

Chapter 5

Computer graphics and visualization

Donald P. Greenberg

The field of computer graphics has made enormous progress during the past decade. It is rapidly approaching the time when we will be able to create images of such realism that it will be possible to 'walk through' nonexistent spaces and to evaluate their aesthetic quality based on the simulations. In this chapter we wish to document the historical development of computer graphics image creation and describe some techniques which are currently being developed. We will try to explain some pilot projects that we are just beginning to undertake at the Program of Computer Graphics and the Center for Theory and Simulation in Science and Engineering at Cornell University.

5.1 How do you make an image?

To create a synthetic scene, it is necessary to perform the following five steps:

- (1) *Three-dimensional model.* The entire geometry of the environment must be mathematically defined as well as the colour of the surfaces.
- (2) *Perspective transformation.* Each vertex of the model is mathematically transformed to generate a true perspective picture on the image plane as well as retaining the correct perspective-depth information.
- (3) *Visible surface determination.* Surfaces remaining within the frustum of vision after the perspective transformation are sorted in depth so that only the elements closest to the observer are displayed.
- (4) *Light-reflection model.* This model predicts the colour and spatial distribution of the light reflected from each surface in the environment.
- (5) *Image display.* The image is rendered by selecting the appropriate red, green and blue intensities for each pixel in the visible scene.

Historically, because of the constraints of processing power, the unavailability of storage and the high cost of memory many short-cuts were taken. Simple environments were used, limited-light-reflection models were incorporated and the separate steps of the image-creation process were combined. Perhaps the most sophisticated of these results were images created using Phong (PHON73) shading.

Today, much of the processing has now been embedded in hardware. Transformations which use a 4×4 multiplication and perform all of the perspective transformations, clipping, viewpoint mapping and windowing are accomplished in hardware. Visible surface algorithms, such as the depth-buffer (Foley and Van Dam, 1982) are becoming standard. Shading models, such as the Gouraud (1971) approach which uses linear interpolation, or even the Phong approach, which can provide specular highlights, are also built in hardware. Furthermore, some parts of the process have been pipelined in more expensive display devices, such as the Iris workstation of Silicon Graphics or the Picture System of Evans and Sutherland.

Raster display manufacturers, such as Lexidata, Tektronix, Adage, Ramtek and Raster Technologies, are all now offering these routines in specialized hardware.

For some time now, very special and expensive hardware has been available for flight simulation and pilot training. These systems take advantage of a predefined environment, a hierarchical database which can be preprocessed, but primarily only the observer or a limited number of specified objects can move. It is difficult - if not impossible - to easily change the description of the environment, and thus these simulators are not generally used for scientific exploration, engineering or design applications.

What is necessary is the ability to design and model an environment, to evaluate it visually and then to interactively redesign the environment based on the evaluations. This process is iterative and must be continuous. The system in its ultimate state should be as easy to use as the yellow tracing paper frequently used in architectural design offices.

We cannot do this yet as many of the ingredients required to accomplish these objectives are still missing. Certainly we need the modelling software. We also need the image-synthesis algorithms to create pictures of sufficient realism, and lastly, we need the processing power to provide the dynamic sequences. In the following pages we will concentrate on the image-synthesis part and discuss our future directions at Cornell University.

5.2 Image synthesis

There are several fundamental flaws in most algorithms or systems which are used today for image synthesis. These deficiencies include poor light-reflection models, incorrect description of material and surface properties, lack of global illumination effects, speed (particularly for dynamic sequences) and limitations on complexity. Let us examine each of these topics more specifically.

5.2.1 Light-reflection models

Light is a very complex wavelength-based phenomenon, and reflections of light from a surface depend on many parameters. Historically, much work has been done in the fields of radiation heat transfer and in radar theory. The most sophisticated models are wavelength-based, and can with some degree of accuracy provide the spatial and spectral distribution of the reflected light (Hall, 1983). Accurate models must preserve the conservation of energy. Comprehensive theories are not yet available and although research groups are working in this direction the problem has not been completely solved.

5.2.2 Surface and light descriptions

For surface descriptions, in addition to describing the macro geometry it is necessary to describe characteristics such as roughness, size and orienta-

tion of the micro facets. Some models which depend on a probabilistic or statistically based description of the microfacets of the surface, such as the Blinn model (Blinn, 1977) or the Cook and Torrance (1982) model, have provided reasonable results. Material properties, including the spectral reflectivity on a wavelength basis, must also be available.

For a light to be completely described it is necessary to give the three-dimensional geometry and orientation of the light, as well as the spectral and spatial distribution of the energy which is being emitted. In the field of illumination engineering this is usually done through goniometric diagrams.

5.2.3 Global illumination effects

It is generally easy to recognize a computer-generated image because of the lack of global illumination. However, it is well known that the effect of intra-environment reflections can substantially enhance the quality of the image. In 1980 Whitted first introduced one technique called ray tracing. Another method was developed at the Program of Computer Graphics at Cornell called the radiosity method (Goral *et al.*, 1984). In ray tracing a ray is sent through each pixel of the image plane into the environment. At each surface intersection reflected and/or transmitted rays must be created. The final pixel colour is obtained by combining the intensity contributions from all of the reflected and/or transmitted rays. When using this approach each image is dependent on the current observer position. On the other hand, the radiosity approach determines the light-energy equilibrium of all the surfaces in a static environment independent of the observer position.

It is not possible within this short chapter to fully describe either of these methods. However, we will illustrate these two types of techniques and show the enormous potential they have for the visualization of threedimensional environments.

5.2.4 Speed

When the techniques described above were first-mentioned very few believed that these simulation endeavours would be worthwhile. Raytracing images on a large minicomputer, such as a Digital Equipment Corporation VAX 11/780, could take from 10 to 40 hours. The radiosity approach, which involves the solution of a large number of simultaneous equations, was something that should not even be tried. In fact, researchers felt that if this problem was ever to be solved, it could only be done for simple environments. However, much progress has been made since those initial reactions. In ray tracing, by using bounding volumes, hierarchical data structures, adaptive tree depths and image coherence which substantially reduces the intersection calculations for the first ray, the times can be reduced by a factor of ten to twenty (Weghorst *et al.*, 1984). This, a ten-hour image of 1980 may only take one-half hour on a standard large minicomputer today.

Perhaps even more significant is the use of the radiosity approach. Since this method is independent of the observer position, once the static environment has been computed in terms of its form factors and intensities

it is only necessary to render the image. This can be shown dramatically by looking at *Figure 5.2*. Although the first image (a) required approximately five hours of computation time, the lights could be turned out without having to recompute the form factors and geometric relationships. Thus the second image (b) required only one-half hour for image generation. Even more dramatic is the third image (c), where the observer position is changed and only the rendering process has to be repeated. The computation time for rendering has been reduced to approximately fifteen minutes on our VAX 11/780.

5.2.5 Complexity limitations

More computer-generated images to date are limited in terms of environmental definition, detail, texture and the number of light sources in the scene. Constraints of hardware, or the available processing power, has restricted this complexity. Unfortunately, for sophisticated lighting models, as the number of surfaces depicting the environment increases, the computational requirements increase exponentially, not to mention the increased difficulty in creating the model. Thus, for realistic simulations a substantial change in hardware and software environments is necessary.

5.3 Recent developments at Cornell University

We are very fortunate at Cornell University to have excellent computer resources. The recently established Center for Theory and Simulation in Engineering and Science has been designated as a national computing centre with a specific mission to experiment with large-scale parallel processing. Initially, we will be using an IBM 3084QX machine with five Floating Point Systems 264 array processors operating in parallel. In a few years we expect to have an array of parallel machines with gigaflop capability. We are also working on the creation of a powerful graphics environment to surround these supercomputers, so that scientists, engineers and designers can interactively monitor their simulations. We are now adding a Floating Point Systems 264 array processor to our graphics laboratory to help with this prototype workstation development. This device will provide a peak capacity of approximately three hundred times the processing power that is available on a large minicomputer. Thus, if we can vectorize the equations properly, it should be possible to get images of the quality shown in *Figures 5.1* and *5.2* in the neighbourhood of ten seconds. Our ultimate goal is to be able to generate the images fast enough for us to simulate walking through the hypothetical space in real time.

To conclude, we strongly believe that these techniques will be used in the design profession before the end of this decade. In order to properly evaluate architecture, one cannot really look at static synthetic images, but one must have the feeling of moving through the space. Architecture is a spatially oriented profession and with today's more sophisticated algorithms which will allow us to model light, shadow, colour, texture and the




Figure 5.2. This high-resolution, ray-traced image of the new Performing Arts Center at Cornell University, designed by Stirling-Wilford Associates but not yet constructed, was created by using Cornell's testbed-modelling system for realistic image synthesis. The light reflections, transparency, shade and shadows, and the contextual information, illustrate how these powerful simulation techniques can be used for architectural design evaluations. (See also the cover illustration)




Figure 5.2. See overleaf




Figure 5.2. These three images were created using the radiosity approach. Although computation for the form factors and surface intensities requires several hours, since the algorithms are independent of observer position subsequent images for dynamic motion can be generated in minutes. The approach is significant in that it correctly models the soft shadows of penumbras and the colour-bleeding effects of diffuse environments, phenomena which can only be modelled using the radiosity approach

affect of intra-environmental surroundings, we should be able to obtain simulations fast enough and which are realistic enough to allow us to make our aesthetic evaluations. The field is extremely exciting and the author hopes to be able to continue to be part of this development process.

Acknowledgements

The author would like to express his appreciation for the enormous support and participation of the entire staff and students of the Program of Computer Graphics. All of the systems development, programming and modelling was performed by this group of dedicated individuals. The support of the National Science Foundation and Digital Equipment Corporation is also gratefully acknowledged.

References

- BLINN, J. (1977). Models of Light Reflection for Computer Synthesized Pictures. *SIGGRAPH 1977 Proceedings, Computer Graphics 11*, No. 2, 192-198
- COOK, L. and TORRANCE, K. E. (1982). A Reflection Model for Computer Graphics *ACM Transactions on Graphics 1*, No. 1, 7-24
- FOLEY, J. D. and VAN DAM, A. (1982). *Fundamentals of Computer Graphics*, Reading, Mass.: Addison-Wesley
- GORAL, M., TORRANCE, K. E., GREENBERG, D. P. and BATTAILE, B. (1984). Modeling the Interaction of Light Between Diffuse Surfaces *ACM Computer Graphics (Proceedings)* 213-222
- GOURAUD, H. (1971). Computer Display of Curved Surfaces. PhD Dissertation, University of Utah
- HALL, R. A. (1971). A Methodology for Realistic Image Synthesis, MS Thesis, Cornell University
- WEGHORST, H., HOPPER, G., and GREENBERG, D. P. (1984). Improved Computational Method for Ray Tracing *ACM Transactions on Graphics* January, 52-69
- WHITTED, T. (1980). An Improved Illumination Model for Shaded Display, *Communications of the ACM*