

Reference Guide



Compute Abstraction Layer
(CAL) Technology

Intermediate Language (IL)

February 2009

v. 2.0

© 2009 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, ATI, the ATI logo, Radeon, FireStream, and combinations thereof are trademarks of Advanced Micro Devices, Inc. Other names are for informational purposes only and general-purpose be trademarks of their respective owners.

The contents of this document are provided in connection with Advanced Micro Devices, Inc. ("AMD") products. AMD makes no representations or warranties with respect to the accuracy or completeness of the contents of this publication and reserves the right to make changes to specifications and product descriptions at any time without notice. The information contained herein general-purpose be of a preliminary or advance nature and is subject to change without notice. No license, whether express, implied, arising by estoppel or otherwise, to any intellectual property rights is granted by this publication. Except as set forth in AMD's Standard Terms and Conditions of Sale, AMD assumes no liability whatsoever, and disclaims any express or implied warranty, relating to its products including, but not limited to, the implied warranty of merchantability, fitness for a particular purpose, or infringement of any intellectual property right.

AMD's products are not designed, intended, authorized or warranted for use as components in systems intended for surgical implant into the body, or in other applications intended to support or sustain life, or in any other application in which the failure of AMD's product could create a situation where personal injury, death, or severe property or environmental damage general-purpose occur. AMD reserves the right to discontinue or make changes to its products at any time without notice.



Advanced Micro Devices, Inc.
One AMD Place
P.O. Box 3453
Sunnyvale, CA 94088-3453
www.amd.com

Preface

About This Document

This document describes the instruction set for the CAL Technology IL compiler.

The document serves two purposes:

1. It specifies the language constructs and behavior, including the organization of each type of instruction in both text syntax and binary format.
2. It provides a reference of instruction operation that compiler writers can use to maximize performance of the processor.

Audience

This document is intended for programmers writing application and system software, including operating systems, compilers, loaders, linkers, device drivers, and system utilities. It assumes an understanding of the R600- and R700-family processor microarchitecture and of programming practices for either graphics or general-purpose computing.

Contact Information

To submit questions or comments about this document, contact our technical documentation staff at: streamcomputing@amd.com.

For questions concerning ATI Stream products, please email: streamcomputing@amd.com.

For questions about developing with ATI Stream, please email: streamdeveloper@amd.com.

You can learn more about ATI Stream at: <http://www.amd.com/stream>.

We also have a growing community of ATI Stream users! Come visit us at the ATI Stream Developer Forum (<http://www.amd.com/streamdevforum>) to find out what applications other users are trying on their ATI Stream products!

Organization

This document begins with an overview summarizing the similarities and differences between ATI Intermediate Language (IL) and general-purpose

computer languages. It describes text and binary formats of the IL program instructions. Then, it describes the types of instructions in detail, presenting a high-level description of the instruction fields, and restrictions that must be observed. It also describes the instruction syntax for text representation. Further, it presents the specification of each type of instruction. A glossary of terms and acronyms ends the document.

Endian Order

The R600 and R700 architectures address memory and registers using little-endian byte-ordering and bit-ordering. Multi-byte values are stored with their least-significant (low-order) byte (LSB) at the lowest byte address; they are illustrated with their LSB at the right side. Byte values are stored with their least-significant (low-order) bit (lsb) at the lowest bit address; they are illustrated with their lsb at the right side.

Conventions

The following conventions are used in this document.

mono-spaced font	A filename, file path, or code.
*	Any number of alphanumeric characters in the name of a microcode format, microcode parameter, or instruction.
< >	Angle brackets denote streams.
[1,2)	A range that includes the left-most value (in this case, 1) but excludes the right-most value (in this case, 2).
[1,2]	A range that includes both the left-most and right-most values (in this case, 1 and 2).
{x y}	One of the multiple options listed. In this case, x or y.
0.0	A single-precision (32-bit) floating-point value.
1011b	A binary value, in this example a 4-bit value.
7:4	A bit range, from bit 7 to 4, inclusive. The high-order bit is shown first.

Related Documents

- AMD, *R600-Family Instruction Set Architecture*, Sunnyvale, CA, 2008. This document includes the RV670 GPU instruction details.
- ISO/IEC 9899:TC2 - *International Standard - Programming Languages - C*
- Kernighan Brian W., and Ritchie, Dennis M., *The C Programming Language*, Prentice-Hall, Inc., Upper Saddle River, NJ, 1978.
- IEEE, *754-1985 IEEE Standard for Binary Floating-Point Arithmetic*, 2003.
- I. Buck, T. Foley, D. Horn, J. Sugerman, K. Fatahalian, M. Houston, and P. Hanrahan, "Brook for GPUs: stream computing on graphics hardware," ACM Trans. Graph., vol. 23, no. 3, pp. 777–786, 2004.

ATI STREAM COMPUTING

- *ATI Compute Abstraction Layer (CAL) Intermediate Language (IL) Reference Manual*. Published by AMD.
- Buck, Ian; Foley, Tim; Horn, Daniel; Sugerman, Jeremy; Hanrahan, Pat; Houston, Mike; Fatahalian, Kayvon. “BrookGPU”
<http://graphics.stanford.edu/projects/brookgpu/>
- Buck, Ian. “Brook Spec v0.2”. October 31, 2003.
<http://merrimac.stanford.edu/brook/brookspec-05-20-03.pdf>
- *OpenGL Programming Guide*, at <http://www.glprogramming.com/red/>
- *Microsoft DirectX Reference Website*, at
[http://msdn.microsoft.com/en-us/library/bb219740\(VS.85\).aspx](http://msdn.microsoft.com/en-us/library/bb219740(VS.85).aspx)
- *Microsoft Programming Guide for HLSL*,
<http://msdn2.microsoft.com/en-us/library/bb509635.aspx>
- *GPGPU*: <http://www.gpgpu.org>, and Stanford BrookGPU discussion forum
<http://www.gpgpu.org/forums/>

Contents

Preface

Contents

Chapter 1 Overview

1.1	Open Design	1-1
1.2	DirectX as a Design Basis	1-1
1.3	Threading Model	1-2
1.4	Access Model for Local Shared Memory	1-2

Chapter 2 Binary Stream Format

2.1	IL Stream	2-1
2.2	IL Token Descriptions	2-1
2.2.1	Language Token	2-2
2.2.2	Version Token	2-2
2.2.3	Opcode Token	2-2
2.2.4	Destination Token	2-3
2.2.5	Destination Modifier Token	2-4
2.2.6	Source Token	2-4
2.2.7	Source Modifier Token	2-6
2.2.8	Relative Address Token	2-8
2.2.9	Source Token Examples	2-9

Chapter 3 Text Instruction Syntax

3.1	Version	3-1
3.2	Registers	3-2
3.3	Relative Addressing	3-2
3.4	Control Specifiers	3-2
3.5	Destination Modifiers	3-3
3.6	Write Mask	3-3
3.7	Source Modifiers	3-4
3.8	Comments	3-4

Chapter 4 Register Types

4.1	INDEX	4-1
4.2	OBJECT_INDEX.....	4-2
4.3	BARYCENTRIC_COORD	4-2
4.4	PRIMITIVE_INDEX	4-3
4.5	QUAD_INDEX.....	4-3
4.6	SPRITE	4-4
4.7	POS.....	4-4
4.8	PINPUT	4-5
4.9	VOUTPUT and OUTPUT	4-5
4.10	LITERAL	4-6
4.11	INTERP	4-6
4.12	TEXCOORD	4-7
4.13	PRICOLOR	4-7
4.14	SECCOLOR	4-8
4.15	PSOUTFOG	4-9
4.16	FOG.....	4-10
4.17	SPRITECOORD	4-11
4.18	PRIMCOORD.....	4-11
4.19	PRIMTYPE	4-12
4.20	vWinCoord	4-12
4.21	FACE.....	4-13
4.22	PCOLOR	4-13
4.23	DEPTH	4-14
4.24	Mask	4-14
4.25	Stencil.....	4-15
4.26	Global	4-15
4.27	Shared Temp.....	4-16
4.28	Thread_ID.....	4-16
4.29	Absolute Thread_ID	4-17
4.30	Thread_Group_ID	4-17
4.31	Generic_Memory	4-18

Chapter 5 Enumerated Types

5.1	ILShader	5-1
5.2	ILLanguageType	5-1
5.3	ILRegType	5-2
5.4	ILTopologyType	5-3
5.5	ILMatrix.....	5-3
5.6	ILComponentSelect.....	5-3

5.7	ILModDstComp.....	5-4
5.8	ILImportUsage	5-4
5.9	ILImportComponent	5-6
5.10	ILDefaultVal	5-7
5.11	ILShiftScale	5-7
5.12	ILDivComp	5-8
5.13	ILRelOp.....	5-8
5.14	ILLogicOp.....	5-8
5.15	ILZeroOp	5-9
5.16	ILCmpValue.....	5-9
5.17	ILTexCoordMode.....	5-10
5.18	ILPixTexUsage	5-10
5.19	ILTexShadowMode	5-11
5.20	ILTexFilterMode	5-11
5.21	ILAnisoFilterMode	5-11
5.22	ILMipFilterMode.....	5-12
5.23	ILNoiseType	5-12
5.24	ILInterpolation	5-12
5.25	ILAddressing	5-12
5.26	IElementFormat	5-13
5.27	ILOpCode	5-13

Chapter 6 Instructions

6.1	Formats	6-1
6.2	Instruction Notes.....	6-1
6.2.1	Notes on Comparison Instructions.....	6-1
6.2.2	Notes on Flow Control Instructions.....	6-2
6.2.3	Notes on Input/Output Instructions	6-2
6.2.4	Notes on Conversion Instructions.....	6-2
6.2.5	Notes on Double Precision Instructions.....	6-3
6.3	Flow Control Instructions	6-3
6.4	Declaration and Initialization Instructions	6-22
6.5	Input/Output Instructions.....	6-48
6.6	Integer Arithmetic Instructions	6-83
6.7	Unsigned Integer Operations	6-90
6.8	Conversion Instructions.....	6-94
6.9	Float Instructions	6-97
6.10	Double-Precision Instructions.....	6-154

Glossary of Terms

Index

Tables

2.1	IL Stream Instruction Packet Ordering	2-1
2.2	IL_Lang: Source Language Information	2-2
2.3	IL_Version: Source Version Information	2-2
2.4	IL_Opcode: Instruction Opcode Details.....	2-2
2.5	IL_Dst: Destination Operand Information	2-3
2.6	IL_Dst_Mod: Destination Modification Information	2-4
2.7	IL_Src: Source Operand Information.....	2-5
2.8	IL_Src_Mod: Source Operand Modification Information	2-7
2.9	IL_Rel_Addr: Relative Addressing Descriptor	2-8
3.1	Destination Modifiers	3-3
3.2	Source Modifiers	3-4
4.1	Registers and Their Restrictions	4-1
5.1	ILShader Enumeration Types	5-1
5.2	ILLanguageType Enumeration Types.....	5-1
5.3	ILRegType Enumeration Types	5-2
5.4	IL_TOPOLOGY Enumeration Types (Input to a Geometry Shader)	5-3
5.5	IL_OUTPUT_TOPOLOGY Enumeration Types (Output from a Geometry Shader)	5-3
5.6	ILMatrix Enumeration Types.....	5-3
5.7	ILComponentSelect Enumeration Types	5-3
5.8	ILModDstComp Enumeration Types.....	5-4
5.9	ILImportUsage Enumeration Types	5-4
5.10	ILImportComponent Enumeration Types.....	5-6
5.11	ILDefaultVal Enumeration Types	5-7
5.12	ILShiftScale Enumeration Types	5-7
5.13	ILDivComp Enumeration Types.....	5-8
5.14	ILRelOp Enumeration Types	5-8
5.15	ILLogicOp Enumeration Types	5-8
5.16	ILZeroOp Enumeration Types	5-9
5.17	ILCmpValue Enumeration Types	5-9
5.18	ILTexCoordMode Enumeration Types.....	5-10
5.19	ILPixTexUsage Enumeration Types.....	5-10
5.20	ILTexShadowMode Enumeration Types	5-11
5.21	ILTexFilterMode Enumeration Types	5-11
5.22	ILAnisoFilterMode Enumeration Types.....	5-11
5.23	ILMipFilterMode Enumeration Types.....	5-12
5.24	ILNoiseType Enumeration Types.....	5-12
5.25	ILInterpolation Enumeration Types	5-12
5.26	ILAddressing Enumeration Types.....	5-12
5.27	ILElementFormat Enumeration Types	5-13
5.28	ILOpCode Enumeration Types	5-13

Chapter 1

Overview

This document defines the format and behavior of the IL. The Intermediate Language (IL) is an abstract representation for hardware vertex, pixel, and geometry shaders that can be taken as input by other modules implementing the IL. An IL compiler uses an IL shader in conjunction with driver state information to translate these shaders into hardware instructions or a software emulation layer.

1.1 Open Design

The IL adopts an open design, where most instructions can appear in any kind of shader.

Note that the IL does not enforce usage restrictions of the input language, but only verifies that the IL, as specified, is properly programmed. For example, in DX, instruction modifiers such as `_x2` general-purpose not be applied to texture address instructions. The IL does not enforce this restriction, since both the IL and current architectures support such operations.

Also, this manual does not describe how to optimally code the IL. By design, the IL has an extensive set of operations, but some of these are not supported by existing hardware. The programmer is advised to adhere in most cases to the syntax restrictions and performance guidelines of the input language.

1.2 DirectX as a Design Basis

Because of changes between DX9 and DX10, many instructions occur in multiple forms:

- unconditional
- conditional (comparing two float values)
- logical (comparing a single integer value with zero)
- Boolean (comparing a Boolean register with true)

To support DX10, many operations required several similar IL opcodes. For example, there are three forms of break: unconditional, break on a Boolean, and break on a logical value.

Other operations were required both in vector and scalar forms. For example, there are two `rsq` instructions: `rsq` corresponds to the IL 1 scalar opcode,

`rsq_vec` corresponds to the DX10 vector form that computes the reciprocal square root on each component.

In DX9, $0 * \text{any value}$ was defined to be 0. DX10 changed this to more closely match IEEE arithmetic which defines $0 * \text{Nan} = \text{Nan}$.

All float operations containing a multiply now take a flag to specify Nan behavior.

1.3 Threading Model

A hierarchical threading model is used. A *kernel* (a shader program running on a GPU) can be launched with a number of thread groups. Threads within this group can communicate through local shared memory. There is no supported communication between thread groups. Threads in a thread group run in units called wavefronts. All threads in a wavefront run in SIMD fashion (in lock steps). All wavefronts within a thread group can be synchronized using a synchronization barrier operation.

1.4 Access Model for Local Shared Memory

Each processor has an amount of local memory that can be shared across the threads in a thread group. IL provides two models of memory access to local data store (LDS).

The first memory access model, called owner-computes, is supported by the HD4000-family of devices. In owner-computes, each thread in a thread group owns a area of LDS memory. The size of the area is declared in the shader. Each thread in a group can write only to the area of memory it owns; however, a thread can read any chunk of memory that is owned by either itself or other threads. An LDS shared memory read is specified by (`owner_thread_ID`, `offset`): read the memory area owned by that `thread_ID` with an offset within the area.

Different from the access model for threads within a wavefront, the access mode for different wavefronts (within a thread group) is specified by the sharing mode, which is either relative or absolute. If it is relative, new and consecutive space is allocated for each wavefront; if it is absolute, all wavefronts are mapped to the same set of memory starting at address 0. In this mode, wavefronts can overwrite each other's data.

The second memory access model is a general read write: each thread can read or write any address in the LDS.

Both models allow threads to read or write memory (video or system), but do not provide synchronization to memory.

Supported inter-thread communication includes:

- SR – Globally shared registers.
- Sharing between all wavefronts in a SIMD.
- Column sharing on the SIMD.

ATI STREAM COMPUTING

- Persistent registers.
- LDS – local data store - read/write.
- Data sharing between all threads in a group.
- Required synchronization.
- Memory - read/write.
- Constant buffers
- Texture cache

The following indexing values are available in the compute shader:

- vTid – ID of thread within a group
- aTid – ID of thread within a kernel

Chapter 2

Binary Stream Format

The following chapter defines the format in which kernels written using the IL are passed to the compiler.

2.1 IL Stream

Clients pass kernels as a stream of 32-bit tokens organized as variable-length instruction packets. These tokens include information about the client language, shader type, and instruction packets that describe the operation of the kernels. Table 2.1 indicates the ordering of packets.

Table 2.1 IL Stream Instruction Packet Ordering

Instruction Packet	Description
1	IL_Lang token. See Section 2.2.1, on page 2-2.
2	IL_Version token. See Section 2.2.2, on page 2-2.
3	IL_OpCode token describing the operation of the first instruction in the stream and the beginning of the first IL instruction packet.
...	... More IL instruction packets. See Chapter 3.
n (number of 32-bit tokens in the stream)	IL instruction packet for an END instruction. See page 6-13.

Instruction Packets contain all the information needed to perform the single instruction specified in the IL_OpCode token. This information can include data about source and destination operands, destination or target labels, and additional data needed to perform the instruction.

There are assorted IL statements that can be used to declare resources, samplers, or registers. Any declaration of an object must appear before all uses of the object. There is no requirement to group all declarations at the start of the program.

Most IL statements and types can be used in any kind of shader. However, as noted below, some statements and types are restricted to specific kinds of shaders.

2.2 IL Token Descriptions

This section describes the generic tokens used in the IL stream. There are additional tokens for use in single specific instruction packets. Those tokens are

described under the instruction packet description for the instruction in which they can be used.

2.2.1 Language Token

This token indicates the type of client generating the IL. This token must be at the beginning of every IL stream passed to the compiler.

Table 2.2 IL_Lang: Source Language Information

Field Name	Bits	Description
client_type	7:0	Specifies the client API. Can be any value of the enumerated type ILLanguage-Type. This value is not used, but can allow IL compilers to make API specific workarounds and optimizations.
reserved	31:8	Must be zero.

2.2.2 Version Token

This token specifies the version of IL used in this IL stream. It also specifies the type of kernel the IL stream represents (pixel or vertex).

Table 2.3 IL_Version: Source Version Information

Field Name	Bits	Description
minor_version	7:0	The minor version.
major_version	15:8	The major version.
shader_type	23:16	Specifies the type of shader described by this binary token stream. See Section 5.1, "ILShader," page 5-1.
multipass	24	0 = Outputs are not ignored. 1 = This shader is for multipass use only (output are ignored).
realtime	25	0 = This shader is not for real-time use. 1 = This shader is for real-time use. This can be set only when the shader type is IL_PIXEL_SHADER.
reserved	31:26	Must be zero.

2.2.3 Opcode Token

This token specifies the current operation and information required to perform the operation.

Table 2.4 IL_Opcode: Instruction Opcode Details

Field Name	Bits	Description
code	15:0	The operation for the current instruction. This value can be any of the enumerated type ILOpCode (see Section 5.27, "ILOpCode," page 5-13).
control	29:16	Opcode specific control. Possible values for this depend on the value of <i>code</i> . Specifies further instruction behavior. This field must be zero for all instructions not using it.

Table 2.4 IL_Opcode: Instruction Opcode Details (Cont.)

Field Name	Bits	Description
sec_modifier_present	30	Specifies whether an opcode-specific token describing further instruction behavior follows the primary modifier token. 0 An opcode specific token does not follow. 1 An opcode specific token follows. This field must be zero for all instructions not using it.
pri_modifier_present	31	Specifies whether an opcode-specific token describing further instruction behavior follows this token. 0 An opcode specific token does not follow. 1 An opcode specific token follows. This field must be zero for all instructions not using it.

2.2.4 Destination Token

This token specifies the register to which the hardware passes the result of the current instruction and other information pertaining to this result. This token can only be issued after an IL_OpCode token has been issued as part of an instruction packet. By default, all components of the register specified are written unless the modifier_present field is 1 and an IL_Dst_Mod token follows.

DX10 allows an additional kind of indexing: some temporary objects can be indexed by a register; thus, the IL now allows an additional modifier (register relative modifier). Two kinds of destinations can be indexed: IL_TEMP_ARRAY and IL_OUTPUT.

The immediate_present field is used for indexed data types: itemp, cb, etc. It can be used if the data type uses absolute, reg-relative, loop, or addr relative addressing. See Section 2.2.6, "Source Token," page 2-4, for examples.

Table 2.5 IL_Dst: Destination Operand Information

Field Name	Bits	Description
register_num	15:0	Register number to which the result of the current instruction is written. See Chapter 4, "Register Types," for acceptable values.
register_type	21:16	This value can be any of the enumerated type ILRegType (see Section 5.3, "ILRegType," page 5-2).
modifier_present	22	Specifies whether an IL_Dst_Mod token follows this token: 0 An IL_Dst_Mod token does not follow. 1 An IL_Dst_Mod token follows.
relative_address	24:23	Specifies whether register_num represents an offset from the address register specified the following IL_Rel_Addr token: 0 Absolute addressing is used. 1 Relative addressing is used (an IL_Rel_Addr token follows). See Section 2.2.8, on page 2-8. 2 Register relative addressing is used (tokens describing the index follows) Dimension specifies the number of source tokens that need to follow.
dimension	25	Number of additional following dimensions (0 == 1 dimension). so v[3][5] is src_token 1: il_regtype_vertex dim = 1 num = 3 src_token 2: il_regtype_vertex dim = 0 num = 5

Table 2.5 IL_Dst: Destination Operand Information (Cont.)

Field Name	Bits	Description
immediate_present	26	0 There is no immediate value. 1 A 32-bit value containing the immediate value follows this token, modifier tokens, and src tokens used in register relative addressing.
reserved	27:30	Must be zero.
extended	31	0 No extended register addressing. 1 The register_number is a 32-bit value. The low 16-bits is represented by the register_num field, and following this token is a 32-bit word containing the high 16-bits and 16-bits reserved.

2.2.5 Destination Modifier Token

This token specifies modifications to the destination operand.

The modifiers precedence is:

1. shift_scale
2. clamp
3. destination comp

An example of all these operations performed on the x-component of a destination operand is:

```
dstCmpMod(clamp(shift_scale(dst.x)))
```

Table 2.6 IL_Dst_Mod: Destination Modification Information

Field Name	Bits	Description
component_x_r	1:0	This value can be any of the enumerated type ILMODDstComp ¹ .
component_y_g	3:2	This value can be any of the enumerated type ILMODDstComp ¹ .
component_z_b	5:4	This value can be any of the enumerated type ILMODDstComp ¹ .
component_w_a	7:6	This value can be any of the enumerated type ILMODDstComp ¹ .
clamp	8	Specifies whether to clamp the value to 0.0 and 1.0: 0 Do not clamp. 1 Clamp. Clamp(NaN) returns 0.
shift_scale	12:9	This value can be any of the enumerated type ILShiftScale ² .
reserved	31:13	Must be zero.

1. See Section 5.7, "ILModDstComp," page 5-4.
2. See Section 5.11, "ILShiftScale," page 5-7.

2.2.6 Source Token

This token specifies the register that the instruction uses as a source operand. This token can only be issued after an IL_OpCode token as part of an instruction packet. If an IL_Src_Mod token does not follow, then:

- the first component is set to the x component,

- the second component is set to the y component,
- the third component is set to the z component, and
- the fourth component is set to the w component.

Starting with IL 2.0, some source tokens are allowed to use register relative indexing. Only one level of indexing is allowed. The type of register that has indexed sources must be one of the following: IL_TEMP_ARRAY, IL_CONST_BUFFER, or IL_INPUT. Since the index must be a scalar value, a modifier field must be used to replicate a single component into four slots.

Table 2.7 lists and briefly describes the IL_src source operands. IL version 2.0 also allows a source token to refer to a literal defined in a dcl_literal statement. DX10 statements such as:

```
add r1, r2, float4 (1.0f, 2.0f, 3.0f, 4.0f)
```

can be translated into:

```
dcl_literal_float4, l1, 1.0f, 2.0f, 3.0f, 4.0f
add r1, r2, l1
```

Table 2.7 IL_Src: Source Operand Information

Field Name	Bits	Description
register_num	15:0	Register number from which to retrieve the operand. See Chapter 4, "Register Types," for acceptable values of this field.
register_type	21:16	This value can be any of the enumerated type ILRegType. See Chapter 4, "Register Types," for a description of these types.
modifier_present	22	Specifies whether an IL_Src_Mod token follows this token: 0 An IL_Src_Mod token does not follow. 1 An IL_Src_Mod token follows.
relative_address	24:23	Specifies whether the register_num represents an offset from the address register specified in the following IL_Rel_Addr token or an absolute address: 0 Absolute addressing is used. 1 Relative addressing is used (an IL_Rel_Addr token follows). See Section 2.2.8, on page 2-8. 2 Register relative addressing is used (source tokens follows). Dimension specifies the number of source tokens that need to follow.
dimension	25	Number of additional following dimensions (0 == 1 dimension). so v[3][5] is src_token 1: il_regtype_vertex dim = 1 num = 3 src_token 2: il_regtype_vertex dim = 0 num = 5
immediate_present	26	0 There is no immediate value. 1 A 32-bit value containing the immediate value follows this token, modifier tokens and src tokens used in register relative addressing.
reserved	27:30	Must be zero.
extended	31	0 There is no extended register addressing. 1 The register_number is a 32-bit value. The low 16-bits is represented by the register_num field and following this token is a 32-bit word containing the high 16-bits and 16-bits reserved.

2.2.7 Source Modifier Token

This token specifies modifications to the source operand. It can be issued only immediately following an IL_Src token, and only if the preceding IL_Src token has the modifier_present field set to 1.

When this token is used in conjunction with an IL_Rel_Addr token, modifiers are applied to the register referenced by the relative address, not to the relative-address register itself. Also, this token always precedes the IL_Rel_Addr token if both are used.

Notes about this token:

- If this token is not present, the x, y, z, w corresponds to the first, second, third, and fourth components, respectively.
 - For floating point arithmetic instructions, the negate modifier simply flips the sign of the number(s) in the source operand, including on INF values. Applying negate on NaN preserves NaN, although the particular NaN bit pattern that results is not defined.
 - For integer instructions, the negate modifier takes the 2's complement of the number(s) in the source operand.
 - For floating point arithmetic instructions: the abs modifier simply forces the sign of the number(s) on the source operand positive, including on INF values. Applying abs on NaN preserves NaN, although the NaN bit pattern that results is not defined.
 - The 1 swizzle inserts a floating point 1.0f, even if the opcode is an integer operation. This can lead to unexpected results. For example, when evaluating the second source of iadd r1, r1, r1.1_neg(xyzw), implementations take a floating point 1.0f and treat it as an integer. To negate an integer, use the INEGATE function.
- The modifiers use the following precedence:
 1. swizzle
 2. invert
 3. bias
 4. x2
 5. sign
 6. divComp
 7. abs
 8. negate (per component)
 9. clamp

An example of all of these operations performed on the x-component of a source operand is:

```
clamp(negate (abs (divComp(sign(x2(bias(invert(swizzle(src's))))))))))
```

The modifiers *bias*, *x2*, *divComp*, and *clamp* cannot be used when the opcode of the instruction specifies an integer or logical operation.

Table 2.8 IL_Src_Mod: Source Operand Modification Information

Field Name	Bits	Description
swizzle_x_r	2:0	This value can be any of the enumerated type ILComponentSelect ¹ . If this token is not present, the IL behaves as if this field is set to IL_COMPSEL_X_R.
negate_x_r	3	Specifies whether to negate the first component: 0 Do not negate. 1 Negate.
swizzle_y_g	6:4	This value can be any of the enumerated type ILComponentSelect ¹ . If this token is not present, the IL behaves as if this field is set to IL_COMPSEL_Y_G.
negate_y_g	7	Specifies whether to negate the second component: 0 Do not negate. 1 Negate.
swizzle_z_b	10:8	This value can be any of the enumerated type ILComponentSelect ¹ . If this token is not present, the IL behaves as if this field is set to IL_COMPSEL_Z_B.
negate_z_b	11	Specifies whether to negate the third component: 0 Do not negate. 1 Negate.
swizzle_w_a	14:12	This value can be any of the enumerated type ILComponentSelect ¹ . If this token is not present, the IL behaves as if this field is set to IL_COMPSEL_W_A.
negate_w_a	15	Specifies whether to negate the fourth component: 0 Do not negate. 1 Negate.
invert	16	Specifies whether to invert each value of the source register prior to the operation. ($s = 1.0 - s$). The 'invert' modifier has higher precedence than bias, x2, negate sign, divComp ops, or abs: 0 Do not invert. 1 Invert.
bias	17	Specifies whether to bias each value of the source register prior to the operation. ($s = s - 0.5$). This takes precedence over x2, negate, sign, divComp ops, and abs: 0 Do not bias. 1 Bias.
x2	18	Specifies whether to multiply each value by 2.0 prior to the operation ($s = 2.0 * s$). This takes precedence over negate, sign, divComp ops, and abs: 0 Do not perform x2. 1 Perform x2.
sign	19	Specifies whether to sign each value of the source register prior to the operation ($s = (s < 0) ? -1 : ((s == 0) ? 0 : 1)$). This takes precedence before divComp ops, and abs: 0 Do not sign. 1 Sign.

Table 2.8 IL_Src_Mod: Source Operand Modification Information (Cont.)

Field Name	Bits	Description
abs	20	Specifies whether to take the absolute value of each value of the source register prior to the operation. (s = (s < 0) ? -s : s): 0 Do not take the absolute value. 1 Take the absolute value.
divComp	23:21	Specifies component divide modifier. This can be any value of the enumerated type ILDivComp. This takes precedence over abs. See Section 5.12, "ILDivComp," page 5-8. This field indicates one of the following: <ul style="list-style-type: none"> No component divide necessary. The first component is divided by the second component. The first and second component is divided by the third component. The first, second, and third component is divided by the fourth component. Any one of these four are performed based on this field.
clamp	31:24	Specifies whether or not to clamp the value to [0.0, 1.0]. 0 Do not clamp 1 Clamp clamp(NaN) returns 0

1. See Section 5.6, "ILComponentSelect," page 5-3.

2.2.8 Relative Address Token

This token specifies data to use for relative addressing. This token can be issued after an IL_Src, IL_Src_Mod, IL_Dst, or IL_Dst_Mod token. It follows an IL_Src token if *relative_address* is 1 and *modifier_present* is 0 in the preceding IL_Src token. It follows an IL_Src_Mod token if *relative_address* is 1 and *modifier_present* is 1 in the preceding IL_Src token. It follows an IL_Dst token if *relative_address* is 1 and *modifier_present* is 0 in the preceding IL_Dst token. It follows an IL_Dst_Mod token if *relative_address* is 1 and *modifier_present* is 1 in the preceding IL_Dst token. See the Section 3.3, "Relative Addressing," page 3-2, for more details on using this token.

Table 2.9 IL_Rel_Addr: Relative Addressing Descriptor

Field Name	Bits	Description
address_register	15:0	Specifies the address register number to use for base relative addressing. (In the future this can specify the loop counter to use for loop relative addressing. In the current version of the IL, only the loop counter for the inner-most loop can be used).
loop_relative	16	Specifies whether loop or base relative addressing is used: 0 Base relative addressing is used. 1 Loop relative addressing is used. This field must be 1 if relative addressing is used on a destination.
component	19:17	Specifies which component of the address register to use as the base address. If <i>loop_relative</i> is 0, this field can be set to any value of the enumerated type ILComponentSelect. If <i>loop_relative</i> is 1, this field must be set to IL_COMPSEL_X_R
reserved	31:20	Must be zero.

2.2.9 Source Token Examples

Some source token encodings can be complicated, as shown in the following examples.

Example 1: *Source Operand:* x5[6].y

- IL_Src
 - register_num = 5
 - register_type = IL_REGTYPE_ITEMP (see Section 6.4, “Declaration and Initialization Instructions,” page 6-22)
 - modifier_present = 1
 - relative_address = IL_ADDR_ABSOLUTE
 - dimension = 0
 - immediate_present = 1
 - extended = 0
- IL_Src_Mod
 - swizzle_x_r = y
 - swizzle_y_g = y
 - swizzle_z_b = y
 - swizzle_w_a = y
- IL_Literal
 - val = 6

Example 2: *Source Operand:* x5[r2.x+6].y

Here are the IL tokens for this:

- IL_Src x5
 - register_num = 5
 - register_type = IL_REGTYPE_ITEMP
 - modifier_present = 1
 - relative_address = IL_ADDR_REG_RELATIVE
 - dimension = 0
 - immediate_present = 1
 - extended = 0
- IL_Src_Mod
 - swizzle_x_r = y
 - swizzle_y_g = y
 - swizzle_z_b = y
 - swizzle_w_a = y

- IL_Src
 - register_num = 2
 - register_type = IL_REGTYPE_TEMP
 - modifier_present = 1
 - relative_address = IL_ADDR_ABSOLUTE
 - dimension = 0
 - immediate_present = 0
 - extended = 0
- IL_Src_Mod .x for r2
 - swizzle_x_r = x
 - swizzle_y_g = x
 - swizzle_z_b = x
 - swizzle_w_a = x
- IL_Literal
 - val = 6

Example 3: *Source Operand:* v[1][2] (all fields set to zero, unless otherwise stated)

- IL_Src
 - Register_num = 1
 - Register_type = IL_REGTYPE_VERTEX
 - Dimension = 1
- IL_Src
 - Register_num = 2
 - Register_type = IL_REGTYPE_VERTEX

Example 4: *Source Operand:* v[1][2].xyxx (all fields set to zero unless otherwise stated)

- IL_Src
 - Register_num = 1
 - Register_type = IL_REGTYPE_VERTEX
 - Dimension = 1
 - Modifier_present = 1
- IL_Src_Mod
 - swizzle_x_r = x
 - swizzle_y_g = y
 - swizzle_z_b = x

ATI STREAM COMPUTING

- swizzle_w_a = x
- IL_Src
 - Register_num = 2
 - Register_type = IL_REGTYPE_VERTEX

Chapter 3

Text Instruction Syntax

IL Text syntax is designed to closely match the IL specification, so that there is an almost complete one-to-one mapping.

Below is a simple vertex and pixel shader pair written in IL Text syntax that renders green stripes.

```

il_vs
dclv_elem(0) v0           ; Declare position
dclv_elem(1) v1           ; Declare color
dclv_elem(2) v2           ; Declare texture coordinates

mmul_matrix(4x4) oPos, v0, c[0] ; Transform position to clip space

mov oPriColor0, v1       ; Export vertex color
mov oT0, v2              ; Export texture coordinates

end

il_ps
dclpi_x(1)_y(1)_z(1)_w(1) vPriColor0 ; Declare primary color import
dclpi_x(1)_y(1)_z(*)_w(*) vT0        ; Declare vs import texture
                                     coordinates

def c0, 0.5, 1, 0, 0
def c1, 0.0, 1.0, 0.0, 1.0           ; Green color

mod r0.x, vT0.x, c0.y                ; x = mod( s, 1.0 )

ifc_relop(lt) r0.x, c0.x              ; if ( x < 0.5 )
    mov oC0, vPriColor0              ; Output surface color
else ; else
    mov oC0.rgb1, c1                 ; Output green color
endif

end

```

3.1 Version

The first two tokens in an IL binary stream are IL_Lang and IL_Version. IL Text syntax combines these into a single version instruction. The IL translator sets the language type to IL_LANG_VERSION and disables the language defaults. The IL_Version token has the following syntax:

```

il_ps_major_minor_mp_rt
il_vs_major_minor_mp_rt

```

If the major and minor version are not specified, the IL translator inserts them based on its own version. Typically, the version information is omitted, unless a specific optimization made in the compiler is wanted. If a shader represents a multipass shader, append `_mp` to the statement. If this shader represents a real-time shader, append `_rt` to the statement.

3.2 Registers

Registers that are prefixed with the letter “v” are read-only (import) buffers, registers prefixed with the letter “o” are write-only (export) buffers. This section only lists the registers. See Chapter 4 for more information on register types.

Common Registers:

`b#`, `c#`, `i#`, `a#`, `aL`, and `r#`

Vertex Shader:

`v#`, `oPos#`, `oPriColor#`, `oSecColor#`, `oT#`, `oInterp#`, `oFog`,
`oSprite`, `vBaryCoord`, `vPrimIndex`, `vQuadIndex`

Pixel Shader:

`vPriColor#`, `vSecColor#`, `vT#`, `vInterp#`, `vFog`, `vSprite`, `vFace`,
`vWinCoord`, `oC#`, `oDepth`

3.3 Relative Addressing

The `[]` syntax can be used on register types that can be relatively addressed.

Absolute Register Index

`reg[m]` Example: `mov r[0], r[1]`
`reg[m0+...+mn]` Example: `mov r0, r[1+2]`

Note that `r[0]` is the same as `r0`.

Base Relative Addressing

The address register is used to address another register file.

`reg[A+m0+...+mj]` Example: `mov r0, r[a0.x+5]`
`reg[m0+...+mn+A+n0+...+nn]` Example: `mov r0, r[1+3+a0.y+2]`

Loop Relative Addressing

The loop register `aL` is used to address another register file.

`reg[A+m0+...+mj]` Example: `mov r0 r[aL+5]`
`reg[m0+...+mn+A+n0+...+nn]` Example: `mov r0, r[1+3+aL+2]`

3.4 Control Specifiers

Control specifiers affect the behavior of the instruction. Binary control specifiers, which typically indicate there are two possible actions, have the following syntax:

`<instr>[_ctrl]`

If the control specifier `_ctrlspec` is left off the instruction, the default action is performed.

Control specifiers that have more than two modes of operation have the following syntax:

```
instruction_ctrl(value)]
```

For these specifiers, the parenthesis is mandatory, and the value specifies the mode of operation. Sections 3.5, 3.6, 3.7, and 3.8 describe the control specifiers in the IL spec. See Chapter 6, "Instructions," to determine which specifiers go with which instructions.

3.5 Destination Modifiers

Instruction modifiers are applied after the instruction runs but before writing the result to the destination register.

Table 3.1 Destination Modifiers

Modifier	Description	Example
<code>_x2</code> <code>_x4</code> <code>_x8</code> <code>_d2</code> <code>_d4</code> <code>_d8</code>	Shift scale modifiers.	<code>add_x2 r0, r1, r2</code>
<code>_sat</code>	Saturate or clamp result to [0,1].	<code>add_sat r0, r1, r2</code> <code>add_x2_sat r0, r1, r2</code>

3.6 Write Mask

Component-wise write masks on the destination have the following syntax in IL Text:

```
reg.{x|_0|1}{y|_0|1}{z|_0|1}{w|_0|1}
reg.{r|_0|1}{g|_0|1}{b|_0|1}{a|_0|1}
```

A letter specifies a vector component to be written, while an underscore specifies that an component is not written. An component can also be forced to zero or one.

Examples:

```
mov r0.x, r1
mov r0.y r1
mov r0.z r1
mov r0.x_zw, r1
mov r0.y_w, r1
mov r0.__w, r1
mov r0._1_w, r1
mov r0.01z, r1
mov r0.01z1, r1
mov r0.11, r1
mov r0._11, r1
mov r0.__11, r1
```

```

mov r0.xyz1, r1
mov r0.l_0, r1
mov r0.ll_0, r1
mov r0.y00, r1
mov r0.yz1, r1
mov r0.y_1, r1

```

3.7 Source Modifiers

Table 3.2 Source Modifiers

Modifier	Description	Example
swizzle	Rearrange and/or replicate components.	mov r0, r1.zyx1 mov r0, r1.xyy
_abs	Takes the absolute value of components.	mov r0, r1_abs
_bias	Components are biased ($x - 0.5$).	add r0, r1, r2_bias
_bx2	Signed scaling. Combined bias and x2 modifiers.	add r0, r1, r2_bx2
_divcomp(<i>type</i>)	Performs division based on <i>type</i> . <i>type</i> : y, z, w, unknown	texId_stage(0) r0, vT0_divcomp(y)
_invert	Invert components ($1.0 - x$).	add r0, r1_invert, r2
_neg(<i>comp</i>)	Provides per component negate.	mov r0 r1_neg(xw)
_sign	Signs Components: Components less than 0 become -1. Components equal to 0 become 0. Components greater than 0 become 1.	mov r0, r1_sign
_x2	Multiply components by 2.0.	add r0, r1, r2_x2

3.8 Comments

Only single-line comments are supported using a semicolon as the delimiter. Other commenting styles, such as `/* */` matching pairs, also are supported. A comment can start from any position on the line.

Example:

```

; The following instruction moves the contents of r1 into r0
mov r0, r1 ; mov instruction

```


Chapter 4

Register Types

This chapter describes the ATI IL valid register types.

Table 4.1 Registers and Their Restrictions

Name	IL Regtype_	Syntax	Num Comps	Read	Write	Relative Address			Notes
						Loop	A0	Index	
Boolean	CONST_BOOL	b#	1	No	Defb only	No	No	No	Can be used only as the source for static control flow.
Prim id	VPRIM	vPrim	1	Yes	No	No	No	No	Used only in geometry shader.
buffer	CONST_BUFF	C#[#]	4	Yes	No	No	No	Yes	Max 4096 elements.
float	CONST_FLOAT	C#	4	Yes	Def only	No	Yes	No	
Int	CONST_INT	I#	4	Loop only	Def only	No	No	No	
Addr	ADDR	A0	4	Yes	Mova	No	No	No	
Temp	TEMP	R#	4	Yes	Write	No	No	No	Max 4096 elements.
Index temp	INDEXED_TEMP	X#	4	Yes	Yes	No	No	Yes	Max 4096 elements.
Vertex	VERTEX	V#	4	Yes	Dclv or initv	No	No	No	Cannot be used in pixel shader.
index	INDEX	vINDEX	4	Yes					
Output	OUTPUT	o#	4	No	Yes	No	No	No	To support DX10.
Input	INPUT	v#	4	Yes	No	No	No	No	To support DX10.

4.1 INDEX

Enum: IL_REGTYPE_INDEX

Text Syntax: vIndex

Common Name: Index Register

Number of Components per Register: 4

Description:

- Vertex shader import for the index from the index buffer of the current vertex processed.

- Not guaranteed to be incremental.
- For non-HOS (high-order surface) rendering, the first component is the index of the current vertex processed.
- For HOS rendering (when HOS is enabled) the first, second, third, and fourth components represent the indices for the superprim vertices for the current tessellated vertex being processed. The superprim indices are output by the tessellation engine based on the relevant HOS state.
- This register cannot be used with relative addressing.
- It is an error to use this register in a pixel shader.
- This is a read-only register.

4.2 OBJECT_INDEX

Enum: IL_REGTYPE_OBJECT_INDEX

Text Syntax: vObjIndex

Common Name: Object Index Register

Number of Components per Register: 1

Description:

- Vertex shader import for the ordered index for the current vertex processed (ordered vertex shader instance).
- Pixel shader import for the ordered index for the current pixel processed (ordered pixel shader instance).
- This value starts at 0 and is incremented for each successive pixel/vertex.
- The first component of this register contains the current vertex processed.
- This register cannot be used with relative addressing.
- This is a read-only register.
- It is an error to use the second, third, and fourth components register.

4.3 BARYCENTRIC_COORD

Enum: IL_REGTYPE_BARYCENTRIC_COORD

Text Syntax: vBaryCoord

Common Name: HOS Barycentric Coordinates

Number of Components per Register: 4

Description:

- This register type is valid only for HOS rendering.

- For non-HOS rendering, this register type is invalid and its contents are undefined.
- For HOS rendering, this is a vertex shader import for the barycentric coordinates of the current, tessellated vertex.
- This register cannot be used with relative addressing.
- This is a read-only register.
- It is an error to use this register in a pixel shader.

4.4 PRIMITIVE_INDEX

Enum: IL_REGTYPE_PRIMITIVE_INDEX

Text Syntax: vPrimIndex

Common Name: HOS Primitive Index

Number of Components per Register: 4

Description:

- This register type is valid only for HOS rendering.
- For normal, non-HOS rendering, this register type is invalid, and its contents are undefined.
- For HOS rendering, this is the incremental index of current tessellation primitive generated by the tessellation engine for certain types of HOS rendering.
- This register cannot be used with relative addressing.
- This is a read-only register.
- It is an error to use this register in a pixel shader.

4.5 QUAD_INDEX

Enum: IL_REGTYPE_QUAD_INDEX

Text Syntax: vQuadIndex

Common Name: HOS QUAD Index

Number of Components per Register: 4

Description:

- This register type is valid only for HOS rendering.
- For normal, non-HOS rendering, this register type is invalid and its contents are undefined.
- For HOS rendering, the first component is a quad index generated by the tessellation engine for certain types of HOS rendering.

- This register cannot be used with relative addressing.
- This is a read-only register. It is an error to use this register in a pixel shader.

4.6 SPRITE

Enum: IL_REGTYPE_SPRITE

Text Syntax: oSprite

Common Name: Sprite Register

Number of Components per Register: 1

Description:

- Vertex shader export for point size.
- The first component contains the point size.
- This is a write-only register. This register cannot be the source of an instruction.
- This register cannot be used with relative addressing.
- It is an error to use this register if a VOUTPUT register is used. You can achieve similar functionality by using the DCLVOUT instruction on a VOUTPUT register and declaring its usage as IL_IMPORTUSAGE_POINTSIZE.
- It is an error to use this register in a pixel shader.
- The second, third and fourth components of this register are undefined.

4.7 POS

Enum: IL_REGTYPE_POS

Text Syntax: oPos

Common Name: Export Position Register

Number of Components per Register: 4

Description:

- Vertex shader export for position data.
- The position of the vertex in clip space.
- This is a write-only register. This register cannot be the source of an instruction.
- It is an error to use this register if a VOUTPUT register is used. You can achieve similar functionality by using the DCLVOUT instruction on a VOUTPUT register and declaring its usage as IL_IMPORTUSAGE_POS.
- This register cannot be used with relative addressing.

- It is an error to use this register in a pixel shader.

4.8 PINPUT

Enum: IL_REGTYPE_PINPUT

Text Syntax: vPixIn#

Common Name: Pixel Shader Input Register

Number of Components per Register: 4

Number Per Shader: 16

- Pixel shader input data.
- It is an error to use this register in a vertex shader.
- These are read and write registers.
- This register can be used with loop-relative addressing as a source only.
- A DCLPIN or DCLPP instruction must be issued on a register of this type before it is used.
- It is an error to use this register if an INTERP, TEXCOORD, PRICOLOR, SECCOLOR, or FOG register is used.

4.9 VOUTPUT and OUTPUT

Enum: IL_REGTYPE_VOUTPUT

Text Syntax: oVtxOut#

Common Name: Vertex Shader Output Register

Number of Components per Register: 4

Number Per Shader: 18

Description:

- Vertex shader output data.
- These are write-only registers. They cannot be the source of an instruction.
- It is an error to use this register in a pixel shader.
- This register can be used with loop-relative addressing as a destination only.
- A DCLVOUT instruction must be issued on a register of this type before it is used.
- It is an error to use this register if a POS, SPRITE, INTERP, TEXCOORD, PRICOLOR, SECCOLOR, or FOG register is used.

4.10 LITERAL

Enum: IL_REGTYPE_LITERAL

Text Syntax: l# where # is the literal register number.

Common Name: Literal Constant Register

Number of Components per Register: 4

Description:

This four-components, constant, typeless, read-only register can be used in place of a GPR. The format of this register can be either integer, floating point, or four-byte hex value.

4.11 INTERP

Enum: IL_REGTYPE_INTERP

Text Syntax (VS): oInterp#

Text Syntax (PS): vInterp#

Common Name: General-Purpose Interpolator Register

Number of Components per Register: 4

Description:

- General-purpose vertex shader export and pixel shader import for interpolated data.
- Perspective correct interpolation is performed on the values of these registers when passed from the vertex shader to the pixel shader.
- A DCLPI instruction must be issued on this shader type before is it used in any other instruction.
- It is an error to use a register of this type in a vertex shader if a VOUTPUT register is used. You can achieve similar functionality by using the DCLVOUT instruction on an VOUTPUT register and declaring its usage as IL_IMPORTUSAGE_GENERIC.
- It is an error to use a register of this type in a pixel shader if a PINPUT register is used. You can achieve similar functionality by using the DCLPIN instruction on a PINPUT register and declaring its usage as IL_IMPORTUSAGE_GENERIC.
- This register can be used only with loop relative addressing.
- In a vertex shader, these are write-only registers. They cannot be the source of an instruction.
- In a pixel shader, these are read-only registers. They cannot be the destination of an instruction.

4.12 TEXCOORD

Enum: IL_REGTYPE_TEXCOORD

Text Syntax (VS): oT#

Text Syntax (PS): vT#

Common Name: Texture Coordinate Interpolator Register

Number of Components per Register: 4

Description:

- Vertex shader export and pixel shader import interpolated for texture coordinate data.
- Perspective correct interpolation is performed on the values of these registers when passed from the vertex shader to the pixel shader.
- A DCLPI instruction must be issued on this shader type before is it used in any other instruction.
- It is an error to use a register of this type in a vertex shader if a VOUTPUT register is used. You can achieve similar functionality by using the DCLVOUT instruction on a VOUTPUT register and declaring its usage as IL_IMPORTUSAGE_GENERIC.
- It is an error to use a register of this type in a pixel shader if a PINPUT register is used. You can achieve similar functionality by using the DCLPIN instruction on a PINPUT register and declaring its usage as IL_IMPORTUSAGE_GENERIC.
- In a vertex shader, these are write-only registers. They cannot be the source of an instruction.
- In a pixel shader, these are read and write registers.
- In a vertex shader, this register can be used with loop-relative addressing as a destination only.
- In a pixel shader, this register can be used with loop-relative addressing as a source only.

4.13 PRICOLOR

Enum: IL_REGTYPE_PRICOLOR

Text Syntax (VS): oPriColor#

Text Syntax (PS): vPriColor#

Common Name: Primary Color Interpolator Register

Number of Components per Register: 4

Description:

TEXCOORD

Copyright © 2009 by Advanced Micro Devices, Inc. All rights reserved.

- Vertex shader export and pixel shader import for interpolated primary color data.
- By convention register number 0 represents the front-facing color, while register number 1 represents the back-facing color.
- When flat shading is used (AS_SHADE_MODE is set to FLAT), no interpolation is performed on the values of these registers when passed from the vertex shader to the pixel shader. Instead, only the value of the provoking vertex is passed to the pixel shader. When smooth shading is used (AS_SHADE_MODE is set to SMOOTH), perspective correct interpolation is performed on the values.
- A DCLPI instruction must be issued on this shader type before is it used in any other instruction.
- This register cannot be used with relative addressing.
- It is an error to use a register of this type in a vertex shader if a VOUTPUT register is used. You can achieve similar functionality by using the DCLVOUT instruction on a VOUTPUT register and declaring its usage as IL_IMPORTUSAGE_COLOR or IL_IMPORTUSAGE_BACKCOLOR.
- It is an error to use a register of this type in a pixel shader if a PINPUT register is used. You can achieve similar functionality by using the DCLPIN instruction on a PINPUT register and declaring its usage as IL_IMPORTUSAGE_COLOR or IL_IMPORTUSAGE_BACKCOLOR.
- In a vertex shader, these are write-only registers. They cannot be the source of an instruction.
- In a pixel shader, these are read and write registers.

4.14 SECCOLOR

Enum: IL_REGTYPE_SECCOLOR

Text Syntax (VS): oSecColor#

Text Syntax (PS): vSecColor#

Common Name: Secondary Color Interpolator Register

Number of Components per Register: 4

Description:

- Vertex shader export and pixel shader import for interpolated secondary color data.
- By convention, register number 0 represents the front-facing color, while register number 1 represents the back-facing color.
- When flat shading is used (AS_SHADE_MODE is set to FLAT), no interpolation is performed on the values of these registers when passed from the vertex shader to the pixel shader. Instead, only the value of the provoking vertex is passed to the pixel shader. When smooth shading is used

(AS_SHADE_MODE is set to SMOOTH), perspective correct interpolation is performed on the values.

- This register cannot be used with relative addressing.
- A DCLPI instruction must be issued on this shader type before is it used in any other instruction.
- It is an error to use a register of this type in a vertex shader if a VOUTPUT register is used. You can achieve similar functionality by using the DCLVOUT instruction on a VOUTPUT register and declaring its usage as IL_IMPORTUSAGE_COLOR or IL_IMPORTUSAGE_BACKCOLOR.
- It is an error to use a register of this type in a pixel shader if a PINPUT register is used. You can achieve similar functionality by using the DCLPIN instruction on a PINPUT register and declaring its usage as IL_IMPORTUSAGE_COLOR or IL_IMPORTUSAGE_BACKCOLOR.
- In a vertex shader, these are write-only registers. They cannot be the source of an instruction.
- In a pixel shader, these are read and write registers.

4.15 PSOUTFOG

Enum: IL_REGTYPE_PS_OUT_FOG

Text Syntax: oPsFog

Common Name: Pixel Shader Fog Output

Number of Components per Register: 1

Description:

- Pixel shader output for fog factor.
- The first component contains a fog factor.
- The second, third, and fourth components of this register are ignored.
- This is a write-only register.
- This register cannot be used with relative addressing.
- It is an error to use this register in a vertex or geometry shader.

4.16 FOG

Enum: IL_REGTYPE_FOG

Text Syntax (VS): $\circ F_{\text{fog}}$

Text Syntax (PS): $\vee F_{\text{fog}}$

Common Name: Fog Interpolator Register

Number of Components per Register: 1

Description:

- Vertex shader export and pixel shader import for interpolated fog data.
- This is a scalar register where the value is contained in the first component.
- In a vertex shader, the second, third, and fourth components must be masked (cannot be written to).
- In a pixel shader the second, third, and fourth components are undefined.
- Perspective correct interpolation is performed on the values of this registers when passed from the vertex shader to the pixel shader.
- Can only use a total of 16 of INTERP, TEXCOORD, PRICOLOR, SECCOLOR, and FOG registers in a single shader.
- A DCLPI instruction must be issued on this shader type before is it used in any other instruction.
- This register cannot be used with relative addressing.
- It is an error to use a register of this type in a vertex shader if a VOUTPUT register is used. You can achieve similar functionality by using the DCLVOUT instruction on a VOUTPUT register and declaring its usage as IL_IMPORTUSAGE_FOG.
- It is an error to use a register of this type in a pixel shader if a PINPUT register is used. You can achieve similar functionality by using the DCLPIN instruction on a PINPUT register and declaring its usage as IL_IMPORTUSAGE_FOG.
- In a vertex shader, this is a write-only register. It cannot be the source of an instruction.
- In a pixel shader, this is a read-only register. It cannot be the destination of an instruction.

4.17 SPRITECOORD

Enum: IL_REGTYPE_SPRITECOORD

Text Syntax: vSpriteCoord

Common Name: Sprite Texture Coordinate Register

Number of Components per Register: 2

Description:

- Pixel shader input for sprite texture coordinate.
- The first and second components contain the pixel's S and T coordinate for the point primitive rendered.
- The values of this register are undefined if the primitive rendered is not a point.
- The third and fourth components of this register are undefined.
- This is a read-only register. It cannot be the destination of an instruction.
- This register cannot be used with relative addressing.
- It is an error to use this register in a vertex shader.

4.18 PRIMCOORD

Enum: IL_REGTYPE_PRIMCOORD

Text Syntax: vPrimCoord

Number of Components per Register: 2

Description:

- This is a graphics-only feature.
- Pixel shader input for point-aa or line-aa texture coordinates.
- SpriteCoord cannot be used in any shader that uses primcoord.
- The first and second components contain the pixel's S and T coordinate for the point/line primitive rendered.
- The values of this register are undefined if the primitive rendered is not a point or a line.
- The third and fourth components of this register are undefined.
- This is a read-only register. It cannot be the destination of an instruction.
- This register cannot be used with relative addressing.
- It is an error to use this register in a vertex or geometry shader.

4.19 PRIMTYPE

Enum: IL_REGTYPE_PRIMTYPE

Text Syntax: vPrimType

Number of Components per Register: 2

Description:

- This is a graphics only feature.
- It is a pixel shader input for a primitive type.
- The first component has sign bit = 1, this is a point. Values in other bits are undefined.
- The second component has sign bit = 1, this is a line. Values in other bits are undefined.
- The third and fourth components of this register are undefined.
- This is a read-only register. It cannot be the destination of an instruction.
- This register cannot be used with relative addressing.
- It is an error to use this register in a vertex or geometry shader.

4.20 vWinCoord

Enum: IL_REGTYPE_WINCOORD

Text Syntax: vWinCoord.xy

Common Name: Window Coordinate Register

Number of Components per Register: 2

Description:

- Pixel shader import for screen position data.
- The first and second components are the X and Y position of the pixel in the domain of execution. The third component is the Z coordinate of the pixel in window space.
- The fourth component is W.
- A DCLPI instruction must be issued on this shader type before is it used in any other instruction.
- The DCLPI instruction specifies whether the X and Y coordinate is relative to the lower-left or upper-left corner of the window and whether it represents the center or upper-left corner of the pixel.

This read-only register cannot be the destination of an instruction. It is an error to use this register in a vertex shader. This register cannot be used with relative

addressing. The second, third, and fourth components of this register are undefined. This register is not valid in compute shaders.

4.21 FACE

Enum: IL_REGTYPE_FACE

Text Syntax: v_{Face}

Common Name: Face Register

Number of Components per Register: 1

Description:

- Pixel shader import for primitive facing.
- The x component is negative if the pixel is the back-face of the primitive. The x component is positive if the pixel is the front-face of the primitive.
- Point and Line primitives are always front-facing. Points and lines rendered as a result of polygons using point or line fill mode inherit the facing of the polygon.
- This is a read-only register. It cannot be the destination of an instruction.
- It is an error to use this register in a vertex shader.
- It is an error to use this register in a real-time pixel shader.
- This register cannot be used with relative addressing.
- The second, third, and fourth components of this register are undefined.

4.22 PCOLOR

Enum: IL_REGTYPE_PCOLOR

Text Syntax: $oC\#$

Common Name: Pixel Color Register

Number of Components per Register: 4

Description:

- Pixel shader export for color data.
- The register number corresponds to the color buffer to which data is output.
- These are write-only registers. They cannot be the source of an instruction.
- It is an error to use this register in a vertex shader.
- This register cannot be used with relative addressing.

4.23 DEPTH

Enum: IL_REGTYPE_DEPTH

Text Syntax: oDepth

Common Name: Pixel Depth Register

Number of Components per Register: 1

Description:

- Pixel shader export for depth data.
- This is a scalar register where the depth values is contained in the first component.
- This is a write-only register. It cannot be the source of an instruction.
- It is an error to use this register in a vertex shader.
- This register cannot be used with relative addressing.
- The second, third, and fourth components of this register are unused and undefined.

4.24 Mask

Enum: IL_REGTYPE_OMASK

Text Syntax: oMask

Common Name: Pixel Coverage Mask

Number of Components per Register: 1

Description:

When the pixel shader runs at sample-frequency, the coverage mask is ANDed with a mask that selects the sample currently being processed. As a result, sample N is always masked by bit N of oMask. This allows a shader to run at either sample-frequency or pixel-frequency with identical oMask behavior. The same rule applies to Alpha and to Coverage when the shader runs at sample-frequency.

- This is a scalar register where the mask value is contained in the first component.
- Values assigned to oMask are treated as integer.
- This is a write-only register. It cannot be the source of an instruction.
- It is an error to use this register in a vertex shader.
- This register cannot be used with relative addressing.
- The second, third, and fourth components of this register are unused and undefined.

4.25 Stencil

Enum: IL_REGTYPE_STENCIL

Text Syntax: oSTENCIL

Common Name: STENCIL ref value

Number of Components per Register: 1

Description:

- This is a scalar register where the stencil value is contained in the first component.
- Values assigned to Stencil are treated as integer.
- This is a write-only register. It cannot be the source of an instruction.
- It is an error to use this register in a vertex shader.
- This register cannot be used with relative addressing.
- The second, third, and fourth components of this register are unused and undefined.

4.26 Global

Enum: IL_REGTYPE_GLOBAL

Text Syntax: g[address]

Note that the address must be a single-component integer.

Common Name: Global Memory Register

Number of Components per Register: 4

Description:

This is a read/write register that can be used to address global memory in the text form of the shader. Global variables are indicated by a `g[address]`.

For example:

```
add g[2].x, r4, g[4]
```

reads address 4, adds r4 to the contents of that address, and scatters the result to address 2. Each address corresponds to a 128-bit, four dword location. Most address locations are indexed.

Global g registers can be indexed using the temp (r) registers. For example,

```
add g [r5.x].x r4, g[r6.y]
```

4.27 Shared Temp

Enum: IL_REGTYPE_SHARED_TEMP

Text Syntax: sr# (where # is any shared register number)

Common Name: Shared General-Purpose Register

Number of Components per Register: 4

Description:

This read/write register is shared by all wavefronts running on a SIMD. Only absolute address mode is allowed; for example: sr2. It is used only in compute shaders.

For example:

```
iadd sr2.xy_w, sr2.xyzw, r0.xxxx
```

4.28 Thread_ID

Enum: IL_REGTYPE_THREAD_ID

Text Syntax: vTid

Common Name: Input Register for a Thread ID Within a Block

Number of Components per Register: 1

Description:

This read-only input register contains the thread ID within a thread block. It is used only in compute shaders as an index or in integer operations.

For example:

```
mov r1.x___, vTid.x
mov g[vTid.x], r1.x
```


4.29 Absolute Thread_ID

Enum: IL_REGTYPE_ABSOLUTE_THREAD_ID

Text Syntax: vaTid

Common Name: Input Register for Absolute Thread ID

Number of Components per Register: 1

Description:

This read-only input register contains an absolute thread ID. It is used only in compute shaders. It is used as an index or in integer operations.

For example:

```
mov r1.x____, vaTid.x
mov g[vaTid.x], r1.x
```

4.30 Thread_Group_ID

Enum: IL_REGTYPE_THREAD_GROUP_ID

Text Syntax: vTGroupid

Common Name: Input Register for Thread Group ID

Number of Components per Register: 1

Description:

This is a read-only input register that contains the thread group ID. It is used only in compute shaders. It is used as an index or in integer operations.

For example:

```
mov r1.x____, vTGroupid.x
mov g[vTGroupid.x], r1.x
```

4.31 Generic_Memory

Enum: IL_REGTYPE_GENERIC_MEM

Text Syntax: mem

Common Name: Generic Memory Type

Number of Components per Register: 4

Description:

This register provides a mask or swizzle. It is used by instructions like `write_lds_vec`, which does not have a dst register but requires a dst mask.

For example:

```
write_lds_vec mem.x_z_, r0.xyzw
```

Chapter 5

Enumerated Types

This chapter lists and briefly describes the ATI IL enumerated types.

5.1 ILShader

Table 5.1 ILShader Enumeration Types

Enumeration	Description
IL_SHADER_PIXEL	This code describes a kernel that uses a thread creation pattern optimized for graphics (pixel kernel).
IL_SHADER_VERTEX	This code describes a vertex shader.
IL_SHADER_GEOMETRY	This code describes a geometry shader.

5.2 ILLanguageType

Table 5.2 ILLanguageType Enumeration Types

Enumeration	Description
IL_LANG_DX10_GS	
IL_LANG_DX10_PS	
IL_LANG_DX10_VS	
IL_LANG_DX8_PS	DX 1.x shader (DX 8)
IL_LANG_DX8_VS	DX vertex shader (DX 8)
IL_LANG_DX9_PS	
IL_LANG_DX9_VS	
IL_LANG_GENERIC	
IL_LANG_OPENGL	

5.3 ILRegType

See Chapter 4, "Register Types," for more information.

Table 5.3 ILRegType Enumeration Types

Register Type	Description
IL_REGTYPE_OBJECT_INDEX	
IL_REGTYPE_ADDR	
IL_REGTYPE_BARYCENTRIC_COORD	
IL_REGTYPE_CLIP	
IL_REGTYPE_CONST_BOOL	
IL_REGTYPE_CONST_BUFF	A four component element of a constant buffer.
IL_REGTYPE_CONST_FLOAT	
IL_REGTYPE_CONST_INT	
IL_REGTYPE_DEPTH	
IL_REGTYPE_FACE	
IL_REGTYPE_FOG	
IL_REGTYPE_INDEX	
IL_REGTYPE_INTERP	
IL_REGTYPE_LITERAL	A four component register containing explicit values.
IL_REGTYPE_PCOLOR	
IL_REGTYPE_PINPUT	
IL_REGTYPE_POS	
IL_REGTYPE_PRICOLOR	
IL_REGTYPE_PRIMITIVE_INDEX	
IL_REGTYPE_QUAD_INDEX	
IL_REGTYPE_SECCOLOR	
IL_REGTYPE_SPRITE	
IL_REGTYPE_SPRITECOORD	
IL_REGTYPE_TEMP	A four component virtual register.
IL_REGTYPE_TEXCOORD	
IL_REGTYPE_VERTEX	
IL_REGTYPE_VOUTPUT	
IL_REGTYPE_VPRIM	
IL_REGTYPE_WINCOORD	A Windows coordinate register.

5.4 ILTopologyType

Table 5.4 IL_TOPOLOGY Enumeration Types (Input to a Geometry Shader)

Enumeration	Description
IL_TOPOLOGY_LINE	
IL_TOPOLOGY_LINE_ADJ	
IL_TOPOLOGY_POINT	
IL_TOPOLOGY_TRIANGLE	
IL_TOPOLOGY_TRIANGLE_ADJ	

Table 5.5 IL_OUTPUT_TOPOLOGY Enumeration Types (Output from a Geometry Shader)

Enumeration	Description
IL_OUTPUT_TOPOLOGY_LINESTRIP	
IL_OUTPUT_TOPOLOGY_POINTLIST	
IL_OUTPUT_TOPOLOGY_TRIANGLE_STRIP	

5.5 ILMatrix

Table 5.6 ILMatrix Enumeration Types

Enumeration	IL Text
IL_MATRIX_3X2	_matrix(3x2)
IL_MATRIX_3X3	_matrix(3x3)
IL_MATRIX_3X4	_matrix(3x4)
IL_MATRIX_4X3	_matrix(4x3)
IL_MATRIX_4X4	_matrix(4x4)

5.6 ILComponentSelect

See Section 2.2.7, "Source Modifier Token," page 2-6 for usage. IL Text details for component selection can be found in Chapter 3, "Text Instruction Syntax."

Table 5.7 ILComponentSelect Enumeration Types

Enumeration	Description
IL_COMPSEL_0	Force this component to 0.0.
IL_COMPSEL_1	Force this component to 1.0.
IL_COMPSEL_W_A	Select the fourth component (w/alpha) for the component.
IL_COMPSEL_X_R	Select the first component (x/red) for the component.
IL_COMPSEL_Y_G	Select the second component (y/green) for the component.
IL_COMPSEL_Z_B	Select the third component (z/blue) for the component.

5.7 IModDstComp

See Section 3.5, “Destination Modifiers,” page 3-3 for usage. IL Text details for write mask can be found in Section 3.6, “Write Mask,” page 3-3.

Table 5.8 IModDstComp Enumeration Types

Enumeration	Description
IL_MODCOMP_0	Write 0.0 to this component.
IL_MODCOMP_1	Write 1.0 to this component.
IL_MODCOMP_NOWRITE	Do not write to this component. The contents of this component of the destination register are not changed.
IL_MODCOMP_WRITE	Write the result of the instruction to this component.

5.8 IImportUsage

See the DCLVOUT and DCLPIN instructions for more information:

Table 5.9 IImportUsage Enumeration Types

Enumeration	Text Syntax	Description
IL_IMPORTUSAGE_BACKCOLOR	<code>_usage(backcolor)</code>	<p>When used to declare a vertex shader output, this usage indicates the register contains a color value to be interpolated across the primitive and passed to the pixel shader once the shader terminates. The processed value can be read from the PINPUT register declared with this usage and the matching <i>usageIndex</i>.</p> <p>When used to declare a pixel shader input, this usage indicates the register contains a color value that has been interpolated across the primitive. The value originates from the VOUTPUT register in the vertex shader declared with this <i>usage</i> and matching <i>usageIndex</i>.</p> <p>Hardware can use lower-precision interpolators for colors.</p>
IL_IMPORTUSAGE_COLOR	<code>_usage(color)</code>	<p>When used to declare a vertex shader output, this usage indicates the register contains a color value to be interpolated across the primitive and passed to the pixel shader once the shader terminates. The processed value can be read from the PINPUT register declared with this usage and the matching <i>usageIndex</i>.</p> <p>When used to declare a pixel shader input, this usage indicates the register contains a color value that has been interpolated across the primitive. The value originates from the VOUTPUT register in the vertex shader declared with this <i>usage</i> and matching <i>usageIndex</i>.</p> <p>Hardware can use lower-precision interpolators for colors.</p>

Table 5.9 ILImportUsage Enumeration Types (Cont.)

Enumeration	Text Syntax	Description
IL_IMPORTUSAGE_FOG	<code>_usage(fog)</code>	<p>When used to declare a vertex shader output, this usage indicates that the x component of the register has the vertices' fog value at shader termination. The value is interpolated across the primitive and passed to the pixel shader. The processed value can be read in the pixel shader from the PINPUT register declared with this usage.</p> <p>When used to declare a pixel shader input, indicates that the x component of the register contains the fog value at shader execution. The value of the VOUTPUT register in the vertex shader declared with the DCLVOUT instruction as having this usage is interpolated across the primitive and passed to the PINPUT register declared with this usage.</p> <p>For any register declared with this usage, the <i>usageIndex</i> must be set to 0.</p>
IL_IMPORTUSAGE_GENERIC	<code>_usage(generic)</code>	<p>When used to declare a vertex shader output, this usage indicates the register contains a generic value to be interpolated across the primitive and passed to the pixel shader once the shader terminates. The processed value can be read from the PINPUT register declared with this usage and the matching <i>usageIndex</i>.</p> <p>When used to declare a pixel shader input, this usage indicates the register contains a generic value that has been interpolated across the primitive. The value originates from the VOUTPUT register in the vertex shader declared with this usage and matching <i>usageIndex</i>.</p>
IL_IMPORTUSAGE_POINTSIZE	<code>_usage(pointsize)</code>	<p>When used to declare a vertex shader output, this usage indicates the x component of the register contains the vertices' point size when the shader terminates. <i>usageIndex</i> must be zero when this usage is set. Only VOUTPUT register 1 can have this usage.</p> <p>This usage cannot be used in a pixel shader input.</p>
IL_IMPORTUSAGE_POS	<code>_usage(pos)</code>	<p>When used to declare a vertex shader output, this usage indicates the register contains the vertices' position when the shader terminates. <i>usageIndex</i> must be zero when this usage is set. Only VOUTPUT register 0 can have this usage.</p> <p>This usage cannot be used in a pixel shader input.</p>
IL_IMPORTUSAGE_WINCOORD	<code>_usage(wincoord)</code>	<p>Can only be used in a pixel shader. The x, y, z, w values correspond to the screen position. This has been added to IL2.0 for OpenGL support.</p>

5.9 ILImportComponent

See the “DCLPI”, “DCLPIN”, “DCLV”, and “DCLVOUT” instructions for usage.

Table 5.10 ILImportComponent Enumeration Types

Enumeration	Text Syntax	Description
IL_IMPORTSEL_DEFAULT0	<code>_<i>comp</i>>(0)</code> where <i>comp</i> is x, y, z, or w	<p>This component is enabled and can be used in a vertex shader. If this register or this component of this register is not written to, the component defaults to 0.0 when used as a source or when the shader terminates. In a pixel shader, if this register or this component of this register is not exported in the vertex shader, the component is set to 0.0.</p> <p>If used in with the DCLVOUT instruction in the vertex shader, this component will default to 0.0 if it is not written to. The component is considered to be exported in the vertex shader in this case; thus, any component in a pixel shader mapped to this component will be set to 0.0 regardless of its default value set by the DCLPIN instruction.</p>
IL_IMPORTSEL_DEFAULT1	<code>_<i>comp</i>>(1)</code> where <i>comp</i> is x, y, z, or w	<p>This component is enabled and can be used. In a vertex shader, if this register or this component of this register is not written to, the component defaults to 1.0 when used as a source or when the shader terminates. In a pixel shader, if this register or this component of this register is not exported in the vertex shader, set the component to 1.0.</p> <p>If used with the DCLVOUT instruction in the vertex shader, this component defaults to 1.0 if it's not written to. The component is considered to be exported in the vertex shader in this case; thus, any component in a pixel shader mapped to this component is set to 1.0, regardless of its default value set by the DCLPIN instruction.</p>
IL_IMPORTSEL_UNDEFINED	<code>_<i>comp</i>>(*)</code> where <i>comp</i> is x, y, z, or w	<p>This component is enabled and can be used. If this register or this component of the register is not exported in the vertex shader, the value of the component is undefined (the value of the component does not matter).</p>
IL_IMPORTSEL_UNUSED	<code>_<i>comp</i>>(-)</code> where <i>comp</i> is x, y, z, or w Example: <code>dclpi_z(-)_w(-)</code>	<p>This component is disabled and cannot be used in the shader. It is an error to reference the component in the shader.</p>

5.10 ILDefaultVal

See the DCLDEF instruction for usage.

Table 5.11 ILDefaultVal Enumeration Types

Enumeration	Text Syntax	Description
IL_DEFVAL_0	<code>_<i>comp</i>>(0)</code> where <i>comp</i> is <i>x</i> , <i>y</i> , <i>z</i> , or <i>w</i>	Indicates that the default value for this component of the register type is 0.0.
IL_DEFVAL_1	<code>_<i>comp</i>>(1)</code> where <i>comp</i> is <i>x</i> , <i>y</i> , <i>z</i> , or <i>w</i>	Indicates that the default value for this component of the register type is 1.0.
IL_DEFVAL_NONE	<code>_<i>comp</i>>(*)</code> where <i>comp</i> is <i>x</i> , <i>y</i> , <i>z</i> , or <i>w</i> Example: <code>dcldef_z(*)_w(*)</code>	Indicates there is no default value for this component of the register type.

5.11 ILShiftScale

See Section 3.5, "Destination Modifiers," page 3-3 for usage.

Table 5.12 ILShiftScale Enumeration Types

Enumeration	Text Syntax	Description
IL_SHIFT_D2	<code>_d2</code>	Shift value right by 1 bit (divide by 2).
IL_SHIFT_D4	<code>_d4</code>	Shift value right by 2 bits (divide by 4).
IL_SHIFT_D8	<code>_d8</code>	Shift value right by 3 bits (divide by 8).
IL_SHIFT_NONE		Do not shift.
IL_SHIFT_X2	<code>_x2</code>	Shift value left by 1 bit (multiply by 2).
IL_SHIFT_X4	<code>_x4</code>	Shift value left by 2 bits (multiply by 4).
IL_SHIFT_X8	<code>_x8</code>	Shift value left by 3 bits (multiply by 8).

5.12 ILDivComp

See Section 3.7, “Source Modifiers,” page 3-4 for usage.

Table 5.13 ILDivComp Enumeration Types

Enumeration	Text Syntax	Description
IL_DIVCOMP_NONE	<code>_divcomp(<i>none</i>)</code>	None
IL_DIVCOMP_UNKNOWN	<code>_divcomp(<i>unknown</i>)</code>	The component used to divide is not known at shader create time. The component is determined by a state setting at shader run time. This value can only be used in a TEXLD, TEXLDD, TEXLDB, TEXLDMS, TEXWEIGHT, PROJECT, or LOD instruction.
IL_DIVCOMP_W	<code>_divcomp(<i>w</i>)</code>	Divide the first, second, and third components by the fourth component. If the fourth component is 0.0, the first, second, and third component become \pm FLT_MAX.
IL_DIVCOMP_Y	<code>_divcomp(<i>y</i>)</code>	Divide the first component by the second component. If the second component is 0.0, the first component becomes \pm FLT_MAX.
IL_DIVCOMP_Z	<code>_divcomp(<i>z</i>)</code>	Divide the first and second components by the third component. If the third component is 0.0, the first and second component become \pm FLT_MAX.

5.13 ILRelOp

See IFC, CONTINUEC, BREAKC, CMP, and SET for usage.

Table 5.14 ILRelOp Enumeration Types

Enumeration	Text Syntax	Description
IL_RELOP_EQ	<code>_relop(<i>eq</i>)</code>	Equal.
IL_RELOP_GE	<code>_relop(<i>ge</i>)</code>	Greater than or equal.
IL_RELOP_GT	<code>_relop(<i>gt</i>)</code>	Greater than.
IL_RELOP_LE	<code>_relop(<i>le</i>)</code>	Less than or equal.
IL_RELOP_LT	<code>_relop(<i>lt</i>)</code>	Less than.
IL_RELOP_NE	<code>_relop(<i>ne</i>)</code>	Not equal.

5.14 ILLogicOp

Can be used in IF_LOGICAL*, CONTINUE_LOGICAL*, and BREAK_LOGICAL*.

Table 5.15 ILLogicOp Enumeration Types

Enumeration	Text Syntax	Description
IL_LOG_EQ	<code>_log(<i>eq</i>)</code>	Equal (32-bit integer compare)
IL_LOG_NE	<code>_log(<i>ne</i>)</code>	Not equal (32-bit integer compare)

5.15 ILZeroOp

See the RSQ, RCP, LOG, LOGP, LN, NRM, and DIV instructions for more information:.

Table 5.16 ILZeroOp Enumeration Types

Enumeration	Text Syntax	Description
IL_ZEROOP_0	<code>_zeroop(<i>zero</i>)</code>	Return 0.0 when the instruction operates on 0.0. Cannot be used with LOG, LOGP, or LN instructions.
IL_ZEROOP_FLTMAX	<code>_zeroop(<i>fltmax</i>)</code>	Return the max floating point value when the instruction operates on 0.0.
IL_ZEROOP_INF_ELSE_MAX	<code>_zeroop(<i>inf_else_max</i>)</code>	Implementation-dependant: Return IEEE infinity if the implementation can support it; otherwise, use FLT_MAX when the instruction operates on 0.0.
IL_ZEROOP_INFINITY	<code>_zeroop(<i>infinity</i>)</code>	Return IEEE Infinity when the instruction operates on 0.0.

5.16 ILCmpValue

See the CMP Instruction for usage.

Table 5.17 ILCmpValue Enumeration Types

Enumeration	Text Syntax	Description
IL_CMPVAL_0_0	<code>_cmpval(<i>0.0</i>)</code>	Compare vs. 0.0.
IL_CMPVAL_0_5	<code>_cmpval(<i>0.5</i>)</code>	Compare vs. 0.5.
IL_CMPVAL_1_0	<code>_cmpval(<i>1.0</i>)</code>	Compare vs. 1.0.
IL_CMPVAL_NEG_0_5	<code>_cmpval(<i>-0.5</i>)</code>	Compare vs. -0.5.
IL_CMPVAL_NEG_1_0	<code>_cmpval(<i>-1.0</i>)</code>	Compare vs. -1.0.

5.17 ILTexCoordMode

See DCLPT instruction for more information:

Table 5.18 ILTexCoordMode Enumeration Types

Enumeration	Text Syntax	Description
IL_TEXCOORDMODE_NORMALIZED	<code>_coordmode(normalized)</code>	The texture coordinates given in the texture load instructions are non-parametric. (the coordinate range [0.0-1.0] spans the entire texture)
IL_TEXCOORDMODE_UNKNOWN	<code>_coordmode(unknown)</code>	At shader create time, it is not known if the texture coordinates given in the texture load instructions are normalized. Instead, this is determined at shader run time based on a state value.
IL_TEXCOORDMODE_UNNORMALIZED	<code>_coordmode(unnormalized)</code>	The texture coordinates given in the texture load instructions are parametric. (the coordinate range [0.0, dimension of the texture] spans the entire dimension of the texture)

5.18 ILPixTexUsage

There are a maximum of eight values. See DCLPT instruction for more information.

Table 5.19 ILPixTexUsage Enumeration Types

Enumeration	Text Syntax
IL_USAGE_PIXTEX_1D	<code>_type(1d)</code>
IL_USAGE_PIXTEX_2D	<code>_type(2d)</code>
IL_USAGE_PIXTEX_2DMS_ARRAY	<code>_type(2dms_array)</code>
IL_USAGE_PIXTEX_2DMSAA	<code>_type(2dmsaa)</code>
IL_USAGE_PIXTEX_3D	<code>_type(3d)</code>
IL_USAGE_PIXTEX_CUBEMAP	<code>_type(cubemap)</code>
IL_USAGE_PIXTEX_CUBEMAPARRAY	<code>_type(cubemaparray)</code>
IL_USAGE_PIXTEX_UNKNOWN	<code>_type(unknown)</code>

5.19 ILTexShadowMode

See the TEXLD (page 6-67), TEXLDB (page 6-70), and TEXLDD (page 6-74) instructions for more information.

Table 5.20 ILTexShadowMode Enumeration Types

Enumeration	Text Value	Description
IL_TEXSHADOWMODE_NEVER	<code>_shadowmode(<i>never</i>)</code>	Never do a shadow map comparison.
IL_TEXSHADOWMODE_UNKNOWN	<code>_shadowmode(<i>unknown</i>)</code>	Do a compare vs. the third component of the texture coordinate.
IL_TEXSHADOWMODE_Z	<code>_shadowmode(<i>z</i>)</code>	Always do a compare vs. the third component of the texture coordinate.

5.20 ILTexFilterMode

See the TEXLD, TEXLDB, and TEXLDD instructions for more information.

Table 5.21 ILTexFilterMode Enumeration Types

Enumeration	Text Syntax	Description
IL_TEXFILTER_ANISO	<code>_<i><filter></i>(<i>aniso</i>)</code> where <i><filter></i> is min, mag, volmag, or volmin	Always use anisotropic filtering if aniso filtering is enabled in state or based on the aniso parameter in the instruction.
IL_TEXFILTER_LINEAR	<code>_<i><filter></i>(<i>linear</i>)</code> where <i><filter></i> is min, mag, volmag, or volmin	Always use linear filtering.
IL_TEXFILTER_POINT	<code>_<i><filter></i>(<i>point</i>)</code> where <i><filter></i> is min, mag, volmag, or volmin	Always use point filtering.
IL_TEXFILTER_UNKNOWN	<code>_<i><filter></i>(<i>unknown</i>)</code> where <i><filter></i> is min, mag, volmag, or volmin	Use the filtering mode specified by state.

5.21 ILAnisoFilterMode

See the TEXLD, TEXLDB, and TEXLDD instructions for more information.

Table 5.22 ILAnisoFilterMode Enumeration Types

Enumeration	Text Syntax	Description
IL_ANISOFILTER_DISABLED	<code>_aniso(<i>disabled</i>)</code>	Never use anisotropic filtering.
IL_ANISOFILTER_MAX_1_TO_1	<code>_aniso(<i>1</i>)</code>	Use a maximum of 1 sample for anisotropic filtering.
IL_ANISOFILTER_MAX_16_TO_1	<code>_aniso(<i>16</i>)</code>	Use a maximum of 16 samples for anisotropic filtering.
IL_ANISOFILTER_MAX_2_TO_1	<code>_aniso(<i>2</i>)</code>	Use a maximum of 2 samples for anisotropic filtering.
IL_ANISOFILTER_MAX_4_TO_1	<code>_aniso(<i>4</i>)</code>	Use a maximum of 4 samples for anisotropic filtering.
IL_ANISOFILTER_MAX_8_TO_1	<code>_aniso(<i>8</i>)</code>	Use a maximum of 8 samples for anisotropic filtering.
IL_ANISOFILTER_UNKNOWN	<code>_aniso(<i>unknown</i>)</code>	Use the anisotropic filtering mode specified by state.

5.22 ILMipFilterMode

See the TEXLD, TEXLDB, and TEXLDD instructions for more information.

Table 5.23 ILMipFilterMode Enumeration Types

Enumeration	Text Syntax	Description
IL_MIPFILTER_BASE	<code>_mip(base)</code>	Always sample the base mipmap only.
IL_MIPFILTER_LINEAR	<code>_mip(linear)</code>	Always use linear filtering across mipmaps.
IL_MIPFILTER_POINT	<code>_mip(point)</code>	Always sample the nearest mipmap only.
IL_MIPFILTER_UNKNOWN	<code>_mip(unknown)</code>	Use the filtering mode specified by state.

5.23 ILNoiseType

Table 5.24 ILNoiseType Enumeration Types

Enumeration	Text Syntax	Description
IL_NOISETYPE_PERLIN1D	<code>_type(perlin1D)</code>	Compute 1D Perlin noise function
IL_NOISETYPE_PERLIN2D	<code>_type(perlin2D)</code>	Compute 2D Perlin noise function
IL_NOISETYPE_PERLIN3D	<code>_type(perlin3D)</code>	Compute 3D Perlin noise function
IL_NOISETYPE_PERLIN4D	<code>_type(perlin4D)</code>	Compute 4D Perlin noise function

5.24 ILInterpolation

Table 5.25 ILInterpolation Enumeration Types

Enumeration	Text Syntax
IL_INTERPMODE_CONSTANT	<code>_interp(constant)</code>
IL_INTERPMODE_LINEAR	<code>_interp(linear)</code>
IL_INTERPMODE_LINEAR_CENTROID	<code>_interp(centroid)</code>
IL_INTERPMODE_LINEAR_NOPERSPECTIVE	<code>_interp(noperspective)</code>
IL_INTERPMODE_LINEAR_NOPERSPECTIVE_CENTROID	<code>_interp(noper_centroid)</code>
IL_INTERPMODE_LINEAR_NOPERSPECTIVE_SAMPLE	<code>_interp(noper_sample)</code>
IL_INTERPMODE_LINEAR_SAMPLE	<code>_interp(sample)</code>
IL_INTERPMODE_NOTUSED = 0	<code>_interp(notused)</code>

5.25 ILAddressing

See IL_Dst and IL_Src for more information.

Table 5.26 ILAddressing Enumeration Types

Enumeration	Description
IL_ADDR_ABSOLUTE (no = 0)	Absolute addressing is used.
IL_ADDR_REG_RELATIVE	Register-relative addressing is used.
IL_ADDR_RELATIVE	Relative addressing is used.

5.26 IElementFormat

Table 5.27 IElementFormat Enumeration Types

Enumeration	Description
IL_ELEMENTFORMAT_FLOAT	Element uses a single-precision, floating point format.
IL_ELEMENTFORMAT_MIXED	Element uses more than one format.
IL_ELEMENTFORMAT_SINT	Element uses a signed, integer format.
IL_ELEMENTFORMAT_SNORM	Element uses a signed, normalized format.
IL_ELEMENTFORMAT_SRGB	Element uses an integer format in the sRGB color space.
IL_ELEMENTFORMAT_UINT	Element uses an unsigned, integer format.
IL_ELEMENTFORMAT_UNKNOWN	Element uses an unknown or user-defined format.
IL_ELEMENTFORMAT_UNORM	Element uses an unsigned, normalized format.

5.27 ILOpCode

Table 5.28 ILOpCode Enumeration Types

Enumeration	Description
IL_OP_ABS	Computes the absolute value of each element of a vector.
IL_DCL_CONST_BUFFER	Declares a constant buffer.
IL_DCL_INDEXED_TEMP_ARRAY	Declares an indexed array.
IL_DCL_INPUT	Declares an input register.
IL_DCL_INPUT_PRIMITIVE	Used to declare the input primitive type.
IL_DCL_LITERAL	Declares a literal value.
IL_DCL_MAX_OUTPUT_VERTEX_COUNT	Used to declare the maximum number of vertices that will be emitted by a shader
IL_DCL_ODEPTH	Used to declare that the pixel shader intends to write to its scalar output oDepth register.
IL_DCL_OUTPUT	Declares an output register.
IL_DCL_OUTPUT_TOPOLOGY	Used to declare the output topology of a primitive.
IL_DCL_PERSIST	Used to declare the amount of persistent storage used by a shader.
IL_DCL_RESOURCE	Declares an input buffer.
IL_DCL_VPRIM	Used to declare an input register.
IL_OP_ACOS	Used to compute the inverse cosine of a component.
IL_OP_ADD	Computes the single-precision, floating point addition of the values in each element of a vector to the values in the corresponding element in another vector.
IL_OP_AND	Performs a logical-AND on two vectors.
IL_OP_ASIN	Used to compute the inverse sine of a component.
IL_OP_ATAN	Used to compute the inverse tangent of a component.
IL_OP_BREAK	Unconditionally terminates a LOOP or SWITCH block.
IL_OP_BREAK_LOGICALNZ	Conditionally terminates a LOOP or SWITCH block if the parameter is not zero.

Table 5.28 ILOpCode Enumeration Types (Cont.)

Enumeration	Description
IL_OP_BREAK_LOGICALZ	Conditionally terminates a LOOP or SWITCH block if the parameter is zero.
IL_OP_BREAKC	Conditionally terminates a LOOP or SWITCH block.
IL_OP_CALL	Unconditionally calls a FUNC block.
IL_OP_CALL_LOGICALNZ	Calls a FUNC block if the parameter is not zero.
IL_OP_CALL_LOGICALZ	Calls a FUNC block if the parameter is zero.
IL_OP_CALLNZ	Used to call a FUNC block if the contents of the Constant Boolean register are not zero.
IL_OP_CASE	Compares <case-value> to <switch-value>, and conditionally executes the CASE instruction block if they are equal.
IL_OP_CLAMP	
IL_OP_CLG	Computes the ceiling of the value in each element of a vector.
IL_OP_CMOV	Performs a single-precision, floating point comparison of the value in each element of a vector to 0.0f, and conditionally moves the value in each element of another vector to the corresponding element of a third vector if the comparison evaluates FALSE.
IL_OP_CMOV_LOGICAL	Performs an integer comparison of the value in each element of a vector to zero, and moves the value in each element of a second vector to the corresponding element of a third vector if the comparison evaluates TRUE; otherwise, moves the corresponding element of a fourth vector to the corresponding element of the third vector.
IL_OP_CMP	Performs a logical comparison of the value in each element of a comparison value, and conditionally moves the value in each element of another vector to the corresponding element of a third vector if the comparison evaluates TRUE; otherwise, the value in the corresponding element of a fourth vector is moved.
IL_OP_COLORCLAMP	
IL_OP_COMMENT	
IL_OP_CONTINUE	Unconditionally continues execution at the next LOOP or WHILELOOP instruction.
IL_OP_CONTINUE_LOGICALNZ	Conditionally continues execution at the next LOOP or WHILELOOP instruction if the parameter is not zero.
IL_OP_CONTINUE_LOGICALZ	Conditionally continues execution at the next LOOP or WHILELOOP instruction if the parameter is zero.
IL_OP_CONTINUEC	Conditionally continues execution at the next LOOP or WHILELOOP instruction.
IL_OP_COS	Computes the cosine of a component.
IL_OP_COS_VEC	Computes the cosine of each element of a vector.
IL_OP_CRS	Used to compute the cross product of two 3-vectors.
IL_OP_CUT	Used to close a primitive topology and open a new one.
IL_OP_D_2_F	Converts a double-precision float value to a single-precision float value.
IL_OP_D_ADD	Performs a double-precision, floating point addition of the values in each element of a vector to the values in the corresponding element in another vector.
IL_OP_D_EQ	Performs a double-precision float equality comparison.

Table 5.28 ILOpCode Enumeration Types (Cont.)

Enumeration	Description
IL_OP_D_FRAC	Returns a double-precision, floating point fraction (mantissa).
IL_OP_D_FREXP	Splits a double-precision, floating point value into fraction (mantissa) and exponent values.
IL_OP_D_GE	Performs a double-precision float greater than or equal comparison.
IL_OP_D_LDEXP	Combines fraction (mantissa) and exponent values into a double-precision, floating point value.
IL_OP_D_LT	Performs a double-precision float less than comparison.
IL_OP_D_MAD	Performs a double-precision, floating point multiplication of two values, then performs a double-precision addition on the result with a third value.
IL_OP_D_MUL	Performs a double-precision, floating point multiplication of two values.
IL_OP_D_NE	Performs a double-precision float inequality comparison.
IL_OP_DCLARRAY	Used to declare a range of registers to be included in an array.
IL_OP_DCLDEF	Not used.
IL_OP_DCLPI	Used to declare the interpolator properties.
IL_OP_DCLPIN	Used to declare an input register.
IL_OP_DCLPP	Not used.
IL_OP_DCLPT	Used to declare a texture's properties.
IL_OP_DCLV	Used to declare the mapping for a vertex shader's inputs.
IL_OP_DCLVOUT	Used to declare the output register for a vertex shader.
IL_OP_DEF	Used to declare the constant integer or float value for a register.
IL_OP_DEFAULT	Starts the default instruction block within a SWITCH instruction block.
IL_OP_DEFB	Used to declare the constant Boolean value for a register.
IL_OP_DET	Used to calculate the determinant of a 4x4 matrix.
IL_OP_DISCARD_LOGICALNZ	Logically compares a parameter to zero, and conditionally stops execution and discards the results of a kernel invocation if the parameter is not zero. <u>When used in the compute kernel, this instruction produces undefined results.</u>
IL_OP_DISCARD_LOGICALZ	Logically compares a parameter to zero, and conditionally stops execution and discards the results of a kernel invocation if the parameter is zero. <u>When used in the compute kernel, this instruction produces undefined results.</u>
IL_OP_DIST	Used to compute the vector distance between two 3-vectors.
IL_OP_DIV	Performs a single-precision, floating point division of the values in each element of a vector by the values in the corresponding element in another vector.
IL_OP_DP2	Computes the dot product of two 2-vectors.
IL_OP_DP2ADD	Used to compute the dot product of two 2-vectors and scalar adds a 32-bit value to the result.
IL_OP_DP3	Computes the dot product of two 3-vectors.
IL_OP_DP4	Computes the dot product of two 4-vectors.
IL_OP_DST	

Table 5.28 ILOpCode Enumeration Types (Cont.)

Enumeration	Description
IL_OP_DSX	Used to compute the single-precision, floating point rate of change of a vector in the x-direction.
IL_OP_DSY	Used to compute the single-precision, floating point rate of change of a vector in the y-direction.
IL_OP_DXSINCOS	
IL_OP_ELSE	Starts the ELSE clause within an IF instruction block.
IL_OP_EMIT	Used to generate a primitive.
IL_OP_EMIT_THEN_CUT	Used to generate and close a primitive then start a new primitive.
IL_OP_END	Indicates the end of an IL stream.
IL_OP_ENDFUNC	Indicates the end of a FUNC instruction block.
IL_OP_ENDIF	Indicates the end of an IF or ELSE instruction block.
IL_OP_ENDLOOP	Indicates the end of a LOOP or WHILELOOP instruction block.
IL_OP_ENDMAIN	Indicates the end of a main program in a kernel.
IL_OP_ENDSWITCH	Indicates the end of a SWITCH instruction block.
IL_OP_EQ	Performs a float equality comparison.
IL_OP_EXN	Used to compute the single-precision, floating point value of e raised to a power.
IL_OP_EXP	Computes the single-precision, floating point value of 2 raised to a power.
IL_OP_EXP_VEC	Computes the single-precision, floating point value of 2 raised to a power for each element in a vector.
IL_OP_EXPP	Used to compute the partial-precision, floating point value of 2 raised to a power.
IL_OP_F_2_D	Converts a single-precision float value to a double-precision float value.
IL_OP_FACEFORWARD	
IL_OP_FLR	Computes the floor of the value in each element of a vector.
IL_OP_FRC	Returns a single-precision, floating point fraction (mantissa).
IL_OP_FTOI	Converts a single-precision float value to an integer value.
IL_OP_FTOU	Converts a single-precision float value to an unsigned integer value.
IL_OP_FUNC	Indicates the start of a FUNC instruction block.
IL_OP_FWIDTH	
IL_OP_GE	Performs a float greater than or equal comparison.
IL_OP_I_ADD	Computes the integer addition of the values in each element of a vector to the values in the corresponding element in another vector.
IL_OP_I_EQ	Performs an integer equality comparison.
IL_OP_I_GE	Performs an integer greater than or equal comparison.
IL_OP_I_LT	Performs an integer less than comparison.
IL_OP_I_MAD	Performs an integer multiplication of the values in each element of a vector by the values in the corresponding elements of another vector; it then adds the value in the corresponding elements of a third vector to the results.

Table 5.28 ILOpCode Enumeration Types (Cont.)

Enumeration	Description
IL_OP_I_MAX	Performs an integer comparison of the value in each element of a vector to the value in the corresponding element of another vector, and returns the larger of the two values in the corresponding element of a third vector.
IL_OP_I_MIN	Performs an integer comparison of the value in each element of a vector to the value in the corresponding element of another vector, and returns the smaller of the two values in the corresponding element of a third vector.
IL_OP_I_MUL	Performs an integer multiplication of the values in each element of a vector by the values in the corresponding element in another vector, and returns the lower 32-bits of the result.
IL_OP_I_MUL_HIGH	Performs an integer multiplication of the values in each element of a vector by the values in the corresponding element in another vector, and returns the upper 32-bits of the result.
IL_OP_I_NE	Performs an integer inequality comparison.
IL_OP_I_NEGATE	Computes the two's complement negation of the values in each element in a vector.
IL_OP_I_NOT	Performs a bit-wise one's complement on a vector.
IL_OP_I_OR	Performs a logical-OR on two vectors.
IL_OP_I_SHL	Shifts integer values in each element of a vector the specified number of bits to the left.
IL_OP_I_SHR	Shifts the integer values in each element of a vector a the specified number of bits to the right through sign extension.
IL_OP_I_XOR	Performs a logical-XOR on two vectors.
IL_OP_IF_LOGICALNZ	Logically compares a parameter to zero, and executes the IF instruction block if the comparison evaluates FALSE.
IL_OP_IF_LOGICALZ	Logically compares a parameter to zero, and executes the IF instruction block if the comparison evaluates TRUE.
IL_OP_IFC	Logically compares two parameters, and conditionally executes the IF instruction block if the comparison evaluates TRUE.
IL_OP_IFNZ	Used to conditionally execute the IF instruction block if the parameter is not zero.
IL_OP_INITV	
IL_OP_ITOF	Converts an integer value to a single-precision float value.
IL_OP_KILL	Used to perform a float comparison on a parameter and conditionally stop execution and discard the results of a pixel shader invocation if any element in the parameter is less than 0.0f.
IL_OP_LEN	Used to compute the length of a 3-vector.
IL_OP_LIT	
IL_OP_LN	Computes the single-precision, floating point natural logarithm of a component.
IL_OP_LOAD	Used to fetch data from a specified Buffer or Texture without filtering.
IL_OP_LOD	
IL_OP_LOG	Computes the single-precision, floating point binary logarithm of a component.

Table 5.28 ILOpCode Enumeration Types (Cont.)

Enumeration	Description
IL_OP_LOG_VEC	Computes the single-precision, floating point binary logarithm of the values in each element of a vector.
IL_OP_LOGP	Used to compute the partial-precision, floating point binary logarithm of a component.
IL_OP_LOOP	Indicates the start of a LOOP instruction block.
IL_OP_LRP	
IL_OP_LT	Performs a float less-than comparison.
IL_OP_MAD	Performs a single-precision, floating point multiplication of the values in each element of a vector by the values in the corresponding elements of another vector; it then adds the value in the corresponding elements of a third vector to the results.
IL_OP_MAX	Performs a single-precision, floating point comparison of the value in each element of a vector to the value in the corresponding element of another vector, and returns the larger of the two values in the corresponding element of a third vector.
IL_OP_MEMEXPORT	Used to export the value in <i>src1</i> to a memory stream.
IL_OP_MEMIMPORT	Used to import a value in memory to a TEMP register.
IL_OP_MIN	Performs a single-precision, floating point comparison of the value in each element of a vector to the value in the corresponding element of another vector, and returns the smaller of the two values in the corresponding element of a third vector.
IL_OP_MMUL	
IL_OP_MOD	Computes the single-precision, floating point division of the values in each element of a vector by the values in the corresponding elements of another vector, and returns the remainder of each division.
IL_OP_MOV	Unconditionally moves the value in each element of a vector to the corresponding element of another vector.
IL_OP_MOVA	Used to unconditionally move the rounded or truncated value in each element of a vector to the corresponding element of another vector.
IL_OP_MUL	Performs a single-precision, floating point multiplication of the values in each element of a vector by the values in the corresponding element in another vector, and returns the lower 32-bits of the result.
IL_OP_NE	Performs a float inequality comparison.
IL_OP_NOISE	
IL_OP_NOP	Used to pad IL streams so that they are the same length in memory without unnecessary computations.
IL_OP_NRM	Computes the single-precision, floating point, normalized value of the value in each element of a vector.
IL_OP_PIREDUCE	
IL_OP_POW	Computes the value in a component raised to the power of the value in another component.
IL_OP_PRECOMP	
IL_OP_PROJECT	
IL_OP_RCP	Computes the single-precision, floating point reciprocal of the value in a component.

Table 5.28 ILOpCode Enumeration Types (Cont.)

Enumeration	Description
IL_OP_REFLECT	
IL_OP_RESINFO	Used to get information about a resource.
IL_OP_RET	Used to unconditionally transfer control from the end of a FUNC instruction block back to the caller.
IL_OP_RET_DYN	Unconditionally transfers control from anywhere in a FUNC instruction block back to the caller.
IL_OP_RET_LOGICALNZ	Compares a parameter to zero, and conditionally transfers control from anywhere in a FUNC instruction block back to the caller if the parameter is not zero.
IL_OP_RET_LOGICALZ	Compares a parameter to zero, and conditionally transfers control from anywhere in a FUNC instruction block back to the caller if the parameter is zero.
IL_OP_RND	Rounds the single-precision, floating point value in each element of a vector to the nearest integer.
IL_OP_ROUND_NEAR	Rounds the single-precision, floating point value in each element of a vector to the nearest even integer.
IL_OP_ROUND_NEG_INF	Rounds the single-precision, floating point value in each element of a vector towards $-\infty$.
IL_OP_ROUND_POS_INF	Rounds the single-precision, floating point value in each element of a vector towards ∞ .
IL_OP_ROUND_ZERO	Rounds the single-precision, floating point value in each element of a vector towards zero.
IL_OP_RSQ	Computes the single-precision, floating point square root of the reciprocal of the value in a component.
IL_OP_RSQ_VEC	Computes the single-precision, floating point square root of the reciprocal of the value in each element of a vector.
IL_OP_SAMPLE	Performs a single-precision, floating point sample of a resource.
IL_OP_SAMPLE_B	Used to perform a single-precision, floating point sample of a resource with filtering and bias.
IL_OP_SAMPLE_C	Used to perform a single-precision, floating point sample of a resource with filtering and comparison.
IL_OP_SAMPLE_C_LZ	Used to perform a single-precision, floating point sample of a resource with filtering and comparison.
IL_OP_SAMPLE_G	Used to perform a single-precision, floating point sample of a resource with filtering and gradient.
IL_OP_SAMPLE_L	Performs a single-precision, floating point sample of a resource with filtering.
IL_OP_SET	Performs a single-precision, floating point comparison of the value in each element of a vector with the value in each element of another vector, and places a 1.0f in the corresponding element of a third vector if the comparison evaluates TRUE; otherwise, it places a 0.0f in the corresponding element of the third vector.
IL_OP_SGN	Computes the single-precision, floating point sign of the value in each element of a vector.
IL_OP_SIN	Computes the sine of the value in a component.
IL_OP_SIN_VEC	Computes the sine of the value in each element of a vector.
IL_OP_SINCOS	Used to compute the sine and cosine of a component.

Table 5.28 ILOpCode Enumeration Types (Cont.)

Enumeration	Description
IL_OP_SQRT	Computes the single-precision, floating point square root of the value in a component.
IL_OP_SQRT_VEC	Computes the single-precision, floating point square root of the value in each element of a vector.
IL_OP_SUB	Computes the single-precision, floating point subtraction of the values in each element of a vector from the values in the corresponding element in another vector.
IL_OP_SWITCH	Indicates the start of a SWITCH instruction block, and provides a <switch-value> to later CASE instructions.
IL_OP_TAN	Used to compute the tangent of a component.
IL_OP_TEXLD	
IL_OP_TEXLDB	
IL_OP_TEXLDD	
IL_OP_TEXLDMS	
IL_OP_TEXWEIGHT	
IL_OP_TRANSPOSE	
IL_OP_TRC	Performs a single-precision, floating point truncation of the value in each element of a vector.
IL_OP_U_DIV	Performs an unsigned integer division of the values in each element of a vector by the values in the corresponding element in another vector.
IL_OP_U_GE	Performs an unsigned integer greater than or equal comparison.
IL_OP_U_LT	Performs an unsigned integer less than comparison.
IL_OP_U_MAD	Performs an unsigned integer multiplication of the values in each element of a vector by the values in the corresponding elements of another vector; it then adds the value in the corresponding elements of a third vector to the results.
IL_OP_U_MAX	Performs an unsigned integer comparison of the value in each element of a vector to the value in the corresponding element of another vector; it then returns the larger of the two values in the corresponding element of a third vector.
IL_OP_U_MIN	Performs an unsigned integer comparison of the value in each element of a vector to the value in the corresponding element of another vector; it then returns the smaller of the two values in the corresponding element of a third vector.
IL_OP_U_MOD	Computes the unsigned integer division of the values in each element of a vector by the values in the corresponding elements of another vector; it then returns the remainder of each division.
IL_OP_U_MUL	Performs an unsigned integer multiplication of the values in each element of a vector by the values in the corresponding element in another vector; it then returns the lower 32-bits of the result.
IL_OP_U_MUL_HIGH	Performs an unsigned integer multiplication of the values in each element of a vector by the values in the corresponding element in another vector; it then returns the upper 32-bits of the result.

Table 5.28 ILOpCode Enumeration Types (Cont.)

Enumeration	Description
IL_OP_U_SHR	Shifts without sign extension unsigned integer values in each element of a vector the specified number of bits to the right.
IL_OP_UTOF	Converts an unsigned integer value to a single-precision float value.
IL_OP_WHILE	Indicates the start of a WHILELOOP instruction block.

Chapter 6

Instructions

An IL stream consists mainly of IL instruction packets. Each packet begins with an IL_OpCode token. The type and number of tokens that follow depends upon the value of the *code* field in the initial IL_OpCode token, as well as the *modifier_present* field in each of the following IL_Dst tokens and IL_Src tokens. This chapter describes each instruction packet, including its operation and usage within the IL stream. Some instructions are defined by pseudo code. In this chapter, the symbol $v[i]$ refers to source *i* post swizzle.

6.1 Formats

Most instructions use a standardized format. Rather than repeat the format in each instruction, the common format is given here. Tokens have the following order:

1. Opcode token.
2. Destination information (zero or one set of tokens: *dst* token, any relative addressing, any modifiers).
3. Source information (zero or more sets of tokens), *src_token*, any relative addressing information, any modifiers.

Formats include: zero or one destination, followed by any number of sources. See the specific opcode to determine the number of destinations/sources allowed.

6.2 Instruction Notes

6.2.1 Notes on Comparison Instructions

Comparison instructions support four comparison types: float, double, integer, and unsigned integer, with different instructions for each type. Instructions return either TRUE (comparison condition met) or FALSE (comparison condition not met). Integer comparison instructions return 0xFFFFFFFF (TRUE) and 0x00000000 (FALSE). Float comparison instructions can be configured to return 1.0f (TRUE) and 0.0f (FALSE).

Integer instructions use signed arithmetic when comparing operands; unsigned integer instructions use unsigned arithmetic. Float instructions use signed arithmetic when comparing operands; however, denorms are flushed before all float instructions (the original source registers remain untouched). Thus, +0 and

-0 are equivalent for float comparisons. Also, float instructions return FALSE when either operand holds NaN. See the IEEE 754 documentation for more information on floating point rules.

6.2.2 Notes on Flow Control Instructions

Flow control instructions determine the order in which instructions are executed by the hardware.

A *flow-control-block* is code within:

- A subroutine: code between FUNC and RET, or between FUNC and ENDFUNC.
- Between IFNZ and ENDIF.
- Between IFC and ENDIF.
- Between LOG_IF and ENDIF.
- Between ELSE and ENDIF.
- Between IFNZ and ELSE.
- Between IFC and ELSE.
- Between LOOP and ENDLOOP.
- Between WHILELOOP and ENDLOOP.

There are two forms of loop: LOOP/ENDLOOP used for DX9-style counted loops, and WHILELOOP/ENDLOOP used for DX10-style while loops.

The following are the restrictions on when particular control flow instructions are allowed.

- LOOPS and WHILELOOPS must terminate in the same *flow-control-block* in which they begin.
- END, ENDMAIN, and ENDFUNC cannot be placed within a *flow-control-block*.

6.2.3 Notes on Input/Output Instructions

Instructions in this section are expected to be used only by graphic clients. Consult the DX10 documentation or other graphics programming reference for details.

6.2.4 Notes on Conversion Instructions

Even though IL is an untyped language, conversion instructions are required to convert between different data formats so that source data for an instruction appears in the proper format. Without format conversion, instructions could provide incorrect results.

6.2.5 Notes on Double Precision Instructions

Double precision values are represented by a pair of registers. Outputs are either the pair *yx* or to the pair *wz*, where the msb is stored in *y/w*. For example:

Idata 3.0 => (0x4008000000000000) in register *r* looks like:

```
r.w = 0x40080000 ;high dword
r.z = 0x00000000 ;low  dword
```

Or by:

```
r.y = 10x40080000 ;high dword
r.x = 10x00000000 ;low  dword
```

All source double inputs must be in *xy* (after swizzle operations). For example:

```
d_add r1.xy, r2.xy, r2.xy
```

Or

```
d_add r1.zw, r2.xy, r2.xy
```

Each computes twice the value in *r2.xy*, and places the result in either *xy* or *zw*.

The user can set the output mask to either *xy* or *zw*. The msb is in *y/w*, so users can test the sign of the result with single precision operations.

All inputs are in the first two components *xy* of each source.

6.3 Flow Control Instructions

Unconditional BREAK out of Loop or Switch Constructs

<i>Instructions</i>	BREAK		
<i>Syntax</i>	BREAK		
<i>Description</i>	BREAK is used only in a loop or switch statement. It terminates the smallest enclosing loop or switch. Control passes to the statement following the terminated statement, if any.		
<i>Format</i>	0-input.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_BREAK
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	BREAKC, BREAK_LOGICALZ, BREAK_LOGICALNZ.		

Conditional BREAK Using Floating Point

<i>Instructions</i>	BREAKC		
<i>Syntax</i>	<code>breakc_relop(<i>op</i>) <i>src0</i>, <i>src1</i></code>		
<i>Description</i>	<p>BREAKC compares two registers using a float comparison on the x component of the sources. It is used only in a loop or switch statement. It terminates the smallest enclosing loop or switch. Control passes to the statement following the terminated statement, if any.</p> <p>Operation:</p> <pre>if (v0.x relop v1.x) { Break; }</pre>		
<i>Format</i>	2-input.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_BREAKC
	control	29:16	relop(<i>op</i>). See Table 5.14 on page 5-8.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	BREAK, BREAK_LOGICALZ, BREAK_LOGICALNZ.		

Conditional BREAK Using Integers

<i>Instructions</i>	BREAK_LOGICALNZ, BREAK_LOGICALZ		
<i>Syntax</i>	Opcode	Syntax	Description
	IL_OP_BREAK_LOGICALNZ	<code>break_logicalnz <i>src0</i></code>	If <i>src0.x</i> ≠ 0 break.
	IL_OP_BREAK_LOGICALZ	<code>break_logicalz <i>src0</i></code>	If <i>src0.x</i> == 0 break.
<i>Description</i>	<p>This form of Break compares a register component using integer compare to zero. The break is done if all bits of <i>src0</i> are zero. <i>src0</i> must have a swizzle that selects a single component. The break statement occurs only in a loop-<i>statement</i> or a switch statement. It causes termination of the smallest enclosing loop or switch statement; control passes to the statement following the terminated statement, if any.</p>		
<i>Format</i>	1-input.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	See <i>Syntax</i> , above.
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	BREAK, BREAKC.		

Unconditional CALL

Instructions **CALL**

Syntax `call <integer label>`

Description Call a subroutine. The subroutine is identified by a label id, corresponding to the ID in a function statement. Subroutines can nest up to 32 levels deep. Direct or indirect recursion is permitted. If there are already 32 entries on the return address stack, and a “call” is issued, the call is skipped.

<u>.Ordinal</u>	<u>Token</u>
1	IL_OpCode token with code set to IL_OP_CALL (control field ignored).
3	Unsigned integer representing the subroutine label.

Format 0-input.

<i>Opcode</i>	Token	Field Name	Bits	Description
	1	code	15:0	IL_OP_CALL
		control	29:16	Must be zero.
		sec_modifier_present	30	Must be zero.
		pri_modifier_present	31	Must be zero.
	2			Must be zero.
	3			Unsigned integer representing label of the subroutine.

Related CALLNZ, CALL_LOGICALZ, CALL_LOGICALNZ.

CALL if Boolean register is not zero

<i>Instructions</i>	CALLNZ																																					
<i>Syntax</i>	<code>callnz src0, <integer label></code>																																					
<i>Description</i>	<p>CALLNZ compares <code>src0.x</code> to zero using an integer comparison. If any bits of <code>src0.x</code> are not zero, CALLNZ pushes the address of the next instruction in the shader onto the return address stack and transfers control to the FUNC block identified by <code><integer label></code>. CALLs can be nested up to 32 levels deep. If all bits of <code>src0.x</code> are zero, or the return address stack already contains 32 addresses, the CALL is skipped and execution continues at the next instruction in the shader. Recursion is permitted either directly, using another CALL instruction, or indirectly using relative addressing.</p> <p><code>src0</code> must be a CONST_BOOL register. See Section 4 on page 4-1 for information on the register types.</p> <p>CALLNZ cannot be used with relative addressing or a source modifier. Both the <code>modifier_present</code> and <code>relative_address</code> fields of the IL_Src token must be zero.</p> <p>Operation:</p> <pre>if (v0.x != 0) { Call label }</pre>																																					
<i>Format</i>	1-input, 0-output.																																					
<i>Opcode</i>	<table> <thead> <tr> <th>Token</th> <th>Field Name</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td rowspan="4">1</td> <td>code</td> <td>15:0</td> <td>IL_OP_CALLNZ</td> </tr> <tr> <td>control</td> <td>29:16</td> <td>Must be zero.</td> </tr> <tr> <td>sec_modifier_present</td> <td>30</td> <td>Must be zero.</td> </tr> <tr> <td>pri_modifier_present</td> <td>31</td> <td>Must be zero.</td> </tr> <tr> <td rowspan="5">2</td> <td colspan="3">IL_Src token (<code>src0</code>) with <code>register_type</code> set to IL_REGTYPE_CONST_BOOL. The <code>modifier_present</code> and <code>relative_address</code> field must be set to 0.</td> </tr> <tr> <th>Field Name</th> <th>Bits</th> <th>Description</th> </tr> <tr> <td>register_num</td> <td>15:0</td> <td>Any available register.</td> </tr> <tr> <td>register_type</td> <td>21:16</td> <td>CONST_BOOL</td> </tr> <tr> <td>Remaining fields</td> <td>31:22</td> <td>Must be zero.</td> </tr> <tr> <td>3</td> <td colspan="3">Unsigned integer representing the label of the subroutine.</td> </tr> </tbody> </table>	Token	Field Name	Bits	Description	1	code	15:0	IL_OP_CALLNZ	control	29:16	Must be zero.	sec_modifier_present	30	Must be zero.	pri_modifier_present	31	Must be zero.	2	IL_Src token (<code>src0</code>) with <code>register_type</code> set to IL_REGTYPE_CONST_BOOL. The <code>modifier_present</code> and <code>relative_address</code> field must be set to 0.			Field Name	Bits	Description	register_num	15:0	Any available register.	register_type	21:16	CONST_BOOL	Remaining fields	31:22	Must be zero.	3	Unsigned integer representing the label of the subroutine.		
Token	Field Name	Bits	Description																																			
1	code	15:0	IL_OP_CALLNZ																																			
	control	29:16	Must be zero.																																			
	sec_modifier_present	30	Must be zero.																																			
	pri_modifier_present	31	Must be zero.																																			
2	IL_Src token (<code>src0</code>) with <code>register_type</code> set to IL_REGTYPE_CONST_BOOL. The <code>modifier_present</code> and <code>relative_address</code> field must be set to 0.																																					
	Field Name	Bits	Description																																			
	register_num	15:0	Any available register.																																			
	register_type	21:16	CONST_BOOL																																			
	Remaining fields	31:22	Must be zero.																																			
3	Unsigned integer representing the label of the subroutine.																																					
<i>Related</i>	CALL, CALL_LOGICALZ, CALL_LOGICALNZ.																																					

Conditional CALL Based on Boolean*Instructions* **CALL_LOGICALZ**

<i>Syntax</i>	Opcode	Syntax	Description
	IL_OP_CALL_LOGICALZ	<code>call_logicalnz src0 <integer label></code>	If <code>src0.x == 0</code> , call <integer label>.

Description Performs conditional call based on a Boolean constant. `Scr0` must be a `CONST_BOOL` register. This instruction does not allow the use of a source modifier. The `modifier_present` field of the `IL_Src` token must be set to 0. This instruction does not allow relative addressing. The `relative_address` field of the `IL_Src` token must be set to 0.

All four forms call a subroutine. The subroutine is identified by a label id, corresponding to the ID in a function statement. Subroutines can nest up to 32 levels deep. Recursion is permitted, directly or indirectly. If there are already 32 entries on the return address stack and a "call" is issued, the call is skipped.

<u>.Ordinal</u>	<u>Token</u>
1	IL_OpCode token with code set to IL_OP_CALL_LOGICALZ (control field ignored).
2	(<code>src0</code>) representing any valid IL Source (unspecified number of tokens).
3	Unsigned integer following all source tokens representing the subroutine label.

Operation:

```
if (v0.x eq 0) {
    Call subroutine;
}
```

Format 1-input.

<i>Opcode</i>	Token	Field Name	Bits	Description
	1	code	15:0	See <i>Syntax</i> , above.
		control	29:16	Must be zero.
		sec_modifier_present	30	Must be zero.
		pri_modifier_present	31	Must be zero.
	2	IL_Src token and IL_Src_Mod token (if required) representing any valid IL Source. See Section 2.2.6, "Source Token," page 2-4.		
	3	Unsigned integer representing the label of the subroutine.		

Related CALL, CALLNZ.

Conditional CALL Based on Integer

Instructions **CALL_LOGICALNZ**

Syntax **Opcode** **Syntax** **Description**

IL_OP_CALL_LOGICALNZ *call_logicalnz, src0* If *src0.x* ≠ 0, call <integer label>.

Description Conditional call subroutine at the label. The 32-bit register component (*src0*) is tested. If any bits are not zero, the instruction performs the call. It checks only the x component (after swizzling) of *src0* is checked. All four forms call a subroutine. The subroutine is identified by a label id, corresponding to the id in a function statement. Subroutines can nest up to 32 levels deep. Recursion is permitted, directly or indirectly. If there are already 32 entries on the return address stack, and a “call” is issued, the call is skipped.

Operation:

```
if (v0.x ne 0) {
    Call subroutine;
}
```

Format 1-input.

<i>Opcode</i>	Token	Field Name	Bits	Description
	1	code	15:0	See <i>Syntax</i> , above.
		control	29:16	Must be zero.
		sec_modifier_present	30	Must be zero.
		pri_modifier_present	31	Must be zero.
	2	IL_Src token and IL_Src_Mod token (if required) representing any valid IL Source. See Section 2.2.6, “Source Token,” page 2-4.		
	3	Unsigned integer representing the label of the subroutine.		

Related CALL, CALLNZ.

CASE Statement

<i>Instructions</i>	CASE															
<i>Syntax</i>	case case-value															
<i>Description</i>	This case statement is used in a switch. This instruction is followed by a 32 integer that identifies the case. Compares are done using integer arithmetic. Falling through cases is valid, as in C. This can be used to implement a DX10 case instruction. Operation: see SWITCH.															
<i>Format</i>	1-input, 0-output.															
<i>Opcode</i>	<table> <thead> <tr> <th>Field Name</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>code</td> <td>15:0</td> <td>IL_OP_CASE</td> </tr> <tr> <td>control</td> <td>29:16</td> <td>Must be zero.</td> </tr> <tr> <td>sec_modifier_present</td> <td>30</td> <td>Must be zero.</td> </tr> <tr> <td>pri_modifier_present</td> <td>31</td> <td>Must be zero.</td> </tr> </tbody> </table>	Field Name	Bits	Description	code	15:0	IL_OP_CASE	control	29:16	Must be zero.	sec_modifier_present	30	Must be zero.	pri_modifier_present	31	Must be zero.
Field Name	Bits	Description														
code	15:0	IL_OP_CASE														
control	29:16	Must be zero.														
sec_modifier_present	30	Must be zero.														
pri_modifier_present	31	Must be zero.														
<i>Related</i>	DEFAULT, ENDSWITCH, SWITCH.															

Unconditional CONTINUE

<i>Instructions</i>	CONTINUE															
<i>Syntax</i>	continue															
<i>Description</i>	CONTINUE causes shader execution to continue at the previous LOOP or WHILELOOP instruction. This is used only in a LOOP - ENDLOOP instruction block. Operation: If this is a counted loop { LoopIterationCount = LoopIterationCount - 1; LoopCounter = LoopCounter + LoopStep; LoopCounter = (LoopCounter > 0) ? LoopCounter : 0; if (LoopIterationCount > 0) Continue execution at the StartLoopOffset; } else { Continue execution at the StartLoopOffset; }															
<i>Format</i>	0-input.															
<i>Opcode</i>	<table> <thead> <tr> <th>Field Name</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>code</td> <td>15:0</td> <td>IL_OP_CONTINUE</td> </tr> <tr> <td>control</td> <td>29:16</td> <td>Must be zero.</td> </tr> <tr> <td>sec_modifier_present</td> <td>30</td> <td>Must be zero.</td> </tr> <tr> <td>pri_modifier_present</td> <td>31</td> <td>Must be zero.</td> </tr> </tbody> </table>	Field Name	Bits	Description	code	15:0	IL_OP_CONTINUE	control	29:16	Must be zero.	sec_modifier_present	30	Must be zero.	pri_modifier_present	31	Must be zero.
Field Name	Bits	Description														
code	15:0	IL_OP_CONTINUE														
control	29:16	Must be zero.														
sec_modifier_present	30	Must be zero.														
pri_modifier_present	31	Must be zero.														
<i>Related</i>	CONTINUEC, CONTINUE_LOGICALZ, CONTINUE_LOGICALNZ.															

CONTINUE Using Floating Point

<i>Instructions</i>	CONTINUEC		
<i>Syntax</i>	continuec_relop(<i>op</i>) <i>src0</i> , <i>src1</i>		
<i>Description</i>	<p>Conditionally continues execution at previous LOOP or WHILELOOP instruction if the condition is true. Can only be used in a LOOP - ENDLOOP instruction block.</p> <p>Operation:</p> <pre>if (v0.x relop v1.x) { Continue; }</pre>		
<i>Format</i>	2-input.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_CONTINUEC
	control	29:16	relop(<i>op</i>). See Table 5.14 on page 5-8 .
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	CONTINUE, CONTINUE_LOGICALZ, CONTINUE_LOGICALNZ.		

Conditional CONTINUE Using Integers

<i>Instructions</i>	CONTINUE_LOGICALNZ, CONTINUE_LOGICALZ		
<i>Syntax</i>	Opcode	Syntax	Description
	IL_OP_CONTINUE_LOGICALNZ	continue_logicalnz <i>src0</i>	If <i>src0.x</i> ≠ 0 call <integer label>.
	IL_OP_CONTINUE_LOGICALZ	continue_logicalnz <i>src0</i>	If <i>src0.x</i> == 0 call <integer label>.
<i>Description</i>	<p>Conditionally continues execution at the beginning of the current loop: CONTINUE_LOGICAL_Z continues the loop if all bits of <i>src0.x</i> are zero. CONTINUE_LOGICAL_NE continues the loop if any bit of <i>src0.x</i> is not zero.</p> <p>Can only be within a LOOP - ENDLOOP switch block.</p>		
<i>Format</i>	1-input, 0-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	See <i>Syntax</i> , above.
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	CONTINUE, CONTINUEC.		

DEFAULT Statement

<i>Instructions</i>	DEFAULT															
<i>Syntax</i>	default															
<i>Description</i>	<p>DEFAULT starts an instruction block within a SWITCH instruction block (see page 6-19). Unlike a CASE label, a DEFAULT label does not provide a value for comparison.</p> <p>This is like the default in C. Falling through or into a DEFAULT section is valid. There can be only one DEFAULT statement in each SWITCH block.</p>															
<i>Format</i>	0-input.															
<i>Opcode</i>	<table> <thead> <tr> <th>Field Name</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>code</td> <td>15:0</td> <td>IL_OP_DEFAULT</td> </tr> <tr> <td>control</td> <td>29:16</td> <td>Must be zero.</td> </tr> <tr> <td>sec_modifier_present</td> <td>30</td> <td>Must be zero.</td> </tr> <tr> <td>pri_modifier_present</td> <td>31</td> <td>Must be zero.</td> </tr> </tbody> </table>	Field Name	Bits	Description	code	15:0	IL_OP_DEFAULT	control	29:16	Must be zero.	sec_modifier_present	30	Must be zero.	pri_modifier_present	31	Must be zero.
Field Name	Bits	Description														
code	15:0	IL_OP_DEFAULT														
control	29:16	Must be zero.														
sec_modifier_present	30	Must be zero.														
pri_modifier_present	31	Must be zero.														
<i>Related</i>	CASE, ENDSWITCH, SWITCH.															

ELSE

<i>Instructions</i>	ELSE															
<i>Syntax</i>	else															
<i>Description</i>	<p>This instruction is the start of the ELSE clause of an IFNZ-ELSE-ENDIF or IFC-ELSE-ENDIF or LOG_IF-ELSE-ENDIF block. ELSE must be after an IFC, IFNZ, or LOG_IF instruction in the stream.</p>															
<i>Format</i>	0-input.															
<i>Opcode</i>	<table> <thead> <tr> <th>Field Name</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>code</td> <td>15:0</td> <td>IL_OP_ELSE</td> </tr> <tr> <td>control</td> <td>29:16</td> <td>Must be zero.</td> </tr> <tr> <td>sec_modifier_present</td> <td>30</td> <td>Must be zero.</td> </tr> <tr> <td>pri_modifier_present</td> <td>31</td> <td>Must be zero.</td> </tr> </tbody> </table>	Field Name	Bits	Description	code	15:0	IL_OP_ELSE	control	29:16	Must be zero.	sec_modifier_present	30	Must be zero.	pri_modifier_present	31	Must be zero.
Field Name	Bits	Description														
code	15:0	IL_OP_ELSE														
control	29:16	Must be zero.														
sec_modifier_present	30	Must be zero.														
pri_modifier_present	31	Must be zero.														
<i>Related</i>	IFNZ, IFC, ENDIF.															

ENDSWITCH

<i>Instructions</i>	ENDSWITCH		
<i>Syntax</i>	endswitch		
<i>Description</i>	ENDSWITCH indicates the end of a SWITCH instruction block.		
<i>Format</i>	0-input.		
<i>Opcodes</i>	Field Name	Bits	Description
	code	15:0	IL_OP_ENDSWITCH
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	CASE, DEFAULT, SWITCH.		

End of Main Program

<i>Instructions</i>	ENDMAIN		
<i>Syntax</i>	endmain		
<i>Description</i>	<p>ENDMAIN indicates the end of the shader source for the main program. Only FUNC instruction blocks and the END statement (see page 6-13) can follow an ENDMAIN statement. An ENDMAIN statement is only required if the shader contains subroutines.</p> <p>Operation:</p> <p>End of shader execution, and beginning of subroutine definitions.</p>		
<i>Format</i>	0-input.		
<i>Opcodes</i>	Field Name	Bits	Description
	code	15:0	IL_OP_ENDMAIN
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	None.		

End of Stream

<i>Instructions</i>	END		
<i>Syntax</i>	end		
<i>Description</i>	END indicates the end of an IL stream and must be the last statement in the stream. All shader programming, including subroutines, must be placed before this instruction.		
<i>Format</i>	0-input.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_END
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	None.		

End of FUNC (subroutine)

<i>Instructions</i>	ENDFUNC		
<i>Syntax</i>	endfunc		
<i>Description</i>	ENDFUNC indicates the end of a shader subroutine. Only FUNC blocks and the END statement can follow an ENDFUNC statement. This instruction is required only if the shader contains subroutines.		
<i>Format</i>	0-input.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_ENDFUNC
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	FUNC.		

End of IF block

<i>Instructions</i>	ENDIF		
<i>Syntax</i>	endif		
<i>Description</i>	Indicates the end of an IFNZ - ENDIF, IFC - ENDIF, IFNZ - ELSE - ENDIF, or IFC - ELSE - ENDIF block. The ENDIF statement must follow an ELSE, IFC, IFNZ or LOG_IF instruction.		
<i>Format</i>	0-input, 0-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_ENDIF
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	ELSE, IFC, IFNZ, IF_LOGICAL*.		

End of a LOOP or WHILELOOP block

<i>Instructions</i>	ENDLOOP		
<i>Syntax</i>	endloop		
<i>Description</i>	<p>Indicates the end of a LOOP - ENDLOOP or WHILELOOP - ENDLOOP instruction block. The instruction must follow a LOOP/WHILELOOP instruction. This instruction cannot be with a FUNC - RET, IFNZ - ENDIF, IFC - ENDIF, IFNZ - ELSE, IFC - ELSE, or ELSE - ENDIF block unless its corresponding LOOP or WHILELOOP is also in that block.</p> <p>Operation:</p> <p>(For counted loops)</p> <pre> LoopIterationCount = LoopIterationCount - 1; LoopCounter = LoopCounter + LoopStep; LoopCounter = (LoopCounter > 0) ? LoopCounter : 0; if (LoopIterationCount > 0) Continue execution at the StartLoopOffset; </pre> <p>For While loops</p> <p>Transfer control back to the test of the while loop.</p>		
<i>Format</i>	0-input.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_ENDLOOP
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	LOOP, WHILELOOP.		

Start of a FUNC Instruction Block

<i>Instructions</i>	FUNC																												
<i>Syntax</i>	<code>func <integer label></code>																												
<i>Description</i>	<p>Indicates the start of a subroutine. Each FUNC statement must have a unique label, <integer-label>. A RET instruction must be before the END instruction. This instruction can be used only after an ENDMAIN instruction. The code between FUNC and RET is executed if a CALL or CALLNZ with the same label value is executed.</p> <p>Operation: Defines the lexical start of a subroutine.</p>																												
<i>Format</i>	0-input, 0-output.																												
<i>Opcode</i>	<table> <thead> <tr> <th>Token</th> <th>Field Name</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>code</td> <td>15:0</td> <td>IL_OP_FUNC</td> </tr> <tr> <td></td> <td>control</td> <td>29:16</td> <td>Must be zero.</td> </tr> <tr> <td></td> <td>sec_modifier_present</td> <td>30</td> <td>Must be zero.</td> </tr> <tr> <td></td> <td>pri_modifier_present</td> <td>31</td> <td>Must be zero.</td> </tr> <tr> <td>2</td> <td colspan="3">Must be zero.</td> </tr> <tr> <td>3</td> <td colspan="3">Unsigned integer representing the label of the subroutine.</td> </tr> </tbody> </table>	Token	Field Name	Bits	Description	1	code	15:0	IL_OP_FUNC		control	29:16	Must be zero.		sec_modifier_present	30	Must be zero.		pri_modifier_present	31	Must be zero.	2	Must be zero.			3	Unsigned integer representing the label of the subroutine.		
Token	Field Name	Bits	Description																										
1	code	15:0	IL_OP_FUNC																										
	control	29:16	Must be zero.																										
	sec_modifier_present	30	Must be zero.																										
	pri_modifier_present	31	Must be zero.																										
2	Must be zero.																												
3	Unsigned integer representing the label of the subroutine.																												
<i>Related</i>	ENDFUNC.																												

Conditional IF Using Floating Point

<i>Instructions</i>	IFC															
<i>Syntax</i>	<code>ifc_relop(op) src0, src1</code>															
<i>Description</i>	<p>This is the start of an IFC - ENDIF or IFC - ELSE - ENDIF block. It skips a block of code based on the value of <i>src0</i> compared to <i>src1</i>. The IFC block must end with an ELSE or ENDIF instruction.</p> <p>Operation:</p> <pre>if (v0.x relop v1.x) { Execute following instructions; } else { Jump to the instruction following the next ELSE or ENDIF instruction; }</pre>															
<i>Format</i>	2-input.															
<i>Opcode</i>	<table> <thead> <tr> <th>Field Name</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>code</td> <td>15:0</td> <td>IL_OP_IFC</td> </tr> <tr> <td>control</td> <td>29:16</td> <td>relop(op). See Table 5.14 on page 5-8.</td> </tr> <tr> <td>sec_modifier_present</td> <td>30</td> <td>Must be zero.</td> </tr> <tr> <td>pri_modifier_present</td> <td>31</td> <td>Must be zero.</td> </tr> </tbody> </table>	Field Name	Bits	Description	code	15:0	IL_OP_IFC	control	29:16	relop(op). See Table 5.14 on page 5-8 .	sec_modifier_present	30	Must be zero.	pri_modifier_present	31	Must be zero.
Field Name	Bits	Description														
code	15:0	IL_OP_IFC														
control	29:16	relop(op). See Table 5.14 on page 5-8 .														
sec_modifier_present	30	Must be zero.														
pri_modifier_present	31	Must be zero.														
<i>Related</i>	IF_LOGICALNZ, IF_LOGICALZ, IFNZ.															

Execute instruction block if Boolean register is not zero

<i>Instructions</i>	IFNZ		
<i>Syntax</i>	ifnz src0		
<i>Description</i>	<p>This starts an IFNZ - ENDIF or IFNZ - ELSE - ENDIF block. IFNZ cannot be used with relative addressing or a source modifier. Both the <code>relative_address</code> and <code>modifier_present</code> fields of the IL_Src token must be zero. An IFNZ block must end with an ELSE or ENDIF instruction.</p> <p>Operation:</p> <pre> if (v0.x) { Execute following instructions; } else { Jump to the instruction following the next ELSE or ENDIF instruction; } </pre>		
<i>Format</i>	1-input.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_IFNZ
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	IFC, IF_LOGICALNZ, IF_LOGICALZ.		

Conditional IF Using Integers

<i>Instructions</i>	IF_LOGICALNZ, IF_LOGICALZ		
<i>Syntax</i>	Opcode	Syntax	Description
	IL_OP_IF_LOGICALNZ	if_logicalnz <i>src0.x</i>	If <i>src0.x</i> \neq 0 execute instruction block.
	IL_OP_IF_LOGICALZ	if_logicalz <i>src0.x</i>	If <i>src0.x</i> == 0 execute instruction block.
<i>Description</i>	<p>These are integer versions of the IF statement. They skip a block of code based on the value of <i>src0.x</i>. The LOG_IF block must end with and ELSE or ENDIF instruction. The source selector must replicate the component to be tested into all four components. The test uses integer tests, so values like Nan or -0 are not equal to zero.</p> <p>Operation for IF_LOGICALNZ:</p> <pre>if (v0.x has any bit non-zero) { Execute following instructions; } else { Jump to the instruction following the next ELSE or ENDIF instruction; }</pre> <p>Operation for IF_LOGICALZ:</p> <pre>if (v0.x has all bits zero) { Execute following instructions; } else { Jump to the instruction following the next ELSE or ENDIF instruction; }</pre>		
<i>Format</i>	1-input, 0-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	See <i>Syntax</i> , above.
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	IFC, IFNZ.		

Start of a Counted LOOP Block

Instructions **LOOP**

Syntax `loop [repeat] src0`

Description LOOP indicates the start of a LOOP instruction block. *src0* must be a register of type IL_REGTYPE_CONST_INT. *src0.x* specifies the iteration count used for the loop. If repeat is set to 0, *src0.y* specifies the initial value for the current *loop-counter* used for relative addressing. If repeat is set to 1, *src0.y* and *src0.z* are not used, and the current auto-increment loop counter is not incremented during this loop. This instruction starts a LOOP - ENDLOOP block.

If this instruction is within a FUNC - RET, IFNX - ENDIF, IFX - ENDIF, IFNZ - ELSE, IFC - ELSE, or ELSE - ENDIF block, its corresponding ENDLOOP must also be within that block.

This instruction is provided for iteration. It only increments the current auto-incremented loop counter if repeat is set to 0.

ENDLOOP must follow the last instruction of a loop block. The ENDLOOP instruction offset must be greater than the LOOP instruction offset.

If repeat is set to 0, the loop initial value and the loop iteration count cannot be negative.

LOOP cannot be used with relative addressing or a source modifier. Both the *relative_address* and *modifier_present* fields of the IL_Src token must be zero.

Format 1-input, 0-output.

<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_LOOP
	control	29:16	<i>repeat</i> flag
		0	<i>src0.y</i> holds the initial value for the current loop-counter used for relative addressing. <i>src0.z</i> holds the loop step. <i>src0.y</i> and <i>src0.z</i> cannot be negative.
		1	<i>src0.y</i> and <i>src0.z</i> are not used and the current auto-increment loop-counter is not incremented during this loop. <i>src0.y</i> and <i>src0.z</i> can be negative.
	<i>sec_modifier_present</i>	30	Must be zero.
	<i>pri_modifier_present</i>	31	Must be zero.

Related ENDLOOP, WHILELOOP.

Start of a WHILELOOP Block

<i>Instructions</i>	WHILELOOP		
<i>Syntax</i>	whileloop		
<i>Description</i>	WHILELOOP indicates the start of a WHILELOOP instruction block. A WHILELOOP block can iterate indefinitely, exiting only when a BREAK instruction is executed. It can be used as the translation of a DX10 loop statement. There is no limit to the amount of loop nesting. Although a WHILELOOP can iterate indefinitely, overall execution of the shader can be terminated after some number of instructions are executed.		
<i>Format</i>	0-input, 0-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_WHILE
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	ENDLOOP, LOOP.		

Start of a SWITCH Block

<i>Instructions</i>	SWITCH		
<i>Syntax</i>	switch <i>src0</i>		
<i>Description</i>	A switch/endswitch construct behaves exactly as a switch construct in the C language. The <i>src0</i> must be a 32-bit register component or immediate quantity. Compares are done using integer arithmetic. Falling through cases are valid, as in C. This instruction can be used to implement DX10 case instruction. Switch statements can be nested without limits. Operation: Same as a C switch statement.		
<i>Format</i>	1-input, 0-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_SWITCH
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	CASE, DEFAULT, ENDSWITCH.		

RETURN from a FUNC Block (end of subroutine)

<i>Instructions</i>	RET															
<i>Syntax</i>	ret															
<i>Description</i>	Returns from a subroutine, and indicates the end of the subroutine. It marks the end of a subroutine and can be present only at the end (it cannot be within a flow-control block). This instruction can implement DX9 returns. Operation: Continue execution after the call statement which executed this subroutine.															
<i>Format</i>	0-input, 0-output.															
<i>Opcod</i>	<table border="0"> <thead> <tr> <th style="text-align: left;">Field Name</th> <th style="text-align: left;">Bits</th> <th style="text-align: left;">Description</th> </tr> </thead> <tbody> <tr> <td>code</td> <td>15:0</td> <td>IL_OP_RET</td> </tr> <tr> <td>control</td> <td>29:16</td> <td>Must be zero.</td> </tr> <tr> <td>sec_modifier_present</td> <td>30</td> <td>Must be zero.</td> </tr> <tr> <td>pri_modifier_present</td> <td>31</td> <td>Must be zero.</td> </tr> </tbody> </table>	Field Name	Bits	Description	code	15:0	IL_OP_RET	control	29:16	Must be zero.	sec_modifier_present	30	Must be zero.	pri_modifier_present	31	Must be zero.
Field Name	Bits	Description														
code	15:0	IL_OP_RET														
control	29:16	Must be zero.														
sec_modifier_present	30	Must be zero.														
pri_modifier_present	31	Must be zero.														
<i>Related</i>	CALL, CALLNZ, CALL_LOGICALZ, CALL_LOGICALNZ, FUNC, RET_DYN, RET_LOGICALNZ, RET_LOGICALZ.															

RETURN from a FUNC Block (not end of subroutine)

<i>Instructions</i>	RET_DYN															
<i>Syntax</i>	ret_dyn															
<i>Description</i>	Returns from a subroutine to the instruction after the call. It can appear anywhere in a subroutine, any number of times.															
<i>Format</i>	0-input, 0-output.															
<i>Opcod</i>	<table border="0"> <thead> <tr> <th style="text-align: left;">Field Name</th> <th style="text-align: left;">Bits</th> <th style="text-align: left;">Description</th> </tr> </thead> <tbody> <tr> <td>code</td> <td>15:0</td> <td>IL_OP_RET_DYN</td> </tr> <tr> <td>control</td> <td>29:16</td> <td>Must be zero.</td> </tr> <tr> <td>sec_modifier_present</td> <td>30</td> <td>Must be zero.</td> </tr> <tr> <td>pri_modifier_present</td> <td>31</td> <td>Must be zero.</td> </tr> </tbody> </table>	Field Name	Bits	Description	code	15:0	IL_OP_RET_DYN	control	29:16	Must be zero.	sec_modifier_present	30	Must be zero.	pri_modifier_present	31	Must be zero.
Field Name	Bits	Description														
code	15:0	IL_OP_RET_DYN														
control	29:16	Must be zero.														
sec_modifier_present	30	Must be zero.														
pri_modifier_present	31	Must be zero.														
<i>Related</i>	CALL, CALLNZ, CALL_LOGICALZ, CALL_LOGICALNZ, FUNC, RET, RET_LOGICALNZ, RET_LOGICALZ.															

Conditional RETURN from FUNC Block Using Integer

Instructions **RET_LOGICALNZ, RET_LOGICALZ**

<i>Syntax</i>	Opcode	Syntax	Description
	IL_OP_IF_LOGICALNZ	if_logicalnz <i>src0.x</i>	If <i>src0.x</i> ≠ 0 execute instruction block.
	IL_OP_IF_LOGICALZ	if_logicalz <i>src0.x</i>	If <i>src0.x</i> == 0 execute instruction block.

Description These instructions conditionally return to the instruction after the call. *src1.x* is tested after swizzle. The instructions can appear anywhere in a subroutine, any number of times.

The 32-bit value supplied by *src1* is tested at the bit level:

For RET_LOGICALNZ, if any bit is non-zero, the statement returns.

For RET_LOGICALZ, if all bits are zero, the statement returns.

Format 1-input, 0-output.

<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	See <i>Syntax</i> , above.
	control	29:16	logic_op
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.

Related CALL, CALLNZ, CALL_LOGICALZ, CALL_LOGICALNZ, FUNC, RET, RET_DYN.

6.4 Declaration and Initialization Instructions

Declare Registers as Part of an Array

<i>Instructions</i>	DCLARRAY		
<i>Syntax</i>	<code>dclarray src0, src1</code>		
<i>Description</i>	<p>Declares a range of registers as part of an array indexed access (to be accessed through base/loop relative addressing). Only the INTERP and TEXCOORD registers can be used. <i>src0</i> and <i>src1</i> must be of the same register type. Base and loop relative addressing cannot be used on TEMP, INTERP, and TEXCOORD registers not declared within the range of <i>src1</i> and <i>src2</i>. If registers outside the range of registers in the array declared with this instruction are accessed, the compiler cannot guarantee the stability of the shader.</p> <p>Base and loop relative addressing perform indexing based on the register number, not on the register's position in the array. (Declaring an array of registers 2 to 6, and indexing when the loop counter is 3, accesses register number 3, not register number 5.)</p> <p>There can be numerous DCLARRAYs for each type, so long as they do not overlap in range.</p> <p>DCLARRAY cannot be used with relative addressing or a source modifier. Both the <code>modifier_present</code> and <code>relative_address</code> fields of the <code>IL_Src</code> tokens must be zero.</p>		
<i>Format</i>	2-input, 0-output.		
<i>Opcode</i>	Field Name	Bits	Description
	<code>code</code>	15:0	IL_OP_DCLARRAY
	<code>control</code>	29:16	Must be zero.
	<code>sec_modifier_present</code>	30	Must be zero.
	<code>pri_modifier_present</code>	31	Must be zero.
<i>Related</i>	DCL_INDEXED_TEMP_ARRAY.		

Declare a Constant Buffer*Instructions* DCL_CB*Syntax* dcl_cb cbm[n]

Description Declares that the shader can use the constant buffer. This instruction must occur before any use of the constant buffer. The source of this instruction must be of type IL_REGTYPE_CONSTANT_BUFFER. The cb index m can be in the range of 0-14. The index n can be 4k.

DCL_CB cannot be used with relative addressing or a source modifier. Both the `modifier_present` and `relative_address` fields of the IL_Src token must be zero.

If the instruction is being used to specify an immediate constant buffer, set the `pri_modifier_present` bit to 1. The subsequent 32 bits are an unsigned integer specifying the number of elements in the constant buffer, followed by one 32-bit floating point value specifying each member of the immediate constant buffer.

Note that the HD4000-family of devices does not support both constant buffer and constant file operations; drivers must select only one mode.

Format 1-input, 0-output.

<i>Opcode</i>	Token	Field Name	Bits	Description
	1	code	15:0	IL_DCL_CONST_BUFFER
		control	29:16	Must be zero.
		sec_modifier_present	30	Must be zero.
		pri_modifier_present	31	0: Instruction declares a constant buffer. 1: Does not declare a constant buffer. Instruction declares an Immediate constant buffer.
	2	pri_modifier_present == 0		IL_Src token (<i>src0</i>).
		pri_modifier_present == 1		Number of elements, n, in an immediate constant buffer.
	3	pri_modifier_present == 0		Not used.
		pri_modifier_present == 1		First (index 0) 32-bit element of the immediate constant buffer.
	4 to n+2	pri_modifier_present == 0		Not used.
		pri_modifier_present == 1		Second (index 1) through nth (index n-1) 32-bit elements of the immediate constant buffer.

Related None.

Declare register defaults

<i>Instructions</i>	DCLDEF																												
<i>Syntax</i>	<code>dcldef_x(comp)_y(comp)_z(comp)_w(comp) dst</code>																												
<i>Description</i>	<p>Declares the default value for register <i>dst</i> and cannot be used more than once per register type in a shader. <i>dst</i> must be a TEMP or ADDR register type (see Section 5.3 on page 5-2). Each component in the register is set individually to an ILDefaultVal enumerated type. See Table 5.11 on page 5-7 for ILDefaultVal values. The xdefault, ydefault, zdefault, and wdefault describe the default value for each component of the register.</p> <p>DCLDEF cannot be used with relative addressing or a destination modifier. Both the modifier_present and relative_address fields of the IL_Dst token must be zero.</p>																												
<i>Format</i>	0-input, 0-output.																												
<i>Opcode</i>	<table border="0"> <thead> <tr> <th style="text-align: left;">Token</th> <th style="text-align: left;">Field Name</th> <th style="text-align: left;">Bits</th> <th style="text-align: left;">Description</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>code</td> <td>15:0</td> <td>IL_OP_DCLDEF</td> </tr> <tr> <td></td> <td>xdefault</td> <td>17:16</td> <td>x-component default value.</td> </tr> <tr> <td></td> <td>ydefault</td> <td>19:18</td> <td>y-component default value.</td> </tr> <tr> <td></td> <td>zdefault</td> <td>21:20</td> <td>z-component default value.</td> </tr> <tr> <td></td> <td>wdefault</td> <td>23:22</td> <td>w-component default value.</td> </tr> <tr> <td></td> <td>reserved</td> <td>31:24</td> <td>Must be zero.</td> </tr> </tbody> </table> <p>The component default types can be any value of the enumerated type ILDefaultVal. See Table 5.11 on page 5-7.</p> <p>2 IL_Dst token (<i>dst</i>) where the register_type field is set to IL_REGTYPE_TEMP or IL_REGTYPE_ADDR. The modifier_present and relative_address fields must be set to 0.</p>	Token	Field Name	Bits	Description	1	code	15:0	IL_OP_DCLDEF		xdefault	17:16	x-component default value.		ydefault	19:18	y-component default value.		zdefault	21:20	z-component default value.		wdefault	23:22	w-component default value.		reserved	31:24	Must be zero.
Token	Field Name	Bits	Description																										
1	code	15:0	IL_OP_DCLDEF																										
	xdefault	17:16	x-component default value.																										
	ydefault	19:18	y-component default value.																										
	zdefault	21:20	z-component default value.																										
	wdefault	23:22	w-component default value.																										
	reserved	31:24	Must be zero.																										
<i>Related</i>	None.																												

Constant Integer or Float Register Definition

Instructions **DEF**

Syntax `def dst, <number>, <number>, <number>, <number>`

Description DEF declares the constant value for register *dst* and cannot be used more than once on the same register. It indicates that a floating point or integer constant contains a given value when the shader is executed. This instruction is followed by four words that contain the bit values of the constant. *dst* must be a CONST_INT or CONST_FLOAT register type (See Section 4 on page 4-1). Each component of a register can be set individually. CONST_INT registers do not use the w-component.

Clients must ensure that the constant values set through state match the values indicated in this instruction.

DEF cannot be used with relative addressing or a destination modifier. Both the `modifier_present` and `relative_address` fields of the IL_Dst token must be zero.

Operation:

```
VECTOR v;
v[0] = x;
v[1] = y;
v[2] = z;
v[3] = w;
WriteResult(v, dst);
```

Format 0-input, 1-output.

Opcode	Token	Field Name	Bits	Description
1	code		15:0	IL_OP_DEF
		control	29:16	Must be zero.
		sec_modifier_present	30	Must be zero.
		pri_modifier_present	31	Must be zero.
	2	IL_Dst token (<i>dst</i>) where the <code>register_type</code> field is set to IL_REGTYPE_CONST_INT or IL_REGTYPE_CONST_FLOAT. The <code>modifier_present</code> and <code>relative_address</code> fields must be set to zero.		
		3	x-component, 32-bit integer or float	
4	y-component, 32-bit integer or float			
5	z-component, 32-bit integer or float			
6	w-component, 32-bit float only (not used for a CONST_INT register)			

Related DEFB.

Constant Boolean Register Definition

Instructions **DEFB**

Syntax `defb dst, <unsigned integer>`

Description Indicates that a boolean constant contains a given value when the shader is executed. Clients must ensure that the constant values set through state match the values indicated in this instruction. A shader cannot use this instruction on the same register more than once.

If the value of the given unsigned integer is 0, the boolean is set to false. If the value the given unsigned integer is non-zero, set the boolean to true.

DEFB cannot be used with relative addressing or a destination modifier. Both the `modifier_present` and `relative_address` fields of the `IL_Dst` token must be zero.

Operation:

```
CONST_BOOL b;
b = TRUE;
if(val == 0) {
    b=FALSE;
}
WriteResult(b, dst);
```

Format	<table border="0"> <thead> <tr> <th style="text-align: left;">Token</th> <th style="text-align: left;">Field Name</th> <th style="text-align: left;">Bits</th> <th style="text-align: left;">Description</th> </tr> </thead> <tbody> <tr> <td rowspan="4">1</td> <td><code>code</code></td> <td>15:0</td> <td><code>IL_OP_DEFB</code></td> </tr> <tr> <td><code>control</code></td> <td>29:16</td> <td>Must be zero.</td> </tr> <tr> <td><code>sec_modifier_present</code></td> <td>30</td> <td>Must be zero.</td> </tr> <tr> <td><code>pri_modifier_present</code></td> <td>31</td> <td>Must be zero.</td> </tr> <tr> <td rowspan="1">2</td> <td colspan="3"> <code>IL_Dst</code> token (<i>dst</i>) where the <code>register_type</code> field is set to <code>IL_REGTYPE_CONST_BOOL</code>. The <code>modifier_present</code> and <code>relative_address</code> fields must be set to zero. </td> </tr> <tr> <td rowspan="1">3</td> <td colspan="3"> Boolean value, 32-bit unsigned integer <ul style="list-style-type: none"> • 0: FALSE • not 0: TRUE </td> </tr> </tbody> </table>	Token	Field Name	Bits	Description	1	<code>code</code>	15:0	<code>IL_OP_DEFB</code>	<code>control</code>	29:16	Must be zero.	<code>sec_modifier_present</code>	30	Must be zero.	<code>pri_modifier_present</code>	31	Must be zero.	2	<code>IL_Dst</code> token (<i>dst</i>) where the <code>register_type</code> field is set to <code>IL_REGTYPE_CONST_BOOL</code> . The <code>modifier_present</code> and <code>relative_address</code> fields must be set to zero.			3	Boolean value, 32-bit unsigned integer <ul style="list-style-type: none"> • 0: FALSE • not 0: TRUE 		
Token	Field Name	Bits	Description																							
1	<code>code</code>	15:0	<code>IL_OP_DEFB</code>																							
	<code>control</code>	29:16	Must be zero.																							
	<code>sec_modifier_present</code>	30	Must be zero.																							
	<code>pri_modifier_present</code>	31	Must be zero.																							
2	<code>IL_Dst</code> token (<i>dst</i>) where the <code>register_type</code> field is set to <code>IL_REGTYPE_CONST_BOOL</code> . The <code>modifier_present</code> and <code>relative_address</code> fields must be set to zero.																									
3	Boolean value, 32-bit unsigned integer <ul style="list-style-type: none"> • 0: FALSE • not 0: TRUE 																									

Related DEF.

Declare an Indexed Array

<i>Instructions</i>	DCL_INDEXED_TEMP_ARRAY		
<i>Syntax</i>	<code>dcl_indexed_temp_array src0[n]</code>		
<i>Description</i>	<p>Declares a temporary array.</p> <p>Each indexed temp array to be used in the Shader must be declared. <i>src0</i> must be of type IL_REGTYPE_ITEMP. The modifier_present and relative_address fields must be zero. <i>n</i> is the maximum size in 128-bit units.</p> <p>DX10 limits the total storage for temps (indexed plus non-indexed) to be ≤ 4096 registers (each a four-component vector).</p>		
<i>Format</i>	1-input, 0-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_DCL_INDEXED_TEMP_ARRAY
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	None. DCLARRAY.		

Declare an Input Register

<i>Instructions</i>	DCL_INPUT												
<i>Syntax</i>	<code>dcl_input[_usage(usage)][_interp(mode)] dst[.mask]</code>												
<i>Description</i>	<p>Declares an input register for the shader. <i>dst</i> must be an IL_REGTYPE_INPUT register type (see Chapter 4, "Register Types") and can have a mask. It is legal to declare a subset of the component mask from what is output from the previous shader in the pipeline; however, mutually exclusive masks are not allowed (for example: VS outputting o3.xy means PS declaring v3.z input is invalid, but v3.x, v3.y, or v3.xy is legal). Interpolation mode, <code>_interp(mode)</code>, is only applicable to pixel shaders; it is an error to use this in GS or VS.</p> <p>Possible values for <code>_interp(mode)</code> are:</p> <table> <tr> <td>IL_INTERPMODE_NOTUSED</td> <td>IL_INTERPMODE_LINEAR_CENTROID</td> </tr> <tr> <td>IL_INTERPMODE_CONSTANT</td> <td>IL_INTERPMODE_LINEAR_NOPERSPECTIVE</td> </tr> <tr> <td>IL_INTERPMODE_LINEAR</td> <td>IL_INTERPMODE_LINEAR_SAMPLE</td> </tr> <tr> <td>IL_INTERPMODE_LINEAR_NOPERSPECTIVE_CENYROID</td> <td></td> </tr> <tr> <td>IL_INTERPMODE_LINEAR_NOPERSPECTIVE_SAMPLE</td> <td></td> </tr> </table>			IL_INTERPMODE_NOTUSED	IL_INTERPMODE_LINEAR_CENTROID	IL_INTERPMODE_CONSTANT	IL_INTERPMODE_LINEAR_NOPERSPECTIVE	IL_INTERPMODE_LINEAR	IL_INTERPMODE_LINEAR_SAMPLE	IL_INTERPMODE_LINEAR_NOPERSPECTIVE_CENYROID		IL_INTERPMODE_LINEAR_NOPERSPECTIVE_SAMPLE	
IL_INTERPMODE_NOTUSED	IL_INTERPMODE_LINEAR_CENTROID												
IL_INTERPMODE_CONSTANT	IL_INTERPMODE_LINEAR_NOPERSPECTIVE												
IL_INTERPMODE_LINEAR	IL_INTERPMODE_LINEAR_SAMPLE												
IL_INTERPMODE_LINEAR_NOPERSPECTIVE_CENYROID													
IL_INTERPMODE_LINEAR_NOPERSPECTIVE_SAMPLE													

Declare an Input Register (Cont.)

Format	Field Name	Bits	Description
	code	15:0	IL_DCL_INPUT
	control	20:16	<i>usage</i> . Unless using system value (SV), must be IL_IMPORTUSAGE_GENERIC (see Table 5.9 on page 5-4). The usage index is specified by the register number.
	interp mode	23:21	Interpolation mode.
	center	24	This field is always ignores. Center vs centroid must be specified through <i>_interp(mode)</i> .
	bias	25	This is ignored, unless the <i>dst</i> type is wincoord. For the window coordinate register, this bit specifies if the compiler generates code that adds a bias supplied in constants that are named SC_CONS_SRC_VIEWPORT_BIAS_{X,Y}. Without a bias, the origin is located at the upper-left corner of the screen (DX Style). A bias can be used to move the origin. For example, to move the origin to the lower-left corner of a window (OGL style), DX must set this to 0, OGL must set this to 1.
	invert	26	This is ignored, unless the <i>dst</i> type is wincoord. If the register_type field of the following IL_DstToken is IL_REGTYPE_WINCOORD, and this bit is set, then the compiler inverts (multiplies by -1) the y coordinate of the position. By default, the y axis goes down (DX style), but this bit can be used to change the direction to up (OGL style).
	centered	27	This is ignored, unless the <i>dst</i> type is wincoord. For the window coordinate register, this bit controls (value 0) defines the origin as the upper-left corner of the RenderTarget. Thus, pixel centers are offset by (0.5f,0.5f) from integer locations on the RenderTarget. This choice of origin makes rendering screen-aligned textures trivial, as the pixel coordinate system is aligned with the texel coordinate system. This choice matches OGL and DX10. The second value (1) defines the origin as the center of the upper-left pixel in the RenderTarget. In other words, the origin is (0.5,0.5) away from the upper left corner of the RenderTarget. Thus, pixel centers are at integer locations. This choice of origin matches DX9. 0 – center of the pixel is 0.5,0.5 1 – center of the pixel is 0,0
	reserved	29:28	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	DCL_OUTPUT, DCLPIN, DCLVPOUT.		

Declare a Primitive

Instructions **DCL_INPUTPRIMITIVE**

Syntax `dcl_input_primitive prim_type`

Description Declare the type of primitive that can be accepted by a geometry shader. Must appear in a geometry shader. Only valid in a geometry shader. The `prim_type` must be an element of the enumeration `IL_TOPOLOGY`.

Format 0-input, 0-output.

<i>Opcode</i>	Token	Field Name	Bits	Description
	1	code	15:0	IL_DCL_INPUTPRIMITIVE
		control	29:16	prim_type
		sec_modifier_present	30	Must be zero.
		pri_modifier_present	31	Must be zero.
	2	IL_Src token (<i>src0</i>) where the <code>register_type</code> field is set to <code>IL_REGTYPE_LITERAL</code> .		
	3	x-bits, 32-bit untyped literal.		
	4	y-bits, 32-bit untyped literal.		
	5	z-bits, 32-bit untyped literal.		
	6	w-bits, 32-bit untyped literal.		

Related None.

Declare a Literal

Instructions **DCL_LITERAL**

Syntax `dcl_literal src0, <x-bits>, <y-bits>, <z-bits>, <w-bits>`

Description DCL_LITERAL declares the literal to be used in the following instruction. The instruction is followed by four words containing the actual bits of the literal, in order x, y, z, w. value for register *src0* and can be used more than once on the same register in a shader. The lexically nearest preceding value is used. The 32-bit component literals (x-bits, y-bits, z-bits, and w-bits) are untyped, so that integer and float literals can be initialized with this instruction. *src0* must be a IL_REGTYPE_LITERAL register type (see Section 4.10, "LITERAL," page 4-6). This instruction cannot be placed in an unreachable code block such as after an unconditional break or return instruction. No modifiers are allowed.

Format 1-input, 0-output.

<i>Opcode</i>	Token	Field Name	Bits	Description
	1	code	15:0	IL_DCL_LITERAL
		control	29:16	Must be zero.
		sec_modifier_present	30	Must be zero.
		pri_modifier_present	31	Must be zero.
	2	IL_Src token (<i>src0</i>) where the <i>register_type</i> field is set to IL_REGTYPE_LITERAL.		
	3	x-bits, 32-bit untyped literal		
	4	y-bits, 32-bit untyped literal		
	5	z-bits, 32-bit untyped literal		
	6	w-bits, 32-bit untyped literal		

Related None.

Declare the Maximum Number of Vertices

<i>Instructions</i>	DCL_MAX_OUTPUT_VERTEX_COUNT		
<i>Syntax</i>	dcl_max_output_vertex_count <i>n</i>		
<i>Description</i>	<p>DCL_MAX_OUTPUT_VERTEX_COUNT declares the maximum number of vertices that a single invocation of a geometry shader can emit. A geometry shader can emit a maximum of 1024 32-bit values. Thus, <i>n</i> can be no larger than floor (1024/number of 32-bit values per vertex). For example, if a geometry shader emits one 4-component vector (four 32-bit values) per vertex then <i>n</i> must be 256 or less.</p> <p>Some implementations may be able to make optimizations by knowing the maximum number of vertices a single geometry shader invocation emits (for a single input primitive). The upper bound on the number of vertices that a geometry shader can produce depends on how large each vertex is. The sum of the number of components in each declared geometry shader output register defines how many 32-bit values are present in a single vertex. For example, if a geometry shader declares that it outputs a single four-component position, plus a three-component color per vertex, then the maximum number of vertices that can be declared for output by a single invocation is floor(1024 / 7). Or, if a Geometry Shader declares that it outputs 32 four-component vectors, the maximum number of vertices that can be declared for output by a single invocation is floor(1024 / 128).</p> <p>DCL_MAX_OUTPUT_VERTEX_COUNT sets an upper limit on the number of vertices that can be emitted and, a geometry shader invocation can terminate after emitting fewer vertices than the maximum number allowed. An invocation terminates when DCL_MAX_OUTPUT_VERTEX_COUNT is reached. There is no requirement on the minimum number of vertices a geometry shader invocation must emit. The amount of vertices generated by a geometry shader invocation is simply the total number of emit* instructions executed in an invocation.</p>		
<i>Format</i>	0-input, 0-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_DCL_MAX_OUTPUT_VERTEX_COUNT
	control	29:16	<i>n</i> , maximum number of vertices.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	EMIT, EMIT_THEN_CUT.		

Declare that the Pixel Shader intends to write to its scalar output oDepth register

<i>Instructions</i>	DCL_ODEPTH		
<i>Syntax</i>	dcl_odepth		
<i>Description</i>	Declare that the pixel shader intends to write to its scalar output oDepth register. DX10 has some rules for what happens if oDepth is declared, but the shader does not write it.		
<i>Format</i>	0-input, 0-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_DCL_ODEPTH
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	None.		

Declare Primitive Topology of Geometry Shader Output

<i>Instructions</i>	DCL_OUTPUT_TOPOLOGY		
<i>Syntax</i>	dcl_output_topology <i>n</i>		
<i>Description</i>	The geometry shader can only emit a single primitive topology from a given shader; the choices are: pointlist, linestrip or trianglestrip. Geometry shaders must contain this declaration. Note that for strip topologies, a single invocation of the geometry shader can emit multiple strips by using the cut instruction. This instruction is required in a geometry shader.		
<i>Format</i>	0-input, 0-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_DCL_OUTPUT_TOPOLOGY
	control	29:16	Any value of the enumerated type ILTopologyType. See Table 5.4 on page 5-3.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	None.		

Declare an Output Register**Instructions** **DCL_OUTPUT****Syntax** `dcl_output[_usage(type)] dst`

Description Declares the usage of a shader output register. *dst* must be of type IL_REGTYPE_OUTPUT and can have a modifier that specifies a component mask (see Section 3.6, "Write Mask," page 3-3). The component mask can be any subset of [xyzw]. DCL_OUTPUT can declare a superset of the component mask declared for input by the next shader stage, indicating the current shader writes more registers than the next shader stage reads; however, mutually exclusive masks are not allowed. For example, a vertex shader that writes o3.xy means the pixel shader reading only v3.z is invalid, but reading v3.x or v3.y or v3.xy would be valid. An output variable can be declared more than once, so that different components can be given different output types. Also, it is possible to have multiple clip distances. DX10 specifies that any component which is of type "none" must appear in xyzw order after components with non-none types. The component mask must be appropriate to the particular type. For example, the current set of system-generated values are all scalars, so the mask must have only one component.

A system-generated value cannot be output from a stage that is before the place in the pipeline where the hardware normally generates the value. For example, a geometry shader cannot output `IsFrontFace`, and VS cannot output `PrimitiveID`.

Output type is:

Type	IL Enumeration (No EXPORTUSAGE in IL Headers)
Generic	IL_IMPORTUSAGE_GENERIC
Position	IL_IMPORTUSAGE_POSITION
ClipDistance	IL_IMPORTUSAGE_CLIPDISTANCE
CullDistance	IL_IMPORTUSAGE_CULLDISTANCE
PrimitiveID	IL_IMPORTUSAGE_PRIMITIVEID
VertexID	IL_IMPORTUSAGE_VERTEXID
InstancelD	IL_IMPORTUSAGE_INSTANCEID
RenderTargetArrayIndex	IL_IMPORTUSAGE_RENDERTARGET_ARRAY_INDEX
ViewportArrayIndex	IL_IMPORTUSAGE_VIEWPORT_ARRAY_INDEX

Format 0-input, 1-output.

Opcode	Field Name	Bits	Description
	<code>code</code>	15:0	IL_DCL_ODEPTH
	<code>control</code>	29:16	Any value of the enumerated type <code>ILImportUsage</code> . See Table 5.9 on page 5-4.
	<code>sec_modifier_present</code>	30	Must be zero.
	<code>pri_modifier_present</code>	31	Must be zero.

Related None.

Declare a Primitive ID

<i>Instructions</i>	DCL_VPRIM (Input Primitive)		
<i>Syntax</i>	dcl_vprim		
<i>Description</i>	Declares that the geometry shader intends to use its scalar input register vPrim. For the geometry shader, input primitive data only comes in the form of a scalar (vPrim, no mask). Also, there is no Primitive Data for adjacent primitives available in a geometry shader invocation.		
<i>Format</i>	0-input, 0-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_DCL_VPRIM
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	None.		

Declare Shared Registers

<i>Instructions</i>	DCL_SHARED_TEMP		
<i>Syntax</i>	dcl_shared_temp src#		
<i>Description</i>	Declares the number of shared GPRs used by this kernel (shared for each SIMD). <i>src0</i> must be of type IL_REGTYPE_SHARED_TEMP, with the number (#) indicating one more than the maximum used index. For example, <code>dcl_shared_temp sr4</code> indicates that sr0, sr1, sr2, and sr3 are used in the shader. The number declared must be smaller than the maximum available in hardware (for example: 128 for the HD4000-family of devices).		
<i>Format</i>	1-input, 0-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_DCL_SHARED_TEMP
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	DCL_OUTPUT.		

Declare LDS Size Used in a Shader

<i>Instructions</i>	DCL_LDS_SIZE_PER_THREAD		
<i>Syntax</i>	<code>dcl_lds_size_per_thread n</code>		
<i>Description</i>	Declare the space or size of LDS memory (in dwords) to be used in a compute shader. The value must be: <ul style="list-style-type: none"> • in dwords • no greater than 64, and • a factor of 4. 		
<i>Format</i>	0-input, 0-output, no additional token.		
<i>Opcode</i>	Field Name	Bits	Description
	<code>code</code>	15:0	Must be set to <code>IL_OP_DCL_LDS_SIZE_PER_THREAD</code> .
	<code>controls</code>	29:16	<code>n</code> = size in dwords
	<code>sec_modifier_present</code>	30	Must be zero.
	<code>pri_modifier_present</code>	31	Must be zero.
<i>Example</i>	<code>dcl_lds_size_per_thread 8</code>		
<i>Related</i>	None.		

Declare the LDS Sharing Mode

<i>Instructions</i>	DCL_LDS_SHARING_MODE		
<i>Syntax</i>	<code>dcl_lds_sharing_mode _wavefrontRel or _wavefrontAbs</code>		
<i>Description</i>	Local data share (LDS) memory has two sharing mode: wavefront relative or absolute. Relative means each wavefront has its private LDS memory. Absolute means all wavefronts share the same piece of LDS memory. Only used in a compute kernel.		
<i>Format</i>	0-input, 0-output, 0 additional token.		
<i>Opcode</i>	Field Name	Bits	Description
	<code>code</code>	15:0	Must be set to <code>IL_OP_DCL_LDS_SHARING_MODE</code> .
	<code>controls</code>	29:16	Mode: 0 <code>_wavefrontRel</code> 1 <code>_wavefrontAbs</code>
	<code>sec_modifier_present</code>	30	Must be zero.
	<code>pri_modifier_present</code>	31	Must be zero.
<i>Examples</i>	<code>dcl_lds_sharing_mode _wavefrontRel</code>		
	<code>dcl_lds_sharing_mode _wavefrontAbs</code>		
<i>Related</i>	None.		

Declare the Thread Group Size

<i>Instructions</i>	DCL_NUM_THREAD_PER_GROUP		
<i>Syntax</i>	<code>dcl_num_thread_per_group n1, n2, n3</code>		
<i>Description</i>	Specifies the number of threads per group. The product of the three sizes can be at most 1024. Only used in a compute kernel. For the HD4000-family of devices, n2 and n3 must be 1 because they allow only one dimension.		
<i>Format</i>	0-input, 0-output, 1 to 3 additional tokens: up to three unsigned integers representing the literal value for n (n1, n2, n3).		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	Must be set to IL_OP_DCL_NUM_THREAD_PER_GROUP.
	controls	29:16	Number of dimensions (1 to 3).
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Examples</i>	<code>dcl_num_thread_per_group 5, 10, 1</code> <code>dcl_num_thread_per_group 5 (which is the same as dcl_num_thread_per_group 5, 1, 1)</code>		
<i>Related</i>	None.		

Declare Interpolator Properties

<i>Instructions</i>	DCLPI
<i>Syntax</i>	<code>dclpi_x(comp)_y(comp)_z(comp)_w(comp)_center_bias_invert_centered] dst</code>
<i>Description</i>	<p>Declares properties of pixel shader named interpolator inputs and the window coordinate register.</p> <p>A shader cannot use this instruction on the same register more than once. There can be at most one DCLPI per register.</p> <p><code>ximport</code>, <code>yimport</code>, <code>zimport</code>, and <code>wimport</code> each describe what components of the register are used by the pixel shader. See the enumerated type for more information.</p> <p>In shaders where PINPUT registers are used, this instruction can only be used to declare a WINCOORD register.</p> <p>You cannot use a destination modifier with this instruction. The <code>modifier_present</code> field of the IL_Dst token must be set to 0.</p> <p>You cannot use relative addressing with this instruction. The <code>relative_address</code> field of the IL_Dst token must be set to 0.</p> <p>To use both bias and centered, set just bias and adjust the bias values: when invert is not set, add 0.5 to both bias.x and bias.y; when invert is set, add 0.5 to bias.x and -0.5 to bias.y. This optimization reduces instructions generated to support these control fields.</p> <p>A DCLPI token can be used to declare interpolated outputs in a vertex shader and interpolated inputs in a pixel shader.</p>

Declare Interpolator Properties (Cont.)

Format	Token	Field Name	Bits	Description
	1	code	15:0	IL_OP_DCLPI
		ximport	17:16	Any value of the enumerated type ILImportComponent. See Table 5.10 on page 5-6.
		yimport	19:18	Any value of the enumerated type ILImportComponent. See Table 5.10 on page 5-6.
		zimport	21:20	Any value of the enumerated type ILImportComponent. See Table 5.10 on page 5-6.
		wimport	23:22	Any value of the enumerated type ILImportComponent. See Table 5.10 on page 5-6.
		center	24	For interpolator registers, this field specifies if center or centroid sampling is used. If the <code>register_type</code> field of the following IL_Dst token is IL_REGTYPE_WINCOORD, this bit determines whether to use the center (1), or the nearest sample (0), called centroid, within the pixel as the interpolation value when multisampling, since the center might be outside the polygon. For DX10 drivers, this value can be zero or one, depending on the shader. For all other drivers this value must be one.
		bias	25	This field has no effect if the <code>register_type</code> field of the following IL_Dst token is not IL_REGTYPE_WINCOORD. For the window coordinate register, specifies if the compiler generates instructions to add a bias supplied in constants named SC_CONS_SRC_VIEWPORT_BIAS_{X,Y}. Without a bias, the origin is located at the upper-left corner of the screen (DX Style). A bias can be used to move the origin. For example, to move the origin to the lower-left corner of a window (OGL style): <ul style="list-style-type: none"> • DX must set this to 0 • OGL must set this to 1.
		invert	26	This field does nothing if the <code>register_type</code> field of the following IL_Dst token is not IL_REGTYPE_WINCOORD. If the <code>register_type</code> field of the following IL_DstToken is IL_REGTYPE_WINCOORD and this bit is set, the compiler inverts (multiplies by -1) the y coordinate of the position. By default, the y axis goes down (DX style), but this bit can be used to change the direction to up (OGL style).

Declare Interpolator Properties (Cont.)

centered	27	<p>This field does nothing if the register_type field of the following IL_Dst token is not IL_REGTYPE_WINCOORD. If the register_type field of the following IL_Dst token is IL_REGTYPE_WINCOORD, this bit specifies the following:</p> <p>0 defines the origin as the upper-left corner of the RenderTarget. Thus, pixel centers are offset by (0.5f,0.5f) from integer locations on the RenderTarget. This origin makes rendering screen-aligned textures trivial, as the pixel coordinate system is aligned with the texel coordinate system. This choice matches OGL and DX10.</p> <p>1 defines the origin as the center of the upper-left pixel in the RenderTarget: the origin is (0.5,0.5) away from the upper-left corner of the RenderTarget. Thus, pixel centers are at integer locations. This choice of origin matches DX9.</p> <p>0 center of the pixel is (0.5,0.5) 1 center of pixel is (0,0)</p>
reserved	31:28	Must be zero.
2		<p>IL_Dst token (<i>dst</i>) where the register_type field is set to IL_REGTYPE_INTERP, IL_REGTYPE_FOG, IL_REGTYPE_PRIMCOORD, IL_REGTYPE_TEXCOORD, IL_REGTYPE_PRICOLOR, or IL_REGTYPE_SECCOLOR, or IL_REGTYPE_WINCOORD. The modifier_present and relative_address fields must be set to 0.</p>

Related None.

Declare Pixel Shader Input Register**Instructions** **DCLPIN****Syntax** `dclpin_usage(op)_usageIndex(n)_x(comp)_y(comp)_z(comp)_w(comp)[_centroid]
dst`**Description** Declares a mapping of a vertex shader output to a pixel shader input. This instruction or a DCLPP instruction must be issued on each PINPUT register before the register is used in a shader. There can be at most one DCLPIN instruction per *usage-usageIndex* pair.An *enabled* component is a component set to IL_IMPORTSEL_UNDEFINED, IL_IMPORTSEL_DEFAULT0, or IL_IMPORTSEL_DEFAULT1.A *disabled* component is a component set to IL_IMPORTSEL_UNUSED.*Packed Registers:* a PINPUT register can be declared multiple times with this instruction; thus, a single register can have multiple unique *usage-usageIndex* pairs. However, the same component of a register cannot be *enabled* in both declarations. Thus, if in one declaration a component is *enabled*, the component must be *disabled* in the other declaration(s). For example, if vIN3 is declared as having *usage(interp)_usageIndex(1)_x (*)*, and vIN3 is also declared as having *usage(interp)_usageIndex(2)*, then that declaration must set x(-).

If an IL_PrimaryDCLPIN_Mod token is not preset, the shader behaves as if ximport, yimport, zimport, and wimport are set to IL_IMPORTSEL_UNDEFINED and centroid is set to 0.

Only one register can be declared to have the usage IL_IMPORTUSAGE_FOG. In this case, usageIndex must be zero.

It is an error to use this instruction in a vertex shader or a real-time pixel shader.

Note that a shader using a PINPUT register cannot use the INTERP, TEXCOORD, PRICOLOR, SECCOLOR, or FOG registers.

A source modifier cannot be used with this instruction. The *modifier_present* field of the IL_Dst token must be set to 0.You cannot use relative addressing with this instruction. The *relative_address* field of the IL_Dst token must be set to 0.**Format**

Token	Field Name	Bits	Description
1	code	15:0	IL_OP_DCLPIN The lower five bits of the control field must be set to an acceptable value of the enumerated type ILImportUsage(usage). The next eight bits of the control field must be a unique number for usage that specifies a mapping of usage type to a vertex shader output (<i>usageIndex</i>).
	usage	20:16	Any value of the enumerated type ILImportUsage. See Table 5.9 on page 5-4.
	usageIndex	28:21	Unique number for <i>usage</i> which specifies a mapping of <i>usage</i> type to a vertex shader output (<i>usageIndex</i>), 0 to 255.
	reserved	30:29	Must be zero.
	pri_instruction _modifier	31	If <i>pri_modifier_present</i> is set to 1, an IL_PrimaryDCLPIN_Mod token immediately follows this token.

Declare Pixel Shader Input Register (Cont.)

2	Primary pixel shader input register declaration modifier. IL_PrimaryDCLPIN_Mod token described below: Note that the IL_PrimaryDCLPIN_Mod is present only if the <code>pri_modifier_present</code> field is 1 in the previous IL_OpCode token.		
	Field Name	Bits	Description
	<code>ximport</code>	1:0	Specifies if the x component is enabled for the usage-usageIndex for this declaration. Also specifies the default value if the component is enabled. Can be any value of the enumerated type <code>ILImportComponent</code> . See Table 5.10 on page 5-6.
	<code>yimport</code>	3:2	Specifies if the y component is enabled for the usage-usageIndex for this declaration. Also specifies the default value if the component is enabled. Can be any value of the enumerated type <code>ILImportComponent</code> . See Table 5.10 on page 5-6.
	<code>zimport</code>	5:4	Specifies if the z component is enabled for the usage-usageIndex for this declaration. Also specifies the default value if the component is enabled. Can be any value of the enumerated type <code>ILImportComponent</code> . See Table 5.10 on page 5-6.
	<code>wimport</code>	7:6	Specifies if the w component is enabled for the usage-usageIndex for this declaration. Also specifies the default value if the component is enabled. Can be any value of the enumerated type <code>ILImportComponent</code> . See Table 5.10 on page 5-6.
	<code>centroid</code>	8	Specifies if this value represents a value at the pixel centroid or the center 0 Value at the center. 1 Value at the centroid.
	<code>origin</code>	9	This field is ignored if the import usage is not <code>IL_REGTYPE_WINCOORD</code> . For the window coordinate register, specifies if the value is relative to the upper-left corner of the window or the lower-left corner of the window. If the import usage is <code>IL_IMPORTUSAGE_WINCOORD</code> , then: 0 The value of the <code>IL_REGTYPE_WINCOORD</code> register is relative to the upper-left-hand corner of the window (DX style). 1 The value of the <code>IL_REGTYPE_WINCOORD</code> register is relative to the lower-left-hand corner of the window (OpenGL style).
	<code>constant</code>	10	Constant interpolation.
	<code>no_perspective</code>	11	Do not perform perspective divide during interpolation.
	<code>reserved</code>	31:12	Must be zero.
3	IL_Dst token (<code>dst</code>) where the <code>register_type</code> field is set to <code>IL_REGTYPE_PINPUT</code> . The <code>modifier_present</code> and <code>relative_address</code> fields must be set to 0.		

Related

None.

Map Interpolator Register to Realtime Interpolator Parameter

Instructions **DCLPP****Syntax** `dclpp_param(n) dst`**Description** Maps a PINPUT register to a realtime interpolator parameter. This instruction can be used only in a real-time pixel shader. It is an error to use this instruction in a normal pixel shader or vertex shader. This instruction must be issued on each PINPUT register before the register is used in a real-time pixel shader. There can be at most one DCLPP per register.

Format	Token	Field Name	Bits	Description
	1	code	15:0	IL_OP_DCLPP
		param	23:16	Specifies the realtime state register to which the register is mapped (<i>param</i>).
		<i>reserved</i>	31:24	Must be zero.
	2	IL_Dst token (<i>dst</i>), where <i>register_type</i> is set to IL_REGTYPE_PINPUT. The <i>modifier_present</i> and <i>relative_address</i> fields must be set to 0.		

Related None.

Declare Texture Properties

<i>Instructions</i>	DCLPT		
<i>Syntax</i>	<code>dclpt_stage(<i>n</i>)_type(<i>op</i>)_coordmode(<i>mode</i>)</code>		
<i>Description</i>	<p>Declares properties of a texture stage. Do not use this instruction for DX10.</p> <p>This instruction must be issued on each stage if a TEXLD, TEXLDD, TEXLDB, TEXLDMS, TEXWEIGHT, PROJECT, or LOD instruction is issued on the stage.</p> <p>There can be at most one DCLPT per <i>stage</i> value.</p> <p>A shader cannot use this instruction on the same stage more than once.</p> <p>If <i>type</i> is IL_USAGE_PIXTEX_UNKNOWN, this instruction indicates that the texture type of the texture on the stage/unit indicated by <i>stage</i> is not known at shader-create time</p> <p>If <i>coordmode</i> is set to IL_TEXCOORDMODE_NORMALIZED, this instruction indicates that the texture coordinate is normalized in any subsequent TEXLD, TEXLDB, or LOD instruction for the texture on <i>stage</i>.</p> <p>If <i>coordmode</i> is set to IL_TEXCOORDMODE_UNNORMALIZED, this instruction indicates that the texture coordinate is not normalized in any subsequent TEXLD, TEXLDD, TEXLDB, TEXLDMS, TEXWEIGHT, PROJECT, or LOD instruction for the texture on <i>stage</i> when the wrap mode set in state for the stage/unit is set to clamp-to-border, clamp-half-way-to-border, or clamp-to-edge. In this case, the x texture coordinate ranges from 0.0 to the width of the texture, the y texture coordinate ranges from 0.0 to the height of the texture, and the z texture coordinate ranges from 0.0 to the depth of the texture.</p> <p>If <i>coordmode</i> is set to IL_TEXCOORDMODE_UNKNOWN, the value of AS_TEX_DENORM_N (<i>stage</i>) determines if the texture coordinate used in any subsequent TEXLD, TEXLDD, TEXLDB, TEXLDMS, TEXWEIGHT, PROJECT, or LOD instructions are normalized at shader run time.</p> <p>If <i>cleartype mode</i> is set, the compiler multiplies the result of any fetch on this texture by 1/(kernel height * width). This is used for DX9 only.</p>		
<i>Format</i>	Field Name	Bits	Description
	<code>code</code>	15:0	IL_OP_DCLPT
	<code>stage</code>	23:16	Stage or unit number.
	<code>type</code>	26:24	Any value of the enumerated type ILPixTexUsage. See Table 5.19 on page 5-10.
	<code>coordmode</code>	28:27	Any value of the enumerated type ILTexCoordMode. See Table 5.18 on page 5-10.
	<code>cleartype mode</code>	29	1 = yes, 0 = no.
	<i>reserved</i>	31:30	Must be zero.
<i>Related</i>	None.		

Declare an Input Buffer

<i>Instructions</i>	DCL_RESOURCE			
<i>Syntax</i>	<code>dcl_resource_id(<i>n</i>)_type(<i>pixtexusage</i> [, <i>unnorm</i>])_fmtx(<i>fmt</i>)_fnty(<i>fmt</i>)_fmtz(<i>fmt</i>)_fmtw(<i>fmt</i>)</code>			
<i>Description</i>	<p>Declares an input buffer, specifies its dimension, and assigns it to a resource number. It identifies the resource type as a Buffer, Texture1D, Texture1DArray, Texture2D, Texture2DArray, Texture3D or TextureCube. Resources of type buffer can be used in an ld instruction. Resources of type texture * can as used in both ld and sample* instructions.</p> <p>Use the ILPixTexUsage enumeration for this field. The instruction can also indicate if a texture resource has been normalized. Return types identify the data type fetched from the input buffer on a per-component basis.</p> <p>Declares an input buffer, assigns it a number, and specifies its dimension.</p>			
<i>Opcode</i>	Token	Field Name	Bits	Description
	1	code	15:0	IL_OP_DCL_RESOURCE
		id	23:16	Resource id number, 0 to 255.
		type	27:24	<p>Must be set to any value of the enumerated type ILPixTexUsage (see Section 5.18, "ILPixTexUsage," page 5-10. Possible types: Buffer, Texture1D, Texture1DArray, Texture2D, Texture2DArray, Texture3D, or TextureCube. Use the ILPixTexUsage enumeration for this field.</p> <p>Resources of type buffer can be used in an ld instruction.</p> <p>Resources of type texture * can as used in both ld and sample* instructions.</p>
		<i>reserved</i>	30:24	Must be zero.
		unnorm	31	<p>0: Texture is normalized.</p> <p>1: Texture is not normalized.</p> <p>The optional <code>_unnorm</code> option in the textural source can be used to specify this bit.</p> <p>The <code>_unnorm</code> option in the textural source can be used to specify this bit.</p>
	2			<p>The return type identifies the data type fetched from the input buffer. Return types are four groups, each of three bits that must be set to any value of the enumerated type ILElementFormat. See Section 5.26, "ILElementFormat," page 5-13. Return-types are specified on a per-component basis, DX10 specification has no need to repeat identical return types; however, IL requires the type to be repeated all four times.</p>
		<i>reserved</i>	19:0	Must be zero.
		fmtx	22:20	x-component format.
		fnty	25:23	y-component format.
		fmtz	28:26	z-component format.
		fmtw	31:29	w-component format.
<i>Example</i>	<code>dcl_resource_id(1)_type(ld,unnorm)_fmtx(float)_fnty(float)_fmtz(float)_fmtw(float)</code>			
<i>Related</i>	None.			

Declare Mapping for Vertex Shader Inputs

<i>Instructions</i>	DCLV
<i>Syntax</i>	<code>dclv_elem(n)_x(comp)_y(comp)_z(comp)_w(comp) dst</code>
<i>Description</i>	<p>Declares a mapping between vertex buffer elements (logical streams) set through state and vertex shader inputs. It is an error to use this instruction in a pixel shader.</p> <p>The INITV and DCLV instruction are mutually exclusive for a given VERTEX register. There can be at most one DCLV per register. Use only registers of type TEMP and VERTEX. Each VERTEX register must either be declared at least once with a DCLV, or initialized with an INITV before it is used as a source in any instruction.</p> <p>The value of <code>elem</code> corresponds to vertex buffer element <i>n</i>.</p> <p>An <i>enabled</i> component is a component set to IL_IMPORTSEL_UNDEFINED, IL_IMPORTSEL_DEFAULT0, or IL_IMPORTSEL_DEFAULT1.</p> <p>A <i>disabled</i> component is a component set to IL_IMPORTSEL_UNUSED.</p> <p><i>Packed Registers:</i> A VERTEX register can be initialized multiple times with this instruction; thus, a single register can be mapped to multiple unique vertex buffer elements. However, the same component of a register cannot be used in both declarations: if in one declaration a component is <i>enabled</i>, the component must be <i>disabled</i> in the other declaration(s) of the register.</p> <p>The <i>j</i>th-enabled component receives the <i>j</i>th piece of data of the vertex buffer element specified by <code>elem</code> if it exists. If the data does not exist (the dimension of the vertex buffer element specified in state is less than <i>i</i>), the <i>i</i>th component receives the default value specified in this instruction.</p> <p>Do not use a destination modifier with this instruction. The <code>modifier_present</code> field of the IL_Dst token must be set to 0.</p> <p>Do not use relative addressing with this instruction. The <code>relative_address</code> field of the IL_Dst token must be set to 0.</p> <p>Operation:</p> <pre> VECTOR v; for (i=0; i < 4; i++) { if(i < ElementDimension(elem)) # AS_VS_DECL_TYPE_DIMENSION_N(elem) v[i]=FetchData(elem,currentIndex,i); else { v[i]=Default(i); } } WriteResult(v, dst); </pre>

Declare Mapping for Vertex Shader Inputs (Cont.)

<i>Format</i>	Token	Field Name	Bits	Description
	1	code	15:0	IL_OP_DCLV
		elem	21:16	Vertex buffer element, 0 to 63. These bits specify the vertex buffer element from which vertex data is loaded/fetched (elem).
		reserved	30:22	Must be zero.
		pri_modifier_present	31	If pri_modifier_present is set to 1, an IL_PrimaryDCLV_Mod token immediately follows this token.
	2	Primary vertex shader input register declaration modifier. IL_PrimaryDCLV_Mod token described below. The IL_PrimaryDCLV_Mod is present only if the pri_modifier_present field is 1 in the previous IL_OpCode token.		
		ximport	1:0	Specifies if the x component is enabled for the vertex buffer element specified by elem for this declaration. Also specifies the default value if the component is enabled. Can be any value of the enumerated type ILImportComponent. See Table 5.10 on page 5-6.
		yimport	3:2	Specifies if the y component is enabled for the vertex buffer element specified by elem for this declaration. Also specifies the default value if the component is enabled. Can be any value of the enumerated type ILImportComponent. See Table 5.10 on page 5-6.
		zimport	5:4	Specifies if the z component is enabled for the vertex buffer element specified by elem for this declaration. Also specifies the default value if the component is enabled. Can be any value of the enumerated type ILImportComponent. See Table 5.10 on page 5-6.
		wimport	7:6	Specifies if the w component is enabled for the vertex buffer element specified by elem for this declaration. Also specifies the default value if the component is enabled. Can be any value of the enumerated type ILImportComponent. See Table 5.10 on page 5-6.
		reserved	31:8	Must be zero.
	3	IL_Dst token with register_type set to IL_REGTYPE_TEMP or IL_REGTYPE_VERTEX (dst). The modifier_present and relative_address field must be set to 0.		
<i>Related</i>	None.			

Declare Vertex Shader Output Register

<i>Instructions</i>	DCLVOUT
<i>Syntax</i>	<code>dclvout_usage(op)_usageIndex(n)_x(comp)_y(comp)_z(comp)_w(comp) dst</code>
<i>Description</i>	<p>Declares the usage of a vertex shader output register, as well as a mapping of a vertex shader output to a pixel shader input.</p> <p>This instruction must be issued on each VOUTPUT register before the register is used in a shader. A shader using a VOUTPUT register cannot use the POS, SPRITE, INTERP, TEXCOORD, PRICOLOR, SECCOLOR, or FOG registers.</p> <p>There can be at most one DCLVOUT instruction per <code>usage-usageIndex</code> pair.</p> <p>An <i>enabled</i> component is a component set to <code>IL_IMPORTSEL_UNDEFINED</code>, <code>IL_IMPORTSEL_DEFAULT0</code>, or <code>IL_IMPORTSEL_DEFAULT1</code>.</p> <p>A <i>disabled</i> component is a component set to <code>IL_IMPORTSEL_UNUSED</code>.</p> <p>Packed Registers: A VOUTPUT register can be declared multiple times with this instruction; thus, a single register can have multiple unique <code>usage-usageIndex</code> pairs. However, the same component of a register cannot be enabled in both declarations: if in one declaration a component is enabled, the component must be disabled in the other declaration(s). For example, if <code>oOUT3</code> is declared as having <code>usage(interp)_usageIndex(1)_x (*)</code>, and <code>oOUT3</code> is also declared as having <code>usage(interp)_usageIndex(2)</code>, then the second declaration must set <code>x(-)</code>.</p> <p>If an <code>IL_PrimaryDCLVOUT_Mod</code> token is not preset, the shader behaves as if <code>xexport</code>, <code>yexport</code>, <code>zexport</code>, and <code>wexport</code> are set to <code>IL_IMPORTSEL_UNDEFINED</code>.</p> <p>Only one register can be declared to have the <code>usage IL_IMPORTUSAGE_POS</code>. If loop relative addressing is used on vertex shader outputs, that register can only be the VOUTPUT register number 0.</p> <p><code>usageIndex</code> must be if <code>IL_IMPORTUSAGE_POINTSIZE</code> is used.</p> <p>Only one register can be declared to have the <code>usage IL_IMPORTUSAGE_POINTSIZE</code>. If loop relative addressing is used on vertex shader outputs, that register can only be the VOUTPUT register number 1. <code>usageIndex</code> must be zero if <code>IL_IMPORTUSAGE_POINTSIZE</code> is used.</p> <p>Only one register can be declared to have the <code>usage IL_IMPORTUSAGE_FOG</code>. <code>usageIndex</code> must be zero in this case if <code>IL_IMPORTUSAGE_FOG</code> is used.</p> <p>It is an error to use this instruction in a pixel shader.</p> <p>You cannot use a source modifier with this instruction. The <code>modifier_present</code> field of the <code>IL_Dst</code> token must be set to 0.</p> <p>You cannot use relative addressing with this instruction. The <code>relative_address</code> field of the <code>IL_Dst</code> token must be set to 0.</p>

<i>Format</i>	<table border="0"> <thead> <tr> <th>Token</th> <th>Field Name</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>1</td> <td><code>code</code></td> <td>15:0</td> <td><code>IL_OP_DCLVOUT</code></td> </tr> <tr> <td></td> <td><code>usage</code></td> <td>20:16</td> <td>Any value of the enumerated type <code>ILImportUsage</code>. See Table 5.9 on page 5-4. The first five bits of the must be set to an acceptable value of the enumerated type <code>ILImportUsage (usage)</code>. The next eight bits must be a unique number for <code>usage</code> which specifies a mapping of <code>usage</code> type to a pixel shader input (<code>usageIndex</code>).</td> </tr> <tr> <td></td> <td><code>usageIndex</code></td> <td>28:21</td> <td>Unique number for <code>usage</code> which specifies a mapping of <code>usage</code> type to a pixel shader input (<code>usageIndex</code>), 0 to 255.</td> </tr> <tr> <td></td> <td><i>reserved</i></td> <td>30:29</td> <td>Must be zero.</td> </tr> <tr> <td></td> <td><code>pri_modifier_present</code></td> <td>31</td> <td>If <code>pri_modifier_present</code> is set to 1, an <code>IL_PrimaryDCLVOUT_Mod</code> token immediately follows this token.</td> </tr> </tbody> </table>	Token	Field Name	Bits	Description	1	<code>code</code>	15:0	<code>IL_OP_DCLVOUT</code>		<code>usage</code>	20:16	Any value of the enumerated type <code>ILImportUsage</code> . See Table 5.9 on page 5-4. The first five bits of the must be set to an acceptable value of the enumerated type <code>ILImportUsage (usage)</code> . The next eight bits must be a unique number for <code>usage</code> which specifies a mapping of <code>usage</code> type to a pixel shader input (<code>usageIndex</code>).		<code>usageIndex</code>	28:21	Unique number for <code>usage</code> which specifies a mapping of <code>usage</code> type to a pixel shader input (<code>usageIndex</code>), 0 to 255.		<i>reserved</i>	30:29	Must be zero.		<code>pri_modifier_present</code>	31	If <code>pri_modifier_present</code> is set to 1, an <code>IL_PrimaryDCLVOUT_Mod</code> token immediately follows this token.
Token	Field Name	Bits	Description																						
1	<code>code</code>	15:0	<code>IL_OP_DCLVOUT</code>																						
	<code>usage</code>	20:16	Any value of the enumerated type <code>ILImportUsage</code> . See Table 5.9 on page 5-4. The first five bits of the must be set to an acceptable value of the enumerated type <code>ILImportUsage (usage)</code> . The next eight bits must be a unique number for <code>usage</code> which specifies a mapping of <code>usage</code> type to a pixel shader input (<code>usageIndex</code>).																						
	<code>usageIndex</code>	28:21	Unique number for <code>usage</code> which specifies a mapping of <code>usage</code> type to a pixel shader input (<code>usageIndex</code>), 0 to 255.																						
	<i>reserved</i>	30:29	Must be zero.																						
	<code>pri_modifier_present</code>	31	If <code>pri_modifier_present</code> is set to 1, an <code>IL_PrimaryDCLVOUT_Mod</code> token immediately follows this token.																						

Declare Vertex Shader Output Register (Cont.)

Primary vertex shader output register declaration modifier. IL_PrimaryDCLVOUT_Mod token¹ described below.

Field Name	Bits	Description
xexport	1:0	Specifies if the x component is enabled for the usage-usageIndex for this declaration. Also specifies the default value if the component is enabled. Can be any value of the enumerated type ILImportComponent. See Table 5.10 on page 5-6.
yexport	3:2	Specifies if the y component is enabled for the usage-usageIndex for this declaration. Also specifies the default value if the component is enabled. Can be any value of the enumerated type ILImportComponent. See Table 5.10 on page 5-6.
zexport	5:4	Specifies if the z component is enabled for the usage-usageIndex for this declaration. Also specifies the default value if the component is enabled. Can be any value of the enumerated type ILImportComponent. See Table 5.10 on page 5-6.
wexport	7:6	Specifies if the w component is enabled for the usage-usageIndex for this declaration. Also specifies the default value if the component is enabled. Can be any value of the enumerated type ILImportComponent. See Table 5.10 on page 5-6.
reserved	31:8	Must be zero.

- 3 IL_Dst token (*dst*) where the `register_type` field is set to `IL_REGTYPE_VOUTPUT`. The `modifier_present` and `relative_address` fields must be set to 0.

Related None.

1. The `IL_PrimaryDCLVOUT_Mod` is present only if the `pri_modifier_present` field is 1 in the previous `IL_OpCode` token.

6.5 Input/Output Instructions

Finish Current Topology

<i>Instructions</i>	CUT		
<i>Syntax</i>	cut		
<i>Description</i>	<p>Completes the current primitive topology (if any vertices have been emitted), and any previous primitive topology, then starts a new topology of the type declared by the geometry shader. No vertices are leftover since a geometry shader can only emit pointlist, linestrip and trianglestrip topologies. When CUT is executed, any previously emitted topology by the Geometry Shader invocation is completed. If not enough vertices were emitted for the previous primitive topology, the extra vertices are discarded.</p> <p>After the previous topology (if any) is completed, CUT causes a new topology to begin, using the topology declared as the geometry shader's output.</p> <p>CUT can be used only in a geometry shader. This instruction can appear any number of times in a geometry shader and is executed implicitly at the end of an invocation of the geometry shader.</p>		
<i>Format</i>	0-input, 0-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_CUT
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	None.		

Pixel Kill

<i>Instructions</i>	KILL																											
<i>Syntax</i>	kill[_stage(n)_sample] <i>src0</i>																											
<i>Description</i>	<p>Component-wise test to terminate current pixel shader execution and discard all results.</p> <p>If sample is set to 1, KILL does not use the actual value of <i>src0</i> to test. Instead, the shader performs a KILL based upon a texture sample at the coordinate specified by <i>src0</i> on the texture stage/unit specified by stage.</p> <p>It is an error to use this instruction in a vertex or geometry shader.</p> <p>To kill based on a subset of the four components, set the swizzle for the component not used for the test to IL_COMPSEL_1.</p> <p>Operation:</p> <pre> VECTOR v; If(sample == 1) { v = Sample(src0, stage); } else { v = EvalSource(src0); } if((v[0] < 0.0) (v[1] < 0.0) (v[2] < 0.0) (v[3] < 0.0)) { Discard outputs; Terminate pixel shader; } </pre>																											
<i>Format</i>	1-input, 0-output.																											
<i>Opcode</i>	<table> <thead> <tr> <th>Field Name</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>code</td> <td>15:0</td> <td>IL_OP_KILL</td> </tr> <tr> <td>control</td> <td>29:16</td> <td> <table> <thead> <tr> <th>Field Name</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>stage</td> <td>23:16</td> <td>Texture stage or unit number, 0 to 255, if sample is 1; otherwise, must be zero.</td> </tr> <tr> <td>sample</td> <td>24</td> <td>0: <i>src0</i> is the test value. 1: Test value is sampled from texture stage or unit <i>stage</i>. <i>src0</i> is the sample coordinate, and the resulting sample is used as the test value.</td> </tr> <tr> <td><i>reserved</i></td> <td>29:25</td> <td>Must be zero.</td> </tr> </tbody> </table> </td> </tr> <tr> <td></td> <td>sec_modifier_present 30</td> <td>Must be zero.</td> </tr> <tr> <td></td> <td>pri_modifier_present 31</td> <td>Must be zero.</td> </tr> </tbody> </table> <p>IL_Src token (<i>src0</i>)</p> <p>IL_Src_Mod token: only present if modifier_present field is 1 in previous IL_Src token.</p> <p>IL_Src token (<i>src0</i>): only present if relative_address field is 1 in the preceding IL_Src or IL_Dst token.</p>	Field Name	Bits	Description	code	15:0	IL_OP_KILL	control	29:16	<table> <thead> <tr> <th>Field Name</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>stage</td> <td>23:16</td> <td>Texture stage or unit number, 0 to 255, if sample is 1; otherwise, must be zero.</td> </tr> <tr> <td>sample</td> <td>24</td> <td>0: <i>src0</i> is the test value. 1: Test value is sampled from texture stage or unit <i>stage</i>. <i>src0</i> is the sample coordinate, and the resulting sample is used as the test value.</td> </tr> <tr> <td><i>reserved</i></td> <td>29:25</td> <td>Must be zero.</td> </tr> </tbody> </table>	Field Name	Bits	Description	stage	23:16	Texture stage or unit number, 0 to 255, if sample is 1; otherwise, must be zero.	sample	24	0: <i>src0</i> is the test value. 1: Test value is sampled from texture stage or unit <i>stage</i> . <i>src0</i> is the sample coordinate, and the resulting sample is used as the test value.	<i>reserved</i>	29:25	Must be zero.		sec_modifier_present 30	Must be zero.		pri_modifier_present 31	Must be zero.
Field Name	Bits	Description																										
code	15:0	IL_OP_KILL																										
control	29:16	<table> <thead> <tr> <th>Field Name</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>stage</td> <td>23:16</td> <td>Texture stage or unit number, 0 to 255, if sample is 1; otherwise, must be zero.</td> </tr> <tr> <td>sample</td> <td>24</td> <td>0: <i>src0</i> is the test value. 1: Test value is sampled from texture stage or unit <i>stage</i>. <i>src0</i> is the sample coordinate, and the resulting sample is used as the test value.</td> </tr> <tr> <td><i>reserved</i></td> <td>29:25</td> <td>Must be zero.</td> </tr> </tbody> </table>	Field Name	Bits	Description	stage	23:16	Texture stage or unit number, 0 to 255, if sample is 1; otherwise, must be zero.	sample	24	0: <i>src0</i> is the test value. 1: Test value is sampled from texture stage or unit <i>stage</i> . <i>src0</i> is the sample coordinate, and the resulting sample is used as the test value.	<i>reserved</i>	29:25	Must be zero.														
Field Name	Bits	Description																										
stage	23:16	Texture stage or unit number, 0 to 255, if sample is 1; otherwise, must be zero.																										
sample	24	0: <i>src0</i> is the test value. 1: Test value is sampled from texture stage or unit <i>stage</i> . <i>src0</i> is the sample coordinate, and the resulting sample is used as the test value.																										
<i>reserved</i>	29:25	Must be zero.																										
	sec_modifier_present 30	Must be zero.																										
	pri_modifier_present 31	Must be zero.																										
<i>Related</i>	DISCARD_LOGICALNZ, DISCARD_LOGICALZ.																											

Discard Results Based on Integer

<i>Instructions</i>	DISCARD_LOGICALNZ, DISCARD_LOGICALZ		
<i>Syntax</i>	Opcode	Syntax	Description
	IL_OP_DISCARD_LOGICALNZ	discard_logicalnz <i>src0.select_component</i>	Discard results if all bits in <i>src0.{x y z w} ≠ 0</i> .
	IL_OP_DISCARD_LOGICALZ	discard_logicalz <i>src0.select_component</i>	Discard results if all bits in <i>src0.{x y z w} == 0</i> .
<i>Description</i>	<p>Conditionally flags results of pixel shader to be discarded when the end of the program is reached. This instruction flags the current pixel as terminated, while continuing execution, so that other pixels executing in parallel can obtain gradients, if necessary. Although execution continues, all pixel shader output writes before or after the discard_* instruction are discarded.</p> <p>The discard_* instruction can be present inside any flow control construct.</p> <p>Multiple discard instructions can be present in a pixel shader, and if any is executed, the pixel is terminated.</p> <p>Can be used only in a pixel shader.</p>		
<i>Format</i>	1-input, 0-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	See Opcode part of <i>Syntax</i> , above.
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	None. KILL.		

Emit a Vertex

<i>Instructions</i>	EMIT		
<i>Syntax</i>	emit		
<i>Description</i>	<p>Causes all declared o# registers to be read out of a geometry shader to generate a vertex. Multiple EMIT instructions are used to generate a primitive. Any number of EMIT instructions can appear in a geometry shader, including within flow control. This instruction can be used only in a geometry shader.</p>		
<i>Format</i>	0-input, 0-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_EMIT
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	None.		

Emit Followed by Cut

<i>Instructions</i>	EMIT_THEN_CUT		
<i>Syntax</i>	emitcut		
<i>Description</i>	Operation is the same as an EMIT instruction immediately followed by a CUT instruction. It has the same restrictions as the union of restrictions for the EMIT and CUT commands.		
<i>Format</i>	0-input, 0-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_EMIT_THEN_CUT
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	CUT, EMIT.		

LOAD from Buffer

<i>Instructions</i>	LOAD
<i>Syntax</i>	load_resource(<i>n</i>)[_aoffimmi(<i>u,v,w</i>)] <i>dst</i> , <i>src0</i>
<i>Description</i>	<p>Simplified alternative to the “sample” instruction. Using the provided signed integer address, LOAD fetches data from the specified buffer or texture without any filtering (for example, point sampling). The source data can come from any resource type except TextureCube.</p> <p><i>src0</i> must specify a single component used as the address.</p> <p>DX10 allows an output swizzle on the ld instruction. IL requires an additional move if the swizzle is used.</p> <p>Unlike SAMPLE, LOAD can fetch data from buffers. The buffer with the data being fetched is identified by the resource id stored in the control field.</p> <p><i>src0</i> provides the set of texture coordinates needed to sample the texture in the form of signed integers.</p> <p>If the <i>srcAddress</i> is out of the range [0, (#texels in dimension -1)], the load returns 0.</p> <p><i>src0.a</i> (post-swizzle) always provides a signed integer mipmap level. If the value is out of range [0, (num miplevels in resource-1)], the load returns 0. If the resource has no mipmaps, <i>src0.a</i> is ignored.</p> <p><i>src0.gb</i> (post-swizzle) are ignored for buffers and Texture1D (non-Array). <i>srcAddress.b</i> (post-swizzle) is ignored for Texture1D Arrays and Texture2Ds.</p> <p>For Texture1D Arrays, <i>src0.g</i> (post-swizzle) provides the array index as a signed integer. If the value is out of range of the available array (indices [0, (Array size-1)]), the load returns 0.</p> <p>For Texture2D Arrays, <i>src0.b</i> (post-swizzle) provides the array index; otherwise, it has the same semantics as Texture1D, described above.</p>

LOAD from Buffer (Cont.)

Address Offset

- If the optional `_aoffimmi` suffix is included (bit 29 of the Opcode token is set), a 32-bit value containing the packed offsets immediately follows the Opcode token (and modifier, if present).
- The `optional` field indicates that the texture coordinates for the load are to be offset by the provided immediate texel space integer constant values. The literal values must be a set of unnormalized (texel space) values represented as signed fixed point with one bit of fraction (7.1) that allows ranges from [-64, 63.5]. `src1` can be used with all resources, including Texture1D/2D Arrays and Texture3D; however, it cannot be used with a resource of type TextureCube.
- The offsets are added to the texture coordinates, in texel space, relative to the miplevel being accessed by the load.
- Address offsets are not applied along the array axis of Texture1D/2D Arrays.
- `_aoffimmi v,w` components are ignored for Buffers, and Texture1Ds.
- `_aoffimmi w` component is ignored for Texture2Ds.

Return Type Control

The data format returned by LOAD to the `dst` register is determined in the same way as described for the sample instruction; it is based on the format bound to the Resource parameter (`n`).

As with the SAMPLE instruction, returned values for load are four-component vectors (with format-specific defaults for components not present in the format).

DX10 allows a swizzle on the resource that is used to determine how to swizzle the four-component result back from the texture load, after which `.mask` on `dst` is used to determine which components in `dst` is updated. The equivalent effect can be achieved in IL by adding an extra move after LOAD.

The resource cannot be a TextureCube or a Constant Buffer.

Out-of-bounds access on any axis always results in 0 as the result of the memory fetch.

Formats

1-input, 1-output.

Opcode

Field Name	Bits	Description															
<code>code</code>	15:0	IL_OP_LOAD															
<code>control</code>	29:16	<table border="0"> <thead> <tr> <th>Field Name</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td><code>resource</code></td> <td>23:16</td> <td><code>resource_id</code>.</td> </tr> <tr> <td><code>sampler</code></td> <td>27:24</td> <td><code>sampler_id</code>.</td> </tr> <tr> <td><code>reserved</code></td> <td>28:24</td> <td>Must be zero.</td> </tr> <tr> <td><code>aoffimmi</code></td> <td>29</td> <td>0 = <code>aoffimmi</code> does not exist. 1 = <code>aoffimmi</code> exists.</td> </tr> </tbody> </table>	Field Name	Bits	Description	<code>resource</code>	23:16	<code>resource_id</code> .	<code>sampler</code>	27:24	<code>sampler_id</code> .	<code>reserved</code>	28:24	Must be zero.	<code>aoffimmi</code>	29	0 = <code>aoffimmi</code> does not exist. 1 = <code>aoffimmi</code> exists.
Field Name	Bits	Description															
<code>resource</code>	23:16	<code>resource_id</code> .															
<code>sampler</code>	27:24	<code>sampler_id</code> .															
<code>reserved</code>	28:24	Must be zero.															
<code>aoffimmi</code>	29	0 = <code>aoffimmi</code> does not exist. 1 = <code>aoffimmi</code> exists.															
<code>sec_modifier_present</code>	30	Must be zero.															
<code>pri_modifier_present</code>	31	Must be zero.															

Related

SAMPLE.

Get Texture Mipmap Level of Detail

Instructions **LOD**

Syntax lod_stage(*n*) *dst*, *src0*

Description Uses the provided texture coordinate to determine the computed mipmap level(s) sampled from at *src0*. The LOD value returned includes any clamping and biasing defined by the texture currently bound to the specified texture unit. This instruction cannot be used on a stage that has not been declared with a DCLPT instruction.

You can project using the `divComp` source modifier on *src0*. If the `divComp` source modifier is used, and the component to divide by is negative, the result of this instruction is undefined (if `IL_DIVCOMP_Y` is used, and the second component of *src0* is negative, the results of this texture load is undefined).

Operation:

src0 is the texture coordinate with which the mipmap LOD is determined. The control field of the `IL_ OpCode` token specifies the texture stage to determine the LOD. Associated with the sampler specified by control are 1) a texture, 2) the texture's dimension, and 3) all filter, wrap, bias, and clamp states. The LOD value is replicated on each component of *dst*.

Format 1-input, 1-output.

Opcode

Token	Field Name	Bits	Description									
1	code	15:0	IL_OP_LOD The first eight bits specify the texture stage/unit from which to determine the texture's level of detail (stage).									
	control	29:16	<table border="0"> <thead> <tr> <th>Field Name</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>stage</td> <td>23:16</td> <td>Texture stage or unit number.</td> </tr> <tr> <td>reserved</td> <td>29:24</td> <td>Must be zero.</td> </tr> </tbody> </table>	Field Name	Bits	Description	stage	23:16	Texture stage or unit number.	reserved	29:24	Must be zero.
Field Name	Bits	Description										
stage	23:16	Texture stage or unit number.										
reserved	29:24	Must be zero.										
	sec_modifier_present	30	Must be zero.									
	pri_modifier_present	31	Must be zero.									
2	IL_Dst token (<i>dst</i>)											
3	IL_Dst_Mod token is present only if the modifier_present field is 1 in previous IL_Dst token.											
4	IL_Rel_Addr token is present only if the relative_address field is 1 in the preceding IL_Src or IL_Dst token.											
5	IL_Src token (<i>src0</i>).											
6	IL_Src_Mod token is present only if the modifier_present field is 1 in previous IL_Src token.											
7	IL_Rel_Addr token is present only if the relative_address field is 1 in the preceding IL_Src or IL_Dst token.											

Related None.

Export Data to Memory Stream

Instructions MEMEXPORT

Syntax memexport_exportStream(*n*)_elemOffset(*i*) *src0*, *src1*

Description Exports a value in *src1* to a memory stream. *exportStream* corresponds to the number of the export stream buffer to which data is written. *elemOffset* specifies the offset in terms of elements (not bytes or dwords) to which data is written. *src0.x* contains the index to which data is written. *src1* contains the data to write to the export stream.

<i>Format</i>	Token	Field Name	Bits	Description
	1	code	15:0	IL_OP_MEMEXPORT
		exportStream	21:16	Export stream number, 0 to 16
		<i>reserved</i>	22	Must be zero.
		elemOffset	28:23	Offset from the start of the export stream in elements.
		<i>reserved</i>	31:29	Must be zero.
	2	IL_Src token (<i>src0</i>).		
	3	IL_Src_Mod token is present only if the <i>modifier_present</i> field is 1 in the previous IL_Src token.		
	4	IL_Rel_Addr token is present only if the <i>relative_address</i> field is 1 in the preceding IL_Src token.		
	5	IL_Src token (<i>src1</i>).		
	6	IL_Src_Mod token is present only if the <i>modifier_present</i> field is 1 in the previous IL_Src token.		
	7	IL_Rel_Addr token is present only if the <i>relative_address</i> field is 1 in the preceding IL_Src token.		

Related MEMEXPORT, SCATTER_QUAD.

Import Memory from Stream

Instructions MEMIMPORT

Syntax memimport_elem(*n*) *dst*, *src0*

Description Imports data from buffers set externally. This instruction provides an alternate means of fetching vertex element data using an arbitrary index. Refer the DCLV instruction for normal operation. This instruction typically is used only for higher-order surface shaders (see the INITV instruction).

In a vertex shader:

- the value of *elem* corresponds to vertex buffer element *n*.
- if the vertex buffer corresponding to *elem* is disabled in state, apply defaults when this instruction is executed.

In a pixel shader

- the value of *elem* corresponds to pixel buffer element *n*.
- if the pixel buffer corresponding to *elem* is disabled in state, apply defaults when this instruction is executed.

The first component of *src0* (*src0.x*) contains the index for the element from which to fetch. The value is floored before it is used.

dst contains the data to be fetched from memory.

dst can only be a TEMP register.

Operation:

```
VECTOR v;
v = Fetch(importElement, FLOOR(src0.x));
WriteResult(v, dst);
```

Format	Token	Field Name	Bits	Description
	1	code	15:0	IL_OP_MEMIMPORT The lower six bits are set to a number between 0 and 16 inclusive, specifying the import element from which memory is loaded/fetched (<i>elem</i>).
		<i>elem</i>	21:16	Element number, 0 to 16
		<i>reserved</i>	31:22	Must be zero.
	2	IL_Dst token (<i>dst</i>) where <i>register_type</i> is set to IL_REGTYPE_TEMP.		
	3	IL_Dst_Mod token is present only if the <i>modifier_present</i> field is 1 in the previous IL_Dst token.		
	4	IL_Rel_Addr token is present only if the <i>relative_address</i> field is 1 in the preceding IL_Src or IL_Dst token.		
	5	IL_Src token (<i>src0</i>).		
	6	IL_Src_Mod token is present only if the <i>modifier_present</i> field is 1 in the previous IL_Src token.		
	7	IL_Rel_Addr token is present only if the <i>relative_address</i> field is 1 in the preceding IL_Src or IL_Dst token.		

Related MEMEXPORT, SCATTER_QUAD.

Query Information from a Resource

Instructions **RESINFO**

Syntax `resinfo_resource(n)[_uint] dst, src0`

Description `src0.x` is the 2's complement integer scalar.

DX10 supports both destination swizzles and repFloat modifiers on this instruction. Clients must insert additional IL instructions to achieve the same effect.

`dst` receives `<width, height, depth, total mip count>`.

If `src0.x` is out of the range of the available number of miplevels in the resource, `[0, (numMipLevels - 1)]`, then `resinfo` returns `[0, 0, 0, total mip count]`.

The returned width, height, and depth values are for the mip-level selected by `src0.x`; they are in number of texels, independent of texel data size.

Returned values are all floating point, unless the control field is negative (`_uint` modifier is used), in which case the returned values are integers.

`dst.w` always receives total-mip-count (if `.w` is included in write mask). Thus, the total-mip-count is independent of `src0.x`.

DX10 allows a destination swizzle. IL, however, requires an extra move to allow the returned values to be swizzled arbitrarily before they are written to the destination.

If the resource is not a Texture3D or Texture1D/2D with Array > 1, then depth is zero.

For a Texture1D/2D with Array > 1 or Texture3D, depth always represents the number of array slices.

If the resource is a Buffer or Texture1D, the height is zero, unless it is a Texture1D Array, in which case height is the number of array slices.

If the resource is a TextureCube, then width and height represent individual cube face dimensions, and depth is zero.

If `src0.x` is out of the range of the available number of miplevels in the resource, then `resinfo` returns `[0, 0, 0, total mip count]`.

Formats 1-input, 1-output.

<i>Opcode</i>	Field Name	Bits	Description												
	<code>code</code>	15:0	IL_OP_RESINFO												
	<code>control</code>	29:16	<table border="0"> <tr> <td>Field Name</td> <td>Bits</td> <td>Description</td> </tr> <tr> <td><code>resource_id</code></td> <td>23:16</td> <td>Resource ID</td> </tr> <tr> <td><code>return_type</code></td> <td>24</td> <td>0x0: Return type is Float. 0x1: Return type is Integer.</td> </tr> <tr> <td><code>reserved</code></td> <td>29:25</td> <td>Must be zero.</td> </tr> </table>	Field Name	Bits	Description	<code>resource_id</code>	23:16	Resource ID	<code>return_type</code>	24	0x0: Return type is Float. 0x1: Return type is Integer.	<code>reserved</code>	29:25	Must be zero.
Field Name	Bits	Description													
<code>resource_id</code>	23:16	Resource ID													
<code>return_type</code>	24	0x0: Return type is Float. 0x1: Return type is Integer.													
<code>reserved</code>	29:25	Must be zero.													
	<code>sec_modifier_present</code>	30	Must be zero.												
	<code>pri_modifier_present</code>	31	Must be zero.												

Related None.

Query Information from a Resource

Instructions **SAMPLEINFO**

Syntax `sampleinfo_resource(n)[_uint] dst, src0`

Description `dst` receives a number in components `x`. The returned value is floating point, unless the control field is negative (`_uint` modifier is used), in which case the returned value is integers. If the resource is not a multi-sample resource and not a render target, the result is 0.

Formats 0-input, 1-output.

Opcode

Field Name	Bits	Description																		
<code>code</code>	15:0	IL_OP_SAMPLEINFO																		
<code>control</code>	29:16	<table border="0"> <thead> <tr> <th style="text-align: left;">Field Name</th> <th style="text-align: left;">Bits</th> <th style="text-align: left;">Description</th> </tr> </thead> <tbody> <tr> <td><code>resource_id</code></td> <td>23:16</td> <td>Resource ID</td> </tr> <tr> <td><code>return_type</code></td> <td>24</td> <td>0x0: Return type is Float. 0x1: Return type is Integer.</td> </tr> <tr> <td><code>reserved</code></td> <td>29:25</td> <td>Must be zero.</td> </tr> <tr> <td><code>sec_modifier_present</code></td> <td>30</td> <td>Must be zero.</td> </tr> <tr> <td><code>pri_modifier_present</code></td> <td>31</td> <td>Must be zero.</td> </tr> </tbody> </table>	Field Name	Bits	Description	<code>resource_id</code>	23:16	Resource ID	<code>return_type</code>	24	0x0: Return type is Float. 0x1: Return type is Integer.	<code>reserved</code>	29:25	Must be zero.	<code>sec_modifier_present</code>	30	Must be zero.	<code>pri_modifier_present</code>	31	Must be zero.
Field Name	Bits	Description																		
<code>resource_id</code>	23:16	Resource ID																		
<code>return_type</code>	24	0x0: Return type is Float. 0x1: Return type is Integer.																		
<code>reserved</code>	29:25	Must be zero.																		
<code>sec_modifier_present</code>	30	Must be zero.																		
<code>pri_modifier_present</code>	31	Must be zero.																		

Related None.

Query Position Information from a Resource

Instructions **SAMPLEPOS**

Syntax `samplepos_resource(n)[_uint] dst, src0`

Description The returned value is a float4 (`x,y,0,0`) indicating where the sample is located. If the resource is not a multi-sample resource and not a render target, the result is 0.

Formats 1-input, 1-output.

Opcode

Field Name	Bits	Description																		
<code>code</code>	15:0	IL_OP_SAMPLEPOS																		
<code>control</code>	29:16	<table border="0"> <thead> <tr> <th style="text-align: left;">Field Name</th> <th style="text-align: left;">Bits</th> <th style="text-align: left;">Description</th> </tr> </thead> <tbody> <tr> <td><code>resource_id</code></td> <td>23:16</td> <td>Resource ID</td> </tr> <tr> <td><code>return_type</code></td> <td>24</td> <td>0x0: Return type is Float. 0x1: Return type is Integer.</td> </tr> <tr> <td><code>reserved</code></td> <td>29:25</td> <td>Must be zero.</td> </tr> <tr> <td><code>sec_modifier_present</code></td> <td>30</td> <td>Must be zero.</td> </tr> <tr> <td><code>pri_modifier_present</code></td> <td>31</td> <td>Must be zero.</td> </tr> </tbody> </table>	Field Name	Bits	Description	<code>resource_id</code>	23:16	Resource ID	<code>return_type</code>	24	0x0: Return type is Float. 0x1: Return type is Integer.	<code>reserved</code>	29:25	Must be zero.	<code>sec_modifier_present</code>	30	Must be zero.	<code>pri_modifier_present</code>	31	Must be zero.
Field Name	Bits	Description																		
<code>resource_id</code>	23:16	Resource ID																		
<code>return_type</code>	24	0x0: Return type is Float. 0x1: Return type is Integer.																		
<code>reserved</code>	29:25	Must be zero.																		
<code>sec_modifier_present</code>	30	Must be zero.																		
<code>pri_modifier_present</code>	31	Must be zero.																		

Related None.

Sample Data from Resource with Filter

<i>Instructions</i>	SAMPLE
<i>Syntax</i>	<code>sample_resource(n)_sampler(m)[_aoffinmi(u, v, w)] dst, src0</code>
<i>Description</i>	<p>Samples data from the specified Element/texture using the filtering mode identified by the given sampler. The source data can come from any resource type other than buffers. <i>src0</i> provides the set of texture coordinates needed sampling as floating point values, referencing normalized space in the texture. Address wrapping modes (wrap, mirror, clamp, border, etc.) are applied to texture coordinates outside [0...1] range, taken from the sampler state, and applied after any address offset is applied to texture coordinates (see Address Offset section, below).</p> <p>For Texture2D Arrays, <i>src0.b</i> (post-swizzle) selects the Array Slice from which to fetch and uses the same semantics described for Texture1D Arrays.</p> <p>Address Offset</p> <p>The optional <i>_aoffinmi</i> suffix is specified by the <i>addroff_present</i> bit in <i>IL_SampleOpCode</i>. If set, a 32 bit value containing the packed offsets immediately follows the opcode token and modifier, if present.</p> <p>The optional <i>address offset by immediate integer</i> indicates that the texture coordinates for the sample are offset by the provided immediate texel space integer constant values. The literal values must be a set of unnormalized (texel space) values represented as Signed fixed point with 1 bit of fraction (7.1) that allows ranges from [-64,63.5]. This modifier is defined for all resources, including Texture1D/2D Arrays and Texture3D; but it is undefined for TextureCube.</p> <p>The offsets are added to the texture coordinates, in texel space, relative to each mip level being accessed. Thus, even though texture coordinates are provided as normalized float values, the offset applies a texel-space integer offset. Address offsets are not applied along the array axis of Texture1D/2D Arrays.</p> <p>The v and w components are ignored for buffers and Texture1Ds.</p> <p>The w component is ignored for Texture2Ds.</p> <p>Address wrapping modes (wrap, mirror, clamp, border, etc.) from the sampler state are applied after any address offset is applied to texture coordinates.</p> <p>Return Type Control</p> <p>The data format returned by sample to the destination register is determined by the resource format (<i>WGFFMT*</i>). For example, if the resource format is <i>WGFFMT_A8B8G8R8_UNORM_SRGB</i>, the sampling operation converts sampled texels from gamma 2.0 to 1.0, applies filtering, and the result are written to the destination register as floating point values in the range [0,1].</p> <p>Returned values are 4-vectors (with format-specific defaults for components not present in the format).</p> <p>DX10 supports an output swizzle in this instruction. To achieve the same effect in IL, use an extra move.</p>

Sample Data from Resource with Filter (Cont.)

LOD Calculation

See the `deriv_rtx` and `deriv_rty` instructions on how derivatives are calculated while determining LOD for filtering. The `sample` instruction implicitly computes derivatives on the texture coordinates using the same definition that the `deriv*` Shader instructions use. This does not apply to `sample_l` or `sample_g` instructions; for these, LOD or gradients are provided directly by the application.

Miscellaneous Details

`src0.w` (post-swizzle) must be zero. `src0` provides the set of texture coordinates to perform the sample as floating point values referencing normalized space within the texture. Address wrapping modes (wrap, mirror, clamp, border, etc.) are applied for texture coordinates outside the [0, 1] range, taken from the sampler state, and applied after any address offset (see subsection, above) is applied to texture coordinates.

The information required for the hardware to perform sampling is split into two orthogonal parts. The first part (texture register), provides source data type information, including if the texture contains SRGB data; also, it references the memory being sampled. The second part (sampler register), defines the filtering mode to apply to the sample operation.

For Texture1D Arrays, `src0.y` (post-swizzle) selects which Array Slice to fetch from. This is treated as a scaled float value, as opposed to the normalized space for standard texture coordinates; and a round-to-nearest even is applied on the value, followed by a clamp to the available BufferArray range.

Formats 1-input, 1-output or 3-input, 1-output.

<i>Opcode</i>	Field Name	Bits	Description															
	<code>code</code>	15:0	IL_OP_SAMPLE															
	<code>control</code>	29:16	<table border="0" style="margin-left: 20px;"> <thead> <tr> <th style="text-align: left;">Field Name</th> <th style="text-align: left;">Bits</th> <th style="text-align: left;">Description</th> </tr> </thead> <tbody> <tr> <td><code>resource</code></td> <td>23:16</td> <td>resource_id, 0 to 255.</td> </tr> <tr> <td><code>sampler</code></td> <td>27:24</td> <td>sampler_id, 0 to 15.</td> </tr> <tr> <td><code>reserved</code></td> <td>28</td> <td>Must be zero.</td> </tr> <tr> <td><code>aoffimmi</code></td> <td>29</td> <td>0 = aoffimmi does not exist. 1 = aoffimmi exists.</td> </tr> </tbody> </table>	Field Name	Bits	Description	<code>resource</code>	23:16	resource_id, 0 to 255.	<code>sampler</code>	27:24	sampler_id, 0 to 15.	<code>reserved</code>	28	Must be zero.	<code>aoffimmi</code>	29	0 = aoffimmi does not exist. 1 = aoffimmi exists.
Field Name	Bits	Description																
<code>resource</code>	23:16	resource_id, 0 to 255.																
<code>sampler</code>	27:24	sampler_id, 0 to 15.																
<code>reserved</code>	28	Must be zero.																
<code>aoffimmi</code>	29	0 = aoffimmi does not exist. 1 = aoffimmi exists.																
	<code>sec_modifier_present</code>	30	Must be zero.															
	<code>pri_modifier_present</code>	31	Must be zero.															

Related LOAD, SAMPLE_B, SAMPLE_C, SAMPLE_C_LZ, SAMPLE_G, SAMPLE_L.SAMPLE_L.

Sample Data from Resource with Filter and Bias

Instructions **SAMPLE_B**

Syntax `sample_b_resource(n)_sampler(m)[_aoffimmi(u, v, w)] dst, src0 , src1`

Description Samples data from the specified element/texture using the filtering mode identified by the given sampler. The source data can come from any resource type except for buffers. An additional bias is applied to the level of detail computed as part of the instruction execution. SAMPLE_B is like the SAMPLE instruction with the addition of applying the specified srcLOBBias (in *src1*) value to the level of detail value prior to selecting the mip map. *src1* must be either a literal value or a replicated single component. SAMPLE_B has the same restrictions as the SAMPLE instruction.

Formats 2-input, 1-output or 4-input, 1-output.

Opcode

Field Name	Bits	Description															
code	15:0	IL_OP_SAMPLE_B															
control	29:16	<table border="0"> <thead> <tr> <th>Field Name</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>resource</td> <td>23:16</td> <td>resource_id.</td> </tr> <tr> <td>sampler</td> <td>27:24</td> <td>sampler_id.</td> </tr> <tr> <td><i>reserved</i></td> <td>28</td> <td>Must be zero.</td> </tr> <tr> <td>aoffimmi</td> <td>29</td> <td>0 = aoffimmi does not exist. 1 = aoffimmi exists.</td> </tr> </tbody> </table>	Field Name	Bits	Description	resource	23:16	resource_id.	sampler	27:24	sampler_id.	<i>reserved</i>	28	Must be zero.	aoffimmi	29	0 = aoffimmi does not exist. 1 = aoffimmi exists.
Field Name	Bits	Description															
resource	23:16	resource_id.															
sampler	27:24	sampler_id.															
<i>reserved</i>	28	Must be zero.															
aoffimmi	29	0 = aoffimmi does not exist. 1 = aoffimmi exists.															
sec_modifier_present	30	Must be zero.															
pri_modifier_present	31	Must be zero.															

Related LOAD, SAMPLE, SAMPLE_C, SAMPLE_C_LZ, SAMPLE_G, SAMPLE_L.

Sample Data from Resource with Filter and Gradient

<i>Instructions</i>	SAMPLE_G														
<i>Syntax</i>	<code>sample_g_resource(n)_sampler(m)[_aoffimimmi(u, v, w)] dst, src0, src1, src2</code>														
<i>Description</i>	<p>Samples data from the specified element/texture using the filtering mode identified by the given sampler. The source data can come from any Resource Type, other than Buffers.</p> <p>SAMPLE_G behaves exactly as the “sample” instruction, except that gradients for the source address in the x direction and the y direction are provided by extra parameters, <i>src1</i> and <i>src2</i>, respectively.</p> <p>The x, y, and z components of <i>src1/src2</i> (post-swizzle) provide du/dx, dv/dx and dw/dx. The w component (post-swizzle) is ignored.</p> <p>The w component (post-swizzle) is ignored.</p> <p>This instruction has the same restrictions as the SAMPLE instruction.</p>														
<i>Formats</i>	3-input, 1-output.														
<i>Opcode</i>	Field Name	Bits	Description												
	code	15:0	IL_OP_SAMPLE_G												
	control	29:16	<table border="0"> <thead> <tr> <th>Field Name</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>resource</td> <td>23:16</td> <td>resource_id.</td> </tr> <tr> <td>sampler</td> <td>27:24</td> <td>sampler_id.</td> </tr> <tr> <td>reserved</td> <td>29:28</td> <td>Must be zero.</td> </tr> </tbody> </table>	Field Name	Bits	Description	resource	23:16	resource_id.	sampler	27:24	sampler_id.	reserved	29:28	Must be zero.
Field Name	Bits	Description													
resource	23:16	resource_id.													
sampler	27:24	sampler_id.													
reserved	29:28	Must be zero.													
	sec_modifier_present	30	Must be zero.												
	pri_modifier_present	31	Must be zero.												
<i>Related</i>	LOAD, SAMPLE, SAMPLE_B, SAMPLE_C, SAMPLE_C_LZ, SAMPLE_L.														

Sample Data from Specified Memory Using Given Sampler. Source Data can come from any Resource Type

<i>Instructions</i>	SAMPLE_L				
<i>Syntax</i>	<code>sample_l_resource(n)_sampler(m) dst, src0, src1</code>				
<i>Description</i>	<p>This is identical to the SAMPLE instruction (6-58), except that the level of detail (LOD) is provided directly by the application as a scalar value, representing no anisotropy. This instruction also is available in all programmable shader stages, not only the pixel shader (as with SAMPLE).</p> <p>It samples the texture using <code>src1.x</code> to choose the LOD. If the LOD value is negative, the 0th (biggest map) is chosen with MAGFILTER applied. Since <code>src1.x</code> is a floating point value, the fractional value is used to interpolate (if MIPFILTER is linear) between two mip levels.</p> <p>This instruction ignores address gradients (filtering is isotropic).</p> <p>See the description of the SAMPLE instruction for operational details of this instruction other than the LOD calculation.</p> <p>Note that when used in the pixel kernel, <code>sample_l</code> implies the choice of LOD is per-pixel, with no effect from neighboring pixels.</p>				
<i>Formats</i>	2-input, 1-output.				
<i>Opcode</i>	Field Name	Bits	Description		
	<code>code</code>	15:0	IL_OP_SAMPLE_L		
	<code>control</code>	29:16	Field Name	Bits	Description
			<code>resource</code>	23:16	<code>resource_id</code> .
			<code>sampler</code>	27:24	<code>sampler_id</code> .
			<code>reserved</code>	29:28	Must be zero.
	<code>sec_modifier_present</code>	30	Must be zero.		
	<code>pri_modifier_present</code>	31	Must be zero.		
<i>Related</i>	LOAD, SAMPLE_B, SAMPLE_C, SAMPLE_C_LZ, SAMPLE_G, SAMPLE_L, SAMPLE.				

Sample Data from Resource with Filter and Comparison

Instructions **SAMPLE_C_LZ**

Syntax `sample_c_lz_resource(m)_sampler(n)_aoffimimmi(u, v, w) dst, src0, src1`

Description Performs a comparison filter. SAMPLE_C mainly provides a building-block for Percentage-Closer Depth filtering. The 'c' in SAMPLE_C stands for comparison.
 src0 is the index. src1.x contains the reference value.
 Same as SAMPLE_C, except LOD is 0, and derivatives are ignored (as if they are 0).
 The LZ stands for level-zero. Because derivatives are ignored, this instruction is available in vertex and geometry shaders. It can also be used inside of control flow.

Formats 2-input, 1-output or 4-input, 1-output.

Opcode

Field Name	Bits	Description												
code	15:0	IL_OP_SAMPLE_B												
control	29:16	<table border="0" style="margin-left: 20px;"> <thead> <tr> <th style="text-align: left;">Field Name</th> <th style="text-align: left;">Bits</th> <th style="text-align: left;">Description</th> </tr> </thead> <tbody> <tr> <td>resource</td> <td>23:16</td> <td>resource_id.</td> </tr> <tr> <td>sampler</td> <td>27:24</td> <td>sampler_id.</td> </tr> <tr> <td>reserved</td> <td>29:28</td> <td>Must be zero.</td> </tr> </tbody> </table>	Field Name	Bits	Description	resource	23:16	resource_id.	sampler	27:24	sampler_id.	reserved	29:28	Must be zero.
Field Name	Bits	Description												
resource	23:16	resource_id.												
sampler	27:24	sampler_id.												
reserved	29:28	Must be zero.												
sec_modifier_present	30	Must be zero.												
pri_modifier_present	31	Must be zero.												

Related LOAD, SAMPLE, SAMPLE_B, SAMPLE_C, SAMPLE_G, SAMPLE_L.

Sample Data from Resource with Filter and Comparison

Instructions **SAMPLE_C**

Syntax `sample_c_resource(m)_sampler(n)_aoffimimmi(u, v, w) dst, src0, src1`

Description Performs a comparison filter. *src0* is the index. *src1* contains the reference value. The primary purpose for SAMPLE_C is to provide a building-block for percentage-closer depth filtering. The C in SAMPLE_C stands for comparison.

The operands to SAMPLE_C are identical to those of SAMPLE, except for an additional float32 source operand, *src1*.

If SAMPLE_C is used with a resource that is not a texture1D/2D/Cube, or an unsupported format, then this instruction produces undefined results. It also produces undefined results if used with texture arrays.

When the SAMPLE_C instruction is executed, the hardware uses the current sampler's comparison function to compare *src1.x* against the corresponding texture value at each filter "tap" location (texel) that the currently configured texture filter covers, based on the provided coordinates (*src0*).

The comparison occurs after the source texel's x-component is converted to float32, prior to filtering texels.

For texels that fall off the resource, the x-component value is determined by applying the address modes (and BorderColorR, if in Border mode) from the sampler.

Each comparison that passes returns 1.0f as the x-component value for the texel; each comparison that fails returns 0.0f as the x value for the texture. Then, filtering occurs as specified by the sampler states, operating only in the x-component, and returning a single scalar filter result to the shader (replicated to all masked *dst* components).

The use of SAMPLE_C is orthogonal to all other general-purpose filtering controls (SAMPLE_C works with the other general-purpose filter modes). What SAMPLE_C changes the behavior of the general-purpose filters so that the values being filtered all become 1.0f or 0.0f (comparison results).

See the SAMPLE instruction description for operation of this instruction other than those specified here.

Formats 2-input, 1-output or 4-input, 1-output.

Opcode

Field Name	Bits	Description												
code	15:0	IL_OP_SAMPLE_C												
control	29:16	<table border="1"> <thead> <tr> <th>Field Name</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>resource</td> <td>23:16</td> <td>resource_id.</td> </tr> <tr> <td>sampler</td> <td>27:24</td> <td>sampler_id.</td> </tr> <tr> <td>reserved</td> <td>29:28</td> <td>Must be zero.</td> </tr> </tbody> </table>	Field Name	Bits	Description	resource	23:16	resource_id.	sampler	27:24	sampler_id.	reserved	29:28	Must be zero.
Field Name	Bits	Description												
resource	23:16	resource_id.												
sampler	27:24	sampler_id.												
reserved	29:28	Must be zero.												
sec_modifier_present	30	Must be zero.												
pri_modifier_present	31	Must be zero.												

Related LOAD, SAMPLE, SAMPLE_B, SAMPLE_C_LZ, SAMPLE_G, SAMPLE_L.

Sample Data From Element/Texture With Gradient Using Filter

Instructions `SAMPLE_C_G`**Syntax** `sample_c_g_resource(n)_sampler(m)_aoffiminmi(u,v,w) dst, src0, src1, src2, src3`**Description** Sample data from the specified element/texture using the filtering mode identified by the given sampler. The source data can come from any resource type other than buffers.Behaves like the `SAMPLE_C`, except that gradients for the source address in the x direction and the y direction are provided through `src2` and `src3`, respectively.The x, y, and z components of `src2/src3` (post-swizzle) provide `du/dx`, `dv/dx`, and `dw/dx`. The w component (post-swizzle) is ignored.This instruction has the same restrictions as the `SAMPLE_C` and `SAMPLE_G` instructions.**Format** 4-input 1-output.

Opcode	Field Name	Bits	Description		
	code	15:0	IL_OP_SAMPLE_C_G		
	control	29:16	Field Name	Bits	Description
			resource	23:16	resource_id.
			sampler	27:24	sampler_id.
			reserved	29:28	Must be zero.
	sec_modifier_present	30	Must be zero.		
	pri_modifier_present	31	Must be zero.		

Related `LOAD`, `SAMPLE`, `SAMPLE_B`, `SAMPLE_C`, `SAMPLE_C_LZ`, `SAMPLE_G`, `SAMPLE_L`.

Sample Data from Element/Texture with LOD

Instructions **SAMPLE_C_L**

Syntax `sample_c_l_resource(n)_sampler(m)_aoffiminmi(u,v,w) dst, src0, src1, src2`

Description Sample data from the specified element/texture using the filtering mode identified by the given sampler. The source data can come from any resource type other than buffers.
 Behaves like the SAMPLE_C instruction, except that LOD is provided directly through src2.x.
 This instruction has the same restrictions as the SAMPLE_C and SAMPLE_L instructions.

Format 3-input 1-output.

<i>Opcode</i>	Field Name	Bits	Description	
	code	15:0	IL_OP_SAMPLE_C_L	
	control	29:16	Field Name	Bits Description
			resource	23:16 resource_id.
			sampler	27:24 sampler_id.
			reserved	29:28 Must be zero.
	sec_modifier_present	30	Must be zero.	
	pri_modifier_present	31	Must be zero.	

Related LOAD, SAMPLE, SAMPLE_B, SAMPLE_C, SAMPLE_C_LZ, SAMPLE_G, SAMPLE_C_G.

Texture LOAD

Instructions **TEXLD**

Syntax `texld_stage(n)_centroid_shadowmode(op)_mag(op)_min(op)_volmag(op)_volmin(op)_mip(op)_aniso(op)_lodbias(f)_xoffset(x)_yoffset(y)_zoffset(z) dst, src0`

Description Samples a texture specified by *stage* at the coordinates specified by the *src0* register.

This instruction cannot be used on a *stage* that has not been declared with a DCLPT instruction. Also, it cannot be used on a *stage* set to IL_USAGE_PIXTEX_2DMSAA by a DCLPT instruction. The value of *stage* corresponds to the stage/unit.

The coordinates can be projected using the divComp source modifier on *src0*.

- If the divComp source modifier is used, and the component to divide by is negative, the result of this instruction is undefined. That is, if IL_DIVCOMP_Y is used, and the second component of *src0* is negative, the results of this instruction are undefined.

If *centroid* is 1, sampling is done based on the pixel centroid, not center.

The *lodbias* value specifies a constant value to bias the mipmap from which to load for this instruction. This value is added to the bias value set in the state (the value set through AS_TEX_LODBIAS_N(*stage*)), and the bias value in the fourth component of *src2*.

The following determines the mipmap level(s) from which to sample:

- The *computed LOD* is the mipmap level-of-detail determined based on the ratio of texels in the base texture to the pixel.
- The *instruction LOD* is the value specified by the *lodbias* parameter in the IL_PrimaryTEXLD_Mod token).
- The *minLOD* is the state-based floating point minimum mipmap LOD value.
- The *maxLOD* is the state-based floating point maximum mipmap LOD value.
- The *minLevel* is the smallest mipmap level specified by state to use.
- The *maxLevel* is the largest mipmap level specified by state to use.
- The mipmap level(s) to sample from are determined by:
 - Adding the *state based LOD* to the *computed LOD*.
 - If LOD clamping is enabled in state, clamping the resulting value to *minLOD* and *maxLOD*.
 - Adding the *instruction LOD*.
 - Clamping the resulting value to *minLevel* and *maxLevel*.

The following pseudocode demonstrates the algorithm used to determine the mipmap level(s) to sample from:

- (initial bias LOD) = (state based bias) + (computed LOD)
- (clamped LOD) = (minLOD) ≤ (initial bias LOD) ≤ (maxLOD)
- (secondary bias LOD) = (clamped LOD) + (instruction LOD)
- (final LOD) = (minLevel) ≤ (secondary bias LOD) ≤ (minLevel)

The *mag*, *min*, *volmag*, *volmin*, *mip*, and *aniso* parameters specify whether (and how) to override filter settings. If the IL_PrimaryTEXLD_Mod token is not present, the filters set through external state are used.

Texture LOAD (Cont.)

The values of *xoffset*, *yoffset*, and *zoffset* are added to the unnormalized values of the first, second, and third components of *src0*, respectively, within the sample mipmap. These values are applied whether or not normalized texture coordinates are used. Clamping policy is obeyed as usual when sampling outside the texture's dimensions using these offset parameters. If the *IL_SecondaryTEXLD_Mod* token is not present, *xoffset*, *yoffset*, and *zoffset* default to 0.0.

The *shadowMode* parameter specifies if this instruction performs a shadow map load (compare the texture value to the z-component of *src0*). *shadowMode* indicates one of the following:

- A shadow load never occurs.
- A shadow load always occurs.

If a shadow load occurs with this instruction, the *mag*, *min*, *volmag*, *volmin*, *aniso*, *xoffset*, *yoffset*, and *zoffset* parameters are ignored.

Operation:

The *src0* provides the texture coordinates for the texture sample. The first eight bits of the control field of the *IL_OpCode* token specify the texture stage from which to sample. Associated with the stage specified in the control field are: 1) a texture image, 2) the texture's dimension, 3) and all format, filter, wrap, bias, and clamp settings specified in state and through the *DCLPT* instruction.

Format	Token	Field Name	Bits	Description
	1	code	15:0	IL_OP_TEXLD
		stage	23:16	Stage or unit number, from which to sample.
		centroid	24	0: Sample on pixel center. 1: Sample on pixel centroid.
		<i>reserved</i>	25	Must be zero.
		shadowmode	27:26	Any value of the enumerated type <i>ILTexShadowMode</i> . See Table 5.20 on page 5-11.
		<i>reserved</i>	29:28	Must be zero.
		sec_modifier_present	30	0: No secondary modifier token is present. 1: <i>IL_SecondaryTEXLD_Mod</i> token immediately follows this token or an <i>IL_PRIMARYTEXLD_Mod</i> token, if bit 31 is set.
		pri_modifier_present	31	0: No primary modifier token is present. 1: <i>IL_PrimaryTEXLD_Mod</i> token immediately follows this token.
	2	Primary Texture Load Instruction Modifier token (is present only if the <i>pri_modifier_present</i> field is 1 in the previous <i>IL_OpCode</i> token).		
		Mag	2:0	Specifies how to filter texture values in the S and T directions when a single texel maps to multiple pixels. Can be any value of the enumerated type <i>ILTexFilterMode</i> . See Table 5.21 on page 5-11.
		Min	5:3	Specifies how to filter texture values in the S and T directions when multiple texels map to a single pixel. Can be any value of the enumerated type <i>ILTexFilterMode</i> .
		volmag	8:6	Specifies how to filter texture values in the R direction when the pixel maps to an area less than one texel. Can be any value of the enumerated type <i>ILTexFilterMode</i> .

Texture LOAD (Cont.)

	Volmin	11:9	Specifies how to filter texture values in the R direction when the pixel maps to an area greater than one texel. Can be any value of the enumerated type ILTexFilterMode.
	Mip	14:12	Specifies how to filter values of multiple mipmaps when multiple texels of the base level maps a single pixel. Can be any value of the enumerated type ILMipFilterMode.
	Aniso	17:15	When anisotropic filtering is enabled in the min or mag filter, specifies the maximum number of samples to use for anisotropic filtering. When anisotropic filtering is used for the min and mag filter, this value is the same. Can be any value of the enumerated type ILAnisoFilterMode. See Table 5.22 on page 5-11.
	lodbias	24:18	Specifies a constant mipmap level of detail bias applied to the texture load. Signed fixed point with 4 bits of fraction (3.4), which allows a bias range from [-4, 3.9375]. This value is added to the computed texture LOD value.
	<i>reserved</i>	31:25	Must be zero.
3	Secondary Texture Load Instruction Modifier token (is present only if the <code>pri_modifier_present</code> field is 1 in the previous IL_OpCode token).		
	Xoffset	7:0	Unnormalized (texel space) values added to the X texture coordinate before sampling. Signed fixed point with one bit of fraction (7.1) that allows ranges from [-64, 63.5].
	Yoffset	15:8	Unnormalized (texel space) values added to the Y texture coordinate before sampling. Signed fixed point with one bit of fraction (7.1) that allows ranges from [-64, 63.5].
	Zoffset	23:16	Unnormalized (texel space) values added to the Z texture coordinate before sampling. Signed fixed point with one bit of fraction (7.1) that allows ranges from [-64, 63.5].
	<i>reserved</i>	31:24	Must be zero.
4	IL_Dst token (<i>dst</i>)		
5	IL_Dst_Mod token is present only if the <code>modifier_present</code> field is 1 in the previous IL_Dst token.		
6	IL_Rel_Addr token is present only if the <code>relative_address</code> field is 1 in the preceding IL_Src or IL_Dst token.		
7	IL_Src token (<i>src0</i>)		
8	IL_Src_Mod token is present only if the <code>modifier_present</code> field is 1 in the previous IL_Src token.		
9	IL_Rel_Addr token is present only if the <code>relative_address</code> field is 1 in the preceding IL_Src or IL_Dst token.		

Related TEXTLDB, TEXTLDD, TEXTLDS.

Biased Texture LOAD

Instructions `TEXLDB`

Syntax `texldb_stage(n)[_centroid][_absolute]_shadowmode(op)_mag(op)_min(op)_volmag(op)_volmin(op)_mip(op)_aniso(op)_lodbias(f)_qualitybias_xoffset(x)_yoffset(y)_zoffset(z) dst, src0, src1`

Description Samples a texture specified by *stage* at coordinates specified by the *src0* register and biased by the value in the fourth component of *src1*.

This instruction cannot be used on a *stage* that has not been declared with a DCLPT instruction. Also, it cannot be used on a *stage* set to `IL_USAGE_PIXTEX_2DMSAA` by a DCLPT instruction.

The value of *stage* corresponds to the stage/unit defined externally.

The fourth component of *src2* (*src2.w*) is a factor in determining the mipmap level from which to sample.

The coordinates can be projected using the *divComp* source modifier on *src0*.

- If the *divComp* source modifier is used, and the component to divide by is negative, the result of this instruction is undefined. That is, if `IL_DIVCOMP_Y` is used, and the second (*y*) component of *src0* is negative, the result of this instruction is undefined.
- *divComp* can be set to `IL_DIVCOMP_UNKNOWN` in this instruction; in this case the component used to divide is specified externally.

If *centroid* is 1, sampling is done based on the pixel centroid, not center.

The *lodbias* value specifies a constant value to bias the mipmap from which to load for this instruction. This value is added to the bias value set in the in state and the bias value in the fourth component (*w*) of *src2*.

The following is used to determine the mipmap level(s) to sample from:

- The *computed LOD* is the mipmap level of detail determined based on the ratio of texels in the base texture to the pixel.
- The *instruction LOD* is the value specified by the *lodbias* parameter in the `IL_PrimaryTEXLDB_Mod` token)
- The *state based bias* is the texture bias value specified externally.
- The *minLOD* is the state based floating point minimum mipmap LOD value.
- The *maxLOD* is the state based floating point maximum mipmap LOD value.
- The *minLevel* is the smallest mipmap level specified by state to use.
- The *maxLevel* is the largest mipmap level specified by state to use.
- When *absolute* is 0, the computed LOD is a factor in determining the mipmap level(s) to sample from. The mipmap level(s) to sample from are determined by:
 - Adding the fourth component (*w*) of *src2*, the *state based LOD*, and the *computed LOD*
 - If LOD clamping is enabled in state, clamping the resulting value to *minLOD* and *maxLOD*.
 - Adding the *instruction LOD*.
 - Clamping the resulting value to *minLevel* and *maxLevel*.

The following pseudo code demonstrates the algorithm used to determine the mipmap level(s) to sample from:

```
(initial bias LOD) = (fourth component of src1) + (state based bias) + (computed LOD)
(clamped LOD) = (minLOD) ≤ (initial bias LOD) ≤ (maxLOD)
(secondary bias LOD) = (clamped LOD) + (instruction LOD)
(final LOD) = (minLevel) ≤ (secondary bias LOD) ≤ (maxLevel)
```

Biased Texture LOAD (Cont.)

- When *absolute* is 1, the computed LOD is not a factor in determining the mipmap level(s) from which to sample. These level(s) are determined by:
 - Adding the fourth component of *src2* and the *state based LOD*.
 - If LOD clamping is enabled in state, clamping the resulting value to *minLOD* and *maxLOD*.
 - Adding the *instruction LOD*.
 - Clamping the resulting value to *minLevel* and *maxLevel*.

The following pseudo code demonstrates the algorithm used to determine the mipmap level(s) to sample from:

```
(initial bias LOD) = (fourth component of src2) + (state based bias)
(clamped LOD) = (minLOD) ≤ (initial bias LOD) ≤ (maxLOD)
(secondary bias LOD) = (clamped LOD) + (instruction LOD)
(final LOD) = (minLevel) ≤ (secondary bias LOD) ≤ (maxLevel)
```

The *mag*, *min*, *volmag*, *volmin*, *mip*, and *aniso* parameters specify whether (and how) to override external filter settings. If the *IL_PrimaryTEXLD_Mod* token is not present, the filters are set externally.

The values of *xoffset*, *yoffset*, and *zoffset* are added to the unnormalized values of the first, second, and third components of *src0*, respectively, within the sample mipmap. These values are applied whether or not normalized texture coordinates are used. Clamping policy is obeyed when sampling outside the texture's dimensions using these offset parameters. If the *IL_SecondaryTEXLD_Mod* token is not present, *xoffset*, *yoffset*, and *zoffset* default to 0.0.

The *shadowmode* parameter specifies if this instruction performs a shadow map load, (compare the texture value to the z-component of *src0*). *shadowMode* indicates if a shadow load never occurs or always occurs.

See shadow texture load appendix for texture load algorithm.

If a shadow load occurs with this instruction, the *mag*, *min*, *volmag*, *volmin*, *aniso*, *xoffset*, *yoffset*, and *zoffset* parameters are ignored.

<i>Format</i>	Token	Field Name	Bits	Description
	1	code	15:0	IL_OP_TEXLDB
		stage	23:16	Stage or unit number.
		centroid	24	0: Sample on pixel center. 1: Sample on pixel centroid.
		absolute	25	0: the fourth component of <i>src2</i> is a relative mipmap. 1: the fourth component of <i>src2</i> is an absolute mipmap.
		shadowmode	27:26	Any value of the enumerated type <i>ILTexShadowMode</i> . See Table 5.20 on page 5-11.
		<i>reserved</i>	29:28	Must be zero.
		sec_modifier_present	30	0: No secondary modifier token is present. 1: <i>IL_SecondaryTEXLD_Mod</i> token immediately follows the <i>IL_OP_TEXLDB</i> token or an <i>IL_PrimaryTEXLDB_Mod</i> token if bit 31 is set.
		pri_modifier_present	31	0: No primary modifier token is present. 1: <i>IL_PrimaryTEXLD_Mod</i> token immediately follows the <i>IL_OP_TEXLDB</i> token.

Biased Texture LOAD (Cont.)

2	Primary Texture Load Instruction Modifier token (is present only if the <code>pri_modifier_present</code> field is 1 in the previous <code>IL_OpCode</code> token).	
	Mag	2:0 Specifies how to filter texture values in the S and T directions when the pixel maps to an area than one texel. Can be any value of the enumerated type <code>ILTexFilterMode</code> .
	Min	5:3 Specifies how to filter texture values in the S and T directions when the pixel maps to an area than one texel. Can be any value of the enumerated type <code>ILTexFilterMode</code> .
	volmag	8:6 Specifies how to filter texture values in the R direction when the pixel maps to an area than one texel. Can be any value of the enumerated type <code>ILTexFilterMode</code> .
	Volmin	11:9 Specifies how to filter texture values in the R direction when the pixel maps to an area than one texel. Can be any value of the enumerated type <code>ILTexFilterMode</code> .
	Mip	14:12 Specifies how to filter values of multiple mipmaps when the pixel maps to an area greater than one texel of the base map. Can be any value of the enumerated type <code>ILMipFilterMode</code> .
	Aniso	17:15 When anisotropic filtering is enabled in the min or mag filter, specifies the maximum number of samples to use for anisotropic filtering. When anisotropic filtering is used for the min and mag filter, this value is the same. Can be any value of the enumerated type <code>ILAnisoFilterMode</code> .
	lodbias	24:18 Specifies a constant mipmap level of detail bias applied to the texture load. Signed fixed point with 4 bits of fraction (3.4), which allows a bias range from [-4, 3.9375].
	qualitybias	25 Specifies if the bias is per pixels or per quad: 0 = Per quad. 1 = Per pixel.
	<i>reserved</i>	31:26 Must be zero.
3	Secondary Texture Load Instruction Modifier token (is present only if the <code>sec_modifier_present</code> field is 1 in the previous <code>IL_OpCode</code> token).	
	Xoffset	7:0 Unnormalized (texel space) values added to the X texture coordinate before sampling. Signed fixed point with 1 bit of fraction (7.1) which allows ranges from [-64, 63.5].
	Yoffset	15:8 Unnormalized (texel space) values added to the Y texture coordinate before sampling. Signed fixed point with 1 bit of fraction (7.1) which allows ranges from [-64, 63.5].
	Zoffset	23:16 Unnormalized (texel space) values added to the Z texture coordinate before sampling. Signed fixed point with 1 bit of fraction (7.1) which allows ranges from [-64, 63.5].
	<i>reserved</i>	31:24 Must be zero.

Biased Texture LOAD (Cont.)

- 4 IL_Dst token (*dst*)
- 5 IL_Dst_Mod token is present only if the `modifier_present` field is 1 in the previous IL_Dst token.
- 6 IL_Rel_Addr token is present only if the `relative_address` field is 1 in the preceding IL_Src or IL_Dst token.
- 7 IL_Src token (*src1*)
- 8 IL_Src_Mod token is present only if the `modifier_present` field is 1 in the previous IL_Src token.
- 9 IL_Rel_Addr token is present only if the `relative_address` field is 1 in the preceding IL_Src or IL_Dst token.
- 10 IL_Src token (*src2*)
- 11 IL_Src_Mod token is present only if the `modifier_present` field is 1 in the previous IL_Src token.
- 12 IL_Rel_Addr token is present only if the `relative_address` field is 1 in the preceding IL_Src or IL_Dst token.

Related TEXLDB, TEXLDD, TEXLDMS.

Gradient Texture LOAD

Instructions **TEXLDD**

Syntax `texldd_stage(n)_centroid_shadowmode(op)_mag(op)_min(op)_volmag(op)_volmin(op)_mip(op)_aniso(op)_lodbias(f)_xoffset(x)_yoffset(y)_zoffset(z) dst, src0, src1, src2`

Description *src0* contains the texture coordinate to sample from the specified texture stage.

This instruction cannot be used on a *stage* set to IL_USAGE_PIXTEX_2DMSAA by a DCLPT instruction. Also, it cannot be used on a *stage* that has not been declared with a DCLPT instruction.

The value of *stage* corresponds to the stage/unit defined in state.

The coordinates can be projected using the divComp source modifier on *src0*.

- If the divComp source modifier is used, and the component to divide by is negative, the result of this instruction is undefined. That is, if IL_DIVCOMP_Y is used, and the second component of *src0* is negative, the result of this instruction is undefined.
- divComp can be set to IL_DIVCOMP_UNKNOWN in this instruction; in this case, the component used to divide is specified externally.

If *centroid* is set to 1, sampling is done based on the pixel centroid, not center.

The *lodbias* value specifies a constant value to bias the mipmap from which to load for this instruction. This value is added to the bias value set in the in state and the bias value in the fourth component of *src2*.

The following is used to determine the mipmap level(s) from which to sample.

- The *computed LOD* is the mipmap level of detail determined based on the ratio of texels in the base texture to the pixel.
- The *instruction LOD* is the value specified by the *lodbias* parameter in the IL_PrimaryTEXLD_Mod token)
- The *state based bias* is the texture bias value specified externally.
- The *minLOD* is the state based floating point minimum mipmap LOD value.
- The *maxLOD* is the state based floating point maximum mipmap LOD value.
- The *minLevel* is the smallest mipmap level specified by state to use.
- The *maxLevel* is the largest mipmap level specified by state to use.
- The mipmap level(s) to sample from are determined by:
 - Adding the *state based LOD* to the *computed LOD*.
 - If LOD clamping is enabled in state, clamping the resulting value to *minLOD* and *maxLOD*.
 - Adding the *instruction LOD*.
 - Clamping the resulting value to *minLevel* and *maxLevel*.

The following pseudo code demonstrates the algorithm used to determine the mipmap level(s) to sample from:

- (initial bias LOD) = (state based bias) + (computed LOD)
- (clamped LOD) = (minLOD) ≤ (initial bias LOD) ≤ (maxLOD)
- (secondary bias LOD) = (clamped LOD) + (instruction LOD)
- (final LOD) = (minLevel) ≤ (secondary bias LOD) ≤ (minLevel)

The *mag*, *min*, *volmag*, *volmin*, *mip*, and *aniso* parameters specify whether (and how) to override filter settings. If the IL_PrimaryTEXLD_Mod token is not present, the external filters settings are used.

Gradient Texture LOAD (Cont.)

The values of *xoffset*, *yoffset*, and *zoffset* are added to the unnormalized values of the first, second, and third components of *src0*, respectively, within the sample mipmap. These values are applied whether or not normalized texture coordinates are used. Clamping policy is obeyed when sampling outside the texture's dimensions using these offset parameters. If the IL_SecondaryTEXLD_Mod token is not present, *xoffset*, *yoffset*, and *zoffset* default to 0.0.

The *shadowmode* parameter specifies if this instruction performs a shadow map load (compare the texture value to the z-component of *src0*). *shadowMode* indicates if a shadow load never occurs or always occurs.

If a shadow load occurs with this instruction, the *mag*, *min*, *volmag*, *volmin*, *aniso*, *xoffset*, *yoffset*, and *zoffset* parameters are ignored.

<i>Format</i>	Token	Field Name	Bits	Description
	1	code	15:0	IL_OP_TEXLDD
		stage	23:16	Stage or unit number, 0 to 255.
		centroid	24	0: Sample on pixel center. 1: Sample on pixel centroid.
		reserved	25	Must be zero.
		shadowmode	27:26	Any value of the enumerated type ILTexShadowMode. See Table 5.20 on page 5-11.
		<i>reserved</i>	29:28	Must be zero.
		sec_modifier_present	30	0: No secondary modifier token is present. 1: IL_SecondaryTEXLD_Mod token is present.
		pri_modifier_present	31	0: No primary modifier token is present. 1: IL_PrimaryTEXLD_Mod token is present.
	2	Primary Texture Load Instruction Modifier token (is present only if the <i>pri_modifier_present</i> field is 1 in the previous IL_OpCode token).		
		Mag	2:0	Specifies how to filter texture values in the S and T directions when the pixel maps to an area than one texel. Can be any value of the enumerated type ILTexFilterMode.
		Min	5:3	Specifies how to filter texture values in the S and T directions when the pixel maps to an area than one texel. Can be any value of the enumerated type ILTexFilterMode.
		volmag	8:6	Specifies how to filter texture values in the R direction when the pixel maps to an area than one texel. Can be any value of the enumerated type ILTexFilterMode.
		Volmin	11:9	Specifies how to filter texture values in the R direction when the pixel maps to an area than one texel. Can be any value of the enumerated type ILTexFilterMode.
		Mip	14:12	Specifies how to filter values of multiple mipmaps when the pixel maps to an area greater than one texel of the base map. Can be any value of the enumerated type ILMipFilterMode.

Gradient Texture LOAD (Cont.)

	Aniso	17:15	When anisotropic filtering is enabled in the min or mag filter, specifies the maximum number of samples to use for anisotropic filtering. When anisotropic filtering is used for the min and mag filter, this value is the same. This can be any value of the enumerated type ILAnisoFilterMode.
	lodbias	24:18	Specifies a constant mipmap level of detail bias applied to the texture load. Signed fixed point with 4 bits of fraction (3.4), which allows a bias range from [-4, 3.9375]. This value is added to the computed texture LOD value.
	reserved	31:25	Must be zero.
3	Secondary Texture Load Instruction Modifier token (is present only if the <code>sec_modifier_present</code> field is 1 in the previous IL_OpCode token.)		
	Xoffset	7:0	Unnormalized (texel space) values added to the X texture coordinate before sampling. Signed fixed point with 1 bit of fraction (7.1) which allows ranges from [-64, 63.5].
	Yoffset	15:8	Unnormalized (texel space) values added to the Y texture coordinate before sampling. Signed fixed point with 1 bit of fraction (7.1) which allows ranges from [-64, 63.5].
	Zoffset	23:16	Unnormalized (texel space) values added to the Z texture coordinate before sampling. Signed fixed point with 1 bit of fraction (7.1) which allows ranges from [-64, 63.5].
	reserved	31:24	Must be zero.
4	IL_Dst token (<i>dst</i>)		
5	IL_Dst_Mod token is present only if the <code>modifier_present</code> field is 1 in the previous IL_Dst token.		
6	IL_Rel_Addr token is present only if the <code>relative_address</code> field is 1 in the preceding IL_Src or IL_Dst token.		
7	IL_Src token (<i>src1</i>)		
8	IL_Src_Mod token is present only if the <code>sec_modifier_present</code> field is 1 in the previous IL_OpCode token.		
9	IL_Rel_Addr token is present only if the <code>relative_address</code> field is 1 in the preceding IL_Src or IL_Dst token.		
10	IL_Src token (<i>src2</i>)		
11	IL_Src_Mod token is present only if the <code>sec_modifier_present</code> field is 1 in the previous IL_OpCode token.		
12	IL_Rel_Addr token is present only if the <code>relative_address</code> field is 1 in the preceding IL_Src or IL_Dst token.		
13	IL_Src token (<i>src3</i>)		
14	IL_Src_Mod token is present only if the <code>sec_modifier_present</code> field is 1 in the previous IL_OpCode token.		
15	IL_Rel_Addr token is present only if the <code>relative_address</code> field is 1 in the preceding IL_Src or IL_Dst token.		

Related TEXLD, TEXLDB, TEXLDMS.

Multisample Texture LOAD

<i>Instructions</i>	TEXLDMS		
<i>Syntax</i>	<code>texld_stage(<i>n</i>)_mag(<i>op</i>)_min(<i>op</i>)_aniso(<i>op</i>)_sample(<i>n</i>)_xoffset(<i>x</i>)_yoffset(<i>y</i>) dst, src0</code>		
<i>Description</i>	<p>Samples a multi-sample texture specified by <i>stage</i> at coordinates specified by the <i>src0</i>.</p> <p>This instruction cannot be used on a <i>stage</i> that has not been declared with a DCLPT instruction.</p> <p>This instruction can only be executed on a stage where the <i>usage</i> of the DCLPT instruction is set to IL_USAGE_PIXTEX_2DMSAA.</p> <p>The value of <i>stage</i> corresponds to the stage/unit defined in state.</p> <p>The first two components of <i>src0</i> (<i>src0.xy</i>) are used as the texture coordinate to sample from the specified texture stage.</p> <p>The sum of the third component of <i>src0</i>, <i>src0.z</i>, and the value specified in the <i>sample</i> parameter specifies the single sample to retrieve from the multi-sample buffer when the state based multi-sample filter is set to point filtering.</p> <ul style="list-style-type: none"> • The sum of the third component of <i>src0</i>, <i>src0.z</i>, and the value specified in the <i>sample</i> parameter specifies the single sample to retrieve when AS_MSAA_TEX_FILTER_FUNCTION_N(<i>stage</i>) is set to POINT. <p>The coordinates can be projected using the divComp source modifier on <i>src0</i>.</p> <ul style="list-style-type: none"> • If the divComp source modifier is used, and the component to divide by is negative, the results of this instruction are undefined. For example, if IL_DIVCOMP_Y is used and the second component of <i>src0</i> is negative, the result of this instruction is undefined. • divComp can be set to IL_DIVCOMP_UNKNOWN in this instruction in which case the component used to divide is specified externally. <p>If the IL_PrimaryTEXLDMS_Mod token is not present, the filters are set externally.</p> <p>The values of <i>xoffset</i> and <i>yoffset</i> are added to the unnormalized values of the first, second, and third components of <i>src0</i> respectively within the sample mipmap. These values are applied whether or not normalized texture coordinates are used. Clamping policy is obeyed when sampling outside the texture's dimensions using these offset parameters. If the IL_SecondaryTEXLDMS_Mod token is not present, <i>xoffset</i> and <i>yoffset</i> default to 0.0</p> <p>Operation:</p> <p><i>src0</i> provides the texture coordinates for the texture sample. The first eight bits of the control field of the IL_OpCode token specify the texture stage from which to sample. Associated with the stage specified in the control field are 1) a texture image, 2) the texture's dimension, 3) and all format, filter, wrap, bias, and clamp settings specified in state and through the DCLPT instruction.</p>		
<i>Format</i>	Token	Field Name	Bits Description
	1	code	15:0 IL_OP_TEXLDMS
		stage	23:16 Stage or unit number.
		<i>reserved</i>	29:24 Must be zero.
		sec_modifier_present	30 0: No secondary modifier token is present. 1: IL_SecondaryTEXLD_Mod token immediately follows the IL_OP_TEXLDMS token or an IL_PrimaryTEXLDMS_Mod token if bit 31 is set.
		pri_modifier_present	31 0: No primary modifier token is present. 1: IL_PrimaryTEXLD_Mod token immediately follows the IL_OP_TEXLDMS token.

Multisample Texture LOAD (Cont.)

2	Primary Multi-sample Texture Load Instruction Modifier token (is present only if the <code>pri_modifier_present</code> field is 1 in the previous <code>IL_OpCode</code> token).
Mag	2:0 Specifies how to filter texture values in the S and T directions when the pixel maps to an area less than one texel. Can be any value of the enumerated type <code>ILTexFilterMode</code> .
Min	5:3 Specifies how to filter texture values in the S and T directions when the pixel maps to an area greater than one texel. Can be any value of the enumerated type <code>ILTexFilterMode</code> .
<i>reserved</i>	14:6 Must be zero.
Aniso	17:15 When anisotropic filtering is enabled in the min or mag filter, specifies the maximum number of samples to use for anisotropic filtering. When anisotropic filtering is used for the min and mag filter, this value is the same. This can be any value of the enumerated type <code>ILAnisoFilterMode</code> .
<i>reserved</i>	31:18 Must be zero.
3	Secondary Multi-sample Texture Load Instruction Modifier token (is present only if the <code>sec_modifier_present</code> field is 1 in the previous <code>IL_OpCode</code> token).
Xoffset	7:0 Unnormalized (texel space) values added to the X texture coordinate before sampling. Signed fixed point with one bit of fraction (7.1) that allows ranges from [-64, 63.5].
Xoffset	15:8 Unnormalized (texel space) values added to the Y texture coordinate before sampling. Signed fixed point with one bit of fraction (7.1) that allows ranges from [-64, 63.5].
<i>reserved</i>	23:16 Must be zero.
Sample	28:24 This number is added to the third component of <code>src0</code> to select the single sample to load in the case that point multi-sample filtering is used. This field only applies if the state based point multi-sample filter is set.
<i>reserved</i>	31:29 Must be zero.
4	<code>IL_Dst</code> token (<i>dst</i>)
5	<code>IL_Dst_Mod</code> token is present only if the <code>modifier_present</code> field is 1 in the previous <code>IL_Dst</code> token.
6	<code>IL_Rel_Addr</code> token is present only if the <code>relative_address</code> field is 1 in the preceding <code>IL_Src</code> or <code>IL_Dst</code> token.
7	<code>IL_Src</code> token (<i>src0</i>)
8	<code>IL_Src_Mod</code> token is present only if the <code>modifier_present</code> field is 1 in the previous <code>IL_Src</code> token.
9	<code>IL_Rel_Addr</code> token is present only if the <code>relative_address</code> field is 1 in the preceding <code>IL_Src</code> or <code>IL_Dst</code> token.

Related `TEXLD`, `TEXLDB`, `TEXLDD`.

Get Texel Weight

Instructions	TEXWEIGHT			
Syntax	<code>texweight_stage(<i>n</i>) <i>dst</i>, <i>src0</i></code>			
Description	<p>Retrieves the weights used by a bilinear filtered fetch based upon the texture coordinate provided in <i>src0</i>.</p> <p>The <i>dst.x</i> represents the horizontal LERP factor; the <i>dst.y</i> represents the vertical LERP factor. The <i>dst.z</i> and <i>dst.w</i> component are not written to by this instruction; however, if the <i>component_z_b</i> or <i>component_w_a</i> field of the IL_Dst_Mod token is set to IL_MODCOMP_0 or IL_MODCOMP_1, the <i>dst.z</i> and <i>dst.w</i> are written (<i>dst.z</i> and <i>dst.w</i> can be set to 0.0 or 1.0 if IL_MODCOMP_0 or IL_MODCOMP_1 is used on the <i>component_z_b</i> or <i>component_w_a</i> field of the IL_Dst_Mod token).</p> <p>The coordinates can be projected using the divComp source modifier on <i>src0</i>.</p> <ul style="list-style-type: none"> • If the divComp source modifier is used, and the component to divide by is negative, the result of this instruction is undefined. That is, if IL_DIVCOMP_Y is used, and the second component of <i>src0</i> is negative, the results of this texture load is undefined. • divComp can be set to IL_DIVCOMP_UNKNOWN in this instruction; in this case, the component used to divide is specified in state. <p>The value of <i>stage</i> corresponds to the stage/unit.</p> <p>This instruction can be used on only a <i>stage</i> that has been declared with a DCLPT instruction.</p> <p>Operation:</p> <p><i>src0</i> provides the texture coordinates for the texture weight. The first eight bits of the control field of the IL_OpCode token specify the texture stage from which to retrieve the weight. Associated with the stage specified in the control field are 1) a texture, 2) the texture's dimension, 3) and all format, filter, wrap, bias, and clamp states defined externally and through the DCLPT instruction.</p> <p>ALU instructions perform arithmetic and relational operations. All take in at least one source operand and output to 1 destination operand.</p>			
Format	Token	Field Name	Bits	Description
	1	<i>code</i>	15:0	IL_OP_TEXWEIGHT
		<i>stage</i>	23:16	Stage or unit number, 0 to 255.
		<i>reserved</i>	31:24	Must be zero.
	2	IL_Dst token (<i>dst</i>)		
	3	IL_Dst_Mod token is present only if the <i>modifier_present</i> field is 1 in the previous IL_Dst token.		
	4	IL_Rel_Addr token is present only if the <i>relative_address</i> field is 1 in the preceding IL_Src or IL_Dst token.		
	5	IL_Src token (<i>src0</i>)		
	6	IL_Src_Mod token is present only if the <i>modifier_present</i> field is 1 in the previous IL_Src token.		
	7	IL_Rel_Addr token is present only if the <i>relative_address</i> field is 1 in the preceding IL_Src or IL_Dst token.		
Related	None.			

Read Local Data Share (LDS) Memory into a Vector

Instructions **LDS_READ_VEC**

Syntax `lds_read_vec (_neighborExch)(_sharingMode) dst, src0.xy`

Description Reads the LDS memory into the *dst* register (can read a vector of up to four components). This uses the ownership accessing model. Each thread owns a part of the LDS memory. The LDS location is specified using the owner thread ID and offset, such as “Read data written/owned by thread Tid at Offset.”

Register *src0* specifies the location with *src0.x* = Tid, *src0.y* = offset.

Dst can be any writable register.

The sharing modes `_sharingMode(rel)` and `_sharingMode(abs)` are relative and absolute, respectively. The flag is optional. If it is not specified, the instruction takes the default sharing mode specified by `dcl_lds_sharing_mode`.

The flag `neighborExch` is optional. If it is specified, the output of lds memory is exchanged with its neighbor threads, so that the first thread receives all values from x-components, the second thread receives all values from y-components, etc. This flag is useful for applications like FFT matrix transpose.

It is used only in a compute shader.

Formats 1-input, 1-output, 0 additional token.

Opcode	Field Name	Bits	Description
	<code>code</code>	15:0	Must be set to <code>IL_OP_LDS_READ_VEC</code> .
	<code>controls</code>	17:16	<code>_sharingMode</code> : 0: <code>IL_LDS_SHARING_MODE_RELATIVE</code> 1: <code>IL_LDS_SHARING_MODE_ABSOLUTE</code> 2-3: reserved
		18	<code>_neighborExch</code> : 0: off 1: on and enabled
		29:19	Not used. Must be 0
	<code>sec_modifier_present</code>	30	Must be zero.
	<code>pri_modifier_present</code>	31	Must be zero.

Examples

```
lds_read_vec r1, r0.xy
lds_read_vec_neighborExch r1, r0.xy
lds_read_vec_sharingMode(rel) r1, r0.xy
lds_read_vec_neighborExch_sharingMode(abs) r1, r0.xy
```

Related **LDS_WRITE_VEC**

Write a Vector to Local Data Share (LDS) Memory

Instructions **LDS_WRITE_VEC**

Syntax `lds_write_vec _lOffset(N) (_sharingMode) dst, src0`

Description Writes data (a vector of up to four components) in *src0* into LDS memory. This uses the ownership accessing model: each thread owns a part of the LDS memory. Each thread can write only to the part it owns. The location is specified by offset.

src0 can be any register.

Dst must be of type `IL_REGTYPE_GENERIC_MEM`. This is used to provide the write mask.

`_lOffset()` is optional. If not specified, the offset is 0. If it is specified, such as `_lOffset(n)`, *n* must be a value of multiples of four in the range of [0,60], and must be smaller than the declared `lds_size_per_thread`.

`_sharingMode(rel)` or `_sharingMode(abs)` are for relative or absolute sharing mode, respectively. The flag is optional. If it is not specified, the instruction takes the default sharing mode specified by `dcl_lds_sharing_mode`.

This instruction is used only in a compute kernel.

Formats 1-input, 1-output, no additional token.

Opcode

Field Name	Bits	Description
code	15:0	Must be set to <code>IL_OP_LDS_READ_VEC</code> .
controls	17:16	<code>_sharingMode</code> : 0: <code>IL_LDS_SHARING_MODE_RELATIVE</code> 1: <code>IL_LDS_SHARING_MODE_ABSOLUTE</code> 2-3: reserved
	23:18	<code>_lOffset</code> : offset within its reserved space, must be one of [0,4, 8, ... 60]
	29:24	Not used. Must be zero.
<code>sec_modifier_present</code>	30	Must be zero.
<code>pri_modifier_present</code>	31	Must be zero.

Examples

```
lds_write_vec mem._y_, r0.yyyy
lds_write_vec_lOffset(4) mem.xy_, r0.xyzw
lds_write_vec_sharingMode(rel) mem.__zw, r0.xyzw
lds_write_vec_lOffset(4)_sharingMode(abs) mem.x_z_w, r0.xyzw
```

Related **LDS_READ_VEC.**

Fence for Synchronization Threads, and/or LDS and/or Global Memory or UAV

Instructions **FENCE**

Syntax fence[_threads][_lds][_memory][_sr]

Description It must include at least one option. It is an error if none is specified. All selected options must complete before the thread continues.

 _threads - synchronize threads in a group such that all threads must have reached this point before any thread can go further. The instruction cannot be used inside of any control flow. It can be used only in a compute shader (except for bit 18; see below).

 _lds - shared memory fence. It ensures that:

- no LDS read/write instructions can be re-ordered or moved across this fence instruction.
- all LDS write instructions are complete (the data has been written to LDS memory, not in internal buffers) and visible to other threads.

 _memory - global/scatter memory fence. It ensures that:

- no memory import/export instructions can be re-ordered or moved across this fence instruction.
- all memory export instructions are complete (the data has been written to physical memory, not in the cache) and is visible to other threads.

 _sr - shared register write/read fence. No shared register writes/reads can be re-ordered or moved across this fence instruction.

All prior writes to shared registers done by this thread become visible to all other threads.

Bit 16 (0 = this is a thread barrier point; 1 = a barrier point is required)

Bit 17 (1 = all prior LDS writes must become visible to other threads in the current thread group)

Bit 18 (1 = all prior writes to global memory and UAV surfaces by this thread in program order become visible to all other threads) Note that both compute and pixel shaders can use this bit in the FENCE instruction.

Bit 19 (1 = all prior writes to shared registers done by this thread become visible to all other threads)

Pixel kernels can only use fence instructions for global memory. Compute kernels can use all of options.

In pixel kernels, use of discard instructions implies a fence_memory.

Use of discard with a fence_threads instruction is undefined in IL, although specific implementations can select an interpolation.

Formats 0-input, 0-output, no additional token.

Examples lds_write_vec mem.y__, r0.yyyy
 fence_threads_lds
 lds_read_vec r1, r0.xy

Opcode	Field Name	Bits	Description
	code	15:0	Must be set to IL_OP_FENCE.
	_threads	16	1 indicates that this is a fence for thread synchronization.
	_lds	17	1 indicates that this is a fence for the LDS.
	_memory	18	1 indicates that this is a fence for global memory or UAV.
	_sr	19	1 indicates that this is a fence for sr.
	Controls	29:20	Must be zero.
	sec_modifier_present	30	Must be zero.

Fence for Synchronization Threads, and/or LDS and/or Global Memory or UAV (Cont.)

`pri_modifier_present` 31 Must be zero.

Examples `lds_write_vec mem._y_, r0.yyyy`
 `fence_threads_lds`
 `lds_read_vec r1, r0.xy`

Related None.

6.6 Integer Arithmetic Instructions**Integer AND**

Instructions **IAND**

Syntax `iand dst, src0, src1`

Description Component-wise logical AND of each pair of 32-bit values from *src0* and *src1*. The 32-bit result is placed in *dst*.

Format 2-input 1-output.

<i>Opcode</i>	Field Name	Bits	Description
	<code>code</code>	15:0	IL_OP_IAND
	<code>control</code>	29:16	Must be zero.
	<code>sec_modifier_present</code>	30	Must be zero.
	<code>pri_modifier_present</code>	31	Must be zero.

Related None.

Integer Negate

Instructions **INOT**

Syntax `inot dst, src0`

Description Performs a bit-wise one's complement on each component of *src0*. The 32-bit results are placed in the corresponding components of *dst*.

Format 1-input, 1-output.

<i>Opcode</i>	Field Name	Bits	Description
	<code>code</code>	15:0	IL_OP_I_NOT
	<code>control</code>	29:16	Must be zero.
	<code>sec_modifier_present</code>	30	Must be zero.
	<code>pri_modifier_present</code>	31	Must be zero.

Related None.

Integer inclusive-OR, Integer Exclusive-OR

<i>Instructions</i>	IOR, IXOR			
<i>Syntax</i>	Function	Opcode	Syntax	Description
	IOR	IL_OP_I_OR	<i>ior dst, src0, src1</i>	Integer inclusive-OR, bit-wise OR of each component in <i>src0</i> with the corresponding component in <i>src1</i>
	IXOR	IL_OP_I_XOR	<i>ixor dst, src0, src1</i>	Integer exclusive-OR, bit-wise XOR of each component in <i>src0</i> with the corresponding component in <i>src1</i>
<i>Description</i>	Performs a bit-wise logical operation of each component of <i>src0</i> with the corresponding component of <i>src1</i> . The 32-bit results are placed in the corresponding components of <i>dst</i> .			
<i>Format</i>	2-input, 1-output.			
<i>Opcode</i>	Field Name	Bits	Description	
	code	15:0	See <i>Syntax</i> , above.	
	control	29:16	Must be zero.	
	sec_modifier_present	30	Must be zero.	
	pri_modifier_present	31	Must be zero.	
<i>Related</i>	None.			

Bitwise Integer Addition

<i>Instructions</i>	I_ADD			
<i>Syntax</i>	<i>iadd dst, src0, src1</i>			
<i>Description</i>	Component-wise add of 32-bit operands <i>src0</i> and <i>src1</i> . The 32-bit result is placed in <i>dst</i> . No carry or borrow is done beyond the 32-bit values of each component; thus, this instruction is not sensitive to the sign of the operand.			
<i>Format</i>	2-input, 1-output.			
<i>Opcode</i>	Field Name	Bits	Description	
	code	15:0	IL_OP_I_ADD	
	control	29:16	Must be zero.	
	sec_modifier_present	30	Must be zero.	
	pri_modifier_present	31	Must be zero.	
<i>Related</i>	DADD.			

Signed Integer Multiply and Add

Instructions **I_MAD***Syntax* `imad dst, src0, src1, src2`*Description* Component-wise integer multiply of 32-bit signed operands *src0* and *src1*, keeping the low 32 bits (per component) of the result, followed by an integer add of *src2* to produce the low 32-bit (per component) result. and placing it in *dst*.*Format* 3-input, 1-output.

<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_I_MAD
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.

Related DMAD.

Integer Maximum, Integer Minimum

Instructions	IMAX, IMIN			
Syntax	Function	Opcode	Syntax	Description
	IMAX	IL_OP_I_MAX	$imax\ dst, src0, src1$	$dst = src0 > src1 ? src0 : src1$
	IMIN	IL_OP_I_MIN	$imin\ dst, src0, src1$	$dst = src0 < src1 ? src0 : src1$
Description	<p>Component-wise integer maximum (I_MAX) and minimum (I_MIN).</p> <p>Compares each component of <i>src0</i> with the corresponding component of <i>src1</i>. If the comparison evaluates true, <i>src0</i> is returned in the corresponding component of <i>dst</i>; otherwise, <i>src1</i> is returned.</p> <p>For the MAX and MIN instructions, if <i>src0</i>.{x y z w} or <i>src1</i>.{x y z w} is NaN, then the other component, <i>src1</i>.{x y z w} or <i>src0</i>.{x y z w}, respectively, is returned. Denorms are flushed before the comparison is performed. If a flushed denorm is the maximum value (for the MAX instruction) or minimum value (for the MIN instruction), then the flushed denorm is returned. The MAX instruction uses a greater-than-or-equal-to comparison. Thus, if $\min(src0, src1) = src0$, then $\max(src0, src1) = src1$, including cases with +0 and -0, such as when denorms are flushed to sign preserve zero.</p> <p>If the <code>_ieee</code> flag is included, then MIN and MAX follow IEEE-754r rules for minimum and maximum (except denorms are flushed before the comparison is made). Also, MIN returns -0, and MAX returns +0, for comparisons between -0 and +0.</p>			
Format	2-input, 1-output.			
Opcode	Field Name	Bits	Description	
	code	15:0	IL_OP_I_MAX, IL_OP_I_MIN	
	control	29:16	Must be zero.	
	sec_modifier_present	30	Must be zero.	
	pri_modifier_present	31	Must be zero.	
Related	None.			

Signed Integer Multiplication (Low 32 Bits)

Instructions	IMUL			
Syntax	$imul\ dst, src0, src1$			
Description	Component-wise multiply of 32-bit operands <i>src0</i> and <i>src1</i> . The lower 32 bits of the 64-bit result, which is placed in <i>dst</i> .			
Format	2-input, 1-output.			
Opcode	Field Name	Bits	Description	
	code	15:0	IL_OP_I_MUL	
	control	29:16	Must be zero.	
	sec_modifier_present	30	Must be zero.	
	pri_modifier_present	31	Must be zero.	
Related	IMUL_HIGH.			

Signed Integer Multiplication (High 32 Bits)

Instructions **IMUL_HIGH***Syntax* `imul_high dst, src0, src1`*Description* Component-wise multiply of 32-bit signed operands *src0* and *src1*. The upper 32 bits of the 64-bit result is placed in the corresponding component of *dst*.*Format* 2-input, 1-output.

<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_I_MUL_HIGH
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.

Related DMUL, IMUL, MUL, UMUL.

Integer Compare

<i>Instructions</i>	IEQ, IGE, ILT, INE			
<i>Syntax</i>	Function	Opcode	Syntax	Description
	==	IL_OP_I_EQ	ieq <i>dst</i> , <i>src0</i> , <i>src1</i>	Compares if integer vector ins <i>src0</i> is equal to the one in <i>src1</i> . If TRUE, 0xFFFFFFFF is returned for that component; otherwise, 0x00000000 is returned.
	≥	IL_OP_I_GE	ige <i>dst</i> , <i>src0</i> , <i>src1</i>	Compares if integer vector in <i>src0</i> is greater or equal to the one in <i>src1</i> . If TRUE, 0xFFFFFFFF is returned for that component; otherwise, 0x00000000 is returned.
	<	IL_OP_I_LT	ilt <i>dst</i> , <i>src0</i> , <i>src1</i>	Compares if integer vector in <i>src0</i> is less than the one in <i>src1</i> . If TRUE, 0xFFFFFFFF is returned for that component; otherwise, 0x00000000 is returned.
	≠	IL_OP_I_NE	ine <i>dst</i> , <i>src0</i> , <i>src1</i>	Compares two integer vectors, one in <i>src0</i> , the other in <i>src1</i> , to check if they are not equal. If TRUE, 0xFFFFFFFF is returned for that component; otherwise, 0x00000000 is returned.
<i>Description</i>	Component-wise compares two vectors using an integer comparison. If <i>src0</i> .{x y z w} and the corresponding component in <i>src1</i> satisfy the comparison condition, the corresponding component of <i>dst</i> is set to TRUE; otherwise, it is set to FALSE. Primary and secondary opcode modifiers are not permitted, and no additional control options are supported. Thus, the control field, <i>pri_modifier_present</i> , and <i>sec_modifier_present</i> fields must be zero.			
<i>Format</i>	2-input, 1-output.			
<i>Opcode</i>	Field Name	Bits	Description	
	code	15:0	See <i>Syntax</i> , above.	
	control	29:16	Must be zero.	
	<i>sec_modifier_present</i>	30	Must be zero.	
	<i>pri_modifier_present</i>	31	Must be zero.	
<i>Related</i>	Double-precision Compare, Float Compare, Unsigned Integer Compare.			

Two's Complement Negate

<i>Instructions</i>	INEGATE		
<i>Syntax</i>	inegate <i>dst</i> , <i>src0</i>		
<i>Description</i>	Computes the two's complement negation of each component in <i>src0</i> . The 32-bit results are placed in the corresponding components of <i>dst</i> .		
<i>Format</i>	1-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_I_NEGATE
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	None.		

Integer Shift Left, Integer Shift Right

<i>Instructions</i>	ISHL, ISHR			
<i>Syntax</i>	Function	Opcode	Syntax	Description
	ISHL	IL_OP_I_SHL	ishl <i>dst</i> , <i>src0</i> , <i>src1</i>	<p><i>src1</i> must be a scalar value; thus, all four swizzles on <i>src1</i> must be the same.</p> <p>Component-wise shift of each 32-bit value in <i>src0</i> left by an unsigned integer bit count provided by the LSB five bits (0-31 range) in <i>src1.selected_component</i>, inserting 0.</p> <p>The 32-bit per component results are placed in <i>dst</i>.</p> <p>The count is a scalar value applied to all components.</p>
	ISHR	IL_OP_I_SHR	ishr <i>dst</i> , <i>src0</i> , <i>src1</i>	<p>Component-wise arithmetic shift of each 32-bit value in <i>src0</i> right by an unsigned integer bit count provided by the LSB five bits (0-31 range) in <i>src1.x</i>, replicating the value of bit 31. The 32-bit per component result is placed in <i>dst</i>. The count is a scalar value applied to all components.</p>
<i>Format</i>	2-input, 1-output.			
<i>Opcode</i>	Field Name	Bits	Description	
	code	15:0	See <i>Syntax</i> , above.	
	control	29:16	Must be zero.	
	sec_modifier_present	30	Must be zero.	
	pri_modifier_present	31	Must be zero.	
<i>Related</i>	None.			

6.7 Unsigned Integer Operations

Integer Shift Left, Integer Shift Right

<i>Instructions</i>	USHR		
<i>Syntax</i>	<code>ushr dst, src0, src1</code>		
<i>Description</i>	Component-wise shift of each 32-bit value in <i>src0</i> right by an unsigned integer bit count provided by the LSB five bits (0-31 range) in <i>src1.selected_component</i> , inserting 0. The 32-bit per component result is placed in <i>dst</i> . The count is a scalar value applied to all components.		
<i>Format</i>	2-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_U_SHR
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	None.		

Floating Point Division, Unsigned Integer Division

<i>Instructions</i>	UDIV		
<i>Syntax</i>	<code>udiv dst, src0, src1</code>		
<i>Description</i>	Component-wise unsigned division of the 32-bit operand in <i>src0</i> by the corresponding component in <i>src1</i> . The 32-bit quotient is placed in the corresponding component of <i>dst</i> .		
<i>Format</i>	2-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_U_DIV
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	None.		

Unsigned Integer Modulo

<i>Instructions</i>	UMOD		
<i>Syntax</i>	<code>umod dst, src0, src1</code>		
<i>Description</i>	Component-wise unsigned division of the 32-bit operand <i>src0</i> by the 32-bit operand <i>src1</i> . The 32-bit remainder is placed in the corresponding component of <i>dst</i> .		
<i>Format</i>	2-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_U_MOD
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	None.		

Unsigned Integer Multiply and Add

<i>Instructions</i>	UMAD		
<i>Syntax</i>	<code>umad dst, src0, src1, src2</code>		
<i>Description</i>	Component-wise unsigned multiply of the 32-bit operand <i>src0</i> by the 32-bit operand <i>src1</i> is added to the corresponding component in <i>src2</i> . The result of the multiply-add operation is placed in the corresponding component of <i>dst</i> .		
<i>Format</i>	3-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_U_MAD
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	DMAD.		

Unsigned Integer Maximum, Unsigned Integer Minimum

Instructions **UMAX, UMIN**

<i>Syntax</i>	Function	Opcode	Syntax	Description
	UMAX	IL_OP_U_MAX	<code>umax dst, src0, src1</code>	unsigned integer maximum, $dst = src0 > src1 ? src0 : src1$
	UMIN	IL_OP_U_MIN	<code>umin dst, src0, src1</code>	unsigned integer minimum, $dst = src0 < src1 ? src0 : src1$

Description Compares each component of *src0* with the corresponding component of *src1*. If the comparison is TRUE, *src0* is returned in the corresponding component of *dst*; otherwise, *src1* is returned.

For the MAX and MIN instructions, if *src0*.{x|y|z|w} or *src1*.{x|y|z|w} is NaN, then the other component, *src1*.{x|y|z|w} or *src0*.{x|y|z|w}, respectively, is returned. Denorms are flushed before the comparison is performed. If a flushed denorm is the maximum value (for the MAX instruction) or minimum value (for the MIN instruction), then the flushed denorm is returned. The MAX instruction uses a greater-than-or-equal-to comparison. Thus, if $\min(src0, src1) = src0$, then $\max(src0, src1) = src1$, including cases with +0 and -0 such as when denorms are flushed to sign preserve zero.

If the `_ieee` flag is included, MIN and MAX follow IEEE-754r rules for minimum and maximum (except denorms are flushed before the comparison is made). In addition, MIN returns -0 and MAX returns +0 for comparisons between -0 and +0.

Format 2-input, 1-output.

<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	See <i>Syntax</i> , above.
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.

Related None.

Unsigned Integer Compare

<i>Instructions</i>	UGE, ULT		
<i>Syntax</i>	Function	Opcode	Syntax
	≥	IL_OP_U_GE	ige <i>dst</i> , <i>src0</i> , <i>src1</i>
	<	IL_OP_U_LT	ilt <i>dst</i> , <i>src0</i> , <i>src1</i>
<i>Description</i>	Component-wise compares two vectors using an unsigned integer comparison. For UGE: (<i>src0</i> ≥ <i>src1</i>); for ULT: (<i>src0</i> < <i>src1</i>). If <i>src0</i> .{x y z w} and the corresponding component in <i>src1</i> satisfy the comparison condition, the corresponding component of <i>dst</i> is set to TRUE and returns 0xFFFFFFFF; otherwise, it is set to FALSE and returns 0x00000000. Primary and secondary opcode modifiers are not permitted, and no additional control options are supported. Thus, the control field, <i>pri_modifier_present</i> , and <i>sec_modifier_present</i> fields must be zero.		
<i>Format</i>	2-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_U_GE, IL_OP_U_LT
	control	29:16	Must be zero.
	<i>sec_modifier_present</i>	30	Must be zero.
	<i>pri_modifier_present</i>	31	Must be zero.
<i>Related</i>	Double-precision Compare, Float Compare, Integer Compare.		

Unsigned Integer Multiplication

<i>Instructions</i>	UMUL		
<i>Syntax</i>	umul <i>dst</i> , <i>src0</i> , <i>src1</i>		
<i>Description</i>	Component-wise multiply of 32-bit unsigned operands <i>src0</i> and <i>src1</i> . The lower 32 bits of the 64-bit result (per component) is placed in the corresponding component of <i>dst</i> .		
<i>Format</i>	2-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_U_MUL
	control	29:16	Must be zero.
	<i>sec_modifier_present</i>	30	Must be zero.
	<i>pri_modifier_present</i>	31	Must be zero.
<i>Related</i>	DMUL, IMUL_HIGH, UMUL_HIGH.		

Unsigned Integer Multiplication

Instructions **UMUL_HIGH**

Syntax `umul_high dst, src0, src1`

Description Component-wise multiply of 32-bit unsigned operands *src0* and *src1*. The upper 32 bits of the 64-bit result (per component) is placed in the corresponding component of *dst*.

Format 2-input, 1-output.

<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_U_MUL_HIGH
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.

Related DMUL, IMUL, MUL, UMUL.

6.8 Conversion Instructions

Float to Signed Integer Conversion

Instructions **FTOI**

Syntax `ftoi dst, src0`

Description Converts each component of *src0* in the range [-2147483648.999f, 2147483647.999f] to a signed integer, and puts the result in the corresponding component of *dst*. Using values outside the specified range produces undefined results.

After the instruction executes, *dst* contains a 32-bit signed integer 4-tuple. The conversion is performed per component.

Rounding is performed towards zero on each component of *src0*, following the C convention for casts from float to int. Applications that require different rounding semantics can invoke the ROUND_* instructions before using FTOI.

Format 1-input, 1-output.

<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_FTOI
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.

Related F2D, FTOU.

Float to Unsigned Integer Conversion

<i>Instructions</i>	F_TO_U		
<i>Syntax</i>	<code>ftou dst, src0</code>		
<i>Description</i>	<p>Converts each component of <i>src0</i> in the range [0.0f, 4294967296.999f] to an unsigned integer and puts the result in the corresponding component of <i>dst</i>. Using values outside the specified range produces undefined results.</p> <p>After the instruction executes, <i>dst</i> contains a 32-bit unsigned integer 4-tuple. The conversion is performed per component.</p> <p>Rounding is performed towards zero on each component of <i>src0</i>, following the C convention for casts from float to unsigned int. Applications that require different rounding semantics can invoke the ROUND_* instructions before using F_TO_U.</p>		
<i>Format</i>	1-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_FT <u>OU</u>
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	F2D, F _T O _I .		

Signed Integer to Float Conversion

<i>Instructions</i>	I_TO_F		
<i>Syntax</i>	<code>itof dst, src0</code>		
<i>Description</i>	<p>Converts each component of <i>src0</i> to a float, and puts the result in the corresponding component of <i>dst</i>. Each component of <i>src0</i> is assumed to contain a signed 32-bit integer 4-tuple. After conversion, <i>dst</i> contains a floating-point 4-tuple.</p>		
<i>Format</i>	1-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_I <u>T</u> O <u>F</u>
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	D2F, U _T O _F .		

Unsigned Integer to Float Conversion

<i>Instructions</i>	UTOF		
<i>Syntax</i>	utof <i>dst</i> , <i>src0</i>		
<i>Description</i>	<p>Converts each component of <i>src0</i> to a float and puts the result in the corresponding component of <i>dst</i>. Each component of <i>src0</i> is assumed to contain an unsigned 32-bit integer 4-tuple.</p> <p>After the instruction executes, <i>dst</i> contains a floating-point 4-tuple.</p> <p>If an integer is not represented exactly, the nearest representable value is used. Rounding is performed towards zero on each component of <i>src0</i>, following the C convention for casts from unsigned int to float. Applications that require different rounding semantics can invoke the ROUND_* instructions before using UTOF.</p>		
<i>Format</i>	1-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_UTOF
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	D2F, ITOF.		

Double to Float Conversion

<i>Instructions</i>	D2F																												
<i>Syntax</i>	d2f <i>dst</i> , <i>src0</i>																												
<i>Description</i>	<p>Abs and negate modifiers can be used on the source. All four output components are set to the float converted value. Nan/inf convert to nan/inf; signs are preserved. If the result does not fit as a float, inf is returned. Uses round to nearest even.</p> <table> <tr> <td>d</td> <td>-NaN</td> <td>-inf</td> <td>-F</td> <td>-1.0</td> <td>-denom</td> <td>-0</td> <td>+0</td> <td>+denom</td> <td>+1.0</td> <td>+F</td> <td>+inf</td> <td>+NaN</td> </tr> <tr> <td>f</td> <td>0xffc00000</td> <td>-inf</td> <td>-F</td> <td>-1.0</td> <td>-0.0</td> <td>-0.0</td> <td>++0.0</td> <td>+0.0</td> <td>+1.0</td> <td>+F</td> <td>+inf</td> <td>0x7fc00000</td> </tr> </table>			d	-NaN	-inf	-F	-1.0	-denom	-0	+0	+denom	+1.0	+F	+inf	+NaN	f	0xffc00000	-inf	-F	-1.0	-0.0	-0.0	++0.0	+0.0	+1.0	+F	+inf	0x7fc00000
d	-NaN	-inf	-F	-1.0	-denom	-0	+0	+denom	+1.0	+F	+inf	+NaN																	
f	0xffc00000	-inf	-F	-1.0	-0.0	-0.0	++0.0	+0.0	+1.0	+F	+inf	0x7fc00000																	
<i>Format</i>	1-input, 1-output.																												
<i>Opcode</i>	Field Name	Bits	Description																										
	code	15:0	IL_OP_D_2_F																										
	control	29:16	Must be zero.																										
	sec_modifier_present	30	Must be zero.																										
	pri_modifier_present	31	Must be zero.																										
<i>Related</i>	ITOF, UTOF.																												

Float to Double Conversion

<i>Instructions</i>	F2D																								
<i>Syntax</i>	<code>f2d dst, src0</code>																								
<i>Description</i>	<p>The source is interpreted as a float in component x (after swizzles). Abs and negate modifiers can be used on the source. Output yx or output wz is set to the double converted value. Nan/inf convert to nan/inf; signs are preserved.</p> <p>Rounding is performed towards zero on each component of <i>src0</i>, following the ANSI C convention for casts from float to double. Applications that require different rounding semantics can invoke the ROUND_* instructions before using F2D.</p> <table> <tr> <td>f</td> <td>-inf</td> <td>-F</td> <td>-1.0</td> <td>-denom</td> <td>-0</td> <td>+0</td> <td>+denom</td> <td>+1.0</td> <td>+F</td> <td>+inf</td> <td>NaN</td> </tr> <tr> <td>d</td> <td>-inf</td> <td>-F</td> <td>-1.0</td> <td>-0.0</td> <td>-0.0</td> <td>+0.0</td> <td>+0.0</td> <td>+1.0</td> <td>+F</td> <td>+inf</td> <td>Nan*</td> </tr> </table>	f	-inf	-F	-1.0	-denom	-0	+0	+denom	+1.0	+F	+inf	NaN	d	-inf	-F	-1.0	-0.0	-0.0	+0.0	+0.0	+1.0	+F	+inf	Nan*
f	-inf	-F	-1.0	-denom	-0	+0	+denom	+1.0	+F	+inf	NaN														
d	-inf	-F	-1.0	-0.0	-0.0	+0.0	+0.0	+1.0	+F	+inf	Nan*														
<i>Format</i>	1-input, 1-output.																								
<i>Opcode</i>	<table> <thead> <tr> <th>Field Name</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>code</td> <td>15:0</td> <td>IL_OP_F_2_D</td> </tr> <tr> <td>control</td> <td>29:16</td> <td>Must be zero.</td> </tr> <tr> <td>sec_modifier_present</td> <td>30</td> <td>Must be zero.</td> </tr> <tr> <td>pri_modifier_present</td> <td>31</td> <td>Must be zero.</td> </tr> </tbody> </table>	Field Name	Bits	Description	code	15:0	IL_OP_F_2_D	control	29:16	Must be zero.	sec_modifier_present	30	Must be zero.	pri_modifier_present	31	Must be zero.									
Field Name	Bits	Description																							
code	15:0	IL_OP_F_2_D																							
control	29:16	Must be zero.																							
sec_modifier_present	30	Must be zero.																							
pri_modifier_present	31	Must be zero.																							
<i>Related</i>	F2OI, F2OU.																								

6.9 Float Instructions**Absolute Value**

<i>Instructions</i>	ABS															
<i>Syntax</i>	<code>abs dst, src0</code>															
<i>Description</i>	<p>Computes the absolute value of each component in a vector (<i>src0</i>).</p> <p>Operation:</p> <pre>VECTOR v1 = EvalSource(src0); VECTOR v; for (=0; i < 4; i++) v[i] = abs(v1[i]); WriteResult(v, dst);</pre>															
<i>Format</i>	1-input, 1-output.															
<i>Opcode</i>	<table> <thead> <tr> <th>Field Name</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>code</td> <td>15:0</td> <td>IL_OP_ABS</td> </tr> <tr> <td>control</td> <td>29:16</td> <td>Must be zero.</td> </tr> <tr> <td>sec_modifier_present</td> <td>30</td> <td>Must be zero.</td> </tr> <tr> <td>pri_modifier_present</td> <td>31</td> <td>Must be zero.</td> </tr> </tbody> </table>	Field Name	Bits	Description	code	15:0	IL_OP_ABS	control	29:16	Must be zero.	sec_modifier_present	30	Must be zero.	pri_modifier_present	31	Must be zero.
Field Name	Bits	Description														
code	15:0	IL_OP_ABS														
control	29:16	Must be zero.														
sec_modifier_present	30	Must be zero.														
pri_modifier_present	31	Must be zero.														
<i>Related</i>	None.															

Floating Point Addition

<i>Instructions</i>	ADD		
<i>Syntax</i>	add <i>dst</i> , <i>src0</i> , <i>src1</i>		
<i>Description</i>	Adds two float vectors: <i>src0</i> .{x y z w} + <i>src1</i> .{x y z w}. Operation: <pre>VECTOR v1 = EvalSource(<i>src1</i>); VECTOR v2 = EvalSource(<i>src2</i>); VECTOR v; for (i=0; i < 4; i++) v[i] = v1[i] + v2[i]; WriteResult(v, <i>dst</i>);</pre>		
<i>Format</i>	2-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_ADD
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	DADD.		

Inverse Cosine (arccos)

<i>Instructions</i>	ACOS		
<i>Syntax</i>	acos <i>dst</i> , <i>src0</i>		
<i>Description</i>	Computes the inverse cosine in radians of <i>src0</i> .w. By default, this instruction operates on <i>src0</i> .w, but can operate on any component by swizzling it into the w component. <i>src0</i> .w must be within the range [-1.0, 1.0]; otherwise, the result is undefined. The maximum absolute error is 0.002. Operation: <pre>VECTOR v1 = EvalSource(<i>src0</i>); VECTOR v v[0] = v[1] = v[2] = v[3] = acos(v1[3]); WriteResult(v, <i>dst</i>);</pre>		
<i>Format</i>	1-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_ACOS
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	COS_VEC, SIN_VEC, SINCOS.		

Logical-AND

<i>Instructions</i>	AND		
<i>Syntax</i>	and <i>dst</i> , <i>src0</i> , <i>src1</i>		
<i>Description</i>	Performs a component-wise logical AND of each pair of 32-bit values from <i>src1</i> and <i>src2</i> . The 32-bit result is placed in <i>dst</i> .		
<i>Format</i>	2-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_AND
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	None.		

Inverse Sine (arcsin)

<i>Instructions</i>	ASIN		
<i>Syntax</i>	asin <i>dst</i> , <i>src0</i>		
<i>Description</i>	<p>Computes the inverse sine in radians of the w component of <i>src0</i>. By default, this instruction operates on <i>src0.w</i>, but can operate on any component by swizzling it into the w component. <i>src0.w</i> must be within the range [-1.0, 1.0]; otherwise, the result is undefined. The maximum absolute error is 0.002.</p> <p>Operation:</p> <pre>VECTOR v1 = EvalSource(<i>src0</i>); VECTOR v v[0] = v[1] = v[2] = v[3] = asin(v1[3]); WriteResult(v, <i>dst</i>);</pre>		
<i>Format</i>	1-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_ASIN
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	COS_VEC, SIN_VEC, SINCOS.		

Inverse Tangent (arctan)

Instructions **ATAN**

Syntax `atan dst, src0`

Description Computes the inverse tangent in radians of `src0.w`. By default, this instruction operates on `src0.w`, but can operate on any component by swizzling it into the `w` component. `src0.w` must be within

the range $[-\frac{\pi}{2}, \frac{\pi}{2}]$; otherwise, the result of this instruction is undefined.

Operation:

```
VECTOR v1 = EvalSource(src0);
VECTOR v
v[0] = v[1] = v[2] = v[3] = atan(v1[3]);
WriteResult(v, dst);
```

Format 1-input, 1-output.

<i>Opcode</i>	Field Name	Bits	Description
	<code>code</code>	15:0	IL_OP_ATAN
	<code>control</code>	29:16	Must be zero.
	<code>sec_modifier_present</code>	30	Must be zero.
	<code>pri_modifier_present</code>	31	Must be zero.

Related COS_VEC, SIN_VEC, SINCOS.

Clamp to Given Range

<i>Instructions</i>	CLAMP		
<i>Syntax</i>	clamp <i>dst</i> , <i>src0</i> , <i>src1</i> , <i>src2</i>		
<i>Description</i>	Clamps <i>src0</i> to <i>src1</i> and <i>src2</i> . Operation: <pre> VECTOR v1 = EvalSource(<i>src0</i>); VECTOR v2 = EvalSource(<i>src1</i>); VECTOR v3 = EvalSource(<i>src2</i>); VECTOR v; for (i=0; i < 4; i++) { v[i] = v1[i]; if(v[i] < v2[i]) { v[i] = v2[i]; } if(v[i] > v3[i]) { v[i] = v3[i]; } } WriteResult(v, <i>dst</i>); </pre>		
<i>Format</i>	3-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_CLAMP
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	None.		

Ceiling

<i>Instructions</i>	CLG		
<i>Syntax</i>	clg <i>dst</i> , <i>src0</i>		
<i>Description</i>	<p>Calculates the ceiling of each component of <i>src0</i>, and places the 32-bit result in the corresponding component of <i>dst</i>. The ceiling of a component is defined as the smallest integer greater than or equal to the value in the component. For example, the ceiling of 2.3 is 3.0, and the ceiling of -3.6 is -3.0. The operation is identical to ROUND_PLUS_INF.</p> <p>Operation:</p> <pre>VECTOR v1 = EvalSource(src0); VECTOR v; for (i=0; i < 4; i++) v[i] = ceiling(v1[i]); WriteResult(v, dst);</pre>		
<i>Format</i>	1-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_CLG
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	FLR.		

Component-Wise Conditional Move

<i>Instructions</i>	CMOV															
<i>Syntax</i>	<code>cmov dst, src0, src1</code>															
<i>Description</i>	<p>Conditionally moves a value from <i>src1</i> to <i>dst</i>. If <i>src0</i>.{x y z w} is not 0.0f, the value of the corresponding component in <i>dst</i> becomes the value of the corresponding component of <i>src1</i>; otherwise, the corresponding component in <i>dst</i> remains unchanged.</p> <p>IL_DSTMOD_1 and IL_DSTMOD_0 in the IL_Dst_Mod token are treated as IL_DSTMOD_NOWRITE for components where <i>src1</i> is 0.0f.</p> <p>Operation:</p> <pre> VECTOR v1 = EvalSource(src0); VECTOR v2 = EvalSource(src1); VECTOR v; for (i=0; i < 4; i++) { v[i] = v2[i]; if(v1[i] == 0.0) v[i].MASK = NOWRITE; } WriteResult(v, dst); </pre>															
<i>Format</i>	2-input, 1-output.															
<i>Opcode</i>	<table> <thead> <tr> <th>Field Name</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>code</td> <td>15:0</td> <td>IL_OP_CMOV</td> </tr> <tr> <td>control</td> <td>29:16</td> <td>Must be zero.</td> </tr> <tr> <td>sec_modifier_present</td> <td>30</td> <td>Must be zero.</td> </tr> <tr> <td>pri_modifier_present</td> <td>31</td> <td>Must be zero.</td> </tr> </tbody> </table>	Field Name	Bits	Description	code	15:0	IL_OP_CMOV	control	29:16	Must be zero.	sec_modifier_present	30	Must be zero.	pri_modifier_present	31	Must be zero.
Field Name	Bits	Description														
code	15:0	IL_OP_CMOV														
control	29:16	Must be zero.														
sec_modifier_present	30	Must be zero.														
pri_modifier_present	31	Must be zero.														
<i>Related</i>	CMOV_LOGICAL.															

Component-wise Conditional Logical Move

<i>Instructions</i>	CMOV_LOGICAL															
<i>Syntax</i>	<code>cmov_logical dst, src0, src1, src2</code>															
<i>Description</i>	<p>For each component in <i>src0</i> (post-swizzle): if the component has any bit set, the component in <i>src1</i> (post-swizzle) is copied to the corresponding component in <i>dst</i>; otherwise, the corresponding component in <i>src2</i> is copied to the corresponding component in <i>dst</i>.</p>															
<i>Format</i>	3-input, 1-output.															
<i>Opcode</i>	<table> <thead> <tr> <th>Field Name</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>code</td> <td>15:0</td> <td>IL_OP_CMOV_LOGICAL</td> </tr> <tr> <td>control</td> <td>29:16</td> <td>Must be zero.</td> </tr> <tr> <td>sec_modifier_present</td> <td>30</td> <td>Must be zero.</td> </tr> <tr> <td>pri_modifier_present</td> <td>31</td> <td>Must be zero.</td> </tr> </tbody> </table>	Field Name	Bits	Description	code	15:0	IL_OP_CMOV_LOGICAL	control	29:16	Must be zero.	sec_modifier_present	30	Must be zero.	pri_modifier_present	31	Must be zero.
Field Name	Bits	Description														
code	15:0	IL_OP_CMOV_LOGICAL														
control	29:16	Must be zero.														
sec_modifier_present	30	Must be zero.														
pri_modifier_present	31	Must be zero.														
<i>Related</i>	CMOV.															

Compare with Value

<i>Instructions</i>	CMP																														
<i>Syntax</i>	<code>cmp_relop(op)_cmpval(cmpval) dst, src0, src1, src2</code>																														
<i>Description</i>	<p>For every component of <i>src0</i> that passes (is TRUE for) its relational comparison with <i>cmpval</i>, that component is set to the corresponding component of <i>src1</i>; otherwise, it is set to the corresponding component of <i>src2</i>.</p> <p>Operation:</p> <pre> VECTOR v1 = EvalSource(src1); VECTOR v2 = EvalSource(src2); VECTOR v3 = EvalSource(src3); VECTOR v; for (i=0; i < 4; i++) v[i] = (v1[i] relop cmpval) ? v2[i] : v3[i]; WriteResult(v, dst); </pre>																														
<i>Format</i>	3-input, 1-output.																														
<i>Opcode</i>	<table border="0"> <thead> <tr> <th style="text-align: left;">Field Name</th> <th style="text-align: left;">Bits</th> <th style="text-align: left;">Description</th> </tr> </thead> <tbody> <tr> <td>code</td> <td>15:0</td> <td>IL_OP_CMP</td> </tr> <tr> <td>control</td> <td>29:16</td> <td> <table border="0"> <thead> <tr> <th style="text-align: left;">Field Name</th> <th style="text-align: left;">Bits</th> <th style="text-align: left;">Description</th> </tr> </thead> <tbody> <tr> <td>relop</td> <td>18:16</td> <td>Any value of the enumerated type ILRelOp(<i>relop</i>). See Table 5.14 on page 5-8.</td> </tr> <tr> <td><i>reserved</i></td> <td>20:19</td> <td>Must be zero.</td> </tr> <tr> <td>cmpval</td> <td>23:21</td> <td>Any value of the enumerated type ILCmpValue(<i>cmpval</i>). See Table 5.17 on page 5-9.</td> </tr> <tr> <td><i>reserved</i></td> <td>29:24</td> <td>Must be zero.</td> </tr> </tbody> </table> </td> </tr> <tr> <td>sec_modifier_present</td> <td>30</td> <td>Must be zero.</td> </tr> <tr> <td>pri_modifier_present</td> <td>31</td> <td>Must be zero.</td> </tr> </tbody> </table>	Field Name	Bits	Description	code	15:0	IL_OP_CMP	control	29:16	<table border="0"> <thead> <tr> <th style="text-align: left;">Field Name</th> <th style="text-align: left;">Bits</th> <th style="text-align: left;">Description</th> </tr> </thead> <tbody> <tr> <td>relop</td> <td>18:16</td> <td>Any value of the enumerated type ILRelOp(<i>relop</i>). See Table 5.14 on page 5-8.</td> </tr> <tr> <td><i>reserved</i></td> <td>20:19</td> <td>Must be zero.</td> </tr> <tr> <td>cmpval</td> <td>23:21</td> <td>Any value of the enumerated type ILCmpValue(<i>cmpval</i>). See Table 5.17 on page 5-9.</td> </tr> <tr> <td><i>reserved</i></td> <td>29:24</td> <td>Must be zero.</td> </tr> </tbody> </table>	Field Name	Bits	Description	relop	18:16	Any value of the enumerated type ILRelOp(<i>relop</i>). See Table 5.14 on page 5-8.	<i>reserved</i>	20:19	Must be zero.	cmpval	23:21	Any value of the enumerated type ILCmpValue(<i>cmpval</i>). See Table 5.17 on page 5-9.	<i>reserved</i>	29:24	Must be zero.	sec_modifier_present	30	Must be zero.	pri_modifier_present	31	Must be zero.
Field Name	Bits	Description																													
code	15:0	IL_OP_CMP																													
control	29:16	<table border="0"> <thead> <tr> <th style="text-align: left;">Field Name</th> <th style="text-align: left;">Bits</th> <th style="text-align: left;">Description</th> </tr> </thead> <tbody> <tr> <td>relop</td> <td>18:16</td> <td>Any value of the enumerated type ILRelOp(<i>relop</i>). See Table 5.14 on page 5-8.</td> </tr> <tr> <td><i>reserved</i></td> <td>20:19</td> <td>Must be zero.</td> </tr> <tr> <td>cmpval</td> <td>23:21</td> <td>Any value of the enumerated type ILCmpValue(<i>cmpval</i>). See Table 5.17 on page 5-9.</td> </tr> <tr> <td><i>reserved</i></td> <td>29:24</td> <td>Must be zero.</td> </tr> </tbody> </table>	Field Name	Bits	Description	relop	18:16	Any value of the enumerated type ILRelOp(<i>relop</i>). See Table 5.14 on page 5-8.	<i>reserved</i>	20:19	Must be zero.	cmpval	23:21	Any value of the enumerated type ILCmpValue(<i>cmpval</i>). See Table 5.17 on page 5-9.	<i>reserved</i>	29:24	Must be zero.														
Field Name	Bits	Description																													
relop	18:16	Any value of the enumerated type ILRelOp(<i>relop</i>). See Table 5.14 on page 5-8.																													
<i>reserved</i>	20:19	Must be zero.																													
cmpval	23:21	Any value of the enumerated type ILCmpValue(<i>cmpval</i>). See Table 5.17 on page 5-9.																													
<i>reserved</i>	29:24	Must be zero.																													
sec_modifier_present	30	Must be zero.																													
pri_modifier_present	31	Must be zero.																													
<i>Related</i>	None.																														

Clamp Color Output Values

<i>Instructions</i>	COLORCLAMP		
<i>Syntax</i>	colorclamp <i>dst</i> , <i>src0</i>		
<i>Description</i>	Clamps (or does not clamp) an output color to values specified by the frame buffer color clamp state. Operation: <pre> VECTOR v1 = EvalSource(<i>src0</i>); int buffer = RegisterNum(<i>src0</i>); VECTOR v; for (i=0; i < 4; i++) { if(AS_CB_CLAMP_MODE_N(buffer) == 0_1) { v[i] = (v1[i] <= 0.0) ? 0.0 : v1[i]; v[i] = (v[i] > 1.0) ? 1.0 : v[i]; } else if (AS_CB_CLAMP_MODE_N(buffer) == NEG_1_1) { v[i] = (v1[i] < -1.0) ? -1.0 : v1[i]; v[i] = (v[i] > 1.0) ? 1.0 : v[i]; } else if (AS_CB_CLAMP_MODE_N(buffer) == NONE) { v[i]=v1[i]; } } WriteResult(v, <i>dst</i>); </pre>		
<i>Format</i>	1-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_COLORCLAMP
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	None.		

Cosine (cos)

<i>Instructions</i>	COS		
<i>Syntax</i>	<i>cos dst, src0</i>		
<i>Description</i>	<p>Computes the cosine of <i>src0.w</i>. By default, this instruction operates on <i>src0.w</i>, but can operate on any component by swizzling it into the fourth component.</p> <p>The fourth component of <i>src0</i> must be in the range $[-\pi, \pi]$; otherwise, the result is of this instruction is undefined.</p> <p>The max absolute error is 0.002.</p>		
<i>Format</i>	1-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_COS
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	COS_VEC, SIN_VEC, SINCOS.		

Cross Product

<i>Instructions</i>	CRS		
<i>Syntax</i>	<i>crs dst, src0, src1</i>		
<i>Description</i>	<p>Computes a cross product. This instruction does not write to the <i>dst.w</i> component; however, the if the component <i>w_a</i> field of the IL_Dst_Mod token is set to IL_MODCOMP_0 or IL_MODCOMP_1, <i>dst.w</i> is written. That is, <i>dst.w</i> can be set to 0.0 or 1.0 if IL_MODCOMP_0 or IL_MODCOMP_1 is used on the component <i>w_a</i> field of the IL_Dst_Mod token.</p> <p>Operation:</p> <pre> VECTOR v1 = EvalSource(src0); VECTOR v2 = EvalSource(src1); VECTOR v; v[0] = v1[1] * v2[2] - v1[2] * v2[1]; v[1] = v1[2] * v2[0] - v1[0] * v2[2]; v[2] = v1[0] * v2[1] - v1[1] * v2[0]; WriteResult(v, dst); </pre>		
<i>Format</i>	2-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_CRS
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	None.		

Determinant of a 4x4 Matrix -- WARNING: This instruction has been deleted.

<i>Instructions</i>	DET		
<i>Syntax</i>	det <i>dst</i> , <i>src0</i>		
<i>Description</i>	<p>Calculates the determinant of a 4x4 matrix. Relative addressing and source modifiers of <i>src0</i> are applied to each row of the matrix beginning with <i>src0</i>. The 32-bit scalar result is placed in all four components of <i>dst</i>.</p> <p>Operation:</p> <pre> VECTOR v0 = EvalSource(src0); VECTOR v1 = EvalSource(src0+1); # Register number of src0 + 1 VECTOR v2 = EvalSource(src0+2); # Register number of src0 + 2 VECTOR v3 = EvalSource(src0+3); # Register number of src0 + 3 VECTOR v; float f; f = v0[0]*(v1[1]*(v2[2]*v3[3]-v2[3]*v3[2])-v1[2]*(v2[1]*v3[3]- v3[1]*v2[3])+v1[3]*(v2[1]*v3[2]-v3[1]*v2[2])) - v0[1]*(v1[0]*(v2[2]*v3[3]-v2[3]*v3[2])-v1[2]*(v2[0]*v3[3]- v3[0]*v2[3])+v1[3]*(v2[0]*v3[2]-v3[0]*v2[2])) + v0[2]*(v1[0]*(v2[1]*v3[3]-v2[3]*v3[1])-v1[1]*(v2[0]*v3[3]- v3[0]*v2[3])+v1[3]*(v2[0]*v3[1]-v3[0]*v2[1])) - v0[3]*(v1[0]*(v2[1]*v3[2]-v2[2]*v3[1])-v1[1]*(v2[0]*v3[2]- v3[0]*v2[2])+v1[2]*(v2[0]*v3[1]-v3[0]*v2[1])); v[0] = v[1] = v[2] = v[3] = f; WriteResult(v, dst); </pre>		
<i>Format</i>	1-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_DET
	control	29:16	Must be the enumerated type ILMatrix(IL_MATRIX_4X4). See Table 5.6 on page 5-3.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	None.		

Vector Distance

Instructions **DIST**

Syntax `dist dst, src0, src1`

Description Computes the vector distance from `src0.[xyz]` to `src1.[xyz]`. The 32-bit scalar result is placed in all four components of `dst`.

Operation:

```
VECTOR v1 = EvalSource(src0;
VECTOR v2 = EvalSource(src1;
VECTOR v;
V[0] = v[1] = v[2] = v[3] =
    sqrt((v1[0]-v2[0])*(v1[0]-v2[0]) +
        (v1[1]-v2[1])*(v1[1]-v2[1]) +
        (v1[2]-v2[2])*(v1[2]-v2[2]) +
        (v1[3]-v2[3])*(v1[3]-v2[3]));
WriteResult(v, dst);
```

Format 2-input, 1-output.

<i>Opcode</i>	Field Name	Bits	Description
	<code>code</code>	15:0	IL_OP_DIST
	<code>control</code>	29:16	Must be zero.
	<code>sec_modifier_present</code>	30	Must be zero.
	<code>pri_modifier_present</code>	31	Must be zero.

Related None.

Floating Point Division

<i>Instructions</i>	DIV		
<i>Syntax</i>	<code>div_zeroop(op) dst, src0, src1</code>		
<i>Description</i>	<p>Floating point division $src0.\{x y z w\} / src1.\{x y z w\}$. The 32-bit quotient is placed in the corresponding component of <i>dst</i>. DX10 requires the DIV instruction to use zeroop = IL_ZEROOP_INFINITY.</p> <p>Operation:</p> <pre> VECTOR v1 = EvalSource(src0); VECTOR v2 = EvalSource(src1); VECTOR v; float f; for (i=0; i < 4; i++) { if(v2[i] == 0.0) { if(zeroop == IL_ZEROOP_0) f = 0.0; else if(zeroop == IL_ZEROOPs_FLT_MAX) f = FLT_MAX; else if(zeroop == IL_ZEROOP_INFINITY) f = INFINITY; else if(zeroop == IL_ZEROOP_INF_ELSE_MAX) f = INFINITY or FLT_MAX; # Depends on IL Implementation if(v1[i] < 0.0) f = -f; v[i] = f; } else v[i] = v1[i] / v2[i]; } WriteResult(v, dst); </pre>		
<i>Format</i>	2-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_DIV
	control	29:16	DIV: Any value of the enumerated type ILZeroOp(<i>zeroop</i>). See Table 5.16 on page 5-9.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	None.		

Two-Component Dot Product and Scalar Add

<i>Instructions</i>	DP2ADD		
<i>Syntax</i>	dp2add <i>dst</i> , <i>src0</i> , <i>src1</i> , <i>src2</i>		
<i>Description</i>	<p>Computes the two-component dot product of <i>src0</i>.[xy] and <i>src1</i>.[xy] and adds scalar <i>src2</i>.z to the result. The 32-bit scalar result is placed in all four components of <i>dst</i>. This instruction corresponds to the DX10 dp2 operation. If the control field is 0, then 0*n equals zero, even if n = NaN.</p> <p>Operation:</p> <pre> VECTOR v1 = EvalSource(<i>src0</i>); VECTOR v2 = EvalSource(<i>src1</i>); VECTOR v3 = EvalSource(<i>src2</i>); VECTOR v; v[0] = v[1] = v[2] = v[3] = v1[0]*v2[0] + v1[1]*v2[1] + v3[2]; WriteResult(v, <i>dst</i>); </pre>		
<i>Format</i>	3-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_DP2ADD
	control	29:16	0 DX9 DP2-style multiple. 1 IEEE-style multiply.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	DP2, DP3, DP4.		

Two-Component Dot Product

<i>Instructions</i>	DP2		
<i>Syntax</i>	dp2[_ieee] <i>dst</i> , <i>src0</i> , <i>src1</i>		
<i>Description</i>	Two-component dot product of vector operands <i>src0</i> . <i>[xy]</i> and <i>src1</i> . <i>[xy]</i> . The 32-bit scalar result is placed in all four components of <i>dst</i> . If the control field is 0, then 0*n = 0, even if n = NaN. Operation: VECTOR v1 = EvalSource(<i>src0</i>); VECTOR v2 = EvalSource(<i>src1</i>); VECTOR v; v[0] = v[1] = v[2] = v[3] = v1[0]*v2[0] + v1[1]*v2[1]; WriteResult(v, <i>dst</i>);		
<i>Format</i>	2-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_DP2
	control	29:16	0 DX9 DP2-style multiple. 1 IEEE-style multiply.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	NoneDP2ADD.		

Three-Component Dot Product

<i>Instructions</i>	DP3		
<i>Syntax</i>	dp3[_ieee] <i>dst</i> , <i>src0</i> , <i>src1</i>		
<i>Description</i>	Three-component dot product vector operands <i>src0</i> . <i>[xyz]</i> and <i>src1</i> . <i>[xyz]</i> . The 32-bit scalar result is placed in all four components of <i>dst</i> . If the control field is 0, then 0*n = 0, even if n = NaN. Operation: VECTOR v1 = EvalSource(<i>src0</i>); VECTOR v2 = EvalSource(<i>src1</i>); VECTOR v; v[0] = v[1] = v[2] = v[3] = v1[0]*v2[0] + v1[1]*v2[1] + v1[2]*v2[2]; WriteResult(v, <i>dst</i>);		
<i>Format</i>	2-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_DP3
	control	29:16	0 DX9 DP2-style multiple. 1 IEEE-style multiply.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	NoneDP2ADD.		

Four-Component Dot Product

<i>Instructions</i>	DP4		
<i>Syntax</i>	dp4[_ieee] <i>dst</i> , <i>src0</i> , <i>src1</i>		
<i>Description</i>	<p>Four-component dot product vector operands <i>src0</i>.[xyzw] and <i>src1</i>.[xyzw]. The 32-bit scalar result is placed in all four components of <i>dst</i>. If the control field is 0, then $0*n = 0$, even if $n = \text{NaN}$.</p> <p>Operation:</p> <pre> VECTOR v1 = EvalSource(src0); VECTOR v2 = EvalSource(src1); VECTOR v; v[0] = v[1] = v[2] = v[3] = v1[0]*v2[0] + v1[1]*v2[1] + v1[2]*v2[2] + v1[3]*v2[3]; WriteResult(v, dst); </pre>		
<i>Format</i>	2-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_DP4
	control	29:16	0 DX9 DP2-style multiple. 1 IEEE-style multiply.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	NoneDP2ADD.		

Vector Distance

<i>Instructions</i>	DST		
<i>Syntax</i>	<code>dst dst, src0, src1</code>		
<i>Description</i>	<p>Computes a distance-related value from operands <i>src0</i> and <i>src1</i>. The result vector is placed in <i>dst</i>. If the control field is 0, then $0^n = 0$, even if $n = \text{NaN}$.</p> <p>Operation:</p> <pre> VECTOR v1 = EvalSource(src0); VECTOR v2 = EvalSource(src1); VECTOR v; V[0] = 1.0; V[1] = v1[1] * v2[1]; V[2] = v1[2]; V[3] = v2[3]; WriteResult(v, dst); </pre>		
<i>Format</i>	2-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_DST
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	None.		

Instantaneous Derivative in X

Instructions **DSX**

Syntax `dsx[_centroid] dst, src0`

Description Instantaneous derivative in the x-direction. Computes the rate of change of each float32 component of *src0* (post-swizzle) in the RenderTarget x direction. The results are undefined if used in a vertex or geometry shader.

The data in the current pixel shader invocation might participate in the calculation of the requested derivative, since the derivative is calculated only once per 2x2 quad.

Format 1-input, 1-output.

<i>Opcode</i>	Field Name	Bits	Description															
	<code>code</code>	15:0	IL_OP_DSX															
	<code>control</code>	29:16	<table border="0"> <tr> <td>Field Name</td> <td>Bits</td> <td>Description</td> </tr> <tr> <td><i>reserved</i></td> <td>22:16</td> <td>Must be zero.</td> </tr> <tr> <td><i>gradient</i></td> <td>23</td> <td>0: Gradients can be computed once per quad. 1: Gradients are computed for each pixel.</td> </tr> <tr> <td><i>centroid</i></td> <td>24</td> <td>0: Pixel center is used to determine the derivative value. 1: Pixel centroid is used to determine the derivative value.</td> </tr> <tr> <td><i>reserved</i></td> <td>29:25</td> <td>Must be zero.</td> </tr> </table>	Field Name	Bits	Description	<i>reserved</i>	22:16	Must be zero.	<i>gradient</i>	23	0: Gradients can be computed once per quad. 1: Gradients are computed for each pixel.	<i>centroid</i>	24	0: Pixel center is used to determine the derivative value. 1: Pixel centroid is used to determine the derivative value.	<i>reserved</i>	29:25	Must be zero.
Field Name	Bits	Description																
<i>reserved</i>	22:16	Must be zero.																
<i>gradient</i>	23	0: Gradients can be computed once per quad. 1: Gradients are computed for each pixel.																
<i>centroid</i>	24	0: Pixel center is used to determine the derivative value. 1: Pixel centroid is used to determine the derivative value.																
<i>reserved</i>	29:25	Must be zero.																
	<code>sec_modifier_present</code>	30	Must be zero.															
	<code>pri_modifier_present</code>	31	Must be zero.															

Related None.

Instantaneous Derivative in Y

<i>Instructions</i>	DSY																	
<i>Syntax</i>	dsy[_centroid] dst, src0																	
<i>Description</i>	<p>Instantaneous derivative in the y-direction. Computes the rate of change of each float32 component of <i>src0</i> (post-swizzle) in the RenderTarget y direction. The results are undefined if used in a vertex or geometry shader.</p> <p>The data in the current pixel shader invocation might participate in the calculation of the requested derivative, since the derivative is calculated only once per 2x2 quad.</p>																	
<i>Format</i>	1-input, 1-output.																	
<i>Opcode</i>	Field Name	Bits	Description															
	code	15:0	IL_OP_DSY															
	control	29:16	<table> <thead> <tr> <th>Field Name</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td><i>reserved</i></td> <td>22:16</td> <td>Must be zero.</td> </tr> <tr> <td>gradient</td> <td>23</td> <td>0: Gradients can be computed once per quad. 1: Gradients are computed for each pixel.</td> </tr> <tr> <td>centroid</td> <td>24</td> <td>0: Pixel center is used to determine the derivative value. 1: Pixel centroid is used to determine the derivative value.</td> </tr> <tr> <td><i>reserved</i></td> <td>29:25</td> <td>Must be zero.</td> </tr> </tbody> </table>	Field Name	Bits	Description	<i>reserved</i>	22:16	Must be zero.	gradient	23	0: Gradients can be computed once per quad. 1: Gradients are computed for each pixel.	centroid	24	0: Pixel center is used to determine the derivative value. 1: Pixel centroid is used to determine the derivative value.	<i>reserved</i>	29:25	Must be zero.
Field Name	Bits	Description																
<i>reserved</i>	22:16	Must be zero.																
gradient	23	0: Gradients can be computed once per quad. 1: Gradients are computed for each pixel.																
centroid	24	0: Pixel center is used to determine the derivative value. 1: Pixel centroid is used to determine the derivative value.																
<i>reserved</i>	29:25	Must be zero.																
	sec_modifier_present	30	Must be zero.															
	pri_modifier_present	31	Must be zero.															
<i>Related</i>	None.																	

Compute Sine and Cosine

<i>Instructions</i>	DXSINCOS		
<i>Syntax</i>	dxsincos <i>dst</i> , <i>src0</i> , <i>src1</i> , <i>src2</i>		
<i>Description</i>	<p>Computes sine and cosine values of <i>src0.w</i> for the legacy DX9 sincos instruction. The results are returned in <i>dst.x</i> and <i>dst.y</i>.</p> <p>This instruction supports the DX legacy instruction <i>sincos dst, src0, src1, src2</i>. <i>src1</i> and <i>src2</i> must be constant float registers and are used in a Taylor series expansion. See the <i>DX SDK</i> to understand how the two constants are used in a Taylor series expansion. Devices with native sincos support ignore <i>src1</i> and <i>src2</i>.</p> <p>By default this instruction operates on <i>src0.w</i>, but can operate on any component by swizzling it into the fourth component. The fourth component of <i>src</i> must be within the range $[-\pi, \pi]$; otherwise, the result of this instruction is undefined.</p> <p>This instruction does not write to the <i>dst.z</i> and <i>dst.w</i> components; however, if the <i>component_z_b</i> or the <i>component_w_a</i> field of the <i>IL_Dst_Mod</i> token is set to <i>IL_MODCOMP_0</i> or <i>IL_MODCOMP_1</i>, the <i>dst.z</i> and <i>dst.w</i> are written. That is, <i>dst.z</i> and <i>dst.w</i> can be set to 0.0 or 1.0 if <i>IL_MODCOMP_0</i> or <i>IL_MODCOMP_1</i> is used on the <i>component_z_b</i> or <i>component_w_a</i> field of the <i>IL_Dst_Mod</i> token.</p> <p>The maximum absolute error is 0.002.</p> <p>Operation:</p> <pre>VECTOR v1 = EvalSource(src0); VECTOR v; v[0] = cos(v1[3]); // approximated by using Taylor series v[1] = sin(v1[3]); // approximated by using Taylor series WriteResult(v, dst);</pre>		
<i>Format</i>	3-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_DXSINCOS
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	None.		

Equality Compare Floats

<i>Instructions</i>	EQ		
<i>Syntax</i>	<i>eq dst, src0, src1</i>		
<i>Description</i>	<p>Performs the float comparison $src0 == src1$ for each component, and writes the result to <i>dst</i>. If the comparison is true, 0xFFFFFFFF is returned for that component; otherwise, 0x00000000 is returned. This instruction follows DX10 Floating Point Rules. Denorms are flushed before comparison (original source registers untouched). +0 equals -0. Comparison with NaN returns false.</p> <p>Primary and secondary opcode modifiers are not permitted, and no additional control options are supported. Thus, the <code>pri_modifier_present</code> and <code>sec_modifier_present</code> fields must be zero.</p>		
<i>Format</i>	2-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_EQ
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	Double-precision Compare, Integer Compare, Unsigned Integer Compare.		

Full-Precision e to the Power of x

<i>Instructions</i>	EXN		
<i>Syntax</i>	<i>exn dst, src0</i>		
<i>Description</i>	<p>Full-precision base e power <i>src0</i>: e^{src0}.</p> <p>By default, this instruction operates on <i>src0.w</i>, but can operate on any component by swizzling it into the fourth component.</p> <p>Operation:</p> <pre>VECTOR v1 = EvalSource(src0); VECTOR v; v[0] = v[1] = v[2] = v[3] = exp2(v1[3] * log2(e)); WriteResult(v, dst);</pre>		
<i>Format</i>	1-input 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_EXN
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	None.		

2 Raised to a Power

<i>Instructions</i>	EXP		
<i>Syntax</i>	<i>exp dst, src0</i>		
<i>Description</i>	<p>Full precision base 2 power <i>src0</i>: 2^{src}.</p> <p>Computes two to the power of <i>src0.w</i>. The floating point result is placed in all four components of <i>dst</i>. By default this instruction operates on <i>src0.w</i>, but can operate on any component by swizzling it into the fourth component.</p>		
<i>Format</i>	1-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_EXP
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	EXP_VEC, POW.		

Component-Wise Full Precision 2 to the Power of X

<i>Instructions</i>	EXP_VEC		
<i>Syntax</i>	<i>exp_vec dst, src0</i>		
<i>Description</i>	<p>EXP computes a scalar version, EXP_VEC computes the same result per component.</p> <p>Computes 2 raised to the power of each component of <i>src0</i>. The result of the operation is accurate to at least 21 bits. The 32-bit floating point results are placed in the corresponding components of <i>dst</i>.</p>		
<i>Format</i>	1-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_EXP_VEC
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	EXN, EXP, EXPP, POW.		

Component-Wise Partial Precision 2 to the Power of X

<i>Instructions</i>	EXPP		
<i>Syntax</i>	<code>expp dst, src0</code>		
<i>Description</i>	<p>Computes partial precision of 2 raised to the power of each component of <i>src0</i>. By default, this instruction operates on <i>src0.w</i>, but can operate on any component by swizzling it into the fourth component. The result of the operation is accurate to at least 10 bits. The 32-bit floating point results are placed in the corresponding components of <i>dst</i>.</p> <p>Operation:</p> <pre> VECTOR v1 = EvalSource(src0); VECTOR v; float f = floor(v1[3]); v[0] = exp₂(f); v[1] = v1[3] - f; v[2] = exp₂(v1[3]); v[3] = 1.0; WriteResult(v, dst); </pre>		
<i>Format</i>	1-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_EXPP
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	EXN, EXP, EXPP, POW.		

Face Forward

<i>Instructions</i>	FACEFORWARD		
<i>Syntax</i>	<code>faceforward dst, src0, src1, src2</code>		
<i>Description</i>	<p>FACEFORWARD performs the following calculation:</p> $d(dst = src2 * \text{sign}(\text{dot}(src0, src1)))$		
<i>Format</i>	3-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_FACEFORWARD
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	None.		

Floor

<i>Instructions</i>	FLR															
<i>Syntax</i>	<code>flr dst, src0</code>															
<i>Description</i>	<p>Performs a component-wise floor operation on the operand to generate a result vector. Calculates the floor of each component of <i>src0</i>, and places the 32-bit result in the corresponding component of <i>dst</i>. The floor of a component is defined as the largest integer less-than-or -equal to the value in the component. For example, the floor of 2.3 is 2.0 and the floor of -3.6 is -4.0. The operation is identical to ROUND_NEG_INF.</p> <p>Operation:</p> <pre>VECTOR v1 = EvalSource(src0); VECTOR v; for (i=0; i < 4; i++) v[i] = floor(v1[i]); WriteResult(v, dst);</pre>															
<i>Format</i>	1-input, 1-output.															
<i>Opcode</i>	<table border="0"> <thead> <tr> <th style="text-align: left;">Field Name</th> <th style="text-align: left;">Bits</th> <th style="text-align: left;">Description</th> </tr> </thead> <tbody> <tr> <td>code</td> <td>15:0</td> <td>IL_OP_FLR</td> </tr> <tr> <td>control</td> <td>29:16</td> <td>Must be zero.</td> </tr> <tr> <td>sec_modifier_present</td> <td>30</td> <td>Must be zero.</td> </tr> <tr> <td>pri_modifier_present</td> <td>31</td> <td>Must be zero.</td> </tr> </tbody> </table>	Field Name	Bits	Description	code	15:0	IL_OP_FLR	control	29:16	Must be zero.	sec_modifier_present	30	Must be zero.	pri_modifier_present	31	Must be zero.
Field Name	Bits	Description														
code	15:0	IL_OP_FLR														
control	29:16	Must be zero.														
sec_modifier_present	30	Must be zero.														
pri_modifier_present	31	Must be zero.														
<i>Related</i>	CLG.															

Fraction

<i>Instructions</i>	FRC		
<i>Syntax</i>	<code>frc dst, src0</code>		
<i>Description</i>	<p>Extracts the fractional portion of each component of <i>src0</i>. The results are returned in the corresponding component of <i>dst</i>. The fractional portion of a component is defined as the result after subtracting the floor of the component from the component (see FLR). The result is always in the range [0.0, 1.0).</p> <p>For negative values, the fractional portion is not the number to the right of the decimal point. For example, the fractional portion of -1.7 is 0.3, not 0.7. In this case, it is produced by subtracting the floor of -1.7 - (-2.0) from 1.7.</p> <p>Operation:</p> <pre>VECTOR v1 = EvalSource(src0); VECTOR v; for (i=0; i < 4; i++) v[i] = v1[i] - (float)floor(v1[i]); WriteResult(v, dst);</pre>		
<i>Format</i>	1-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_FRC
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	DFRAC, TRC.		

Filter Width

<i>Instructions</i>	FWIDTH		
<i>Syntax</i>	<code>fwidth dst, src0</code>		
<i>Description</i>	<p>Computes the sum of the absolute derivative in x and y using local differencing for each component of <i>src0</i>. The result is returned in the corresponding component of <i>dst</i>. If used in a vertex shader, the results are undefined.</p> <p>Operation:</p> <pre>VECTOR v1 = EvalSource(src0); for (i=0; i < 4; i++) v[i] = abs(dPdx(v1[i])) + abs(dPdy(v1[i]));</pre>		
<i>Format</i>	1-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_FWIDTH
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	None.		

Greater or Equal Floats

<i>Instructions</i>	GE		
<i>Syntax</i>	<i>ge dst, src0, src1</i>		
<i>Description</i>	<p>Compares two float vectors <i>src0</i> and <i>src1</i> for each component, and writes the result to <i>dst</i>. If the comparison is true, 0xFFFFFFFF is returned for that component; otherwise, 0x00000000 is returned. This instruction follows DX10 Floating Point Rules. Denorms are flushed before comparison (original source registers untouched). +0 equals -0. Comparison with NaN returns false.</p> <p>Primary and secondary opcode modifiers are not permitted, and no additional control options are supported. Thus, the <code>pri_modifier_present</code> and <code>sec_modifier_present</code> fields must be zero.</p>		
<i>Format</i>	2-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_EQ
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	Double-precision Compare, Integer Compare, Unsigned Integer Compare.		

Vector Length

<i>Instructions</i>	LEN		
<i>Syntax</i>	<i>len dst, src0</i>		
<i>Description</i>	<p>Computes the length of a vector. Computes the vector length of the three component vector in <i>src0</i>[.xyz]. The scalar 32-bit floating point result is placed in all four components of <i>dst</i>.</p> <p>Operation:</p> <pre>VECTOR v1 = EvalSource(src0); VECTOR v; V[0] = v[1] = v[2] = v[3] = sqrt(v1[0]*v1[0]+ v1[1]*v1[1] + v1[2]*v1[2] + v1[3]*v1[3]); WriteResult(v, dst);</pre>		
<i>Format</i>	1-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_LEN
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	None.		

Lighting Coefficient

<i>Instructions</i>	LIT		
<i>Syntax</i>	lit <i>dst</i> , <i>src0</i>		
<i>Description</i>	<p>Calculates lighting coefficients for ambient, diffuse, and specular light contributions.</p> <p>Operation:</p> <pre> VECTOR v1 = EvalSource(<i>src0</i>); VECTOR v; float epsilon = 1.192092896e-07F; float g = v1[3]; if(g < -(128.0-epsilon)) g = -(128.0-epsilon); else if (g > (128.0-epsilon)) g = 128.0-epsilon; v[0] = 1.0; v[1] = (v1[0] > 0.0) ? v1[0] : 0.0; v[2] = ((v1[0] > 0.0) && (v1[1] > 0)) ? EXP2(g * LOG2(v1[1])) : 0.0; v[3] = 1.0; WriteResult(v, <i>dst</i>); </pre>		
<i>Format</i>	1-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_LIT
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	None.		

Natural Logarithm

Instructions **LN**

Syntax `ln_zeroop(op) dst, src0`

Description Computes the natural logarithm of *src0.w*. The result of the operation is accurate to at least bits. The 32-bit floating point result is placed in all four components of *dst*. By default this instruction operates on *src0.w*, but can operate on any component by swizzling it into the fourth component. *zeroop* can be any value of the enumerated type `ILZeroOp(zeroop)` except `IL_ZEROOP_0`.

Operation:

```
VECTOR v1 = EvalSource(src0);
VECTOR v;
float f;
if (v1[3] == 0.0)
{
    if(zeroop == IL_ZEROOP_FLT_MAX)
        f = -FLT_MAX;
    else if(zeroop == IL_ZEROOP_INFINITY)
        f = -INFINITY;
    else if(zeroop == IL_ZEROOP_INF_ELSE_MAX)
        f = -INFINITY or -FLT_MAX; # Depends on IL Implementation
}
else if (v1[3] < 0.0)
{
    f = undefined;
}
else
{
    f = (float)(log10(v1[3])/log10(e));
}
v[0] = v[1] = v[2] = v[3] = f;
WriteResult(v, dst);
```

Format 1-input, 1-output.

Opcode	Field Name	Bits	Description
	code	15:0	IL_OP_LN
	control	29:16	Any value of the enumerated type <code>ILZeroOp(zeroop)</code> . See Table 5.16 on page 5-9.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.

Related LOG_VEC, LOGP.

Base-2 Logarithm

<i>Instructions</i>	LOG		
<i>Syntax</i>	<code>log_zeroop(op) dst, src0</code>		
<i>Description</i>	<p>Computes the base-2 logarithm of <code>src0.w</code>. The result of the operation is accurate to at least 21 bits. The 32-bit floating point result is placed in all four components of <code>dst</code>. By default this instruction operates on <code>src0.w</code>, but can operate on any component by swizzling it into the fourth component. <code>zeroop</code> can be any value of the enumerated type <code>ILZeroOp(zeroop)</code> except <code>IL_ZEROOP_0</code>.</p> <p>Operation:</p> <pre> VECTOR v1 = EvalSource(src0); VECTOR v; float f; if (v1[3] == 0) { if(zeroop == IL_ZEROOP_FLT_MAX) f = -FLT_MAX; else if(zeroop == IL_ZEROOP_INFINITY) f = -INFINITY; else if(zeroop == IL_ZEROOP_INF_ELSE_MAX) f = -INFINITY or -FLT_MAX; # Depends on IL Implementation } else if (v1[3] < 0.0) { f = undefined; } else { f = (float)(log10(v1[3])/log10(2)); } v[0] = v[1] = v[2] = v[3] = f; WriteResult(v, dst); </pre>		
<i>Format</i>	1-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	<code>code</code>	15:0	<code>IL_OP_LOG</code>
	<code>control</code>	29:16	Any value of the enumerated type <code>ILZeroOp(zeroop)</code> . See Table 5.16 on page 5-9.
	<code>sec_modifier_present</code>	30	Must be zero.
	<code>pri_modifier_present</code>	31	Must be zero.
<i>Related</i>	<code>LOG_VEC</code> , <code>LOGP</code> .		

Component-wise Base-2 Logarithm

<i>Instructions</i>	LOG_VEC		
<i>Syntax</i>	log_vec dst, src0		
<i>Description</i>	Computes the base-2 logarithm of each component of src0. The result of the operation is accurate to at least 21 bits. This is the vector version of LOG with zeroop set to il_zeroop_infinity.		
<i>Format</i>	1-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_LOG_VEC
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	LN, LOG, LOGP.		

Base-2 Logarithm (partial precision)

<i>Instructions</i>	LOGP		
<i>Syntax</i>	<code>logp_zeroop(op) dst, src0</code>		
<i>Description</i>	<p>Computes the base-2 logarithm of <code>src0.w</code> using partial precision. The result of this computation is accurate to at least 10-bits. The 32-bit floating point result is placed in <code>dst.z</code>. By default, this instruction operates on <code>src0.w</code>, but can operate on any component by swizzling it into the fourth component.</p> <p>Operation:</p> <pre> VECTOR v1 = EvalSource(src0); VECTOR v; float f; if (v1[3] == 0) { if(zeroop == IL_ZEROOP_FLT_MAX) f = -FLT_MAX; else if(zeroop == IL_ZEROOP_INFINITY) f = -INFINITY; else if(zeroop == IL_ZEROOP_INF_ELSE_MAX) f = -INFINITY or -FLT_MAX; # Depends on IL Implementation } else if (v1[3] < 0.0) { f = undefined; } else { f = log2(v1[3]); } v[0] = floor(f); v[1] = v1[3] / exp2(floor(f)); v[2] = f; v[3] = 1.0; WriteResult(v, dst); </pre>		
<i>Format</i>	1-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	<code>code</code>	15:0	IL_OP_LOGP
	<code>control</code>	29:16	Any value of the enumerated type <code>ILZeroOp(zeroop)</code> except <code>IL_ZEROOP_0</code> . See Table 5.16 on page 5-9.
	<code>sec_modifier_present</code>	30	Must be zero.
	<code>pri_modifier_present</code>	31	Must be zero.
<i>Related</i>	LN, LOG, LOG_VEC.		

Linear Interpolation

<i>Instructions</i>	LRP															
<i>Syntax</i>	<code>lrp dst, src0, src1, src2</code>															
<i>Description</i>	<p>Computes the linear interpolation between two vectors.</p> <p>Operation:</p> <pre> VECTOR v1 = EvalSource(src0); VECTOR v2 = EvalSource(src1); VECTOR v3 = EvalSource(src2); VECTOR v; for (i=0; i < 4; i++) v[i] = v2[i] * v1[i] + v3[i] * (1 - v1[i]); WriteResult(v, dst); </pre>															
<i>Format</i>	3-input, 1-output.															
<i>Opcode</i>	<table> <thead> <tr> <th>Field Name</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>code</td> <td>15:0</td> <td>IL_OP_LRP</td> </tr> <tr> <td>control</td> <td>29:16</td> <td>Must be zero.</td> </tr> <tr> <td>sec_modifier_present</td> <td>30</td> <td>Must be zero.</td> </tr> <tr> <td>pri_modifier_present</td> <td>31</td> <td>Must be zero.</td> </tr> </tbody> </table>	Field Name	Bits	Description	code	15:0	IL_OP_LRP	control	29:16	Must be zero.	sec_modifier_present	30	Must be zero.	pri_modifier_present	31	Must be zero.
Field Name	Bits	Description														
code	15:0	IL_OP_LRP														
control	29:16	Must be zero.														
sec_modifier_present	30	Must be zero.														
pri_modifier_present	31	Must be zero.														
<i>Related</i>	None.															

Less Than

<i>Instructions</i>	LT															
<i>Syntax</i>	<code>lt dst, src0, src1</code>															
<i>Description</i>	<p>Compares two float vectors, <i>src0</i> and <i>src1</i>, for each component, and writes the result to <i>dst</i>. If the comparison is true, 0xFFFFFFFF is returned for that component; otherwise, 0x00000000 is returned. This instruction follows DX10 Floating Point Rules. Denorms are flushed before comparison (original source registers untouched). +0 equals -0. Comparison with NaN returns false.</p> <p>Primary and secondary opcode modifiers are not permitted, and no additional control options are supported. Thus, the <code>pri_modifier_present</code> and <code>sec_modifier_present</code> fields must be zero.</p>															
<i>Format</i>	2-input, 1-output.															
<i>Opcode</i>	<table> <thead> <tr> <th>Field Name</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>code</td> <td>15:0</td> <td>IL_OP_EQ</td> </tr> <tr> <td>control</td> <td>29:16</td> <td>Must be zero.</td> </tr> <tr> <td>sec_modifier_present</td> <td>30</td> <td>Must be zero.</td> </tr> <tr> <td>pri_modifier_present</td> <td>31</td> <td>Must be zero.</td> </tr> </tbody> </table>	Field Name	Bits	Description	code	15:0	IL_OP_EQ	control	29:16	Must be zero.	sec_modifier_present	30	Must be zero.	pri_modifier_present	31	Must be zero.
Field Name	Bits	Description														
code	15:0	IL_OP_EQ														
control	29:16	Must be zero.														
sec_modifier_present	30	Must be zero.														
pri_modifier_present	31	Must be zero.														
<i>Related</i>	Double-precision Compare, Integer Compare, Unsigned Integer Compare.															

Floating Point Multiply and Add

<i>Instructions</i>	MAD		
<i>Syntax</i>	<code>mad[_ieee] dst, src0, src1, src2</code>		
<i>Description</i>	<p>Multiplies a component in <i>src0</i> by the corresponding component in <i>src1</i>. The lower 32 bits of the multiplication is added to the corresponding component in <i>src2</i>. The result of the multiply-add operation is placed in the corresponding component of <i>dst</i>.</p> <p>Operation:</p> <pre> VECTOR v1 = EvalSource(src0); VECTOR v2 = EvalSource(src1); VECTOR v3 = EvalSource(src2); VECTOR v; for (i=0; i < 4; i++) v[i] = v1[i] * v2[i] + v3[i]; WriteResult(v, dst); </pre>		
<i>Format</i>	3-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	<code>code</code>	15:0	IL_OP_MAD
	<code>control</code>	29:16	MAD: 0 0*n = o, even if n is NaN. 1 IEEE-style multiply.
	<code>sec_modifier_present</code>	30	Must be zero.
	<code>pri_modifier_present</code>	31	Must be zero.
<i>Related</i>	DMAD.		

Floating Point Maximum

<i>Instructions</i>	MAX		
<i>Syntax</i>	<pre>max[_ieee] dst, src0, src1 MAX IL_OP_MAX max[_ieee] dst, floating point maximum, src0, src1 src0.{x y z w} ≥ src1.{x y z w} ? src0.{x y z w} : src1.{x y z w}</pre>		
<i>Description</i>	<p>Computes maximum value per component. Both max(+0, -0) and max(-0, +0) return +0. NaN has special handling: If one source operand is NaN, then the other source operand is returned (choice made per-component). If both are NaN, any NaN representation is returned.</p> <p>Denorms are flushed to sign preserved 0s before comparison; if it is the maximum, the flushed denorm is written to <i>dst</i>.</p> <p>If the <i>_ieee</i> flag is included then MIN and MAX follow IEEE-754r rules for minimum and maximum except denorms are flushed before the comparison is made. In addition, MIN returns -0 and MAX returns +0 for comparisons between -0 and +0.</p> <p>Operation:</p> <pre>VECTOR v1 = EvalSource(src0); VECTOR v2 = EvalSource(src1); VECTOR v; for (i=0; i < 4; i++) v[i] = (v1[i] > v2[i]) ? v1[i] : v2[i]; WriteResult(v, dst);</pre>		
<i>Format</i>	2-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_MAX
	control	29:16	MAX: 0 No special NaN rules. 1 IEEE-style NaN rules.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	None.		

Floating Point Minimum

<i>Instructions</i>	MIN		
<i>Syntax</i>	min[_ieee] <i>dst</i> , <i>src0</i> , <i>src1</i>		
<i>Description</i>	<p>Computes minimum value per component. Both min(+0, -0) and min(-0, +0) return -0. NaN has special handling: If one source operand is NaN, then the other source operand is returned (choice made per-component). If both are NaN, any NaN representation is returned. Denorms are flushed to sign preserved 0s before comparison; if it is the minimum, the flushed denorm is written to <i>dst</i>.</p> <p>If the <i>_ieee</i> flag is included then MIN and MAX follow IEEE-754r rules for minimum and maximum except denorms are flushed before the comparison is made. In addition, MIN returns -0 and MAX returns +0 for comparisons between -0 and +0.</p> <p>Operation:</p> <pre> VECTOR v1 = EvalSource(<i>src0</i>); VECTOR v2 = EvalSource(<i>src1</i>); VECTOR v; for (i=0; i < 4; i++) v[i] = (v1[i] < v2[i]) ? v1[i] : v2[i]; WriteResult(v, <i>dst</i>); </pre>		
<i>Format</i>	2-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_MIN
	control	29:16	MIN: 0 No special NaN rules. 1 IEEE-style NaN rules.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	None.		

Matrix Multiply by Vector

Instructions	MMUL		
Syntax	<code>mmul_matrix(<i>op</i>) <i>dst</i>, <i>src0</i>, <i>src1</i></code>		
Description	<p>Performs a matrix multiply of the generic matrices <i>src0</i> and <i>src1</i>. Data must be arranged in one of the formats indicated by the enumerated type <code>ILMatrix(<i>matrix</i>)</code>. <i>src1</i> must be of type <code>IL_REGTYPE_CONST_FLOAT</code> or <code>IL_REGTYPE_TEMP</code>. Relative addressing and source modifiers of <i>src1</i> are applied to each row of the matrix beginning with <i>src1</i>.</p> <p>Certain components are not returned by this instruction, depending on the value of <i>matrix</i>. Each component is returned only if the component mask field (<i>component_z_b</i> or <i>component_w_a</i>) of the <code>IL_Dst_Mod</code> token for that component is set to <code>IL_MODCOMP_0</code> or <code>IL_MODCOMP_1</code>, in which case 0.0 or 1.0 is returned.</p> <p>Operation:</p> <pre> v[2] = v1[0]*v4[0] + v1[1]*v4[1] + v1[2]*v4[2] + v1[3]*v4[3]; } else if(matrix == IL_MATRIX_3X4) { VECTOR v1 = EvalSource(<i>src1</i>); VECTOR v2 = EvalSource(<i>src2</i>); VECTOR v3 = EvalSource(<i>src2</i>+1); # Register number of <i>src2</i> + 1 VECTOR v4 = EvalSource(<i>src2</i>+2); # Register number of <i>src2</i> + 2 VECTOR v5 = EvalSource(<i>src2</i>+3); # Register number of <i>src2</i> + 3 v[0] = v1[0]*v2[0] + v1[1]*v2[1] + v1[2]*v2[2]; v[1] = v1[0]*v3[0] + v1[1]*v3[1] + v1[2]*v3[2]; v[2] = v1[0]*v4[0] + v1[1]*v4[1] + v1[2]*v4[2]; v[3] = v1[0]*v5[0] + v1[1]*v5[1] + v1[2]*v5[2]; } else if(matrix == IL_MATRIX_3X3) { VECTOR v1 = EvalSource(<i>src1</i>); VECTOR v2 = EvalSource(<i>src2</i>); VECTOR v3 = EvalSource(<i>src2</i>+1); # Register number of <i>src2</i> + 1 VECTOR v4 = EvalSource(<i>src2</i>+2); # Register number of <i>src2</i> + 2 v[0] = v1[0]*v2[0] + v1[1]*v2[1] + v1[2]*v2[2]; v[1] = v1[0]*v3[0] + v1[1]*v3[1] + v1[2]*v3[2]; v[2] = v1[0]*v4[0] + v1[1]*v4[1] + v1[2]*v4[2]; } else if(matrix == IL_MATRIX_3X2) { VECTOR v1 = EvalSource(<i>src1</i>); VECTOR v2 = EvalSource(<i>src2</i>); VECTOR v3 = EvalSource(<i>src2</i>+1); # Register number of <i>src2</i> + 1 v[0] = v1[0]*v2[0] + v1[1]*v2[1] + v1[2]*v2[2]; v[1] = v1[0]*v3[0] + v1[1]*v3[1] + v1[2]*v3[2]; } WriteResult(v, <i>dst</i>); </pre>		
Format	2-input, 1-output.		
Opcode	Field Name	Bits	Description
	<code>code</code>	15:0	<code>IL_OP_MMUL</code>
	<code>control</code>	29:16	Any value of the enumerated type <code>ILMatrix(<i>matrix</i>)</code> . See Table 5.6 on page 5-3.
	<code>sec_modifier_present</code>	30	Must be zero.
	<code>pri_modifier_present</code>	31	Must be zero.
Related	None.		

Floating Point Modulo

<i>Instructions</i>	MOD		
<i>Syntax</i>	mod <i>dst</i> , <i>src0</i> , <i>src1</i>		
<i>Description</i>	<p>Performs the modulo operation. The sign remains intact (for example, mod(-1.7,1.0)=-0.7).</p> <p>Operation:</p> <pre> VECTOR v1 = EvalSource(src0); VECTOR v2 = EvalSource(src1); VECTOR v; for (i=0; i < 4; i++) { if(v2[i] == 0.0) v[i] = v1[i]; else v[i] = v1[i] - v2[i] * trunc(v1[i]/v2[i]); } WriteResult(v, dst); </pre>		
<i>Format</i>	2-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_MOD
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	None.		

Move Data

<i>Instructions</i>	INV_MOV		
<i>Syntax</i>	invariant_move <i>dst</i> , <i>src0</i>		
<i>Summary</i>	Move data between registers.		
<i>Description</i>	<p>Moves value from <i>src0</i> to <i>dst</i>. As a result of using this instruction, the compiler ensures that any other shader that computes this source using the same instructions gets the identical answer. Generally, use of INV_MOVE prevents many compiler optimizations and lowers performance.</p> <p>Operation:</p> <pre> VECTOR v = EvalSource(src0); WriteResult(v, dst); </pre>		
<i>Format</i>	1-input 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_INVARIANT_MOV
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	None.		

Move Data Between Registers

Instructions **MOV***Syntax* `mov dst, src0`*Description* Places all four components of *src0* into the corresponding components of *dst*.

Operation:

VECTOR *v* = EvalSource(*src0*);
WriteResult(*v*, *dst*);*Format* 1-input, 1-output.

<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_MOV
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.

Related INV_MOV, MOVA.

Set Address Register Value

<i>Instructions</i>	MOVA			
<i>Syntax</i>	mov _a [_{round}] <i>dst</i> , <i>src0</i>			
<i>Description</i>	<p>Sets the values of an ADDR register to the truncated or rounded value of <i>src0</i>. Relative addressing is not possible on the destination register. The <i>relative_address</i> field of the IL_Dst token must be 0.</p> <p>See Base Relative Addressing for more information, in Section 3.3, "Relative Addressing," page 3-2, and Table 2.9 on page 2-8.</p> <p>If the rounding mode (bit 16) is 1, the values of <i>src0</i> are rounded to the nearest integer before they are loaded into the address register. If the rounding mode (bit 16) is 0, the values of <i>src0</i> are floored before they are loaded into the address register.</p> <p>Operation:</p> <pre> VECTOR v1 = EvalSource(src0); VECTOR v; for(i = 0; i < 4; i++) { if(round == 1) { v[i] = round(v1[i]); } else { v[i] = floor(v1[i]); } } WriteResult(v, dst); </pre>			
<i>Format</i>	1-input, 1-output.			
<i>Opcode</i>	Field Name	Bits	Description	
	code	15:0	IL_OP_MOVA	
	control	29:16	Field Name	Bits
			rounding_mode	16
				0: Truncate <i>src0</i> before moving to <i>dst</i> .
				1: Round <i>src0</i> before moving to <i>dst</i> .
			<i>reserved</i>	29:25 Must be zero.
	sec_modifier_present	30	Must be zero.	
	pri_modifier_present	31	Must be zero.	
<i>Related</i>	INV_MOV, MOV.			

Floating Point Multiplication

<i>Instructions</i>	MUL		
<i>Syntax</i>	mul[_ieee] <i>dst</i> , <i>src0</i> , <i>src1</i>		
<i>Description</i>	<p>Multiplies two vectors.</p> <p>Operation:</p> <pre>VECTOR v1 = EvalSource(<i>src0</i>); VECTOR v2 = EvalSource(<i>src1</i>); VECTOR v; for (i=0; i < 4; i++) v[i] = v1[i] * v2[i]; WriteResult(v, <i>dst</i>);</pre>		
<i>Format</i>	2-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_MUL
	control	29:16	MUL: 0 DX9 DP2-style multiple. 1 IEEE-style multiply.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	DMUL, IMUL_HIGH, UMUL_HIGH.		

Not Equal

<i>Instructions</i>	NE		
<i>Syntax</i>	nt <i>dst</i> , <i>src0</i> , <i>src1</i>		
<i>Description</i>	<p>Compares two float vectors, <i>src0</i> and <i>src1</i>, for each component, and writes the result to <i>dst</i>. If the comparison is true, 0xFFFFFFFF is returned for that component; otherwise, 0x00000000 is returned. This instruction follows DX10 Floating Point Rules. Denorms are flushed before comparison (original source registers untouched). +0 equals -0. Comparison with NaN returns false.</p> <p>Primary and secondary opcode modifiers are not permitted, and no additional control options are supported. Thus, the <code>pri_modifier_present</code> and <code>sec_modifier_present</code> fields must be zero.</p>		
<i>Format</i>	2-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_NE
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	Double-precision Compare, Integer Compare, Unsigned Integer Compare.		

Compute Noise Value

<i>Instructions</i>	NOISE																	
<i>Syntax</i>	noise_type(<i>op</i>) <i>dst</i> , <i>src0</i>																	
<i>Description</i>	<p>Uses <i>src0</i> as the seed of a pseudo-random number generator to compute a noise value and places it in <i>dst</i>.</p> <p>Operation:</p> <p>The values in <i>dst</i> has the following characteristics:</p> <ul style="list-style-type: none"> • They are in the range [-1.0, 1.0] • Over many iterations, their average value is 0.0. • The value is repeatable. That is, <i>dst</i> is always the same value given the same value of <i>src0</i>. • The value is statistically invariant under rotation (no matter how the domain is rotated, it has the same statistical character). • The value is statistically invariant under translation (no matter how the domain is translated, it has the same statistical character). • The value is typically different under translation. • Over many iterations, the values have a narrow band pass limit in frequency (the values have no visible features larger or smaller than a certain narrow-size range). • Over many iterations, the values are C₁ continuous everywhere (the first derivative is continuous). 																	
<i>Format</i>	1-input, 1-output.																	
<i>Opcode</i>	Field Name	Bits	Description															
	code	15:0	IL_OP_NOISE															
	control	29:16	<table> <thead> <tr> <th>Field Name</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>noise</td> <td>19:16</td> <td>Any value of the enumerated type ILNoiseType(type). See Table 5.24 on page 5-12.</td> </tr> <tr> <td>reserved</td> <td>29:20</td> <td>Must be zero.</td> </tr> <tr> <td>sec_modifier_present</td> <td>30</td> <td>Must be zero.</td> </tr> <tr> <td>pri_modifier_present</td> <td>31</td> <td>Must be zero.</td> </tr> </tbody> </table>	Field Name	Bits	Description	noise	19:16	Any value of the enumerated type ILNoiseType(type). See Table 5.24 on page 5-12.	reserved	29:20	Must be zero.	sec_modifier_present	30	Must be zero.	pri_modifier_present	31	Must be zero.
Field Name	Bits	Description																
noise	19:16	Any value of the enumerated type ILNoiseType(type). See Table 5.24 on page 5-12.																
reserved	29:20	Must be zero.																
sec_modifier_present	30	Must be zero.																
pri_modifier_present	31	Must be zero.																
<i>Related</i>	None.																	

Three or Four Component Normalization

Instructions **NRM**

Syntax `nrm_zeroop(op) dst, src0`

Description Normalize a 4D vector using three or four components. Each component of the vector in *src0* is divided by the square root of the sum of squares of the components in *src0*. This operation has the result of limiting each component to the range [-1.0, 1.0]. The result is placed in *dst*. If *nrm4* is one, all four components of *src0* are used in normalization. If *nrm4* is zero, only x, y, and z are used in the normalization calculation.

Operation:

```

VECTOR v1 = EvalSource(src0);
VECTOR v;
Float f;

If( nrm4 == 0)
{
    f = v1[0]*v1[0] + v1[1]*v1[1] + v1[2]*v1[2];
}
else
{
    f = v1[0]*v1[0] + v1[1]*v1[1] + v1[2]*v1[2] + v1[3]*v1[3];
}
if (f == 0.0)
{
    if(zeroop == IL_ZEROOP_0)
        f = 0.0;
    else if(zeroop == IL_ZEROOP_FLT_MAX)
        f = FLT_MAX;
    else if(zeroop == IL_ZEROOP_INFINITY)
        f = INFINITY;
    else if(zeroop == IL_ZEROOP_INF_ELSE_MAX)
        f = INFINITY or FLT_MAX; # Depends on IL Implementation
}
else
{
    f = 1.0/sqrt(f);
}
for (i=0; i < 4; i++)
    v[i] = v1[i] * f;
WriteResult(v, dst);
    
```

Format 1-input, 1-output.

Opcode	Field Name	Bits	Description																		
	code	15:0	IL_OP_NRM																		
	control	29:19	<table border="0" style="width: 100%;"> <thead> <tr> <th style="text-align: left;">Field Name</th> <th style="text-align: left;">Bits</th> <th style="text-align: left;">Description</th> </tr> </thead> <tbody> <tr> <td><i>reserved</i></td> <td>17:16</td> <td>Any value of the enumerated type ILZeroOp(<i>zeroop</i>). See Table 5.16 on page 5-9 when the first three components of <i>src0</i> are 0.0.</td> </tr> <tr> <td><i>nrm4</i></td> <td>18</td> <td>0: Normalize <i>src0</i>.xyz 1: Normalize <i>src0</i>.xyzw</td> </tr> <tr> <td><i>reserved</i></td> <td>29:19</td> <td>Must be zero.</td> </tr> <tr> <td></td> <td>18</td> <td>Any value of the enumerated type ILZeroOp(<i>zeroop</i>). See Table 5.16 on page 5-9.</td> </tr> <tr> <td></td> <td>29:19</td> <td>Must be zero.</td> </tr> </tbody> </table>	Field Name	Bits	Description	<i>reserved</i>	17:16	Any value of the enumerated type ILZeroOp(<i>zeroop</i>). See Table 5.16 on page 5-9 when the first three components of <i>src0</i> are 0.0.	<i>nrm4</i>	18	0: Normalize <i>src0</i> .xyz 1: Normalize <i>src0</i> .xyzw	<i>reserved</i>	29:19	Must be zero.		18	Any value of the enumerated type ILZeroOp(<i>zeroop</i>). See Table 5.16 on page 5-9.		29:19	Must be zero.
Field Name	Bits	Description																			
<i>reserved</i>	17:16	Any value of the enumerated type ILZeroOp(<i>zeroop</i>). See Table 5.16 on page 5-9 when the first three components of <i>src0</i> are 0.0.																			
<i>nrm4</i>	18	0: Normalize <i>src0</i> .xyz 1: Normalize <i>src0</i> .xyzw																			
<i>reserved</i>	29:19	Must be zero.																			
	18	Any value of the enumerated type ILZeroOp(<i>zeroop</i>). See Table 5.16 on page 5-9.																			
	29:19	Must be zero.																			
	<i>sec_modifier_present</i>	30	Must be zero.																		

Three or Four Component Normalization (Cont.)

	pri_modifier_present	31	Must be zero.
--	----------------------	----	---------------

Related None.

Reduce Vector to $[-\pi, \pi]$

<i>Instructions</i>	PIREDUCE		
<i>Syntax</i>	pireduce <i>dst</i> , <i>src0</i>		
<i>Description</i>	All four components of the vector in <i>src0</i> is reduced to the range $[-\pi, \pi]$. Operation: <pre>VECTOR v1 = EvalSource(src0); VECTOR v; for (i=0; i < 4; i++) v[i] = (frac((v1[i]/(2*Pi))+0.5)* 2 * PI) - PI; WriteResult(v, dst);</pre>		
<i>Format</i>	1-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_PIREDUCE
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	NRM.		

X to the Power of Y

<i>Instructions</i>	POW															
<i>Syntax</i>	<code>pow dst, src0, src1</code>															
<i>Description</i>	<p>Computes $src0.w$ to the power of $src1.w$ ($src0.w^{src1.w}$). By default, this instruction operates on $src0.w$ and $src1.w$, but can operate on any component of either operand by swizzling it into the fourth component.</p> <p>Operation:</p> <pre> VECTOR v1 = EvalSource(src0); VECTOR v2 = EvalSource(src1); VECTOR v; If(v1[3] < 0.0) { v[0] = v[1] = v[2] = v[3] = undefined; } else { v[0] = v[1] = v[2] = v[3] = exp₂(v2[3] * log₂(v1[3])); } WriteResult(v, dst); </pre>															
<i>Format</i>	2-input, 1-output.															
<i>Opcode</i>	<table border="0"> <thead> <tr> <th style="text-align: left;">Field Name</th> <th style="text-align: left;">Bits</th> <th style="text-align: left;">Description</th> </tr> </thead> <tbody> <tr> <td>code</td> <td>15:0</td> <td>IL_OP_POW</td> </tr> <tr> <td>control</td> <td>29:16</td> <td>Must be zero.</td> </tr> <tr> <td>sec_modifier_present</td> <td>30</td> <td>Must be zero.</td> </tr> <tr> <td>pri_modifier_present</td> <td>31</td> <td>Must be zero.</td> </tr> </tbody> </table>	Field Name	Bits	Description	code	15:0	IL_OP_POW	control	29:16	Must be zero.	sec_modifier_present	30	Must be zero.	pri_modifier_present	31	Must be zero.
Field Name	Bits	Description														
code	15:0	IL_OP_POW														
control	29:16	Must be zero.														
sec_modifier_present	30	Must be zero.														
pri_modifier_present	31	Must be zero.														
<i>Related</i>	EXN, EXP, EXP_VEC, EXPP.															

Reciprocal

<i>Instructions</i>	RCP		
<i>Syntax</i>	<code>rcp_zeroop(op) dst, src0</code>		
<i>Description</i>	<p>Computes the reciprocal of <code>src0.w</code>. By default, this instruction operates on <code>src0.w</code> and <code>src1.w</code>, but can operate on any component of either operand by swizzling it into the fourth component.</p> <p>Operation:</p> <pre> VECTOR v1 = EvalSource(src0); VECTOR v; float f = v1[3]; if (f == 0.0) { if(zeroop == IL_ZEROOP_0) f = 0.0; else if(zeroop == IL_ZEROOP_FLT_MAX) f = FLT_MAX; else if(zeroop == IL_ZEROOP_INFINITY) f = INFINITY; else if(zeroop == IL_ZEROOP_INF_ELSE_MAX) f = INFINITY or FLT_MAX; # Depends on IL Implementation } else if (f != 1.0) f = 1/f; v[0] = v[1] = v[2] = v[3] = f; WriteResult(v, dst); </pre>		
<i>Format</i>	1-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_RCP
	control	17:16	The value of the enumerated type <code>ILZeroOp(zeroop)</code> , which controls how this instruction behaves when the value of <code>src0</code> is 0.0. See Table 5.16 on page 5-9.
		29:18	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	RSQ, RSQ_VEC.		

Compute Reflection Vector

<i>Instructions</i>	REFLECT		
<i>Syntax</i>	reflect_normalize <i>dst</i> , <i>src0</i> , <i>src1</i>		
<i>Description</i>	<p>Computes the reflection direction of the vector in <i>src0</i> using the source vector in <i>src1</i>, and placing the direction vector in <i>dst</i>. <i>_normalize</i> specifies if <i>src1</i> is normalized before computing the reflection vector.</p> <p>Operation:</p> <pre> VECTOR v1 = EvalSource(<i>src0</i>); VECTOR v2 = EvalSource(<i>src1</i>); VECTOR v; float f = 2 * (v1[0]*v2[0] + v1[1]*v2[1] + v1[2]*v2[2] + v1[3]*v2[3]); if(normalize) { float fnorm = (v2[0]*v2[0] + v2[1]*v2[1] + v2[2]*v2[2] + v2[3]*v2[3]); for (i=0; i < 4; i++) v[i] = fnorm * v1[i] - f * v2[i]; } else { for (i=0; i < 4; i++) v[i] = v1[i] - f * v2[i]; } WriteResult(v, <i>dst</i>); </pre>		
<i>Format</i>	2-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_REFLECT
	control	16	0: Do not normalize <i>src1</i> . 1: Normalize <i>src1</i> before computing the reflection vector.
		29:17	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	None.		

Round

<i>Instructions</i>	RND															
<i>Syntax</i>	<code>rnd dst, src0</code>															
<i>Description</i>	<p>Rounds the float value in each component of a floating point vertex (<i>src0</i>) to the nearest integer.</p> <p>Operation:</p> <pre>VECTOR v1 = EvalSource(src0); VECTOR v; for (i=0; i < 4; i++) v[i] = floor(v1[i] + 0.5); WriteResult(v, dst);</pre>															
<i>Format</i>	1-input, 1-output.															
<i>Opcode</i>	<table> <thead> <tr> <th>Field Name</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>code</td> <td>15:0</td> <td>IL_OP_RND</td> </tr> <tr> <td>control</td> <td>29:16</td> <td>Must be zero.</td> </tr> <tr> <td>sec_modifier_present</td> <td>30</td> <td>Must be zero.</td> </tr> <tr> <td>pri_modifier_present</td> <td>31</td> <td>Must be zero.</td> </tr> </tbody> </table>	Field Name	Bits	Description	code	15:0	IL_OP_RND	control	29:16	Must be zero.	sec_modifier_present	30	Must be zero.	pri_modifier_present	31	Must be zero.
Field Name	Bits	Description														
code	15:0	IL_OP_RND														
control	29:16	Must be zero.														
sec_modifier_present	30	Must be zero.														
pri_modifier_present	31	Must be zero.														
<i>Related</i>	ROUND_NEAREST, ROUND_NEG_INF, ROUND_PLUS_INF, ROUND_Z.															

Float Round to Nearest Even Float Integral

<i>Instructions</i>	ROUND_NEAREST															
<i>Syntax</i>	<code>round_nearest dst, src0</code>															
<i>Description</i>	Rounds the float value in each component of <i>src0</i> to the nearest even integral floating-point.															
<i>Format</i>	1-input, 1-output.															
<i>Opcode</i>	<table> <thead> <tr> <th>Field Name</th> <th>Bits</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>code</td> <td>15:0</td> <td>IL_OP_ROUND_NEAR</td> </tr> <tr> <td>control</td> <td>29:16</td> <td>Must be zero.</td> </tr> <tr> <td>sec_modifier_present</td> <td>30</td> <td>Must be zero.</td> </tr> <tr> <td>pri_modifier_present</td> <td>31</td> <td>Must be zero.</td> </tr> </tbody> </table>	Field Name	Bits	Description	code	15:0	IL_OP_ROUND_NEAR	control	29:16	Must be zero.	sec_modifier_present	30	Must be zero.	pri_modifier_present	31	Must be zero.
Field Name	Bits	Description														
code	15:0	IL_OP_ROUND_NEAR														
control	29:16	Must be zero.														
sec_modifier_present	30	Must be zero.														
pri_modifier_present	31	Must be zero.														
<i>Related</i>	RND, ROUND_NEG_INF, ROUND_PLUS_INF, ROUND_Z.															

Float Round to $-\infty$

<i>Instructions</i>	ROUND_NEG_INF		
<i>Syntax</i>	round_neginf <i>dst</i> , <i>src0</i>		
<i>Description</i>	Rounds the float values in each component of <i>src0</i> towards $-\infty$. This is sometimes called a floor() instruction.		
<i>Format</i>	1-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_ROUND_NEG_INF
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	FLR, RND, ROUND_NEAREST, ROUND_PLUS_INF, ROUND_Z.		

Float Round to $+\infty$

<i>Instructions</i>	ROUND_POS_INF		
<i>Syntax</i>	round_posinf <i>dst</i> , <i>src0</i>		
<i>Description</i>	Rounds the float values in each component of <i>src0</i> towards ∞ . This is sometimes called a ceil() instruction.		
<i>Format</i>	1-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_ROUND_POS_INF
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	CLG, RND, ROUND_NEAREST, ROUND_NEG_INF, ROUND_Z.		

Float Round to Zero

Instructions **ROUND_ZERO***Syntax* `round_z dst, src0`*Description* Rounds the float values in each component of *src0* towards zero.*Format* 1-input, 1-output.

<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_ROUND_ZERO
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.

Related RND, ROUND_NEAREST, ROUND_NEG_INF, ROUND_PLUS_INF.

Component-wise Reciprocal Square Root

Instructions **RSQ_VEC***Syntax* `rsq_vec dst, src0`*Description* Computes the reciprocal of the square root of each component of *src0*.*Format* 1-input, 1-output.

<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_RSQ_VEC
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.

Related RCP, RSQ.

Reciprocal Square Root

Instructions **RSQ**

Syntax `rsq_zeroop(op) dst, src0`

Description Computes the reciprocal of the square root (positive only) of `src0.w`. By default, this instruction operates on `src0.w`, but can operate on any component by swizzling it into the fourth component of `src0`.

Operation:

```
VECTOR v1 = EvalSource(src0);
VECTOR v;
float f = v1[3];
if (f == 0.0)
{
    if(zeroop == IL_ZEROOP_0)
        f = 0.0;
    else if(zeroop == IL_ZEROOP_FLT_MAX)
        f = FLT_MAX;
    else if(zeroop == IL_ZEROOP_INFINITY)
        f = INFINITY;
    else if(zeroop == IL_ZEROOP_INF_ELSE_MAX)
        f = INFINITY or FLT_MAX; # Depends on IL Implementation
}
else if (f < 0.0)
{
    f = undefined;
}
else
{
    f = 1.0/(float)sqrt(f);
}
v[0] = v[1] = v[2] = v[3] = f;
WriteResult(v, dst);
```

Format 1-input, 1-output.

Opcode	Field Name	Bits	Description
	code	15:0	IL_OP_RSQ
	control	29:16	Any value of the enumerated type ILZeroOp(zeroop). See Table 5.16 on page 5-9.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.

Related RCP, RSQ_VEC.

Set on Comparison

<i>Instructions</i>	SET		
<i>Syntax</i>	<code>set_relop(<i>op</i>) <i>dst</i>, <i>src0</i>, <i>src1</i></code>		
<i>Description</i>	<p>Compares each component of <i>src0</i> with the corresponding component of <i>src1</i>. The type of comparison performed is dictated by <i>relop(op)</i>. If the comparison <i>src0</i>.{<i>x</i> <i>y</i> <i>z</i> <i>w</i>} <i>relop(op)</i> <i>src1</i>.{<i>x</i> <i>y</i> <i>z</i> <i>w</i>} evaluates TRUE, the result is 1.0; otherwise, the result is 0.0.</p> <p>Operation:</p> <pre> VECTOR v1 = EvalSource(<i>src0</i>); VECTOR v2 = EvalSource(<i>src1</i>); VECTOR v; for (i=0; i < 4; i++) v[i] = (v1[i] <i>relop</i> v2[i]) ? 1.0 : 0.0; WriteResult(v, <i>dst</i>); </pre>		
<i>Format</i>	2-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_SET
	control	29:16	Any value of the enumerated type ILRelOp(<i>relop</i>). See Table 5.14 on page 5-8.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	None.		

Compute Sign

<i>Instructions</i>	SGN		
<i>Syntax</i>	<code>sgn dst, src0</code>		
<i>Description</i>	<p>Computes the sign of each component of <i>src0</i>.</p> <p>Operation:</p> <pre> VECTOR v = EvalSource(src0); for (i=0; i < 4; i++) { if (v[i] < 0.0) v[i] = -1.0; else if (v[i] == 0.0) v[i] = 0.0; else v[i] = 1.0; } WriteResult(v, dst); </pre>		
<i>Format</i>	1-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_SGN
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	None.		

Sine (sin)

<i>Instructions</i>	SIN		
<i>Syntax</i>	<code>sin dst, src0</code>		
<i>Description</i>	<p>Computes the sine of <i>src0.w</i>. <i>src0.w</i> is in radians for trigonometric functions. <i>src0.w</i> must be within the range $[-\pi, \pi]$ for each function; otherwise, the results are undefined. By default, this instruction operates on <i>src0.w</i>, but can operate on any component by swizzling it into the fourth component. The maximum absolute error is 0.002.</p>		
<i>Format</i>	1-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_SIN
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	COS_VEC, SIN_VEC, SINCOS.		

Compute Sine and Cosine

<i>Instructions</i>	SINCOS		
<i>Syntax</i>	<code>sincos dst, src0</code>		
<i>Description</i>	<p>Computes sine and cosine values of <code>src0.w</code>. By default, this instruction operates on <code>src0.w</code>, but can operate on any component by swizzling it into the fourth component. The maximum absolute error is 0.002. The 32-bit floating point results are returned in <code>dst.x</code> (COS) and <code>dst.y</code> (SIN). <code>src0.w</code> must be in the range $[-\pi, \pi]$; otherwise, the results are undefined. <code>dst.z</code> and <code>dst.w</code> are not written by this instruction; however, the if component <code>_z_b</code> or component <code>_w_a</code> field of the <code>IL_Dst_Mod</code> token is set to <code>IL_MODCOMP_0</code> or <code>IL_MODCOMP_1</code>, the <code>dst.z</code> and <code>dst.w</code> are written. That is, <code>dst.z</code> and <code>dst.w</code> can be set to 0.0 or 1.0 if <code>IL_MODCOMP_0</code> or <code>IL_MODCOMP_1</code> is used on the component <code>_z_b</code> or component <code>_w_a</code> field of the <code>IL_Dst_Mod</code> token.</p> <p>Operation:</p> <pre>VECTOR v1 = EvalSource(src0); VECTOR v; v[0] = cos(v1[3]); v[1] = sin(v1[3]); WriteResult(v, dst);</pre>		
<i>Format</i>	1-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	<code>code</code>	15:0	IL_OP_SINCOS
	<code>control</code>	29:16	Must be zero.
	<code>sec_modifier_present</code>	30	Must be zero.
	<code>pri_modifier_present</code>	31	Must be zero.
<i>Related</i>	COS, SIN.		

Component-Wise Sine

<i>Instructions</i>	SIN_VEC		
<i>Syntax</i>	<code>sin_vec dst, src0</code>		
<i>Description</i>	<p>Computes the sine of each component of <code>src0</code> in radians. The maximum absolute error is 0.0008 in the range $[-100*\pi, 100*\pi]$.</p>		
<i>Format</i>	1-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	<code>code</code>	15:0	IL_OP_SIN_VEC
	<code>control</code>	29:16	Must be zero.
	<code>sec_modifier_present</code>	30	Must be zero.
	<code>pri_modifier_present</code>	31	Must be zero.
<i>Related</i>	COS, SIN.		

Component-Wise Cosine

<i>Instructions</i>	COS_VEC			
<i>Syntax</i>	Function	Opcode	Syntax	Description
	COS_VEC	IL_OP_COS_VEC	<i>cos_vec dst, src0</i>	cosine, $\cos(\text{src0.xyzw})$
<i>Description</i>	Computes the cosine of each component of <i>src0</i> in radians. The maximum absolute error is 0.0008 in the range $[-100*\pi, 100*\pi]$.			
<i>Format</i>	1-input, 1-output.			
<i>Opcode</i>	Field Name	Bits	Description	
	code	15:0	IL_OP_COS_VEC	
	control	29:16	Must be zero.	
	sec_modifier_present	30	Must be zero.	
	pri_modifier_present	31	Must be zero.	
<i>Related</i>	COS, SIN.			

Square Root

<i>Instructions</i>	SQRT			
<i>Syntax</i>	<i>sqrt dst, src0</i>			
<i>Description</i>	Computes the square root of <i>src0.w</i> . By default, this instruction operates on <i>src0.w</i> , but can operate on any component by swizzling it into the fourth component. If <i>src0.w</i> is less than zero, the result is undefined. The result is approximate.			
<i>Format</i>	1-input, 1-output.			
<i>Opcode</i>	Field Name	Bits	Description	
	code	15:0	IL_OP_SQRT	
	control	29:16	Must be zero.	
	sec_modifier_present	30	Must be zero.	
	pri_modifier_present	31	Must be zero.	
<i>Related</i>	SQRT_VEC.			

Component-Wise Square Root

<i>Instructions</i>	SQRT_VEC		
<i>Syntax</i>	<code>sqrt_vec dst, src0</code>		
<i>Description</i>	Computes the square root of each component of <i>src0</i> . If a component of <i>src0</i> is less than zero, the result for that component is undefined. The result is approximate.		
<i>Format</i>	1-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_SQRT_VEC
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	SQRT.		

Floating Point Subtraction

<i>Instructions</i>	SUB		
<i>Syntax</i>	<code>sub dst, src0, src1</code>		
<i>Description</i>	Subtracts each component of <i>src0</i> from the corresponding component of <i>src1</i> . No carry or borrow beyond the 32-bit values of each component is performed. Operation: <pre> VECTOR v1 = EvalSource(src0); VECTOR v2 = EvalSource(src1); VECTOR v; for (i=0; i < 4; i++) v[i] = v1[i] - v2[i]; WriteResult(v, dst); </pre>		
<i>Format</i>	2-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_SUB
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	DADD.		

Tangent (tan)

<i>Instructions</i>	TAN		
<i>Syntax</i>	tan <i>dst</i> , <i>src0</i>		
<i>Description</i>	<p>Computes the tangent of <i>src0.w</i>. <i>src0.w</i> is in radians. <i>src0.w</i> must be within the range $[-\pi, \pi]$ for each function; otherwise, the results are undefined. By default, this instruction operates on <i>src0.w</i>, but can operate on any component by swizzling it into the fourth component. The maximum absolute error is 0.002.</p> <p>Operation:</p> <pre>VECTOR v1 = EvalSource(src0); VECTOR v v[0] = v[1] = v[2] = v[3] = tan(v1[3]); WriteResult(v, dst);</pre>		
<i>Format</i>	1-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_TAN
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	COS_VEC, SIN_VEC, SINCOS.		

Transpose a 4x4 Matrix**Instructions** **TRANPOSE****Syntax** transpose *dst*, *src0***Description** Transposes the rows and columns of the 4x4 matrix indicated by *src0*.

Operation:

```

VECTOR vsrc0 = EvalSource(src0);
VECTOR vsrc1 = EvalSource(src0+1); # Register number of src0 + 1
VECTOR vsrc2 = EvalSource(src0+2); # Register number of src0 + 2
VECTOR vsrc3 = EvalSource(src0+3); # Register number of src0 + 3
VECTOR vdst0;
VECTOR vdst1;
VECTOR vdst2;
VECTOR vdst3;
vdst0[0] = vsrc0[0];
vdst0[1] = vsrc1[0];
vdst0[2] = vsrc2[0];
vdst0[3] = vsrc3[0];
vdst1[0] = vsrc0[1];
vdst1[1] = vsrc1[1];
vdst1[2] = vsrc2[1];
vdst1[3] = vsrc3[1];
vdst2[0] = vsrc0[2];
vdst2[1] = vsrc1[2];
vdst2[2] = vsrc2[2];
vdst2[3] = vsrc3[2];
vdst3[0] = vsrc0[3];
vdst3[1] = vsrc1[3];
vdst3[2] = vsrc2[3];
vdst3[3] = vsrc3[3];
WriteResult(vdst0, dst);
WriteResult(vdst1, dst+1); # Register number of dst + 1
WriteResult(vdst2, dst+2); # Register number of dst + 2
WriteResult(vdst3, dst+3); # Register number of dst + 3

```

Format 1-input, 1-output.

Opcode	Field Name	Bits	Description
	code	15:0	IL_OP_TRANPOSE
	control	29:16	Must be the enumerated type ILMatrix(IL_MATRIX_4X4). See Table 5.6 on page 5-3.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.

Related None.

Truncate

<i>Instructions</i>	TRC		
<i>Syntax</i>	<code>trc dst, src0</code>		
<i>Description</i>	<p>Extracts the integer portion (the number to the left of the decimal point) of each component of <i>src0</i>. For example, truncate(-1.7) is -1.0, not -2.0 as it is with floor(-1.7). Note, the sign remains intact.</p> <p>Operation:</p> <pre>VECTOR v1 = EvalSource(src0); VECTOR v; for (i=0; i < 4; i++) v[i] = trunc(v1[i]); WriteResult(v, dst);</pre>		
<i>Format</i>	1-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_TRC
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	DFRAC, FRC.		

6.10 Double-Precision Instructions

Double Not Equal Compare

<i>Instructions</i>	D_NE		
<i>Syntax</i>	<code>dne dst, src0, src1</code>		
<i>Description</i>	<p>Component-wise compares two vectors using a float comparison that follows floating point rules. If <i>src0</i>{xy zw} and the corresponding component pair in <i>src1</i> satisfy the comparison condition, the corresponding component of <i>dst</i> is set to TRUE and returns 0xFFFFFFFF; otherwise, the corresponding component of <i>dst</i> is set to FALSE and returns 0x00000000. Primary and secondary opcode modifiers are not permitted, and no additional control options are supported. Thus, the <code>pri_modifier_present</code> and <code>sec_modifier_present</code> fields must be zero.</p>		
<i>Format</i>	2-input, 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_D_NE
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	Float Compare, Integer Compare, Unsigned Integer Compare.		

Double Equal Compare

Instructions **D_EQ***Syntax* `deq dst, src0, src1`*Description* Component-wise compares two vectors using a float comparison that follows floating point rules. If *src0*.{xy|zw} and the corresponding component pair in *src1* satisfy the comparison condition, the corresponding component of *dst* is set to TRUE and returns 0xFFFFFFFF; otherwise, the corresponding component of *dst* is set to FALSE and returns 0x00000000. Primary and secondary opcode modifiers are not permitted, and no additional control options are supported. Thus, the `pri_modifier_present` and `sec_modifier_present` fields must be zero.*Format* 2-input, 1-output.

<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_D_EQ
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.

Related Float Compare, Integer Compare, Unsigned Integer Compare.

Double Greater-Than or Equal Compare

Instructions D_GE

Syntax dge dst, src0, src1

Description Component-wise compares two vectors using a float comparison that follows floating point rules. If *src0*.{xy|zw} and the corresponding component pair in *src1* satisfy the comparison condition, the corresponding component of *dst* is set to TRUE and returns 0xFFFFFFFF; otherwise, the corresponding component of *dst* is set to FALSE and returns 0x00000000. Primary and secondary opcode modifiers are not permitted, and no additional control options are supported. Thus, the *pri_modifier_present* and *sec_modifier_present* fields must be zero.

<u>src0</u>	<u>src1</u>								
	-inf	-F	-denorm	-0	+0	+denorm	+F	+inf	NaN
-inf	T	F	F	F	F	F	F	F	F
-F	T	T/F	F	F	F	F	F	F	F
-denorm	T	T	T	T	T	T	F	F	F
-0	T	T	T	T	T	T	F	F	F
+0	T	T	T	T	T	T	F	F	F
+denorm	T	T	T	T	T	T	F	F	F
+F	T	T	T	T	T	T	T/F	F	F
+inf	T	T	T	T	T	T	T	T	F
NaN	F	F	F	F	F	F	F	F	F

Format 2-input, 1-output.

Opcode	Field Name	Bits	Description
	code	15:0	IL_OP_D_GE
	control	29:16	Must be zero.
	<i>sec_modifier_present</i>	30	Must be zero.
	<i>pri_modifier_present</i>	31	Must be zero.

Related Float Compare, Integer Compare, Unsigned Integer Compare.

Double Less-Than Compare

Instructions **D_LT**

Syntax `dlt dst, src0, src1`

Description Component-wise compares two vectors using a float comparison that follows floating point rules. If *src0*.{xy|zw} and the corresponding component pair in *src1* satisfy the comparison condition, the corresponding component of *dst* is set to TRUE and returns 0xFFFFFFFF; otherwise, the corresponding component of *dst* is set to FALSE and returns 0x00000000. Primary and secondary opcode modifiers are not permitted, and no additional control options are supported. Thus, the *pri_modifier_present* and *sec_modifier_present* fields must be zero.

<u><i>src0</i></u>	<u><i>src1</i></u>								
	-inf	-F	-denorm	-0	+0	+denorm	+F	+inf	NaN
-inf	F	F	F	F	F	F	F	F	F
-F	T	T/F	F	F	F	F	F	F	F
-denorm	T	T	F	F	F	F	F	F	F
-0	T	T	F	F	F	F	F	F	F
+0	T	T	F	F	F	F	F	F	F
+denorm	T	T	F	F	F	F	F	F	F
+F	T	T	T	T	T	T	T/F	F	F
+inf	T	T	T	T	T	T	T	F	F
NaN	F	F	F	F	F	F	F	F	F

Format 2-input, 1-output.

<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_D_LT
	control	29:16	Must be zero.
	<i>sec_modifier_present</i>	30	Must be zero.
	<i>pri_modifier_present</i>	31	Must be zero.

Related Float Compare, Integer Compare, Unsigned Integer Compare.

Split Double into Fraction and Exponent

Instructions **D_FREXP**

Syntax `dfrexp dst, src0`

Description Output x is zero. Output y is the exponent as an integer. Output wz is the mantissa in range (-1.0,-0.5][0.5,1.0). Abs or negate modifiers can be used on the source. If the input is +-0, all four results are zero. The size is preserved for the wz output. If the input is a NaN, y is set to -1, and wz is set to a NaN = {s, 0x7ff, 1'b1, mant[50:0]} // QNaN. If the input is an +-iff, y is set to -1, and wz is set to a NaN = 0xfff8000000000000.

d	-inf/+inf	-0/+0	-denorm/+denorm	NaN
f	NAN	{sign,0}	{sign,0}	d
e	0xFFFFFFFF	0	0	0xFFFFFFFF

Format 1-input 1-output.

<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_D_FREXP
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.

Related LDEXP.

Add Two Doubles

Instructions **D_ADD**

Syntax `dadd dst, src0, src1`

Description Abs and negate modifiers can be used on the source. Output is either wz or yx. Output is the correct IEEE result.

	<i>src1</i>						
	-inf	-F	-0	+0	+F	+inf	NaN
<i>src0</i>							
-inf	-inf	-inf	-inf	-inf	-inf	NaN	NaN
-F	-inf	-F	<i>src0</i>	<i>src0</i>	+-F or +-0	+inf	NaN
-0	-int	<i>src1</i>	-0	+0	<i>src1</i>	+inf	NaN
+0	-inf	<i>src1</i>	+0	+0	<i>src1</i>	+inf	NaN
+F	-inf	+-F or +-0	<i>src0</i>	<i>src0</i>	-F	+inf	NaN
+inf	NaN	+inf	+inf	+inf	+inf	+inf	NaN
NaN	NaN	NaN	NaN	NaN	NaN	NaN	NaN

Format 2-input 1-output.

<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_D_ADD
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.

Related None.

Multiply Two Doubles

Instructions `D_MUL`

Syntax `dmul dst, src0, src1`

Description Abs and negate modifiers can be used on the source. Output is either wz or yx. Output is the correct IEEE result.

	<i>src1</i>								
	-inf	-F	-1.0	-0	+0	+1.0	+F	+inf	NaN
<i>src0</i>									
-inf	+inf	+inf	+inf	NaN	NaN	inf	-inf	-inf	NaN
-F	+inf	+F	-src0	+0	-0	src0	-F	-inf	NaN
-1.0	+int	-src1	+1.0	+0	-0	-1.0	-src1	-inf	NaN
-0	NaN	+0	+0	+0	-0	-0	-0	NaN	NaN
+0	NaN	-0	-0	-0	+0	+0	+0	NaN	NaN
+1.0	-inf	src1	-1.0	-0	+0	+1.0	src1	+inf	NaN
+F	-inf	-F	-src0	-0	+0	src0	+F	+inf	NaN
+inf	-inf	-inf	-inf	NaN	NaN	+inf	+inf	+inf	NaN
NaN	NaN	NaN	NaN	NaN	NaN	NaN	NaN	NaN	NaN

Format 2-input 1-output.

Opcode	Field Name	Bits	Description
	code	15:0	IL_OP_D_MUL
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.

Related None.

Divide Two Doubles

<i>Instructions</i>	D_DIV		
<i>Syntax</i>	ddiv <i>dst</i> , <i>src0</i> , <i>src1</i>		
<i>Description</i>	Abs and negate modifiers can be used on the source. Output is either wz or yx. Output is the correct IEEE result.		
<i>Format</i>	2-input 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_D_DIV
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	None.		

Pack an EXP and Mantissa into a Double

<i>Instructions</i>	D_LDEXP		
<i>Syntax</i>	dlldexp <i>dst</i> , <i>src0</i> , <i>src1</i>		
<i>Description</i>	Puts an exp and a mantissa into a double. The computation is result = source1 * 2 source. Abs and negate modifiers can be used. Output.yx or output.wz are set to the double result. NaN in either input produces a NaN. Overflow produces inf. Underflow produces 0.		
<i>Format</i>	2-input 1-output.		
<i>Opcode</i>	Field Name	Bits	Description
	code	15:0	IL_OP_D_LDEXP
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	FREXP.		

Return Fractional Part of a Double

<i>Instructions</i>	D_FRAC		
<i>Syntax</i>	dfrac <i>dst</i> , <i>src0</i>		
<i>Description</i>	Returns the fractional part of the source in range [0.0 – 1.0). Abs and negate modifiers can be used. Output.yx or output.wz is set to the double result.		
	x	-inf -F -1.0 -denorm -0 +0 +denorm +1.0 +F +inf NaN	
	result	NAN [+0.0,+1.0)* +0 +0 d+0 +0 +0 +0 [+0.0,+1.0)* NAN NaN	
<i>Format</i>	1-input 1-output.		
<i>Opcod</i>	Field Name	Bits	Description
	code	15:0	IL_OP_D_FRAC
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	None.		

Double Multiply and Add

<i>Instructions</i>	D_MULADD		
<i>Syntax</i>	dmad <i>dst</i> , <i>src0</i> , <i>src1</i> , <i>src2</i>		
<i>Description</i>	Computes: <i>src0</i> * <i>src1</i> + <i>src2</i> . Abs and negate can be used on the sources. The output in either yz or wz is the result. The compiler determines if this is identical to separate MULs and ADD based on the chipset. The operation truncates internally, so the result is not identical to doing a MUL followed by an ADD.		
<i>Format</i>	3-input 1-output.		
<i>Opcod</i>	Field Name	Bits	Description
	code	15:0	IL_OP_D_MULADD
	control	29:16	Must be zero.
	sec_modifier_present	30	Must be zero.
	pri_modifier_present	31	Must be zero.
<i>Related</i>	None.		

Glossary of Terms

Term	Description
*	Any number of alphanumeric characters in the name of a microcode format, microcode parameter, or instruction.
< >	Angle brackets denote streams.
[1,2)	A range that includes the left-most value (in this case, 1) but excludes the right-most value (in this case, 2).
[1,2]	A range that includes both the left-most and right-most values (in this case, 1 and 2).
{BUF, SWIZ}	One of the multiple options listed. In this case, the string <i>BUF</i> or the string <i>SWIZ</i> .
{x y}	One of the multiple options listed. In this case, x or y.
0.0	A single-precision (32-bit) floating-point value.
0x	Indicates that the following is a hexadecimal number.
1011b	A binary value, in this example a 4-bit value.
29'b0	29 bits with the value 0.
7:4	A bit range, from bit 7 to 4, inclusive. The high-order bit is shown first.
ABI	Application Binary Interface.
<i>absolute</i>	A displacement that references the base of a code segment, rather than an instruction pointer. See <i>relative</i> .
<i>active mask</i>	A 1-bit-per-pixel mask that controls which pixels in a "quad" are really running. Some pixels general-purpose not be running if the current "primitive" does not cover the whole quad. A mask can be updated with a <code>PRED_SET*</code> ALU instruction, but updates do not take effect until the end of the ALU clause.
<i>address stack</i>	A stack that contains only addresses (no other state). Used for flow control. Popping the address stack overrides the instruction address field of a flow control instruction. The address stack is only modified if the flow control instruction decides to jump.
ACML	AMD Core Math Library. Includes implementations of the full BLAS and LAPACK routines, FFT, Math transcendental and Random Number Generator routines, stream processing backend for load balancing of computations between the CPU and stream processor.
<i>aL (also AL)</i>	Loop register. A three-component vector (x, y and z) used to count iterations of a loop.
<i>allocate</i>	To reserve storage space for data in an output buffer ("scratch buffer," "ring buffer," "stream buffer," or "reduction buffer") or for data in an input buffer ("scratch buffer" or "ring buffer") before exporting (writing) or importing (reading) data or addresses to, or from that buffer. Space is allocated only for data, not for addresses. After allocating space in a buffer, an "export" operation can be done.

ATI STREAM COMPUTING

Term	Description
<i>ALU</i>	Arithmetic Logic Unit. Responsible for arithmetic operations like addition, subtraction, multiplication, division, and bit manipulation on integer and floating point values. In stream computing, these are known as <i>stream cores</i> . ALU.[X,Y,Z,W] - an ALU that can perform four vector operations in which the four operands (integers or single-precision floating point values) do not have to be related. It performs "SIMD" operations. Thus, although the four operands need not be related, all four operations execute the same instruction. ALU.Trans - An ALU unit that can perform one ALU.Trans (transcendental, scalar) operation, or advanced integer operation, on one integer or single-precision floating-point value, and replicate the result. A single instruction can co-issue four ALU.Trans operations to an ALU.[X,Y,Z,W] unit and one (possibly complex) operation to an ALU.Trans unit, which can then replicate its result across all four component being operated on in the associated ALU.[X,Y,Z,W] unit.
<i>ATI Stream™ SDK</i>	A complete software development suite from ATI for developing applications for ATI Stream Processors. Currently, the ATI Stream SDK includes Brook+ and CAL.
<i>AR</i>	Address register.
<i>aTid</i>	Absolute thread id. It is the ordinal count of all threads being executed (in a draw call).
<i>b</i>	A bit, as in <i>1Mb</i> for one megabit, or <i>lsb</i> for least-significant bit.
<i>B</i>	A byte, as in <i>1MB</i> for one megabyte, or <i>LSB</i> for least-significant byte.
<i>BLAS</i>	Basic Linear Algebra Subroutines.
<i>border color</i>	Four 32-bit floating-point numbers (XYZW) specifying the border color.
<i>branch granularity</i>	The number of threads executed during a branch. For ATI, branch granularity is equal to wavefront granularity.
<i>brcc</i>	Source-to-source meta-compiler that translates Brook programs (.br files) into device-dependent kernels embedded in valid C++ source code that includes CPU code and stream processor device code, which later are linked into the executable.
<i>Brook+</i>	A high-level language derived from C which allows developers to write their applications at an abstract level without having to worry about the exact details of the hardware. This enables the developer to focus on the algorithm and not the individual instructions run on the stream processor. Brook+ is an enhancement of Brook, which is an open source project out of Stanford. Brook+ adds additional features available on ATI Stream Processors and provides a CAL backend.
<i>brt</i>	The Brook runtime library that executes pre-compiled kernel routines invoked from the CPU code in the application.
<i>burst mode</i>	The limited write combining ability. See write combining.
<i>byte</i>	Eight bits.
<i>cache</i>	A read-only or write-only on-chip or off-chip storage space.
<i>CAL</i>	Compute Abstraction Layer. A device-driver library that provides a forward-compatible interface to ATI Stream processor devices. This lower-level API gives users direct control over the hardware: they can directly open devices, allocate memory resources, transfer data and initiate kernel execution. CAL also provides a JIT compiler for ATI IL.
<i>CF</i>	Control Flow.
<i>cfile</i>	Constant file or constant register.
<i>channel</i>	A component in a vector.

ATI STREAM COMPUTING

Term	Description
<i>clamp</i>	To hold within a stated range.
<i>clause</i>	A group of instructions that are of the same type (all stream core, all fetch, etc.) executed as a group. A clause is part of a CAL program written using the stream processor ISA. Executed without pre-emption.
<i>clause size</i>	The total number of slots required for an stream core clause.
<i>clause temporaries</i>	Temporary values stored at GPR that do not need to be preserved past the end of a clause.
<i>clear</i>	To write a bit-value of 0. Compare "set".
<i>command</i>	A value written by the host processor directly to the stream processor. The commands contain information that is not typically part of an application program, such as setting configuration registers, specifying the data domain on which to operate, and initiating the start of data processing.
<i>command processor</i>	A logic block in the R700 (HD4000-family of devices) that receives host commands, interprets them, and performs the operations they indicate.
<i>component</i>	(1) A 32-bit piece of data in a "vector". (2) A 32-bit piece of data in an array. (3) One of four data items in a 4-component register.
<i>compute shader</i>	Similar to a pixel shader, but exposes data sharing and synchronization.
<i>constant buffer</i>	Off-chip memory that contains constants. A constant buffer can hold up to 1024 four-component vectors. There are fifteen constant buffers, referenced as cb0 to cb14. An immediate constant buffer is similar to a constant buffer. However, an immediate constant buffer is defined within a kernel using special instructions. There are fifteen immediate constant buffers, referenced as icb0 to icb14.
<i>constant cache</i>	A constant cache is a hardware object (off-chip memory) used to hold data that remains unchanged for the duration of a kernel (constants). "Constant cache" is a general term used to describe constant registers, constant buffers or immediate constant buffers.
<i>constant file</i>	Same as constant register.
<i>constant index register</i>	Same as "AR" register.
<i>constant registers</i>	On-chip registers that contain constants. The registers are organized as four 32-bit component of a vector. There are 256 such registers, each one 128-bits wide.
<i>constant waterfaling</i>	Relative addressing of a constant file. See waterfaling.
<i>context</i>	A representation of the state of a CAL device.
<i>core clock</i>	See engine clock. The clock at which the stream processor stream core runs.
<i>CPU</i>	Central Processing Unit. Also called host. Responsible for executing the operating system and the main part of the application. The CPU provides data and instructions to the stream processor.
<i>CRs</i>	Constant registers. There are 512 CRs, each one 128 bits wide, organized as four 32-bit values.
<i>CS</i>	Compute shader. A shader type, analogous to VS/PS/GS/ES.
<i>CTM</i>	Close-to-Metal. A thin, HW/SW interface layer. This was the predecessor of the ATI CAL.

ATI STREAM COMPUTING

Term	Description
<i>DC</i>	Data Copy Shader.
<i>device</i>	A <i>device</i> is an entire ATI Stream processor.
<i>DMA</i>	Direct-memory access. Also called DMA engine. Responsible for independently transferring data to, and from, the stream processor's local memory. This allows other computations to occur in parallel, increasing overall system performance.
<i>double word</i>	Dword. Two words, or four bytes, or 32 bits.
<i>double quad word</i>	Eight words, or 16 bytes, or 128 bits. Also called "octword."
<i>domain of execution</i>	A specified rectangular region of the output buffer to which threads are mapped.
<i>DPP</i>	Data-Parallel Processor.
<i>dst.X</i>	The X "slot" of a destination operand.
<i>dword</i>	Double word. Two words, or four bytes, or 32 bits.
<i>element</i>	A component in a vector.
<i>engine clock</i>	The clock driving the stream core and memory fetch units on the stream processor stream processor core.
<i>enum(7)</i>	A seven-bit field that specifies an enumerated set of decimal values (in this case, a set of up to 27 values). The valid values can begin at a value greater than, or equal to, zero; and the number of valid values can be less than, or equal to, the maximum supported by the field.
<i>event</i>	A token sent through a pipeline that can be used to enforce synchronization, flush caches, and report status back to the host application.
<i>export</i>	To write data from GPRs to an output buffer (scratch, ring, stream, frame or global buffer, or to a register), or to read data from an input buffer (a "scratch buffer" or "ring buffer") to GPRs. The term "export" is a partial misnomer because it performs both input and output functions. Prior to exporting, an allocation operation must be performed to reserve space in the associated buffer.
<i>FFT</i>	Fast Fourier Transform.
<i>flag</i>	A bit that is modified by a CF or stream core operation and that can affect subsequent operations.
<i>FLOP</i>	Floating Point Operation.
<i>flush</i>	To writeback and invalidate cache data.
<i>frame</i>	A single two-dimensional screenful of data, or the storage space required for it.
<i>frame buffer</i>	Off-chip memory that stores a frame.
<i>FS</i>	Fetch subroutine. A global program for fetching vertex data. It can be called by a "vertex shader" (VS), and it runs in the same thread context as the vertex program, and thus is treated for execution purposes as part of the vertex program. The FS provides driver independence between the process of fetching data required by a VS, and the VS itself. This includes having a semantic connection between the outputs of the fetch process and the inputs of the VS.
<i>function</i>	A subprogram called by the main program or another function within an ATI IL stream. Functions are delineated by <code>FUNC</code> and <code>ENDFUNC</code> .

ATI STREAM COMPUTING

Term	Description
<i>gather</i>	Reading from arbitrary memory locations by a thread.
<i>gather stream</i>	Input streams are treated as a memory array, and data elements are addressed directly.
<i>global buffer</i>	Memory space containing the arbitrary address locations to which uncached kernel outputs are written. Can be read either cached or uncached. When read in uncached mode, it is known as mem-import. Allows applications the flexibility to read from and write to arbitrary locations in input buffers and output buffers, respectively.
<i>GPGPU</i>	General-purpose stream processor. A stream processor that performs general-purpose calculations.
<i>GPR</i>	General-purpose register. GPRs hold vectors of either four 32-bit IEEE floating-point, or four 8-, 16-, or 32-bit signed or unsigned integer or two 64-bit IEEE double precision data components (values). These registers can be indexed, and consist of an on-chip part and an off-chip part, called the "scratch buffer," in memory.
<i>GPU</i>	Graphics Processing Unit. An integrated circuit that renders and displays graphical images on a monitor. Also called Graphics Hardware, Stream Processor, and Data Parallel Processor.
<i>GPU engine clock frequency</i>	Also called 3D engine speed.
<i>GS</i>	Geometry Shader.
<i>HAL</i>	Hardware Abstraction Layer.
<i>host</i>	Also called CPU.
<i>iff</i>	If and only if.
<i>IL</i>	Intermediate Language. In this manual, the ATI version: ATI IL. A pseudo-assembly language that can be used to describe kernels for stream processors. ATI IL is designed for efficient generalization of stream processor instructions so that programs can run on a variety of platforms without having to be rewritten for each platform.
<i>in flight</i>	A thread currently being processed.
<i>instruction</i>	A computing function specified by the <i>code</i> field of an IL_OpCode token. Compare "opcode", "operation", and "instruction packet".
<i>instruction packet</i>	A group of tokens starting with an IL_OpCode token that represent a single ATI IL instruction.
<i>int(2)</i>	A 2-bit field that specifies an integer value.
<i>ISA</i>	Instruction Set Architecture. The complete specification of the interface between computer programs and the underlying computer hardware.
<i>kcache</i>	A memory area containing "waterfall" (off-chip) constants. The cache lines of these constants can be locked. The "constant registers" are the 256 on-chip constants.
<i>kernel</i>	A small, user-developed program that is run repeatedly on a stream of data. A parallel function that operates on every element of input streams. A device program is one type of kernel. Unless otherwise specified, an ATI Stream processor program is a kernel composed of a main program and zero or more functions. Also called Shader Program. This is not to be confused with an OS kernel, which controls hardware.
<i>LAPACK</i>	Linear Algebra Package.

ATI STREAM COMPUTING

Term	Description
<i>LERP</i>	Linear Interpolation.
<i>local memory fetch units</i>	Dedicated hardware that a) processes fetch instructions, b) requests data from the memory controller, and c) loads registers with data returned from the cache. They are run at stream processor stream core or engine clock speeds. Formerly called texture units.
<i>LOD</i>	Level Of Detail.
<i>loop index</i>	A register initialized by software and incremented by hardware on each iteration of a loop.
<i>lsb</i>	Least-significant bit.
<i>LSB</i>	Least-significant byte.
<i>MAD</i>	Multiply-Add. A fused instruction that both multiplies and adds.
<i>mask</i>	(1) To prevent from being seen or acted upon. (2) A field of bits used for a control purpose.
<i>MBZ</i>	Must be zero.
<i>mem-export</i>	An ATI IL term random writes to the global buffer.
<i>mem-import</i>	Uncached reads from the global buffer.
<i>memory clock</i>	The clock driving the memory chips on the stream processor.
<i>microcode format</i>	An encoding format whose fields specify instructions and associated parameters. Microcode formats are used in sets of two or four. For example, the two mnemonics, <code>CF_DWORD[0,1]</code> indicate a microcode-format pair, <code>CF_DWORD0</code> and <code>CF_DWORD1</code> .
<i>MIMD</i>	Multiple Instruction Multiple Data. – Multiple SIMD units operating in parallel (Multi-Processor System) – Distributed or shared memory
<i>MRT</i>	Multiple Render Target. One of multiple areas of local stream processor memory, such as a “frame buffer”, to which a graphics pipeline writes data.
<i>MSAA</i>	Multi-Sample Anti-Aliasing.
<i>msb</i>	Most-significant bit.
<i>MSB</i>	Most-significant byte.
<i>neighborhood</i>	A group of four threads in the same wavefront that have consecutive thread IDs (Tid). The first Tid must be a multiple of four. For example, threads with Tid = 0, 1, 2, and 3 form a neighborhood, as do threads with Tid = 12, 13, 14, and 15.
<i>normalized</i>	A numeric value in the range [a, b] that has been converted to a range of 0.0 to 1.0 using the formula: $normalized\ value = value / (b - a + 1)$
<i>oct word</i>	Eight words, or 16 bytes, or 128 bits. Same as “double quad word”.
<i>opcode</i>	The numeric value of the <i>code</i> field of an “instruction”. For example, the opcode for the CMOV instruction is decimal 16 (0x10).
<i>opcode token</i>	A 32-bit value that describes the operation of an instruction.
<i>operation</i>	The function performed by an “instruction”.
<i>PaC</i>	Parameter Cache.

ATI STREAM COMPUTING

Term	Description
<i>page</i>	A program-controlled cache, backing up processor-accessible memory.
<i>PCI Express</i>	A high-speed computer expansion card interface used by modern graphics cards, stream processors and other peripherals needing high data transfer rates. Unlike previous expansion interfaces, PCI Express is structured around point-to-point links. Also called PCIe.
<i>PoC</i>	Position Cache.
<i>pop</i>	Write “stack” entries to their associated hardware-maintained control-flow state. The POP_COUNT field of the CF_DWORD1 microcode format specifies the number of stack entries to pop for instructions that pop the stack. Compare “push.”
<i>pre-emption</i>	The act of temporarily interrupting a task being carried out on a computer system, without requiring its cooperation, with the intention of resuming the task at a later time.
<i>processor</i>	Unless otherwise stated, the ATI Stream Processor.
<i>program</i>	Unless otherwise specified, a program is a set of instructions that can run on the ATI Stream Processor. A device program is a type of kernel.
<i>PS</i>	Pixel Shader.
<i>push</i>	Read hardware-maintained control-flow state and write their contents onto the stack. Compare pop.
<i>PV</i>	Previous vector register. It contains the previous four-component vector result from a ALU.[X,Y,Z,W] unit within a given clause.
<i>quad</i>	Group of 2x2 threads in the domain. Always processed together.
<i>rasterization</i>	The process of mapping threads from the domain of execution to the SIMD engine. This term is a carryover from graphics, where it refers to the process of turning geometry, such as triangles, into pixels.
<i>rasterization order</i>	The order of the thread mapping generated by rasterization.
<i>RB</i>	Ring Buffer.
<i>register</i>	A 128-bit address mapped memory space consisting of four 32-bit components.
<i>relative</i>	Referencing with a displacement (also called offset) from an index register or the loop index, rather than from the base address of a program (the first control flow [CF] instruction).
<i>render backend unit</i>	The hardware units in a stream processor stream processor core responsible for writing the results of a kernel to output streams by writing the results to an output cache and transferring the cache data to memory.
<i>resource</i>	A block of memory used for input to, or output from, a kernel.
<i>ring buffer</i>	An on-chip buffer that indexes itself automatically in a circle.
<i>Rsvd</i>	Reserved.
<i>sampler</i>	A structure that contains information necessary to access data in a resource. Also called Fetch Unit.
<i>SC</i>	Shader Compiler.
<i>scalar</i>	A single data component, unlike a vector which contains a set of two or more data elements.

ATI STREAM COMPUTING

Term	Description
<i>scatter</i>	Writes (by uncached memory) to arbitrary locations.
<i>scatter write</i>	Kernel outputs to arbitrary address locations. Must be uncached. Must be made to a memory space known as the global buffer.
<i>scratch buffer</i>	A variable-sized space in off-chip-memory that stores some of the “GPRs”.
<i>set</i>	To write a bit-value of 1. Compare “clear”.
<i>shader processor</i>	Also called thread processor.
<i>shader program</i>	User developed program. Also called kernel.
<i>SIMD</i>	Single instruction multiple data. – Each SIMD receives independent stream core instructions. – Each SIMD applies the instructions to multiple data elements.
<i>SIMD Engine</i>	A collection of thread processors, each of which executes the same instruction per cycle.
<i>SIMD pipeline</i>	A hardware block consisting of five stream cores, one stream core instruction decoder and issuer, one stream core constant fetcher, and support logic. All parts of a SIMD pipeline receive the same instruction and operate on different data elements. Also known as “slice.”
<i>Simultaneous Instruction Issue</i>	Input, output, fetch, stream core, and control flow per SIMD engine.
<i>SKA</i>	Stream KernelAnalyzer. A performance profiling tool for developing, debugging, and profiling stream kernels using high-level stream computing languages.
<i>slot</i>	A position, in an “instruction group,” for an “instruction” or an associated literal constant. An ALU instruction group consists of one to seven slots, each 64 bits wide. All ALU instructions occupy one slot, except double-precision floating-point instructions, which occupy either two or four slots. The size of an ALU clause is the total number of slots required for the clause.
<i>SPU</i>	Shader processing unit.
<i>src0, src1, etc.</i>	In floating-point operation syntax, a 32-bit source operand. Src0_64 is a 64-bit source operand.
<i>stage</i>	A sampler and resource pair.
<i>stream</i>	A collection of data elements of the same type that can be operated on in parallel.
<i>stream buffer</i>	A variable-sized space in off-chip memory that stores an instruction stream. It is an output-only buffer, configured by the host processor. It does not store inputs from off-chip memory to the processor.
<i>stream core</i>	The fundamental, programmable computational units, responsible for performing integer, single, precision floating point, double precision floating point, and transcendental operations. They execute VLIW instructions for a particular thread. Each stream processor stream core handles a single instruction within the VLIW instruction.
<i>stream operator</i>	A node that can restructure data.
<i>stream processor</i>	A parallel processor capable of executing multiple threads of a kernel in order to process streams of data.
<i>swizzling</i>	To copy or move any component in a source vector to any element-position in a destination vector. Accessing elements in any combination.

ATI STREAM COMPUTING

Term	Description
<i>thread</i>	One invocation of a kernel corresponding to a single element in the domain of execution. An instance of execution of a shader program on an ALU.
<i>thread group</i>	It contains one or more thread blocks. Threads in the same thread-group but different thread-blocks might communicate to each through global per-SIMD shared memory. This is a concept mainly for global data share (GDS). A thread group can contain one or more wavefronts, the last of which can be a partial wavefront. All wavefronts in a thread group can run on only one SIMD engine; however, multiple thread groups can share a SIMD engine, if there are sufficient resources.
<i>thread processor</i>	The hardware units in a SIMD engine responsible for executing the threads of a kernel. It executes the same instruction per cycle. Each thread processor contains multiple stream cores. Also called shader processor.
<i>thread-block</i>	A group of threads which might communicate to each other through local per SIMD shared memory. It can contain one or more wavefronts (the last wavefront can be a partial wavefront). A thread-block (all its wavefronts) can only run on one SIMD engine. However, multiple thread blocks can share a SIMD engine, if there are enough resources to fit them in.
<i>Tid</i>	Thread id within a thread block. An integer number from 0 to Num_threads_per_block-1
<i>token</i>	A 32-bit value that represents an independent part of a stream or instruction.
<i>uncached read/write unit</i>	The hardware units in a stream processor responsible for handling uncached read or write requests from local memory on the stream processor.
<i>vector</i>	(1) A set of up to four related values of the same data type, each of which is an element. For example, a vector with four elements is known as a "4-vector" and a vector with three elements is known as a "3-vector". (2) See "AR". (3) See ALU.[X,Y,Z,W].
<i>VLIW design</i>	Very Long Instruction Word. <ul style="list-style-type: none"> – Co-issued up to 6 operations (5 stream cores + 1 FC) – 1.25 Machine Scalar operation per clock for each of 64 data elements – Independent scalar source and destination addressing
<i>waterfall</i>	To use the address register (AR) for indexing the GPRs. Waterfall behavior is determined by a "configuration registers."
<i>wavefront</i>	Group of threads executed together on a single SIMD engine. Composed of quads. A full wavefront contains 64 threads; a wavefront with fewer than 64 threads is called a partial wavefront. For the HD4000-family of devices, there are 64, 32, 16 threads in a full wavefront.
<i>write combining</i>	Combining several smaller writes to memory into a single larger write to minimize any overhead associated with write commands.

Index

Symbols

[] syntax	3-2
_abs	
definition	3-4
source modifier	3-4
_bias	
definition	3-4
source modifier	3-4
_bx2	
definition	3-4
source modifier	3-4
_divcomp	
definition	3-4
source modifier	3-4
_invert	
definition	3-4
source modifier	3-4
_neg	
definition	3-4
source modifier	3-4
_sign	
definition	3-4
source modifier	3-4
_x2	
definition	3-4
source modifier	3-4

Numerics

32-bit token	2-1
--------------	-----

A

abs	
definition	2-8
IL_Src_Mod	2-8
modifier	2-6
ABS instruction	6-97
absolute register index	3-2
Absolute Thread_ID	
register type	4-17
access mode wavefront	1-2
ACOS instruction	6-98
ADD instruction	6-98

ADDR	4-1
address	2-8
base relative	3-2
loop relative	3-2
relative token	2-8
address_register	
definition	2-8
IL_Rel_Addr	2-8
AND instruction	6-99
ASIN instruction	6-99
ATAN instruction	6-100
aTid	1-3

B

BARYCENTRIC_COORD	
register type	4-2
base relative address	3-2
bias	
definition	2-7
IL_Src_Mod	2-7
binary specifier	3-2
BREAK instruction	6-3
BREAK_LOGICALNZ instruction	6-4
BREAK_LOGICALZ instruction	6-4
BREAKC instruction	6-4

C

CALL instruction	6-5
CALL_LOGICALNZ instruction	6-8
CALL_LOGICALZ instruction	6-7
CALLNZ instruction	6-6
CASE instruction	6-9
clamp	
definition	2-4, 2-8
IL_Dst_Mod	2-4
IL_Src_Mod	2-8
CLAMP instruction	6-101
CLG instruction	6-102
client language	2-1
client_type	2-2
CMOV instruction	6-103
CMOV_LOGICAL instruction	6-103
CMP instruction	6-104

ATI STREAM COMPUTING

code		D_FRAC instruction	6-162
definition	2-2	D_FREXP instruction	6-158
IL_Opcode	2-2	D_GE instruction	6-156
COLORCLAMP instruction	6-105	D_LDEXP instruction	6-161
column sharing	1-2	D_LT instruction	6-157
comment delimiter	3-4	D_MUL instruction	6-160
common register	3-2	D_MULADD instruction	6-162
comparison instruction notes	6-1	D_NE instruction	6-154
component		D2F instruction	6-96
definition	2-8	data	
first	2-4	interpolated	4-6
forced to zero	3-3	texture coordinate	4-7
fourth	2-5	DCL_CB instruction	6-23
IL_Rel_Addr	2-8	DCL_INDEXED_TEMP_ARRAY instruction	6-27
second	2-5	DCL_INPUT instruction	6-27
third	2-5	DCL_INPUTPRIMITIVE instruction	6-29
vector written	3-3	DCL_LDS_SHARING_MODE instruction	6-35
component_w_a		DCL_LDS_SIZE_PER_THREAD instruction	6-35
definition	2-4	DCL_LITERAL instruction	6-30
IL_Dst_Mod	2-4	dcl_literal statement	2-5
component_x_r		DCL_MAX_OUTPUT_VERTEX_COUNT	
definition	2-4	instruction	6-31
IL_Dst_Mod	2-4	DCL_NUM_THREAD_PER_GROUP instruction	6-36
component_y_g		DCL_ODEPTH instruction	6-32
definition	2-4	DCL_OUTPUT instruction	6-33
IL_Dst_Mod	2-4	DCL_OUTPUT_TOPOLOGY instruction	6-32
component_z_b		DCL_RESOURCE instruction	6-43
definition	2-4	DCL_SHARED_TEMP instruction	6-34
IL_Dst_Mod	2-4	DCL_VPRIM instruction	6-34
component-wise write mask	3-3	DCLARRAY instruction	6-22
conjunction with IL_Rel_Addr token	2-6	DCLDEF instruction	6-24
CONST_BOOL	4-1	DCLP	
CONST_BUFF	4-1	I instruction	6-36
CONST_FLOAT	4-1	DCLPI	
CONST_INT	4-1	instruction	4-6, 4-7, 4-8, 4-9, 4-12
CONTINUE instruction	6-9	DCLPIN	
CONTINUE_LOGICALNZ instruction	6-10	instruction	4-5, 6-39
CONTINUE_LOGICALZ instruction	6-10	DCLPP	
CONTINUEC instruction	6-10	instruction	4-5, 6-41
control		DCLPT instruction	6-42
definition	2-2	DCLV instruction	6-44
IL_Opcode	2-2	DCLVOUT	4-4
control specifier	3-2	instruction	4-5, 4-6, 4-7, 4-8, 4-9, 4-10, 6-46
conversion instruction notes	6-2	declare an object	2-1
COS instruction	6-106	declare resources, register	2-1
COS_VEC instruction	6-150	declare resources, sampler	2-1
CRS instruction	6-106	DEF instruction	6-25
CUT instruction	6-48	DEFAULT instruction	6-11
D		DEFB instruction	6-26
D_ADD instruction	6-159		
D_DIV instruction	6-161		
D_EQ instruction	6-155		

ATI STREAM COMPUTING

definition		
_abs	3-4	
_bias	3-4	
_bx2	3-4	
_divcomp	3-4	
_invert	3-4	
_neg	3-4	
_sign	3-4	
_x2	3-4	
abs	2-8	
address_register	2-8	
bias	2-7	
clamp	2-4, 2-8	
code	2-2	
component	2-8	
component_w_a	2-4	
component_x_r	2-4	
component_y_g	2-4	
component_z_b	2-4	
control	2-2	
dimension	2-3, 2-5	
divComp	2-8	
extended	2-4, 2-5	
immediate_present	2-4, 2-5	
Intermediate Language	1-1	
invert	2-7	
kernel	1-2	
loop_relative	2-8	
modifier_present	2-3, 2-5	
negate_w_a	2-7	
negate_x_r	2-7	
negate_y_g	2-7	
negate_z_b	2-7	
pri_modifier_present	2-3	
register_num	2-3, 2-5	
register_type	2-3, 2-5	
relative_address	2-3, 2-5	
sec_modifier_present	2-3	
shift_scale	2-4	
sign	2-7	
swizzle	3-4	
swizzle_w_a	2-7	
swizzle_x_r	2-7	
swizzle_y_g	2-7	
swizzle_z_b	2-7	
x2	2-7	
delimiter comment	3-4	
DEPTH register type	4-14	
destination		
index	2-3	
modifier	3-3	
modifier token	2-4	
operand	2-4	
token	2-3	
DET instruction	6-107	
dimension		
definition	2-3, 2-5	
IL_Dst	2-3	
IL_Src	2-5	
DISCARD_LOGICALNZ instruction	6-50	
DISCARD_LOGICALZ instruction	6-50	
DIST instruction	6-108	
DIV instruction	6-109	
divComp		
definition	2-8	
IL_Src_Mod	2-8	
double precision instruction notes	6-3	
DP2 instruction	6-111	
DP2ADD instruction	6-110	
DP3 instruction	6-111	
DP4 instruction	6-112	
DST instruction	6-113	
DSX instruction	6-114	
DSY instruction	6-115	
DXSINCOS instruction	6-116	
E		
ELSE instruction	6-11	
EMIT instruction	6-50	
EMIT_THEN_CUT instruction	6-51	
END instruction	6-13	
ENDFUNC instruction	6-13	
ENDIF instruction	6-14	
ENDLOOP instruction	6-14	
ENDMAIN instruction	6-12	
ENDSWITCH instruction	6-12	
enumeration type		
IL_OUTPUT_TOPOLOGY	5-3	
IL_TOPOLOGY	5-3	
ILAddressing	5-12	
ILAnisoFilterMode	5-11	
ILCmpValue	5-9	
ILComponentSelect	5-3	
ILDefaultVal	5-7	
ILDivComp	5-8	
ILElementFormat	5-13	
ILImportComponent	5-6	
ILImportUsage	5-4	
ILInterpolation	5-12	
ILLanguageType	5-1	
ILLogicOp	5-8	
ILMatrix	5-3	
ILMipFilterMode	5-12	
ILModDstComp	5-4	
ILNoiseType	5-12	
ILOpCode	5-13	

ILPixTexUsage	5-10	H	
ILRegType	5-2	hierarchical threading model	1-2
ILRelOp	5-8	HOS rendering	4-2, 4-3
ILShader	5-1	I	
ILShiftScale	5-7	I_ADD instruction	6-84
ILTexCoordMode	5-10	IAND instruction	6-83
ILTexFilterMode	5-11	IEQ instruction	6-88
ILTexShadowMode	5-11	IF_LOGICALNZ instruction	6-17
ILTopologyType	5-3	IF_LOGICALZ instruction	6-17
ILZeroOp	5-9	IFC instruction	6-15
EQ instruction	6-117	IFNZ instruction	6-16
example		IGE instruction	6-88
source token	2-9	IL	
EXN instruction	6-117	binary stream	3-1
EXP instruction	6-118	compiler	1-1
EXP_VEC instruction	6-118	memory access model	1-2
export vertex shader	4-7	opcode	1-1
EXPP instruction	6-119	restriction	1-1
extended		shader	1-1
definition	2-4, 2-5	text syntax	3-1
IL_Dst	2-4	translator	3-1
IL_Src	2-5	IL statements used in any shader	2-1
F		IL types used in any shader	2-1
F2D instruction	6-97	IL_Dst	
FACE register type	4-13	dimension	2-3
FACEFORWARD instruction	6-119	extended	2-4
FENCE instruction	6-82	immediate_present	2-4
first component	2-4	modifier_present	2-3
flat shading	4-8	register_num	2-3
floating point arithmetic instructions	2-6	register_type	2-3
flow control instruction notes	6-2	relative_address	2-3
FLR instruction	6-120	IL_Dst_Mod	
FOG		clamp	2-4
register	4-5, 4-10	component_w_a	2-4
register type	4-10	component_x_r	2-4
forced to zero, component	3-3	component_y_g	2-4
format instruction packet	6-1	component_z_b	2-4
format kernel	2-1	shift_scale	2-4
fourth component	2-5	IL_IMPORTUSAGE_BACKCOLOR ...	4-8, 4-9
FRC instruction	6-121	IL_IMPORTUSAGE_COLOR	4-8, 4-9
FTOI instruction	6-94	IL_IMPORTUSAGE_FOG	4-10
FTOU instruction	6-95	IL_IMPORTUSAGE_GENERIC	4-6, 4-7
FUNC instruction	6-15	IL_IMPORTUSAGE_POINTSIZE	4-4
FWIDTH instruction	6-121	IL_IMPORTUSAGE_POS	4-4
G		IL_OpCode	
GE instruction	6-122	token 2-1, 2-3, 2-4, 6-1, 6-5, 6-7, 6-40, 6-45,	
generic token	2-1	6-47, 6-53, 6-68, 6-69, 6-72, 6-75, 6-76,	
Generic_Memory register type	4-18	6-77, 6-78, 6-79	
Global register type	4-15	IL_Opcode	2-2
globally shared register	1-2	code	2-2
group thread	1-2	control	2-2

ATI STREAM COMPUTING

pri_modifier_present	2-3	divComp	2-8
sec_modifier_present	2-3	invert	2-7
IL_OUTPUT_TOPOLOGY		negate_w_a	2-7
enumeration type	5-3	negate_x_r	2-7
IL_REGTYPE_ABSOLUTE_THREAD_ID	4-17	negate_y_g	2-7
IL_REGTYPE_BARYCENTRIC_COORD	4-2	negate_z_b	2-7
IL_REGTYPE_DEPTH	4-14	sign	2-7
IL_REGTYPE_FACE	4-13	swizzle_w_a	2-7
IL_REGTYPE_FOG	4-10	swizzle_x_r	2-7
IL_REGTYPE_GENERIC_MEM	4-18	swizzle_y_g	2-7
IL_REGTYPE_GLOBAL	4-15	swizzle_z_b	2-7
IL_REGTYPE_INDEX	4-1	x2	2-7
IL_REGTYPE_INTERP	4-6	IL_TOPOLOGY	
IL_REGTYPE_LITERAL	4-6	enumeration type	5-3
IL_REGTYPE_OBJECT_INDEX	4-2	ILAddressing	
IL_REGTYPE_OMASK	4-14	enumeration type	5-12
IL_REGTYPE_PCOLOR	4-13	ILAnisoFilterMode	
IL_REGTYPE_PINPUT	4-5	enumeration type	5-11
IL_REGTYPE_POS	4-4	ILCmpValue	
IL_REGTYPE_PRICOLOR	4-7	enumeration type	5-9
IL_REGTYPE_PRIMCOORD	4-11	ILComponentSelect	
IL_REGTYPE_PRIMITIVE_INDEX	4-3	enumeration type	5-3
IL_REGTYPE_PRIMTYPE	4-12	ILDefaultVal	
IL_REGTYPE_PS_OUT_FOG	4-9	enumeration type	5-7
IL_REGTYPE_QUAD_INDEX	4-3	ILDivComp	
IL_REGTYPE_SECCOLOR	4-8	enumeration type	5-8
IL_REGTYPE_SHARED_TEMP	4-16	ILElementFormat	
IL_REGTYPE_SPRITE	4-4	enumeration type	5-13
IL_REGTYPE_SPRITECOORD	4-11	ILImportComponent	
IL_REGTYPE_STENCIL	4-15	enumeration type	5-6
IL_REGTYPE_TEXCOORD	4-7	ILImportUsage	
IL_REGTYPE_THREAD_GROUP_ID	4-17	enumeration type	5-4
IL_REGTYPE_THREAD_ID	4-16	ILInterpolation	
IL_REGTYPE_VOUTPUT	4-5	enumeration type	5-12
IL_REGTYPE_WINCOORD	4-12	ILLanguageType	
IL_Rel_Addr		enumeration type	5-1
address_register	2-8	ILLogicOp	
component	2-8	enumeration type	5-8
loop_relative	2-8	ILMatrix	5-3
IL_Rel_Addr token, conjunction with	2-6	enumeration type	5-3
IL_Rel_Addr token, precedes	2-6	ILMipFilterMode	
IL_Src		enumeration type	5-12
dimension	2-5	ILModDstComp	
extended	2-5	enumeration type	5-4
immediate_present	2-5	ILNoiseType	
modifier_present	2-5	enumeration type	5-12
register_num	2-5	ILOpCode	
register_type	2-5	enumeration type	5-13
relative_address	2-5	ILPixTexUsage	
IL_Src_Mod		enumeration type	5-10
abs	2-8	ILRegType	
bias	2-7	enumeration type	5-2
clamp	2-8		

ATI STREAM COMPUTING

ILRelOp		CALL_LOGICALZ	6-7
enumeration type	5-8	CALLNZ	6-6
ILShader		CASE	6-9
enumeration type	5-1	CLAMP	6-101
ILShiftScale		CLG	6-102
enumeration type	5-7	CMOV	6-103
ILT instruction	6-88	CMOV_LOGICAL	6-103
ILTexCoordMode		CMP	6-104
enumeration type	5-10	COLORCLAMP	6-105
ILTexFilterMode		comparison notes	6-1
enumeration type	5-11	CONTINUE	6-9
ILTexShadowMode		CONTINUE_LOGICALNZ	6-10
enumeration type	5-11	CONTINUE_LOGICALZ	6-10
ILTopologyType		CONTINUEC	6-10
enumeration type	5-3	conversion notes	6-2
ILZeroOp		COS	6-106
enumeration type	5-9	COS_VEC	6-150
IMAD instruction	6-85	CRS	6-106
IMAX instruction	6-86	CUT	6-48
IMIN instruction	6-86	D_ADD	6-159
immediate_present		D_DIV	6-161
definition	2-4, 2-5	D_EQ	6-155
IL_Dst	2-4	D_FRAC	6-162
IL_Src	2-5	D_FREXP	6-158
import shader	4-7	D_GE	6-156
IMUL instruction	6-86	D_LDEXP	6-161
IMUL_HIGH instruction	6-87	D_LT	6-157
INDEX	4-1	D_MUL	6-160
index		D_MULADD	6-162
absolute register	3-2	D_NE	6-154
destination	2-3	D2F	6-96
register relative	2-5	DCL_CB	6-23
source	2-5	DCL_INDEXED_TEMP_ARRAY	6-27
INDEX register type	4-1	DCL_INPUT	6-27
INDEXED_TEMP	4-1	DCL_INPUTPRIMITIVE	6-29
INE instruction	6-88	DCL_LDS_SHARING_MODE	6-35
INEGATE instruction	6-89	DCL_LDS_SIZE_PER_THREAD	6-35
INOT instruction	6-83	DCL_LITERAL	6-30
INPUT	4-1	DCL_MAX_OUTPUT_VERTEX_COUNT	6-31
input/output instruction notes	6-2	DCL_NUM_THREAD_PER_GROUP	6-36
instruction		DCL_ODEPTH	6-32
ABS	6-97	DCL_OUTPUT	6-33
ACOS	6-98	DCL_OUTPUT_TOPOLOGY	6-32
ADD	6-98	DCL_RESOURCE	6-43
AND	6-99	DCL_SHARED_TEMP	6-34
ASIN	6-99	DCL_VPRIM	6-34
ATAN	6-100	DCLARRAY	6-22
BREAK	6-3	DCLDEF	6-24
BREAK_LOGICALNZ	6-4	DCLPI	4-6, 4-7, 4-8, 4-9, 4-12, 6-36
BREAK_LOGICALZ	6-4	DCLPIN	4-5, 6-39
BREAKC	6-4	DCLPP	4-5, 6-41
CALL	6-5	DCLPT	6-42
CALL_LOGICALNZ	6-8	DCLV	6-44

ATI STREAM COMPUTING

DCLVOUT	4-5, 4-6, 4-7, 4-8, 4-9, 4-10, 6-46	IMAX	6-86
DEF	6-25	IMIN	6-86
DEFAULT	6-11	IMUL	6-86
DEFB	6-26	IMUL_HIGH	6-87
DET	6-107	INE	6-88
DISCARD_LOGICALNZ	6-50	INEGATE	6-89
DISCARD_LOGICALZ	6-50	INOT	6-83
DIST	6-108	input/output notes	6-2
DIV	6-109	integer	2-6
double precision notes	6-3	INV_MOV	6-133
DP2	6-111	IOR	6-84
DP2ADD	6-110	ISHL	6-89
DP3	6-111	ISHR	6-89
DP4	6-112	ITOF	6-95
DST	6-113	IXOR	6-84
DSX	6-114	KILL	6-49
DSY	6-115	LDS_READ_VEC	6-80
DXSINCOS	6-116	LDS_WRITE_VEC	6-81
ELSE	6-11	LEN	6-122
EMIT	6-50	LIT	6-123
EMIT_THEN_CUT	6-51	LN	6-124
END	6-13	LOAD	6-51
ENDFUNC	6-13	LOD	6-53
ENDIF	6-14	LOG	6-125
ENDLOOP	6-14	LOG_VEC	6-126
ENDMAIN	6-12	LOGP	6-127
ENDSWITCH	6-12	LOOP	6-18
EQ	6-117	LRP	6-128
EXN	6-117	MAD	6-129
EXP	6-118	MAX	6-130
EXP_VEC	6-118	MEMEXPORT	6-54
EXPP	6-119	MEMIMPORT	6-55
F2D	6-97	MIN	6-131
FACEFORWARD	6-119	MMUL	6-132
FENCE	6-82	MOD	6-133
flow control notes	6-2	MOV	6-134
FLR	6-120	MOVA	6-135
FRC	6-121	MUL	6-136
FTOI	6-94	NE	6-136
FTOU	6-95	NOISE	6-137
FUNC	6-15	NRM	6-138
FWIDTH	6-121	packet	2-1
GE	6-122	PIREDUCE	6-139
I_ADD	6-84	POW	6-140
IAND	6-83	RCP	6-141
IEQ	6-88	REFLECT	6-142
IF_LOGICALNZ	6-17	RESINFO	6-56
IF_LOGICALZ	6-17	RET	6-20
IFC	6-15	RET_DYN	6-20
IFNZ	6-16	RET_LOGICALNZ	6-21
IGE	6-88	RET_LOGICALZ	6-21
ILT	6-88	RND	6-143
IMAD	6-85	ROUND_NEAREST	6-143

ATI STREAM COMPUTING

ROUND_NEG_INF	6-144	interpolated fog data	4-10
ROUND_POS_INF	6-144	interpolated primary color data	4-8
ROUND_ZERO	6-145	inter-thread communication	1-2
RSQ	6-146	INV_MOV instruction	6-133
rsq	1-1	invert	
RSQ_VEC	6-145	definition	2-7
SAMPLE	6-58	IL_Src_Mod	2-7
SAMPLE_B	6-60	IOR instruction	6-84
SAMPLE_C	6-64	ISHL instruction	6-89
SAMPLE_C_G	6-65	ISHR instruction	6-89
SAMPLE_C_L	6-66	ITOF instruction	6-95
SAMPLE_C_LZ	6-63	IXOR instruction	6-84
SAMPLE_G	6-61		
SAMPLE_L	6-62	K	
SAMPLEINFO	6-57	kernel	
SAMPLEPOS	6-57	definition	1-2
SET	6-147	format	2-1
SGN	6-148	KILL instruction	6-49
SIN	6-148		
SIN_VEC	6-149	L	
SINCOS	6-149	language token	2-2
SQRT	6-150	LDS	1-2
SQRT_VEC	6-151	shared memory	1-2
SUB	6-151	LDS_READ_VEC instruction	6-80
SWITCH	6-19	LDS_WRITE_VEC instruction	6-81
TAN	6-152	LEN instruction	6-122
TEXLD	6-67	line-aa texture coordinate	4-11
TEXLDB	6-70	LIT instruction	6-123
TEXLDD	6-74	LN instruction	6-124
TEXLDMS	6-77	LOAD instruction	6-51
TEXWEIGHT	6-79	local data store (LDS)	1-2
TRANPOSE	6-153	local memory	1-2
TRC	6-154	LOD instruction	6-53
UDIV	6-90	LOG instruction	6-125
UGE	6-93	LOG_VEC instruction	6-126
ULT	6-93	LOGP instruction	6-127
UMAD	6-91	LOOP instruction	6-18
UMAX	6-92	loop relative address	3-2
UMIN	6-92	loop_relative	
UMOD	6-91	definition	2-8
UMUL	6-93	IL_Rel_Addr	2-8
UMUL_HIGH	6-94	LRP instruction	6-128
USHR	6-90		
UTOF	6-96	M	
WHILELOOP	6-19	MAD instruction	6-129
instruction packet	6-1	major_version	2-2
format	6-1	mask	
integer instruction	2-6	register type	4-14
Intermediate Language - definition	1-1	write	3-3
INTERP		MAX instruction	6-130
register	4-5, 4-10	MEMEXPORT instruction	6-54
register type	4-6	MEMIMPORT instruction	6-55
interpolated data	4-6		

memory access model, IL	1-2	OUTPUT	4-1
MIN instruction	6-131	owner-computes	1-2
minor_version	2-2	P	
MMUL instruction	6-132	packet instruction	2-1
MOD instruction	6-133	PCOLOR register type	4-13
modifier		perspective correct interpolation	4-8, 4-10
abs	2-6	PINPUT	
register relative	2-3	register	4-9
modifier_present		register type	4-5
definition	2-3, 2-5	PIREDUCE instruction	6-139
IL_Dst	2-3	pixel shader import	4-2, 4-8, 4-10
IL_Src	2-5	point-aa texture coordinate	4-11
modify destination	3-3	POS	
MOV instruction	6-134	register	4-5
MOVA instruction	6-135	register type	1-1, 4-4
MUL instruction	6-136	POW instruction	6-140
multipass	2-2	precedes the IL_Rel_Addr token	2-6
multipass shader	3-2	prefix	
N		o	3-2
NE instruction	6-136	v	3-2
negate modifier	2-6	pri_modifier_present definition	2-3
negate_w_a		PRICOLOR	
definition	2-7	register	4-5, 4-10
IL_Src_Mod	2-7	register type	4-7
negate_x_r		PRIMCOORD register type	4-11
definition	2-7	PRIMITIVE_INDEX register type	4-3
IL_Src_Mod	2-7	PRIMTYPE register type	4-12
negate_y_g		PSOUTFOG register type	4-9
definition	2-7	Q	
IL_Src_Mod	2-7	QUAD_INDEX register type	4-3
negate_z_b		R	
definition	2-7	RCP instruction	6-141
IL_Src_Mod	2-7	realtime	2-2
NOISE instruction	6-137	real-time shader	3-2
non-HOS rendering	4-2, 4-3	REFLECT instruction	6-142
notes		register	
comparison instruction	6-1	absolute index	3-2
conversion instruction	6-2	common	3-2
double precision instruction	6-3	declare resources	2-1
flow control instruction	6-2	FOG	4-5, 4-10
input/output instruction	6-2	INTERP	4-5, 4-10
NRM instruction	6-138	PINPUT	4-9
O		POS	4-5
o prefix	3-2	PRICOLOR	4-5, 4-10
object declaration	2-1	relative indexing	2-5
OBJECT_INDEX register type	4-2	relative modifier	2-3
opcode		SECCOLOR	4-5
IL	1-1	SPRITE	4-5
token	2-2	TEXCOORD	4-5, 4-10
open design	1-1		
operand destination	2-4		

VOUTPUT	4-6, 4-7, 4-8, 4-9, 4-10	restriction	
register restriction	4-1	IL	1-1
register type		register	4-1
Absolute	4-17	RET instruction	6-20
BARYCENTRIC_COORD	4-2	RET_DYN instruction	6-20
DEPTH	4-14	RET_LOGICALNZ instruction	6-21
FACE	4-13	RET_LOGICALZ instruction	6-21
FOG	4-10	RND instruction	6-143
Generic_Memory	4-18	ROUND_NEAREST instruction	6-143
Global	4-15	ROUND_NEG_INF instruction	6-144
INDEX	4-1	ROUND_POS_INF instruction	6-144
INTERP	4-6	ROUND_ZERO instruction	6-145
mask	4-14	RSQ instruction	6-146
OBJECT_INDEX	4-2	rsq instruction	1-1
PCOLOR	4-13	rsq_vec	1-2
PINPUT	4-5	RSQ_VEC instruction	6-145
POS	1-1, 4-4		
PRICOLOR	4-7	S	
PRIMCOORD	4-11	SAMPLE instruction	6-58
PRIMITIVE_INDEX	4-3	SAMPLE_B instruction	6-60
PRIMTYPE	4-12	SAMPLE_C instruction	6-64
PSOUTFOG	4-9	SAMPLE_C_G instruction	6-65
QUAD_INDEX	4-3	SAMPLE_C_L instruction	6-66
SECCOLOR	4-8	SAMPLE_C_LZ instruction	6-63
Shared Temp	4-16	SAMPLE_G instruction	6-61
SPRITE	4-4	SAMPLE_L instruction	6-62
SPRITECOORD	4-11	SAMPLEINFO instruction	6-57
Stencil	4-15	SAMPLEPOS instruction	6-57
TEXCOORD	4-7	sampler declare resources	2-1
Thread_Group_ID	4-17	sec_modifier_present	
Thread_ID	4-16	definition	2-3
VOUTPUT	4-5	IL_Opcode	2-3
vWinCoord	4-12	SECCOLOR	
register types	4-1	register	4-5, 4-10
register_num		register type	4-8
definition	2-3, 2-5	second component	2-5
IL_Dst	2-3	secondary color data	4-8
IL_Src	2-5	SET instruction	6-147
register_type		SGN instruction	6-148
definition	2-3, 2-5	shader	
IL_Dst	2-3	IL	1-1
IL_Src	2-5	IL types used in any	2-1
relative		import	4-7
address token	2-8	multipass	3-2
base address	3-2	real-time	3-2
loop address	3-2	type	2-1
relative_address		vertex export	4-7
definition	2-3, 2-5	shader_type	2-2
IL_Dst	2-3	shared memory, LDS	1-2
IL_Src	2-5	shared register	1-2
RESINFO instruction	6-56	Shared Temp register type	4-16

sharing		swizzle_x_r	
column	1-2	definition	2-7
wavefront	1-2	IL_Src_Mod	2-7
shift_scale		swizzle_y_g	
definition	2-4	definition	2-7
IL_Dst_Mod	2-4	IL_Src_Mod	2-7
sign		swizzle_z_b	
definition	2-7	definition	2-7
IL_Src_Mod	2-7	IL_Src_Mod	2-7
SIN instruction	6-148	syntax	
SIN_VEC instruction	6-149	[]	3-2
SINCOS instruction	6-149	IL Text	3-1
smooth shading	4-8		
source		T	
index	2-5	TAN instruction	6-152
source modifier		TEMP	4-1
_abs	3-4	TEXCOORD	
_bias	3-4	register	4-5, 4-10
_bx2	3-4	register type	4-7
_divcomp	3-4	TEXLD instruction	6-67
_invert	3-4	TEXLDB instruction	6-70
_neg	3-4	TEXLDD instruction	6-74
_sign	3-4	TEXLDM instruction	6-77
_x2	3-4	texture coordinate	
swizzle	3-4	data	4-7
token	2-6	line-aa	4-11
source token	2-4	point-aa	4-11
example	2-9	sprite	4-11
specify		TEXWEIGHT instruction	6-79
binary	3-2	third component	2-5
control	3-2	thread	1-2
SPRITE		thread group	1-2
register	4-5	Thread_Group_ID register type	4-17
register type	4-4	Thread_ID register type	4-16
sprite texture coordinate	4-11	token	
SpriteCoord	4-11	32-bit	2-1
SPRITECOORD register type	4-11	destination	2-3
SQRT instruction	6-150	destination modifier	2-4
SQRT_VEC instruction	6-151	generic	2-1
SR - globally shared register	1-2	IL_OpCode	2-1, 2-3, 2-4, 6-1, 6-5, 6-7, 6-40, 6-45, 6-47, 6-53, 6-68, 6-69, 6-72, 6-75, 6-76, 6-77, 6-78, 6-79
statement		language	2-2
dcl_literal	2-5	opcode	2-2
used in any shader	2-1	relative address	2-8
Stencil register type	4-15	source	2-4
stream, IL binary	3-1	source example	2-9
SUB instruction	6-151	source modifier	2-6
SWITCH instruction	6-19	version	2-2
swizzle	2-6	translator, IL	3-1
definition	3-4	TRANSPPOSE instruction	6-153
source modifier	3-4	TRC instruction	6-154
swizzle_w_a		type of register	4-1
definition	2-7		
IL_Src_Mod	2-7		

types used in any shader 2-1

U

UDIV instruction 6-90
 UGE instruction 6-93
 ULT instruction 6-93
 UMAD instruction 6-91
 UMAX instruction. 6-92
 UMIN instruction 6-92
 UMOD instruction 6-91
 UMUL instruction. 6-93
 UMUL_HIGH instruction 6-94
 USHR instruction. 6-90
 UTOF instruction 6-96

V

v prefix. 3-2
 vector component written 3-3
 version token. 2-2
 VERTEX 4-1
 vertex shader export 4-7, 4-8, 4-10
 vertex shader import 4-2
 VOUTPUT 4-4
 register 4-6, 4-7, 4-8, 4-9, 4-10
 register type 4-5
 VPRIM 4-1
 vTid 1-3
 vWinCoord register type 4-12

W

wavefront 1-2
 access mode. 1-2
 sharing 1-2
 WHILELOOP instruction 6-19
 write mask 3-3
 write mask, component-wise. 3-3
 written vector component 3-3

X

x2
 definition 2-7
 IL_Src_Mod. 2-7

Z

zero force component 3-3