TimesNewRoman–BoldItalic standard font name **795**
- TimesNewRoman,BoldItalic standard font name **795**
- TimesNewRoman–Italic standard font name **795**
- TimesNewRoman,Italic standard font name **795**
- TimesNewRomanPS standard font name **795**
- TimesNewRomanPS–Bold standard font name **795**
- TimesNewRomanPS–BoldItalic standard font name **795**
- TimesNewRomanPS–BoldItalicMT standard font name **795**
- TimesNewRomanPS–BoldMT standard font name **795**
- TimesNewRomanPS–Italic standard font name **795**
- TimesNewRomanPS–ItalicMT standard font name **795**
- TimesNewRomanPSMT standard font name **795**
- Tint** entry (version 1.3 OPI dictionary) **696**
- tint transformation functions 115, 190, 195, **204**
 - and color separations 684
 - for **DeviceN** color spaces 207, 216, 236
 - for **Separation** color spaces 204, 207, 236
 - and spot color components 457

tints

- All colorant name 203
- CS operator 217
- DeviceN color spaces 206
- in halftones 383
- and overprint mode 214, 215
- and overprint parameter 214
- Separation color spaces 201, 202, 203
- for spot color components 457
- subtractive color representation 202, 457, 465
- tint transformation function 207

title bar

- document window 472
- pop-up window 491, 492

Title entry

- document information dictionary 576
- outline item dictionary 478

TJ operator 134, 298, 302, 304, 311, 312, 314, 701, 794

Tj operator 134, 294, 311, 701, 793

- character spacing 302
- and CMaps 354
- glyph positioning 297, 298
- with multiple glyphs 312
- with patterns 219, 224, 228, 232
- special graphical effects 295, 296
- tiling patterns, emulating with 219
- word spacing 302

TK entry (graphics state parameter dictionary) 159, 301, 307

TL operator 134, 302, 701

TM entry (field dictionary) 531, 552

Tm operator 134, 310, 701

- in default appearance strings 535

“to CIE” information (ICC color profile) 192, 686, 807

TOC standard structure type 628, 632

TOCI standard structure type 628

tokens, lexical 24, 25, 26, 27, 62, 89

tool bars, hiding and showing 472

tool tips

- See pop-up help systems

topmost object 467

TopSecret annotation icon 507

ToUnicode CMaps 368–372, 796

- beginbfrange and endbfrange operators in 352
- CMap format, based on 348
- for content extraction 16
- syntax 369, 372
- in Tagged PDF 620, 621
- for Type 0 fonts 353
- for Type 1 fonts 318
- for Type 3 fonts 324

ToUnicode entry 369, 620

- Type 0 font dictionary 353

- Type 1 font dictionary 318

- Type 3 font dictionary 324

ToUnicode Mapping File Tutorial (Adobe Technical Note #5411) 369, 813

TP entry (appearance characteristics dictionary) 537

TR entry

- graphics state parameter dictionary 158, 218, 381
- soft-mask dictionary 445, 446

Tr operator 134, 296, 302, 701

TR standard structure type 629, 631

- stacking direction 642

TR2 entry (graphics state parameter dictionary) 159, 218, 381

trailer, file

- See file trailer

trailer keyword 67, 560

Trans entry (page object) 90, 485

Trans object type 486

transfer functions 373, 380–382, 392

- additive colors produced by 381, 383

- for CMYK devices 380

- color inversion with 381

- current. See current transfer function

- for DeviceGray color space 382

- for halftone screens 393, 395, 398, 400, 401

- halftones, applied before 380, 382

- for soft masks 440, 441, 445, 446

- and transparency 467–468

TransferFunction entry

- type 1 halftone dictionary 393

- type 6 halftone dictionary 395

- type 10 halftone dictionary 398

- type 16 halftone dictionary 400

transformation matrices 137, 141, 144–147

- CIE-based color space 183, 184, 185, 186, 188

- reference XObject 288

- specification 142, 145

- type 1 (function-based) shading 237

See also

- current transformation matrix (CTM)

- font matrix

- form matrix

- pattern matrix

- text line matrix

- text matrix

- text rendering matrix

transformations, coordinate

- See coordinate transformations

- transition dictionaries 90, **486–487**
 - D entry **486, 487**
 - Di entry **486, 487**
 - Dm entry **486, 487**
 - M entry **486, 487**
 - S entry **486**
 - Type entry **486**
- transition duration 486, 487
- transition style **486**
 - Blinds **486, 487**
 - Box **486, 487**
 - Dissolve **486**
 - Glitter **486, 487**
 - R **486**
 - Split **486, 487**
 - Wipe **486, 487**
- translation
 - images 266
 - order of transformations 143
 - text space 297, 299, 309, 311
 - transformation matrices 137, 141, **142**
- transparency
 - See transparent imaging model
- Transparency** entry (version 1.3 OPI dictionary) **696**
- transparency group attributes dictionaries **450–451**
 - CS entry 446, **450, 452**
 - I entry 450, **451**
 - K entry **451, 452**
 - S entry **450**
- Transparency** group subtype 286, 287, 449, **450, 452**
- transparency group XObjects **133, 286, 449–452**
 - as appearance streams 496
 - graphics state parameters, initialization of 149
 - soft masks, definition of 445, 446, 447
 - version compatibility 786
- transparency groups 133, 286, **411–412, 425–439**
 - alpha. See group alpha
 - as annotation appearances 451, 496
 - backdrop. See group backdrop
 - backdrop alpha 429
 - backdrop color 429
 - and black-generation functions 469
 - blend mode 411, 425, 431, 433, 434
 - blending color space. See blending color space
 - bounding box 452
 - clipping to bounding box 452
 - color. See group color
 - color space. See group color space
 - compositing computations. See compositing computations
 - compositing in 411, 412, 425, 427, 428, 440, 450, 452, 454
 - transparency groups (*continued*)
 - compositing of 411, 412, 425, 427, 428, 429, 445, 450, 451, 452, 454, 462, 463
 - constant opacity 425
 - constant shape 425
 - elements **428, 430, 431, 433, 434, 435**
 - group attributes dictionary 89
 - hierarchy 411, 425
 - immediate backdrop (group elements). See immediate backdrop
 - initial backdrop. See initial backdrop
 - mask opacity 425
 - mask shape 425
 - nested 411, 425, 428, 468–469
 - and nonstroking alpha constant 444
 - notation **429**
 - opacity. See group opacity
 - and overprinting 462–463, 466
 - page group 89, 412, **436–437, 449–450**
 - painting 451–452
 - in pattern cells 453
 - and rendering intents 468
 - rendering parameters ignored for 467
 - shape. See group shape
 - soft clipping 444
 - soft masks, deriving from 425, **445, 446**
 - group alpha **439–440, 445, 446**
 - group luminosity 412, **440–441, 445, 446**
 - spot color components, painting of 457–458
 - stack. See group stack
 - structure and nomenclature **427–428**
 - and undercolor-removal functions 469
 - See also
 - isolated groups
 - knockout groups
 - non-isolated groups
 - non-knockout groups
 - transparency group XObjects
- transparency stack 133, **410**
 - graphics objects and 441
 - group. See group stack
 - opacity 411
 - page group 412, 467
 - shape 411
 - and transparency groups 411, 425, 426
- transparent imaging model 2, 4, **11, 133, 409–470**
 - and alternate color space 204, 415
 - appearance streams and 496
 - backdrop. See backdrop
 - halftones and 467–468
 - opaque overprinting, compatibility with **460–462, 463**
 - overprinting **458–466**
 - overview **410–412**
 - patterns and **453–454**

- transparent imaging model (*continued*)
 - PostScript, compatibility with 469–470
 - rendering intents and 468–469
 - rendering parameters and 466–469
 - spot colors and 456–458
 - text, showing of 305
 - text knockout parameter 159, 301, 307–308, 463
 - transfer functions and 467–468
 - version compatibility 786–787
 - See also*
 - alpha
 - blend modes
 - blending color space
 - compositing
 - opacity
 - shape
 - soft masks
 - transparency groups
 - transparency stack
- trap network annotation dictionaries 690–691
 - AnnotStates** entry 690, 691
 - Contents** entry 691
 - FontFauxing** entry 691
 - LastModified** entry 690, 691
 - Subtype** entry 513, 691
 - Version** entry 690, 691, 693, 807
- trap network annotations 499, 513, 689–691, 807
 - appearance streams. *See* trap network appearances
 - See also*
 - trap network annotation dictionaries
- trap network appearance stream dictionaries 692
 - PCM** entry 692
 - SeparationColorNames** entry 692
 - TrapRegions** entry 692
 - TrapStyles** entry 692
- trap network appearances 689, 692–693
 - See also* trap network appearance stream dictionaries
- trap networks 689–693
 - current 690
 - modification date 691
 - regeneration 690
 - validation 690, 807
- TrapNet** annotation type 499, 513, 562, 690, 691
 - opacity inapplicable to 492
- Trapped** entry (document information dictionary) 576
- trapping 459, 513, 574, 676, 688–693
 - document status 576
 - instructions 689
 - parameters 689, 692
 - and pre-separated files 693
 - zones 689, 692
- trapping (*continued*)
 - See also*
 - trap network annotations
 - trap network appearances
 - trap networks
 - TrapRegion** objects (PJTF) 692
 - TrapRegions** entry (trap network appearance stream dictionary) 692
 - TrapStyles** entry (trap network appearance stream dictionary) 692
- trees
 - balanced 86, 813
 - of chained actions 514
 - interactive form 530
 - name. *See* name trees
 - number. *See* number trees
 - page 81, 83, 86–92, 557, 757, 765
 - structural parent 90, 269, 285, 492
 - structure. *See* structure hierarchy
- TRef** entry (FDF template dictionary) 567, 568
- triangle meshes, Gouraud-shaded 253
 - free-form. *See* type 4 shadings
 - lattice-form. *See* type 5 shadings
- trigger events 4, 85, 513, 514–517, 800
 - for annotations 492, 515
 - Bl** (annotation) 515
 - C** (form field) 516, 517
 - C** (page) 515
 - D** (annotation) 515, 517
 - DC** (document) 516
 - for documents 516
 - DP** (document) 516
 - DS** (document) 516
 - E** (annotation) 515, 517, 526
 - F** (form field) 516, 517
 - for FDF fields 565
 - Fo** (annotation) 515
 - for form fields 516, 517, 532
 - K** (form field) 516, 517
 - mouse-related 515, 517
 - O** (page) 515
 - for pages 515
 - and pop-up help systems 526
 - U** (annotation) 515, 517
 - V** (form field) 516, 517
 - WP** (document) 516
 - WS** (document) 516
 - X** (annotation) 515, 517, 526
- trim box 677
 - clipping to 679
 - display of 5, 681
 - in page object 89
 - page placement, ignored in 679

- trim box (*continued*)
 - for page positioning 679
 - printer's marks excluded from 680
 - in printing 679
- TrimBox** entry
 - box color information dictionary 681
 - page object 89, 677, 806
- Trinitron® display, Sony 186
- tristimulus values 181, 183, 185, 186, 188
- true** (boolean object) 28
- true** operator (PostScript) 116, 704
 - TrueType 1.0 Font Files Technical Specification* (Microsoft Corporation) 321, 360, 816
- TrueType font dictionaries 321–322, 367
 - BaseFont** entry 321–322
 - Encoding** entry 321, 332, 333, 334
 - Subtype** entry 321
- TrueType font programs
 - “cmap” table 332, 333–334, 339, 340, 367
 - “cvt_” table 367
 - embedded 16, 357, 365, 366, 367–368
 - “fpgm” table 367
 - “glyf” table 367
 - “head” table 367
 - “hhea” table 367
 - “hmtx” table 367
 - “loca” table 367
 - “maxp” table 367
 - “name” table 321
 - “OS/2” table 360
 - “post” table 332, 334
 - “prep” table 367
 - sFamilyClass field 360
 - “vhea” table 368
 - “vmtx” table 368
- TrueType** font type 36, 315, 321, 365
- TrueType® fonts 5, 321–322
 - built-in encoding 332
 - character encodings 332–334, 367, 710
 - font descriptors for 357
 - format 292, 321
 - glyph descriptions 332, 333, 334, 339, 345, 367
 - glyph indices 339, 345
 - in Linearized PDF 738
 - PostScript name 321–322
 - in PostScript XObjects 290
 - subsets 322
 - synthesized styles 321
 - and Type 2 CIDFonts 337, 339, 345, 353
 - vertical metrics 368
 - See also*
 - TrueType font dictionaries
 - TrueType font programs
- TrueType Reference Manual* (Apple Computer, Inc.) 321, 365, 813
- truncate** operator (PostScript) 116, 703
- TS** entry
 - source information dictionary 670, 671
 - Web Capture content set 668
- Ts** operator 134, 302, 701
- TU** entry (field dictionary) 531, 657
- Tw** operator 134, 302, 701
- TwoColumnLeft page layout 84
- TwoColumnRight page layout 84
- Tx** field type 531
- Type 0 CIDFont programs
 - compact 365, 366, 367
- Type 0 CIDFonts 337
 - glyph selection 339
 - PostScript name 338, 353
 - See also*
 - Type 0 CIDFont programs
- Type 0 font dictionaries 334, 342, 352–353
 - BaseFont** entry 353
 - DescendantFonts** entry 336, 353, 354
 - Encoding** entry 336, 353, 354
 - Subtype** entry 334, 353
 - ToUnicode** entry 353
 - Type** entry 353
- Type 0 fonts 314, 334–355
 - character identification 369
 - CID-keyed fonts as 336
 - CIDFonts 314
 - CIDFonts as descendants of 338, 353, 355, 369
 - CMap mapping 354, 796
 - descendant fonts 334, 353
 - font descriptors lacking in 356
 - glyph selection 354–355
 - Identity–H predefined CMap 345
 - Identity–V predefined CMap 345
 - parent fonts 334
 - root font 334
 - undefined characters 355
 - See also*
 - Type 0 font dictionaries
- type 0 function dictionaries (sampled) 110
 - BitsPerSample** entry 110, 111
 - Decode** entry 109, 110, 112
 - Encode** entry 109, 110
 - Order** entry 110, 113, 792
 - Size** entry 109, 110, 111, 113
- type 0 functions (sampled) 107, 108, 109–113, 792
 - clipping to domain 110
 - clipping to range 111
 - clipping to sample table 110

- type 0 functions (sampled) (*continued*)
 - decoding of output values 110
 - dimensionality 109
 - encoding of input values 110
 - interpolation 110
 - sample data 109, 110, 111
 - and smoothness tolerance 405
 - See also*
 - type 0 function dictionaries
- Type 1 font dictionaries 317–318, 320, 321, 794
 - BaseFont** entry 34, 290, 317
 - Encoding** entry 318, 332
 - FirstChar** entry 317, 318, 319, 795
 - FontDescriptor** entry 317, 318, 319, 795
 - LastChar** entry 317, 318, 319, 795
 - Name** entry 317, 794
 - Subtype** entry 317
 - ToUnicode** entry 318
 - Type** entry 317
 - Widths** entry 317, 318, 319, 795
- Type 1 Font Format Supplement* (Adobe Technical Note #5015) 5, 320, 811, 812
- Type 1 font programs
 - clear-text portion 366
 - compact 365, 367
 - embedded 16, 357, 365, 366, 794
 - encrypted portion 366
 - fixed-content portion 366
 - PaintType** entry 367
- Type 1 fonts 5, 16, 316–321, 794–795
 - built-in encoding 318, 332, 710
 - character encodings 331–332, 710
 - character identification 368
 - compact 366
 - font descriptors for 357
 - FontName** entry 317
 - format 292
 - glyph descriptions 316
 - glyph widths 317, 794–795
 - hints 316
 - in Linearized PDF 738
 - multiple master 319–321, 796
 - in PostScript files 22
 - in PostScript XObjects 290
 - standard. *See* standard 14 fonts
 - subsets 322
 - and Type 0 CIDFonts 337
 - and Type 3 fonts, compared 323
 - See also*
 - Type 1 font dictionaries
 - Type 1 font programs
- type 1 form dictionaries 284–285
 - Group** entry 285, 286, 288, 449, 452
 - LastModified** entry 284
 - Matrix** entry 283, 284, 288, 446, 496
 - Metadata** entry 285
 - Name** entry 284, 793
 - OPI** entry 285, 693
 - PieceInfo** entry 284, 285, 581
 - Ref** entry 285, 287
 - Resources** entry 284, 534, 535
 - StructParent** entry 285, 601
 - StructParents** entry 285, 601
 - Subtype** entry 284
 - Type** entry 284
- type 1 halftone dictionaries 393
 - AccurateScreens** entry 393, 394
 - Angle** entry 393
 - Frequency** entry 393
 - HalftoneName** entry 393
 - HalftoneType** entry 393
 - SpotFunction** entry 393
 - TransferFunction** entry 393
 - Type** entry 393
- type 1 halftones 391, 392–394
 - See also* type 1 halftone dictionaries
- type 1 pattern dictionaries (tiling) 221–222
 - BBox** entry 222, 223
 - Matrix** entry 222, 223, 283
 - PaintType** entry 221, 454
 - PatternType** entry 221
 - Resources** entry 222
 - TilingType** entry 222
 - Type** entry 221
 - XStep** entry 222, 223
 - YStep** entry 222, 223
- type 1 patterns (tiling)
 - See* tiling patterns
- type 1 shading dictionaries (function-based) 237
 - Domain** entry 237
 - Function** entry 237
 - Matrix** entry 237
- type 1 shadings (function-based) 233, 237–238
 - coordinate system 237
 - See also*
 - type 1 shading dictionaries
- Type 2 Charstring Format, The* (Adobe Technical Note #5177) 5, 813
- Type 2 CIDFonts 337
 - encoding 353
 - glyph selection 339–340
 - and Identity–H predefined CMap 345
 - and Identity–V predefined CMap 345
 - PostScript name 338, 353

- type 2 function dictionaries (exponential interpolation)
 - 113
 - C0** entry 113
 - C1** entry 113
 - N** entry 113
- type 2 functions (exponential interpolation) 107, 108, 113, 792
 - See also* type 2 function dictionaries
- type 2 pattern dictionaries (shading) 231, 232, 233
 - ExtGState** entry 231, 453
 - Matrix** entry 231, 283
 - PatternType** entry 231
 - Shading** entry 231
 - Type** entry 231
- type 2 patterns (shading)
 - See* shading patterns
- type 2 shading dictionaries (axial) 238
 - Coords** entry 238
 - Domain** entry 238
 - Extend** entry 238
 - Function** entry 238
- type 2 shadings (axial) 233, 236, 238–239
 - parametric variable 238, 239
 - See also*
 - type 2 shading dictionaries
- Type 3 font dictionaries 314, 323–324
 - CharProcs** entry 324, 325, 326, 332
 - Encoding** entry 324, 332, 796
 - FirstChar** entry 324
 - FontBBox** entry 323
 - FontMatrix** entry 298, 323, 324, 325
 - LastChar** entry 324
 - Name** entry 323
 - Resources** entry 324
 - Subtype** entry 323
 - ToUnicode** entry 324
 - Type** entry 323
 - Widths** entry 324, 326
- Type 3 font operators 134, 325–326
 - d0** 134, 324, 325, 326, 700
 - d1** 134, 218, 324, 325, 326, 700
- Type 3 fonts 96, 140, 314, 323–328, 796
 - bounding box 323, 325, 326, 700
 - character encodings 328, 332
 - font descriptors lacking in 356, 316
 - font matrix 298
 - glyph descriptions 323, 324, 325, 789
 - glyph widths 324, 325, 326, 700
 - hints unavailable in 323
 - metrics 298
 - PostScript and PDF, compared 325
 - resource dictionary 324, 796
 - in text objects 309
 - Type 3 fonts (*continued*)
 - text rendering mode, unaffected by 296, 305
 - and Type 1 fonts, compared 323
 - See also*
 - Type 3 font dictionaries
 - Type 3 font operators
- type 3 function dictionaries (stitching) 114
 - Bounds** entry 114
 - Encode** entry 114
 - Functions** entry 114
- type 3 functions (stitching) 108, 113–115, 792
 - for inverting function domains 115
 - See also*
 - type 3 function dictionaries
- type 3 shading dictionaries (radial) 240
 - Coords** entry 240
 - Domain** entry 240
 - Extend** entry 240
 - Function** entry 240
- type 3 shadings (radial) 233, 236, 239–243
 - blend circles 240–242
 - parametric variable 240
 - See also*
 - type 3 shading dictionaries
- type 4 functions (PostScript calculator) 3, 108, 115–118, 792
 - error detection and reporting 117–118
 - language limitations 115–116
 - null operands and results 28
 - operand stack 117
 - operand syntax 116
 - operators 3, 116, 703–704
 - predefined spot functions, definitions of 385–389
 - as spot functions 797
 - as transfer functions 381
- type 4 shading dictionaries (free-form Gouraud-shaded triangle mesh) 244
 - BitsPerComponent** entry 244
 - BitsPerCoordinate** entry 244
 - BitsPerFlag** entry 244
 - Decode** entry 244
 - Function** entry 244
- type 4 shadings (free-form Gouraud-shaded triangle meshes) 233, 243–247
 - data format 244–247
 - edge flags 244, 245–247
 - parametric variable 244, 245, 247
 - See also*
 - type 4 shading dictionaries
- type 5 halftone dictionaries 393, 395, 398, 400–401
 - Default** entry 401
 - HalftoneName** entry 401
 - HalftoneType** entry 401

- type 5 halftone dictionaries (*continued*)
 - keys 392
 - Type** entry 401
- type 5 halftones 391, 400–403
 - default halftone 401
 - transfer functions required for components 393, 395, 398
 - type 5 halftones, prohibited as components of 401
 - See also*
 - type 5 halftone dictionaries
- type 5 shading dictionaries (lattice-form Gouraud-shaded triangle mesh) 249
 - BitsPerComponent** entry 249
 - BitsPerCoordinate** entry 249
 - Decode** entry 249
 - Function** entry 249
 - VerticesPerRow** entry 249
- type 5 shadings (lattice-form Gouraud-shaded triangle meshes) 233, 248–250
 - data format 248, 249–250
 - parametric variable 249
 - See also*
 - type 5 shading dictionaries
- type 6 halftone dictionaries 394, 395
 - HalftoneName** entry 395
 - HalftoneType** entry 395
 - Height** entry 394, 395, 398
 - TransferFunction** entry 395
 - Type** entry 395
 - Width** entry 394, 395, 398
- type 6 halftones 391, 394–395
 - and type 10 halftones, compared 394, 398
 - See also*
 - type 6 halftone dictionaries
- type 6 shading dictionaries (Coons patch mesh) 253, 256
 - BitsPerComponent** entry 253
 - BitsPerCoordinate** entry 253
 - BitsPerFlag** entry 253
 - Decode** entry 253
 - Function** entry 253
- type 6 shadings (Coons patch meshes) 233, 250–256
 - data format 253–256
 - edge flags 253, 254–256
 - parametric variables 250, 251, 253, 255
 - See also*
 - type 6 shading dictionaries
- type 7 shading dictionaries (tensor-product patch mesh) 256
- type 7 shadings (tensor-product patch meshes) 233, 256–260
 - data format 259–260
 - edge flags 259–260
 - parametric variables 258
- type 7 shadings (tensor-product patch meshes) (*continued*)
 - See also*
 - type 7 shading dictionaries
- type 10 halftone dictionaries 398
 - HalftoneName** entry 398
 - HalftoneType** entry 398
 - TransferFunction** entry 398
 - Type** entry 398
 - Xsquare** entry 398
 - Ysquare** entry 398
- type 10 halftones 391, 394–398
 - and type 6 halftones, compared 394, 398
 - and type 16 halftones, compared 399
 - See also*
 - type 10 halftone dictionaries
- type 16 halftone dictionaries 399, 400
 - HalftoneName** entry 400
 - HalftoneType** entry 400
 - Height** entry 399, 400
 - Height2** entry 399, 400
 - TransferFunction** entry 400
 - Type** entry 400
 - Width** entry 399, 400
 - Width2** entry 399, 400
- type 16 halftones 391, 399–400
 - and type 10 halftones, compared 399
 - See also*
 - type 16 halftone dictionaries
- Type 42 fonts 22
 - inapplicable to PDF 321
- Type** entry 35–36
 - action dictionary 514
 - annotation dictionary 490
 - bead dictionary 484
 - border style dictionary 495
 - CIDFont dictionary 338
 - CMap dictionary 349
 - document catalog 83
 - embedded file stream dictionary 124
 - encoding dictionary 330
 - external object (XObject) 261
 - file specification dictionary 122, 123, 128
 - font descriptor 356
 - font dictionary 36
 - graphics state parameter dictionary 157
 - group attributes dictionary 287
 - halftone dictionary 392
 - image dictionary 267, 279, 448
 - marked-content reference dictionary 594
 - metadata stream dictionary 578
 - object reference dictionary 599
 - outline dictionary 478
 - page label dictionary 483
 - page object 88

Type entry (*continued*)

- page tree node **86**
- PDF/X output intent dictionary **685**
- PostScriptXObject dictionary **290**
- property list (Tagged PDF artifact) **616**
- signature dictionary **549**
- soft-mask dictionary **446**
- sound object **569**
- stream dictionary 261
- structure element dictionary **591**
- structure tree root **590**
- thread dictionary **484**
- transition dictionary **486**
- Type 0 font dictionary **353**
- Type 1 font dictionary **317**
- type 1 form dictionary **284**
- type 1 halftone dictionary **393**
- type 1 pattern dictionary **221**
- type 2 pattern dictionary **231**
- Type 3 font dictionary **323**
- type 5 halftone dictionary **401**
- type 6 halftone dictionary **395**
- type 10 halftone dictionary **398**
- type 16 halftone dictionary **400**
- version 1.3 OPI dictionary **694**
- version 2.0 OPI dictionary **697**
- Web Capture content set **668**
- Type Technology Forum (Adobe Systems Incorporated) 812
- Type0** font type **315, 334, 353**
- Type1** font type **36, 315, 317, 365**
- Type1C** compact font subtype **365, 366**
- Type3** font type **315, 323**
- typefaces
 - Adobe Garamond™ 318
 - Courier 16, 709
 - Helvetica* 16, 292, 293, 294, 709, 760
 - ITC Zapf Dingbats® 16, 224, 709
 - New York 321
 - Poetica® 323
 - Symbol 16
 - Times* 16, 292, 709
- types, object
 - See object types
- Tz** operator 134, **302, 701**

U

- U border style (underline) **495**
- U entry
 - additional-actions dictionary **515**
 - encryption dictionary **76, 79–80**
 - URL alias dictionary **671, 672**

- U trigger event (annotation) **515, 517**
- UCR entry (graphics state parameter dictionary) **158, 218, 379**
- UCR2 entry (graphics state parameter dictionary) **158, 218, 379**
- UCS-2 character encoding 343, 344, 345
- UHC (Unified Hangul Code) character encoding 345, 803
- unbalanced parentheses 29, 30, 31
- uncolored tiling patterns 217, **221–222, 227–230**
 - color value for painting 227
 - content stream 218
 - in transparent imaging model 454
- undefined characters (Type 0 fonts) 355
- undercolor-removal function **150, 378, 379, 380**
 - and transparency 466, 468, 469
 - UCR entry (graphics state parameter dictionary) 158
 - UCR2 entry (graphics state parameter dictionary) 158
- underline annotation dictionaries
 - See markup annotation dictionaries
- Underline** annotation type 499, **506**
- underline annotations
 - See markup annotations
- Underline text decoration type **647**
- underlying color space (**Pattern** color space) 195, 217, **227, 456**
- underscore (_) character
 - in file specifications 121
 - in multiple master font names 320
- undoing changes 18
- UniCNS–UCS2–H predefined CMap **344, 346**
- UniCNS–UCS2–V predefined CMap **344, 346**
- Unicode and Glyph Names* (Adobe Systems Incorporated) **811**
- Unicode® character encoding 16
 - for alternate descriptions 657
 - byte order marker 98
 - CMaps 336
 - for field names 801
 - for field values 4, 540, 542, 562
 - hard hyphen character 617
 - JavaScript 1.2 incompatible with 802
 - for JavaScript scripts 556, 802
 - list numbering 650
 - Microsoft Unicode 333
 - natural language escape 99–100, 652, 657, 658, 659
 - soft hyphen character 617, 621, 714
 - Tagged PDF, mapping from 613, 614, 620–621
 - for text strings 98–100, 710
 - TrueType character names, mapping to 333
 - UCS-2 343, 344, 345
 - UTF-8 34, 788
 - UTF-16 98

- Unicode[®] character encoding (*continued*)
 - vendor space 620
 - word breaks 623
 - See also*
 - ToUnicode** CMaps
- Unicode Consortium 816
 - Bidirectional Algorithm, The* (Standard Annex #9) 642
 - Line Breaking Properties* (Standard Annex #14) 623
 - Unicode Standard, The* 98, 368, 620
- Unicode Standard, The* (Unicode Consortium) 98, 368, 620, **816**
- uniform resource identifiers (URIs) **523**
 - for production condition registry 686
- Uniform Resource Locators* (Internet RFC 1738) 119, 127, 664, **815**
- uniform resource locators (URLs)
 - file specifications 119, 122, 127, 551
 - and Linearized PDF 726, 727
 - redirection 671–672
 - in submit-form actions 550, 802
 - “unsafe” characters in 664
 - Web Capture 93, 660, 661, 662, **664**, 670, 672, 673, 674, 806
 - See also* URL alias dictionaries
- UniGB–UCS2–H predefined CMap **344**, 346
- UniGB–UCS2–V predefined CMap **344**, 346
- UniJIS–UCS2–H predefined CMap **345**, 347
- UniJIS–UCS2–HW–H predefined CMap **345**, 347
- UniJIS–UCS2–HW–V predefined CMap **345**, 347
- UniJIS–UCS2–V predefined CMap **345**, 347
- UniKS–UCS2–H predefined CMap **345**, 347
- UniKS–UCS2–V predefined CMap **345**, 347
- Union function **423**, 424, 426, 430, 432, 435, 438
- Unisys Corporation 48
- Universal Time (UT) 100
- Unix** entry
 - file specification dictionary 122, **123**
 - launch action dictionary 520, **521**
- UNIX[®] operating system
 - Acrobat “Print As Image” feature unavailable in 790
 - application launch parameters 520, 521
 - conversion from PostScript to PDF 20
 - file names 121, 558
 - file system 122
- Up predictor function (LZW and Flate encoding) **50**, 51
- updates, incremental
 - See* incremental updates
- UpperAlpha list numbering style **650**
- UpperRoman list numbering style **650**
- URI action dictionaries **523**
 - IsMap** entry **523**
 - S** entry **523**
 - URI** entry **523**
- URI** action type 518, **523**
- URI actions 85, 518, **523–524**, 800
 - for annotations 523
 - base URI **524**
 - and go-to actions 501
 - and link annotations 501
 - OpenAction** entry (document catalog), ignored for **523**
 - outline items, ignored for **523**
 - See also*
 - URI action dictionaries
 - URI dictionaries
- URI dictionaries 85, **524**
 - Base** entry **524**
- URI** entry
 - document catalog **85**, 524
 - URI action dictionary **523**
- URIs. *See* uniform resource identifiers
- URL alias dictionaries 670, **671–672**, 806
 - C** entry **671**
 - implementation limits 708
 - U** entry **671**, 672
- URL** entry (Web Capture command dictionary) **672**
- URL file specifications 119, **127**, 551
- URL** file system **122**, 127
- URLs. *See* uniform resource locators
- URLS** entry (name dictionary) **93**, 661, 662, 663, 664, 668
- UseCMap** entry (CMap dictionary) **349**, **351**, 369
- usecmap** operator (PostScript) **351**
- usefont** operator (PostScript) **352**
- UseNone page mode **84**, **472**
- UseOutlines page mode **84**, **472**, 730, 734, 737
- user interface
 - controller bars (movies) 572
 - field names 531, 657
 - icons 488, 491, 499, 500, 507, 509, 510, 799
 - menu bar 472
 - menu items 530, 801
 - navigation controls 472
 - Print Setup dialog 806
 - pushbuttons 530
 - scroll bars 472
 - tool bars 472
 - windows
 - See*
 - document windows
 - floating windows
 - pop-up windows

user interface (*continued*)*See also*

- actions
 - annotations
 - document outline
 - interactive forms
 - mouse
 - movies
 - page mode
 - presentations
 - sounds
 - thumbnail images
- user password 74, 76, 78
 computing (algorithm) 79–80
- user space 138–139
 current transformation matrix (CTM) 131, 141, 148
 default. *See* default user space
 form space, mapping from 283, 284
 glyph space, mapping from 325
 glyphs in 294
 and graphics state parameters 305
 image space, relationship with 263, 266
 rotation 494
 scaling 494
 and **sh** operator 232
 shadings, target coordinate space for 233
 soft-mask images in 448
 text position in 294
 text space, relationship with 309
 URI actions, mouse position for 523
- UseThumbs page mode 84, 472
- UT (Universal Time) 100
- UTF-8 character encoding 34, 788
- UTF-16 character encoding 98

V**V** entry

- additional-actions dictionary 516
 - bead dictionary 484
 - encryption dictionary 71, 72, 73, 76, 79, 790
 - FDF field dictionary 544, 562, 564, 803
 - field dictionary 531, 532, 533, 539, 540, 541, 542, 543, 544, 546, 547, 548, 551, 554
 - file specification dictionary 123
 - hint stream dictionary 736
 - Web Capture information dictionary 660
- v** operator 134, 163, 165, 166, 701
- V** predefined CMap 345, 347
- V** trigger event (form field) 516, 517

values

- in dictionaries 35
 - in name trees 101, 102
 - in number trees 106
- variable-pitch fonts 297, 358
- variable text (form fields) 533–537, 801
 default font 534
 default resource dictionary 534, 535, 801
 See also
 default appearance strings
 dynamic appearance streams
- vendor space (Unicode) 620
- VeriSign.PPKVS** signature handler 549
- version compatibility, PDF 1–2, 3, 783–809
 border styles (annotations) 491
 compatibility sections 95
 default color spaces 194
 document outline 479
 extensibility 18
 go-to actions 519
 logical structure 479
 named destinations 477
 procedure sets 574
 version specification 63
- Version** entry
 document catalog 63, 70, 83, 784, 785, 791
 FDF catalog 559, 561, 802
 trap network annotation dictionary 690, 691, 693, 807
 version 1.3 OPI dictionary 694
 version 2.0 OPI dictionary 697
- versions, PDF 2
 character collections 346–347, 369
 CID-keyed fonts 336, 342
 CIDSystemInfo dictionaries 349
 CMap operators 352
 color spaces 178, 181, 190, 194, 374
 compatibility. *See* version compatibility, PDF
 composite fonts 334, 353
 and FDF files 558, 559, 561, 802
 font numbers 352
 form XObjects, resources for 284–285
 ICC profiles 190–191
 imaging models 11, 133
 linearization independent of 731
 masked images 275–276
 PostScript XObjects 290
 specification 4, 61, 63, 70–71, 83, 784
 TrueType font encodings 333
 version numbers 783–786
 major 783, 784
 minor 783, 784, 785

- vertical writing
 - character spacing 303
 - in CIDFonts 338, 341–342, 362, 363
 - glyph displacement 313
 - R2L reading order 472
 - word spacing 303
 - writing mode 1 299–300
- VerticesPerRow** entry (type 5 shading dictionary) 249
- “vhea” table (TrueType font) 368
- ViewArea** entry (viewer preferences dictionary) 473
- ViewClip** entry (viewer preferences dictionary) 473
- viewer applications, PDF 2, 316
 - alternate color space, handling of 190, 203, 206
 - annotation handlers, standard 489
 - annotation icons, predefined appearances for 500, 507, 509, 510
 - annotations, scaling and rotation of 494
 - application appearances, rendering of 496
 - articles, navigation facilities for 483
 - blend modes 442
 - chained actions, execution of 514
 - character encodings, standard 709
 - character sets, standard 709
 - characters, identification of 368–369
 - color conversion 376
 - color mapping function 375
 - color separations, previewing of 684
 - compatibility, cross-platform 20, 21
 - compatibility, version 2, 3, 18, 63, 783–809
 - content extraction 16, 368
 - content streams, processing of 11
 - cross-reference table, processing of 69
 - date strings 490
 - device profiles 375
 - encrypted documents, handling of 74, 75
 - file content, processing of 17
 - file identifiers, use of 123, 288
 - file systems, platform-specific 122
 - font management 16, 291, 292, 315, 317, 320, 794
 - font substitution 356, 358
 - form fields, variable text in 533, 534, 535, 536
 - form XObjects, caching of 281, 452
 - gamut mapping function 375
 - glyph selection, TrueType fonts 332, 333, 340
 - glyphs, positioning of 297
 - graphics state, maintenance of 147
 - ICC profile rendering intents ignored 192
 - image data, handling of 264
 - images, rendering of 263
 - implementation limits 705–707
 - implicit color conversion 196
 - Indexed** color spaces, use of 199
 - launch actions, response to 521
 - viewer applications, PDF (*continued*)
 - Linearized PDF, processing of 725, 728, 738, 739, 740, 751, 752, 753, 754, 755
 - metadata, use of 575
 - mouse, responding to 517
 - movies, asynchronous playing of 572
 - named actions 527
 - named pages, handling of 557
 - numbers, range and precision of 28
 - outline items, responding to 477
 - page boundaries, display of 89, 679–680
 - page tree, handling of 86
 - passwords, handling of 543
 - predefined character sets and encodings 3
 - predefined CMaps, support for 347
 - presentations 90, 485
 - private data ignored by 581
 - procedure sets, compatibility with 574
 - reference XObjects, handling of 287
 - reference XObjects, printing of 288
 - remote go-to actions, response to 520
 - rendering intents, handling of 197
 - resolution of output device, adjusting for 139
 - scan conversion 406
 - shading patterns, interpolation of color values in 235
 - signatures, digital 530
 - Sort choice field flag ignored by 546
 - sound formats 570
 - sounds, synchronous playing of 525
 - standard fonts 16, 319
 - standard security handler 71
 - text annotations, font and size for 499
 - thumbnail images, display of 480
 - tint transformation functions, use of 204
 - transparency, representing in PostScript 469–470
 - transparent objects, rasterization of 410
 - TrueType fonts, treatment of 368
 - Type 3 fonts, showing of text with 324, 325
 - unknown annotation types, handling of 498
 - user interface 472–473
 - volatile files, handling of 123
- viewer preferences 4, 471–473, 798
- viewer preferences dictionary 84, 471–473
 - CenterWindow** entry 472
 - Direction** entry 472
 - DisplayDocTitle** entry 472
 - FitWindow** entry 472
 - HideMenubar** entry 472
 - HideToolbar** entry 472
 - HideWindowUI** entry 472
 - NonFullScreenPageMode** entry 472
 - PageLayout** entry erroneously documented in 791, 798
 - PrintArea** entry 473
 - PrintClip** entry 473
 - ViewArea** entry 473
 - ViewClip** entry 473

ViewerPreferences entry (document catalog) **84**, 471, 734

viewing of documents 1, 17
 document window, size and positioning of 472
 embedded fonts, copyright restrictions on 365
 glyph widths in 794
 output medium, dialog with user on 374
 page boundaries 473
 page mode **84**, **472**
 version compatibility 783

vignettes 421, 439

“vmtx” table (TrueType font) 368

Volume entry

movie activation dictionary 572
 sound action dictionary 525

W

W entry

border style dictionary 495
 box style dictionary **681**
 CIDFont dictionary **338**, **340–341**
 inline image object **279**

W operator 134, 162, 169, 171, **172**, 585, 586, 702

w operator 134, 151, **156**, 296, 702

W* operator 134, 162, 169, 171, **172**, 585, 586, 702

W2 entry (CIDFont dictionary) **338**, 341–342, 368

Warnock, John xx

Web Capture command dictionaries 660, 670, **672–674**

CT entry **672**, 673–674

F entry **672**

H entry **672**, 674

implementation limits 708

L entry **672**

P entry **672**, 673

S entry **672**, 674

URL entry **672**

Web Capture command flags **672–673**

SamePath **673**

SameSite **673**

Submit **673**

Web Capture command settings dictionaries 672, **674–675**

C entry **674**, 708

G entry **674**

implementation limits 708

Web Capture content database **660**

metadata inapplicable to 579

Web Capture content sets **661–663**, **667–669**

creation date 668, 671

CT entry **668**

digital identifier 90, 269, 675

expiration date 670, 671

Web Capture content sets (*continued*)

ID entry **668**

implementation limits 708

modification date 670, 671

in name dictionary 93

O entry 667, **668**, 669, 675, 805

S entry **668**

SI entry **668**, 669, 670

source information **670**, **671**

subtype 668, 669

TS entry **668**

Type entry **668**

URLs for 664

See also

Web Capture image sets

Web Capture page sets

Web Capture image sets 661, 662, 663, **669**, 805

digital identifier **665**

R entry **669**

reference counts 669, 805

S entry **669**

Web Capture information dictionary 85, **660**

C entry **660**, 708

implementation limits 708

V entry **660**

Web Capture page sets 661, 662, 663, **667–668**

digital identifier **665**

form submission type 670

S entry **668**

T entry **668**

text identifier **665**, 667, 668

TID entry 667, **668**

title 668

Web Capture plug-in extension (AcroSpider) 85, 573, **659–676**

commands

See

Web Capture command dictionaries

Web Capture command settings dictionaries

content database **661–663**

content sets. *See* Web Capture content sets

digital identifiers 660, 661, **664–665**, 668, 670, 805

implementation limits 660, **708**

information dictionary. *See* Web Capture information dictionary

and link annotations 501

object attributes related to 675–676

source information. *See* source information dictionaries

unique name generation **665–667**

URLs (uniform resource locators) 93, 660, 661, 662,

664, 670, 672, 673, 674, 806

See also URL alias dictionaries

version number 660

- Web Content Accessibility Guidelines* (World Wide Web Consortium) 652, **816**
- WebLink plug-in extension 800
- Western writing systems
 - character encodings 710
 - glyph displacements 299, 619
 - page content order 618
 - progression directions 624, 625, 632, 642
 - shift direction 647
 - word breaks 623
 - writing mode 642
- white point
 - diffuse 183, 184, 185, 188
 - of output medium 198
- white-preserving blend modes **460**
 - in isolated groups 460
 - for spot colors 460
- white-space characters 24, **25–26**, 808
 - ASCII85Decode** filter, ignored by 44
 - ASCIIHexDecode** filter, ignored by 44
 - in hexadecimal strings 32
 - in inline images 279
 - in name objects 32, 33
- WhitePoint** entry
 - CalGray** color space dictionary **183**, 184
 - CalRGB** color space dictionary **185**, 186
 - Lab** color space dictionary **188**
- widget annotation dictionaries **512**
 - Contents** entry **512**
 - and dynamic appearance streams 534
 - FDF field dictionaries, compared with 564
 - field dictionaries, merged with 511, 528, 548
 - H** entry **512**
 - MK** entry **512**, 536
 - Subtype** entry **512**
- Widget** annotation type 499, **512**, 562
 - opacity inapplicable to 492
- widget annotations 499, **511–512**, 528
 - additional-actions dictionary 492
 - appearance dictionaries for 529, 539
 - appearance streams for 529, 534, 539
 - border style 491, 496
 - contents 512
 - for FDF fields 565, 566
 - and Form standard structure type 637
 - and hide actions 527
 - highlighting mode **512**
 - icon 565, 566
 - for radio button fields 541, 542, 543
 - and ReadOnly field flag 493, 532
 - rotation 536
 - scaling 566
 - and submit-form actions 552, 554
- widget annotations (*continued*)
 - trigger events for 515
 - See also*
 - fields, interactive form
 - widget annotation dictionaries
- Width** entry
 - image dictionary **267**, 277, 448, 480
 - inline image object **279**
 - type 6 halftone dictionary 394, **395**, 398
 - type 16 halftone dictionary 399, **400**
- Width** standard structure attribute 637, 640, **645**, 648, 649
- Width2** entry (type 16 halftone dictionary) 399, **400**
- Widths** entry
 - font dictionary 357
 - Type 1 font dictionary **317**, 318, 319, 795
 - Type 3 font dictionary **324**, 326
- Win** entry (launch action dictionary) 520, **521**, 800
- WinAnsiEncoding** predefined character encoding **329**, 709, **710**
 - as base encoding 330
 - euro character 714
 - for TrueType fonts 333
 - for Type 1 fonts 318
 - and Unicode mapping 368, 620
- windows
 - See*
 - document windows
 - floating windows
 - pop-up windows
- Windows launch parameter dictionaries **521**, 800
 - D** entry **521**
 - F** entry **521**
 - O** entry **521**
 - P** entry **521**
- Windows[®] operating system
 - application launch parameters 520, 521
 - character encoding 328, 329, 710
 - Code Page 932 344
 - Code Page 936 343
 - Code Page 949 345
 - Code Page 950 344
 - Code Page 1252 710
 - directories 521, 802
 - file names 121, 521, 558
 - file system 122
 - font names 321
 - Graphics Device Interface (GDI) 19
 - LOGFONT structure 321
 - PATH environment variable 802
 - PDF Writer printer driver 20
 - ShellExecute function 800
 - TrueType[®] font format 321, 332
 - WNetGetConnection function 120

Wipe transition style **486**, 487

WMode entry (CMap dictionary) 341, **349**

WNetGetConnection function (Windows) 120

word spacing (T_w) parameter 298, 301, **303**
 and horizontal scaling 304
 and quotation mark (") operator 311
 text matrix, updating of 313–314
Tw operator 302, 701, 702

workflow 10, 195, 676, 685, 686, 693, 806

World Wide Web
 accessibility guidelines for 652
 document distribution on 577, 684
 form submission 550
 Linearized PDF and 726
 PDF specification available on xix
See also
 Web Capture plug-in extension

World Wide Web Consortium 50, 816
Extensible Markup Language (XML) 1.0 554, 653
Extensible Stylesheet Language (XSL) 1.0 622
Web Content Accessibility Guidelines 652

WP entry (additional-actions dictionary) **516**

WP trigger event (document) **516**

writing mode (fonts) **299–300**, 340, 341
 and character spacing 302
 CMaps, specified by 343, 349
 horizontal scaling independent of 304
 leading independent of 305
 simple fonts, horizontal only in 316
 text matrix, adjustment of 313
 text rise independent of 307
 and **TJ** operator 311
 and word spacing 303

writing mode (Tagged PDF) **642**
LrTb **642**
RIT**b** **642**
TbRl **642**

writing systems
 Arabic 619, 642
 Asian 299, 335
 Chinese 642
 Hebrew 619, 642
 Japanese 642
 Latin 329, 359
 non-Latin 329, 540
 progression directions 624
 right-to-left 619
 Western. *See* Western writing systems

WritingMode standard structure attribute 624, 640, **642**,
 647, 651

WS entry (additional-actions dictionary) **516**

WS trigger event (document) **516**

X

X entry (additional-actions dictionary) **515**, 517

X trigger event (annotation) **515**, 517, 526

XHeight entry (font descriptor) **357**

XML (Extensible Markup Language)
 file-select controls, not supported for 544
 language identifiers 653
 for metadata streams 577–578
 PDF logical structure compared with 589
 standard attribute owner 639
 strongly structured document organization 632
 in submit-form actions 4, 552, 554
 Tagged PDF, conversion from 613, 627

XML field flag (submit-form field) **552**, 553, 554

XML metadata subtype **578**

XML-1.00 standard attribute owner **639**

XMP™ (Extensible Metadata Platform) framework 578

XMP: Extensible Metadata Platform (Adobe Systems
 Incorporated) 578, **812**

XObject entry (resource dictionary) **97**, 261, 266, 269, 282,
 284

XObject object type 261, **267**, **284**, **290**, 448

XObject operator **134**, **261**
Do 134, 220, 224, 228, 232, **261**, 262, 267, 282, 283,
 290, 451, 452, 467, 468, 469, 586, 595, 596, 599, 700

XObject resource type **97**, 261, 266, 269, 282, 284

XObject subtypes
Form 261, **284**
Image 261, **267**, **448**, 480
PS 261, **290**

XObjects
See external objects

xor operator (PostScript) **116**, **704**

xref keyword **64**, 67, 733

XSL (Extensible Stylesheet Language) file format 622, 624

Xsquare entry (type 10 halftone dictionary) **398**

XStep entry (type 1 pattern dictionary) **222**, 223

XX name prefix **724**

XYZ color representation 191

Y

y operator 134, **163**, 165, 166, 702

yellow color component
 blue, complement of 378
DeviceCMYK color space 178, 180
DeviceN color spaces 206
 grayscale conversion 377, 382
 halftones for 401

yellow color component (*continued*)
initialization 180
overprinting 464
RGB conversion 377, 378
transfer function 380, 381
transparent overprinting 464
undercolor removal 150, 378, 379

yellow colorant
overprinting 464
PANTONE Hexachrome system 205
printing ink 201
process colorant 178, 180
subtractive primary 178, 180
transparent overprinting 464

Yes appearance state (checkbox fields) 539

Ysquare entry (type 10 halftone dictionary) 398

YStep entry (type 1 pattern dictionary) 222, 223

yuan symbol (¥) character 343

YUV color representation 60

YUVK color representation 60

Z

Zapf Dingbats typeface

See ITC Zapf Dingbats® typeface

ZapfDingbats standard font 319, 329, 795, 797

character encoding, built-in 709, 721–722

character set 709, 721–722

ZLIB Compressed Data Format Specification (Internet RFC 1950) 46, 815

zlib/deflate compression

See Flate compression

zone theory of color vision 181

zoom factor

See magnification factor

Colophon

THIS BOOK WAS PRODUCED using Adobe® FrameMaker®, Adobe Illustrator®, Adobe Photoshop®, Adobe Acrobat® Distiller®, and other application software packages that support the PostScript® page description language and Type 1 fonts. The type used is from the Adobe Minion® and Myriad® families. Heads are set in Myriad MM 565 Semibold, 600 Normal, and the body text is set in 10.5-on-13-point Minion.

Authors—Jim Meehan, Ed Taft, Stephen Chernicoff, Caroline Rose

Key Contributors—Nabeel Al-Shamma, Steven Kelley Amerige, Bob Ayers, Tim Bienz, Richard Cohn, Michelle Dalton, Stephen Deach, Matt Foley, Ron Gentile, Jim King, Bennett Leeds, Pierre Louveaux, Carl Orthlieb, Ajay Pande, Roberto Perelman, Jim Pravetz, Dan Rabin, Loretta Guarino Reid, Paul Rovner, Ed Rowe, Craig Rublee, Mike Schuster, Steve Schiller, John Warnock, Bob Wulff, Steve Zilles

Reviewers—Parviz Banki, Xintai Chang, L. Peter Deutsch, Mark Donohoe, Martin Fox, David Gelpman, Brian Havlin, Raph Levien, Ken Lunde, John Nash, Terry O'Donnell, Scott Petersen, Dick Sites, Lydia Stang, Koichi Yoshimura, and numerous others at Adobe Systems and elsewhere

Editing and Book Production—Stephen Chernicoff, Caroline Rose

Index—Stephen Chernicoff

Illustrations—Kim Arney, Wendy Bell, Peter Constable, Lisa Ferdinandsen, Carol Keller, Pierre Louveaux, Jim Meehan, Dayna Porterfield, Carl Yoshihara

Book Design—Sharon Anderson

Publication Management—Robert Morrish, Christine Yarrow




PLATE 1 *Additive and subtractive color (Section 4.5.3, “Device Color Spaces,” page 178)*




PLATE 2 *Uncalibrated color (Section 4.5.4, “CIE-Based Color Spaces,” page 181)*




PLATE 3 *Lab color space* (“*Lab Color Spaces*,” page 187)

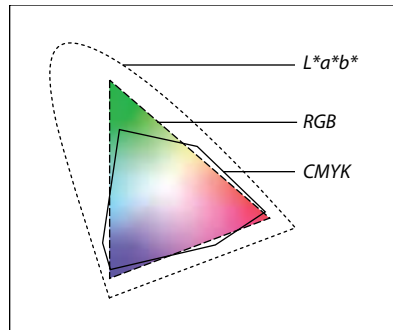


PLATE 4 *Color gamuts* (“*Lab Color Spaces*,” page 187)



AbsoluteColorimetric



RelativeColorimetric



Saturation



Perceptual

PLATE 5 *Rendering intents* (“*Rendering Intents*,” page 197)




PLATE 6 Duotone image (*“DeviceN Color Spaces,”* page 205)




PLATE 7 Quadtone image (*“DeviceN Color Spaces,”* page 205)

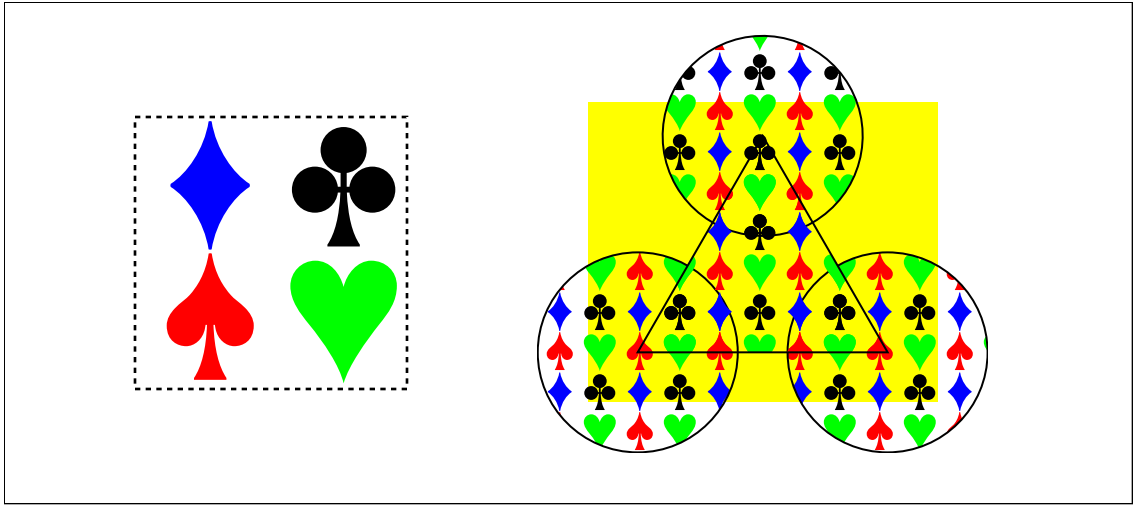


PLATE 8 *Colored tiling pattern* (“Colored Tiling Patterns,” page 224)




PLATE 9 *Uncolored tiling pattern* (“Uncolored Tiling Patterns,” page 228)



Extend = [false false], Background not specified



Extend = [true true], Background not specified



Extend = [false false], Background specified

PLATE 10 Axial shading (“Type 2 (Axial) Shadings,” page 239)




PLATE 11 Radial shadings depicting a cone (“Type 3 (Radial) Shadings,” page 241)




PLATE 12 *Radial shadings depicting a sphere* (“Type 3 (Radial) Shadings,” page 241)




PLATE 13 *Radial shadings with extension* (“Type 3 (Radial) Shadings,” page 242)




PLATE 14 *Radial shading effect* (“Type 3 (Radial) Shadings,” page 242)




PLATE 15 *Coons patch mesh ("Type 6 Shadings (Coons Patch Meshes)," page 250)*




PLATE 16 *Transparency groups (Section 7.1, “Overview of Transparency,” page 411)*


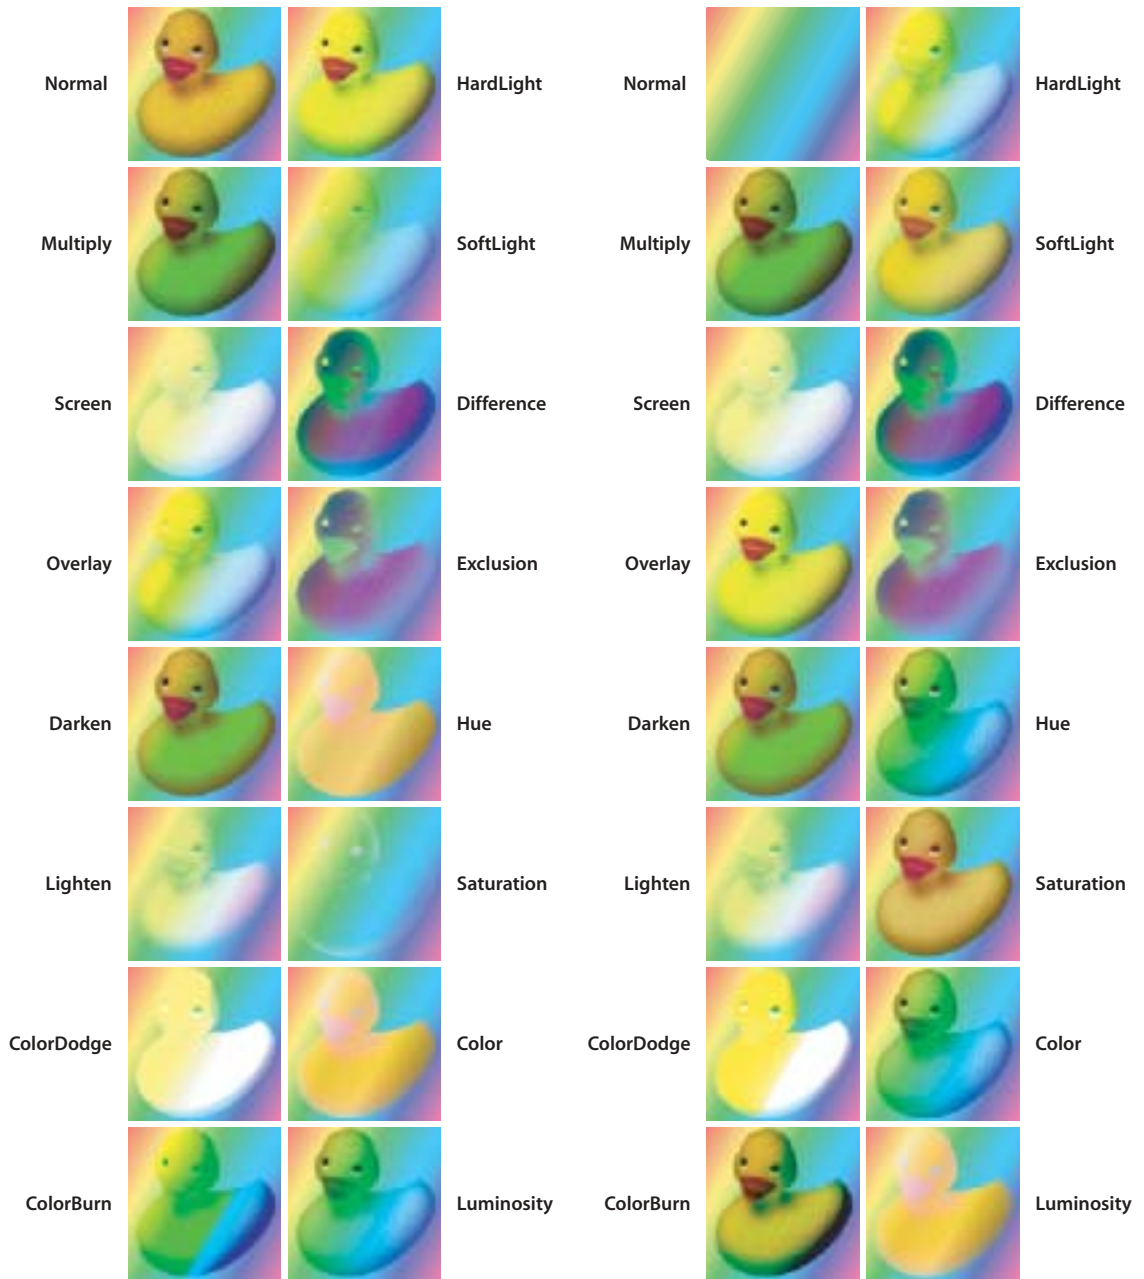



PLATE 17 *Isolated and knockout groups (Sections 7.3.4, “Isolated Groups,” page 433 and 7.3.5, “Knockout Groups,” page 434)*



Duck in foreground, rainbow in background

Rainbow in foreground, duck in background

PLATE 18 *RGB blend modes (Section 7.2.4, "Blend Mode," page 416)*




Duck in foreground, rainbow in background

Rainbow in foreground, duck in background

PLATE 19 CMYK blend modes (Section 7.2.4, "Blend Mode," page 416)

Opacity = 1.0



Opacity = 0.5




PLATE 20 *Blending and overprinting* (“Compatibility with Opaque Overprinting,” page 462)