

30. Visualising ancient Greece: computer graphics in the Sacred Way Project

J. Cornforth

C. Davidson

Mechanical Intelligence, 922, Grange Hall Road, Cardiff-by-the-sea, CA 92007, U.S.A.

C.J. Dallas

Benaki Museum, Athens, Greece.

G.R. Lock

Institute of Archaeology, University of Oxford, 36, Beaumont Street, Oxford, OX1 2PG, U.K.

30.1 Background

The Sacred Way project is a cooperative international initiative for multimedia education in archaeology and Classical civilisation. The project is currently supported by the European Community COMETT II programme. Planned activities include seminars and courses, intended to introduce archaeologists, museum professionals and humanities educators to the techniques and potential of digital multimedia as a tool for training, education and dissemination of information. The project aims to produce a multimedia application, combining the technology of Compact Disc-Interactive (CD-I) with state-of-the-art photorealistic graphics and with advanced authoring and hypermedia techniques. This paper presents work in progress for the Sacred Way application, and discusses some of the issues involved, in the field of computer graphics.¹

CD-I, a state-of-the art digital multimedia format based on the familiar CD-Audio physical medium, belongs to a class of digital technologies that are expected to capture a large part of the market of the analog Interactive Videodiscs (IV). Besides, CD-I in particular is poised to expand interactive multimedia to the general consumer market. Although this is not the place to discuss the details of CD-I technology (cf. Preston 1988; Bastiaens 1989), a summary of its advantages over competing technologies, especially IV, is of interest:

- CD-I is being developed to a world standard so that any disc will be playable anywhere in the world;
- CD-I operates in a digital environment integrating video, audio and other digital data on the same medium;
- like CD-Audio, CD-I is a self-contained system which does not rely on external computing power.

CD-I is to be launched on the consumer electronics market in the very near future² as value-added CD-Audio at an initial price of around US\$ 1,000. It has the potential, therefore, of penetrating commercial, educational and domestic markets that have been denied to IV because of its prohibitive cost.

Another important aspect of CD-I is its place in the growing arena of electronic publishing with discs being described as the 'books of the future'. The projected scenario is that within a few years, High Street bookshops will stock CD-I discs covering a whole range of subject areas. These can be played on a television set in a truly interactive way to maximise the integration of the moving video, still pictures, computer graphics, text and audio stored on the disc. Due to the seamless integration of different types of material, including text, images, video and data, and provided that it is commercially successful, this technology has great potential for archaeology. Possible applications include training discs in areas such as excavation methods, point-of-information systems in museums and sites, and electronic publication of excavation reports (Lock & Dallas 1990).

The Sacred Way application may be seen as a tutorial training system in archaeological method and theory, as a hypertext-based knowledge base on Classical Greek archaeology, society and history, or as a photorealistic computer simulation of a complex Classical site. Sacred Way CD-I materials will be aimed at a range of different potential audiences, including educational institutions, from secondary schools to undergraduate education, museums and heritage centres, and, not least, tourism and the domestic 'armchair travel' market.

The subject-matter of the application is the Sacred Way, the main road to Athens, cultural and political centre of the Greek world in Classical times, from the important sanctuary of Demeter at Eleusis, site of the Eleusinian mysteries (Mylonas 1961). The CD-I discs will incorporate textual information, including original Greek sources, drawings and plans, photographs, audio narrative, video and photorealistic animation, to illustrate relevant aspects of:

- a) the architectural development of the Sanctuary of Eleusis and selected parts of other sites, such as the Athenian agora and the Kerameikos cemetery,
- b) the history and culture of Classical Athens, illustrated by subjects such as the Panathenaic procession, the building programme of Pericles and the Eleusinian mysteries, the major Panhellenic cult up to the advent of Christianity, and

1. An earlier version of this paper was published in the *Proceedings of the Seminar on Data Visualisation and Archaeology: Visual Reconstruction from Partial Data*, British Computer Society Conference, Documentation Displays Group (London, 5th December 1990).

2. The U.S. launch took place in October 1991, the Japanese and European launches are scheduled for 1992.

- c) the methods and techniques of contemporary archaeology, illustrated by means of actual problems encountered in a complex archaeological site.

The basic user experience will be that of a 'surrogate walk' through a computer simulation of the reconstructed site of Eleusis following the route that a modern visitor takes through the ruins of today. Eleusis is a large sanctuary site associated with the goddess Demeter, her daughter Persephone and the Mysteries, a series of unknown ceremonies that took place in the temple. The walk can be interrupted in order to access information on any particular aspect of the site: comparisons with other sites, details of artefacts and architectural methods, reconstructions of aspects of everyday life such as music, drama, religious and burial practices, information on archaeological methods that have resulted in the acquisition of much existing knowledge concerning the past.

30.2 Computer graphics and Eleusis

The computer graphics goal in the Sacred Way Project is to produce images for a guided tour of the sanctuary of Eleusis, possibly at different time periods. As with many other projects involving computer graphics and other disciplines (e.g. art, film) this will be a cooperative effort, requiring the expertise of archaeologists, computer scientists and CD media designers. To produce the tour it will be necessary to construct a 3D solid model of selected areas of the site then to render and animate realistic images.

30.2a The 3D model

There are many considerations involved in the construction of a 3D solid model of a complex site like the Sanctuary of Demeter at Eleusis. The site witnessed many stages of construction, was occupied over many centuries and covered a wide area, so it is important to determine what exactly is to be modeled. A natural inclination is to model everything in the site down to the detail of single stone blocks, and perhaps make one complete model for each of the major phases of the site. While this would result in a spectacular model which could be used for a variety of purposes, it is not necessary for our guided tour. Even if we knew exactly what everything looked like and had enough computer power to store the data and reconstruct and render images of it, such a model would contain data we would not want to include in our tour, as well as detail so minute that a viewer could not possibly see it all.

Practically, the hardware we plan to use could not handle such a complex model. The simple model of room N7 and the associated courtyard at the Roman villa of Fishbourne (Cornforth & Davidson 1989; Haggerty 1990) required about 100 Kbytes for the model and an additional 15 Mbytes for textures. As an estimate, a single brick requires 256 bytes to specify its size, shape, location, texture and colour. This means that 4000 bricks would use 1 Mbyte of memory to specify the model. The hardware we are going to use will be based on the Inmos T-9000 transputer processor. Each processor shall have between 8 and 32 Mbytes of memory. With four processors we then have

an upper limit on model complexity of 128,000 to 512,000 objects, less the memory required for textures.

The easiest way to simplify the modelling procedure is to restrict the model to the minimum structures required for the tour. The archaeologists will choose what is to be represented through an iterative process. To start, everything in the site will be reduced to blocked bounding boxes. The general path of the tour will be chosen and those areas visible along this route will be refined further. This process will continue until the exact path is defined and the visible areas are modelled (Fig. 30.1).

These visible areas will probably be too complex to be all included in a single model, so smaller models, or 'sets', can be created and used much like scenes in a film. Modelling only what is visible to the viewer and using a series of sets increases the efficiency of the computer and enables us to spend more time generating images.

Another way to simplify is to represent the model with as few objects as possible. For instance, a wall made up of bricks would be one object, as opposed to each brick in the wall being a single object. Detail, such as wall decoration can be added using a technique called texture mapping which does not add to the overall complexity of the model. Reducing the number of objects not only makes it easier to construct and manipulate the model, it also speeds up the rendering of the images.

Several commercial modellers are available to construct the model and we want to find one that meets the following criteria for the user interface: an architectural interface as opposed to engineering, the use of object-oriented hierarchical data structures and a sculpture tool.

Most commercial modellers are aimed at the mechanical CAD/CAM market and as such are optimized for the creation of objects which can be machined using automated tools. The few systems that are used for architectural rendering are oriented more towards the modern, mechanical look of steel and glass, while the Classical buildings of ancient Eleusis look anything but mechanical since they are mostly made of stone, terracotta and wood.

An object-oriented interface is one in which the user can create and manipulate objects in a more intuitive, natural way. For instance, the user works with walls and columns not with rectangular prisms, hyperboloids and tori. Also, an object, once defined as a basic shape, can be copied and modified later. In this way common features such as walls, doorways, and columns can be stored as basic objects in libraries. Starting with a basic object, such as a column, it is easy to then change it to include the appropriate base and capital, material, paint colour, fluting, etc.

The hierarchical organization of data allows spatially-related objects to be grouped and stored together and for the user this means it is easier to work with the model during the refinement process. Also, data stored in this way reduces the time it takes to

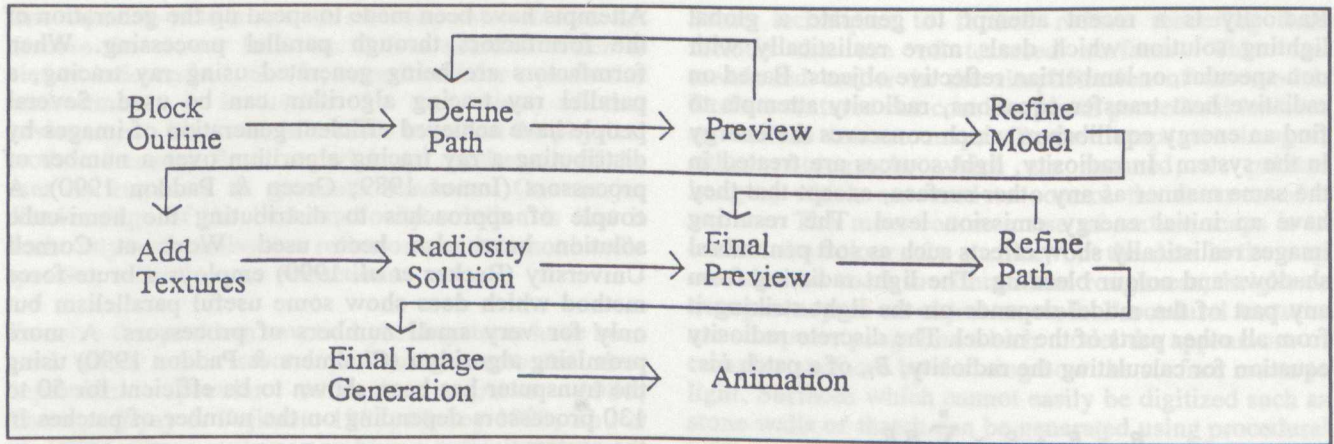


Figure 30.1: Process of defining a tour.

render the images.

What is most lacking for computer modelling at this point is a suitable 3D sculpting tool. Archaeological sites throughout the world have important sculptures, and such non-geometric objects are difficult to model. The current approaches to creating complex forms include digitizing objects (Reeves 1989) and 3D modelling (Williams 1990; Pentland *et al.* 1990).

A problem involved with the digitizing method is that objects may not exist or be available to scan. Also, once a sculpture is built and digitized, as in Pixar's 'Tin Toy', there is still extensive fine tuning of the resulting patch mesh (Reeves 1989). Manipulating the patch mesh is difficult and the final image may not have a satisfactory, realistic look. Scanning by laser is a more accurate way to reproduce a 3D object but the laser scanner is quite expensive. Laser scanners have been used to reproduce the images of people quite successfully and is a method used in film special effects.³

The current interface for 3D paint (Williams 1990) and modelling systems (Pentland *et al.* 1990) requires painting a distortion map as a grey-scale rectangular image which is then applied to an irregular surface. The user feedback is not fast enough or direct enough for this approach to be widely accepted. The other problem is that in representing surfaces which are undercut, the map only allows for displacement from a surface norm. We are currently looking at creating our own sculpting tool.

30.2b Rendering

Once the models, or sets, are built they must be rendered realistically and quickly enough to make the animation sequence possible. All rendering techniques have advantages and disadvantages. The three we will consider here are polygon rendering, ray tracing and radiosity.

In polygon rendering, all objects are decomposed into polygons representing the surfaces of the objects and

the polygons are then rendered. Polygon rendering has been popular for years and there are many algorithms available to quickly render the polygon surfaces. To smooth the polygon edges, Phong and Gouraud shading are used, and can be accelerated using hardware. The drawback of polygon rendering techniques is the difficulty of generating realistic shadows and reflections. These details must be added as textures (Pixar 1989).

Ray tracing is a method which generates impressive computer graphic images. It combines hidden surface removal, shading due to direct and indirect illumination, shadow generation and both reflection and refraction of light (Kay & Greenberg 1979; Whitted 1980). To determine the colour of a pixel of an image, a ray is traced from the eye, through the pixel, until it intersects with an object in the model being rendered. Once an object is intersected, calculations are performed to determine the colour of the object due to direct illumination. If the object is partially reflective, partially refractive or both, additional rays are then generated starting from the surface, and the colour contributions due to those rays are added to the pixel colour of the image. Since the rays observe the same physical laws as light rays, such as reflection and refraction, striking images can be generated.

Simple ray tracing programs use lines, or rays, to test for object intersections. These rays originate from the centre of the display pixels, and are infinitely fine. With only one ray per pixel, several anomalies are introduced by the low sampling frequency. These anomalies result in the image having an unnatural crispness with sharp edges to all the objects including the shadows. Small objects can also fall in the spaces between sample rays and so not be counted in the image. When such an image is animated, sometimes a small object will be found by the rays and will appear occasionally. The result is a scintillation of small objects, similar to the twinkling of stars. Several approaches have been taken to 'soften' the image (Cook *et al.* 1984), but all are computationally very expensive, relying on over sampling to reduce aliasing or introduce a depth-of-field.

3. The film *Star Trek IV* incorporates laser scanned special effects. Work is being done on computer generated images of humans, see ACM SIGGRAPH 88 Course Notes, "Synthetic Actors: The Impact of Artificial Intelligence & Robotics on Animation", Thalmann, D, Magnenat Thalmann, N.

Radiosity is a recent attempt to generate a global lighting solution which deals more realistically with non-specular, or lambertian reflective objects. Based on radiative heat transfer equations, radiosity attempts to find an energy equilibrium which conserves the energy in the system. In radiosity, light sources are treated in the same manner as any other surface, except that they have an initial energy emission level. The resulting images realistically show effects such as soft penumbral shadows and colour bleeding. The light radiating from any part of the model depends on the light striking it from all other parts of the model. The discrete radiosity equation for calculating the radiosity, B_i , of a patch i is:

$$B_i = E_i + R_i \times \sum_{j=1}^n B_j F_{ij}$$


Where E_i is the emission of the patch, R_i is the reflectivity of patch i , B_j is the radiosity of patch j and F_{ij} is the Formfactor between patches i and j , (Cohen 1989). The formfactor between two patches depends on distance between the patches, the orientation of the patches, and the percentage of the patches that is visible from each other. The formfactor takes into account the inverse square law of energy radiation and the fall off in energy transfer between the patches as the line joining the patches approaches a tangent to the surfaces. There is no easy solution for the generation of the formfactor, although several methods have been used including the Hemi-cube (Cohen & Greenberg 1985) and ray tracing by Wallace (Wallace *et al.* 1989) and Sillion (Sillion & Puech 1989) (Fig. 30.2).

Once the model is divided into N patches, it can be translated into N equations as above. These equations can also be represented in matrix form as a set of linear equations. Once the equations are in matrix form, it is a matter of solving the N equations in N unknowns using any number of standard equation solvers to get the resulting equilibrium radiosity of the model.

Attempts have been made to speed up the generation of the formfactors through parallel processing. When formfactors are being generated using ray tracing, a parallel ray tracing algorithm can be used. Several people have achieved efficient generation of images by distributing a ray tracing algorithm over a number of processors (Inmos 1989; Green & Paddon 1990). A couple of approaches to distributing the hemi-cube solution have also been used. Work at Cornell University (Recker *et al.* 1990) employs a brute-force method which does show some useful parallelism but only for very small numbers of processors. A more promising algorithm (Chalmers & Paddon 1990) using the transputer has been shown to be efficient for 50 to 130 processors depending on the number of patches in the system.

In addition to just adding processing power to generate the formfactors, a progressive refinement method of generating the image is also used (Cohen *et al.* 1988). Traditional radiosity requires that all the formfactors be calculated first, and then the system of equations be solved. The solution for the patch i is determined by 'gathering' all the light reaching the patch from all other patches in the system. If the scene could be viewed as each patch is solved, it would change one patch at a time. The refinement method uses the reciprocity of formfactors to 'shoot' light from patch i to all other patches. These patches add the light to their solution combining the calculation of formfactors with the matrix solution. This approach has the advantage that intermediate solutions can be generated and viewed as the algorithm progresses. When both 'shooting' and 'sorting' are used, an image close to the final image can be generated after only a small number of patches have been processed. Because of the amount of data being exchanged at each step of the solution, this has been shown to be less tractable to parallel processing (Chalmers & Paddon 1990).

Figure 30.2: The brightness of the patch on the vase is determined by the amount of light reaching the patch from all other visible surfaces and light sources in the model.



One aspect to note is that the radiosity solution does not involve the viewer location. It is a solution which is viewer independent, that is, once found for a particular environment, it is usable for any view of that environment. Since the solution is based on the physical model of energy flow, no 'trickery' on the part of the user is required to generate shadows. One major disadvantage of the basic radiosity solution is that it deals only with lambertian reflectors, and cannot handle specular reflections or refractions.

To solve this problem several hybrid systems have been used which use a second pass to generate the view dependent specular reflections (Arvo 1986; Chattopadhyay & Fujimoto 1987; Malley 1988; Zhu *et al.* 1988). These systems combine a radiosity first pass with a ray tracing view-dependent second pass. The radiosity solution of the first pass can also be generated using ray tracing for a distributable uniform solution. The advantage of the two pass approaches is that a global solution for the lighting can be calculated and a quick reflection pass can be run for each frame of the animation. Currently our choice for the animation sequence is a two pass radiosity approach. At first glance this may seem an odd choice, given the computational expense of the radiosity approach. There are several reasons for this choice including realism, user interface and the view-independent global lighting solution.

The realism possible using radiosity will help to make the animation sequence as believable as possible. As mentioned above, radiosity allows for the full range of effects such as inter-reflection between surfaces and soft shadows. It also makes possible atmospheric effects such as fog and smoke by introducing what is called participatory media (Rushmier & Torrance 1987). The participatory media introduce volumes of matter which transmit as well as absorb and re-radiate selected colours of light. Although it is very time-consuming to generate a radiosity lighting solution, since the solution is view-independent, it is possible to render the resulting images in real-time (Airey *et al.* 1990). The restriction of this is the requirement that only the point of view changes, exactly what is needed for the tour. This means it is possible to generate a high resolution preview of the tour in near real-time. The preview will be generated without the ray trace post pass, but otherwise will look like the final animation.

The user interface of a radiosity or a ray trace system is more natural since the actual generation of the image more closely follows the physical laws. The laws of energy transfer dictate where the shadows and reflections appear, they are not generated using special effects such as shadow or environmental maps (Pixar 1989).

Given our goals of generating animation of a viewpoint through a fixed model, the only disadvantage of the radiosity method is that reflective objects must be handled using a view-dependent rendering method such as ray tracing. For this reason we plan to use a ray trace pass to generate the images only for the final animation.

Some techniques to further reduce rendering time include the use of textured surfaces. This was mentioned earlier in the simplification of the model. Surface textures can consist of both procedural textures as well as texture maps. Texture maps are images, either photographs which are scanned, or painted images, which are used to control the features of a surface. The most common use of texture maps is to control the colour of a surface, but the data in the map can also be used to control the roughness, height or transparency of an object as well. Procedural textures are computer programs which model the appearance of complex surfaces by defining how the surface reacts to light. Surfaces which cannot easily be digitized such as stone walls or thatch can be generated using procedural textures. Procedural textures and texture maps are commonly used together to define complex surfaces.

30.2c Animation

The goal of the graphics work in the Sacred Way Project is to produce about 5 minutes of animation of walking through Classical Eleusis. An animated sequence consists of a series of images, each slightly different, which are displayed in rapid succession. The rate the images must be displayed is called the frame rate and is either 25 or 30 frames per second (fps) depending on the television standard used. To simulate movement, such as walking through the model, the viewpoint between any two frames must change by an amount which depends on the frame rate and rate of motion. For instance, to simulate movement at 1 meter per second for a 25 fps animation, the viewpoint should move 4cm.

The CD-I disc can generate images for all TV standards currently in use (i.e. PAL, SECAM and NTSC), but the images must be generated for the NTSC frame rate of 30 fps at the PAL or SECAM resolution. This translates to 9000 images to be rendered for the animation. Rather than having to specify the exact view location for each frame, it is possible to specify 'key' frames and have the computer interpolate movement between the frames. This interpolation is called *inbetweening* — the generation of frames in between the key frames.

To simulate walking through the site, the camera, or point of view, is the only thing that will be moving. This allows us to use radiosity economically, since the global lighting solution need be calculated only once. To animate the camera we will use seven control variables: X, Y, Z location, pan and tilt angles, focal length and the time. This allows us to place the camera at a specific time and place, pointed at a specific object. Each of the points defined by the seven variables will be joined by a Catmull-Rom B-spline. The time parameter is used by the system to generate an interpolative curve to access the remaining 6 variables. The Catmull-Rom spline is an interpolative curve, it passes through the control points, and is continuous in the second derivative which assures the curve is smooth. Using this control sequence it will be possible to specify the camera location and view in a natural method and then have the computer generate a smooth path through the points required.

Our current estimate is that it will require between 1 and 5 minutes to generate each frame of final animation, once the radiosity solution has been calculated. The radiosity solution should require 1 to 2 days based on our estimate of the complexity of the model. Once in production we will be able to produce one second of animation every 30 minutes to 2.5 hours. The total time for rendering the animation will be 150 to 750 hours, probably the upper limit, spread over 10 weeks. This schedule and time budget is typical of a five minute animation project.

30.3 Conclusion

For the majority of users, the 'surrogate walk' through a photorealistic model of the site of Eleusis will be the primary access point for information retrieval from the Sacred Way CD-I application. In computer graphics work related to the project, we are striving to produce as believable a simulation of an actual trail through the site, in its original state, as possible. Although we are limited by the capabilities of the reproduction medium and the practicalities of storage and animation efficiency, the combination of some of the techniques illustrated above can give results of striking visual realism. This realism will make the application more attractive to use, and thus more effective in its training, educational and recreational mission.

A word of caution is, however, in order at this point. 'An image is worth a thousand words' not least because it is readily accepted by viewers as indisputable proof of reality. The problem caused by the application of very effective graphical reconstructions for historical-archaeological ethics is related to the fact that very rarely is the form and history of ancient buildings beyond scholarly dispute; in fact, the sanctuary of Demeter presents archaeologists with a number of intriguing problems of phasing and architectural form. The 'innocent' user, who is more likely to access information through the 'surrogate walk' model, is also the most likely prey to this deceptive power of images.

We would like to avoid the 'pink cement' approach to visual reconstruction, which would detract from the appeal and immediacy of navigating within the graphic model. We intend, however, to provide the necessary caution to the users of the Sacred Way discs by exploiting the capabilities of CD-I: using a concurrent and windowed environment, whereby the 'surrogate walk' shares the screen with contextual information and is supplemented by audio commentary, providing where appropriate alternative reconstructions of major architectural monuments, and pre-empting the appearance of tentative reconstruction elements using appropriate visual cues. This issue is, however, intimately linked with the overall user interface and interaction model that we envisage for the Sacred Way application. Its full discussion falls, therefore, outside the scope of the current presentation.

The production of a CD-I application in the archaeology of a complex Classical site involves special problems for the computer graphics researcher. These problems range from the interactivity in all stages of graphics work necessitated by the interdisciplinary nature of the project, to the special solutions required

to deal with the complex form and the texture of Greek building materials, monumental masonry and sculpture. The present paper is an overview of the computer graphics procedures and technical solutions adopted for the Sacred Way Project, from the construction of the computer model, to the final stages of photorealistic rendering and animation of images.

References

- AIREY, J.M., J.H. ROHLF & F.P. BROOKS JR. 1990. "Towards Image Realism with Interactive Update Rates in Complex Virtual Building Environments", *Computer Graphics*, 1990 Symposium on Interactive 3D Graphics Proceedings, 24(2, March): 41-50.
- ARVO, J. 1986. *Backwards Ray Tracing, Developments in Ray Tracing*, SIGGRAPH Course Notes, Vol. 12.
- BASTIAENS, G. 1989. "Compact Disc-Interactive: a multimedia system for entertainment, education and information in the Nineties", unpublished paper presented to the Interactive Media Conference, October.
- CHALMERS, A.G. & D.J. PADDON 1990. "Parallel Radiosity Methods", *Transputer Research and Applications*, 4, NATUG Proceedings, (October): 183-193.
- CHATTOPADHYAY, S. & A. FUJIMOTO 1987. "Bi-Directional Ray Tracing", *Computer Graphics*, Proceedings of CG International '87, Tokyo, Springer-Verlag.
- COHEN, M.F. 1989. *Basic Radiosity Formulation for Diffuse Reflections, Radiosity*, SIGGRAPH Course Notes, 18.
- COHEN, M.F., S.E. CHEN, J.R. WALLACE & D.P. GREENBERG 1988. "A Progressive Refinement Approach to Fast Radiosity Image Generation", *Computer Graphics*, SIGGRAPH '88 Proceedings, 22(3, August): 75-84.
- COHEN, M.F. & D.P. GREENBERG 1985. "A Radiosity Solution for Complex Environments", *Computer Graphics*, SIGGRAPH '85 Proceedings, 19(3, July): 31-40.
- COOK R.L., T. PORTER & T. CARPENTER 1984. "Distributed Ray Tracing", *Computer Graphics*, SIGGRAPH '84 Proceedings, 18(3, July): 137-145.
- CORNFORTH, J. & C. DAVIDSON 1989. "Picturing the Past", *Archaeological Computing Newsletter*, 19 (June): 6-10.
- GREEN, S. & D. PADDON 1990. "A Highly Flexible Multiprocessor Solution for Ray Tracing", *The Visual Computer*, (June): 62-73.
- INMOS, 1989, *The Transputer Applications Notebook, Architecture and Software*.
- HAGGERTY, M. 1990, "Remaking History One Pixel at a Time", *IEEE Computer Graphics and Applications*, July 1990: 3-6.
- KAY, D.S. & D. GREENBERG 1979. "Transparency for Computer Synthesized Objects", *Computer Graphics*, 13(2): 158-164.

- LOCK, G.R. & C.J. DALLAS 1990. "Compact Disc-Interactive: a new technology for archaeology? in Communication in Archaeology: a global view of Information Technology", *Science and Archaeology*, 32: 5-14.
- MALLEY, T.J.V. 1988. "A Shading Method for Computer Generated Images", unpublished Master's Thesis, University of Utah, USA.
- MYLONAS, G.E. 1961. *Eleusis and the Eleusinian Mysteries*, Princeton NJ, Princeton University Press.
- PIXAR 1989. *The Renderman Interface, Version 3.1*, (September).
- PENTLAND, A., I. ESSA, M. FRIEDMANN, B. HOROWITZ & S. SCLAROFF 1990. "The ThingWorld Modelling System: Virtual Sculpting By Modal Forces", *Computer Graphics*, 24(2, March): 143-144.
- PRESTON, J.M. (ed), 1988. *Compact Disc-Interactive. A designer's overview*. 2nd Edition, Deventer, The Netherlands, Kluwer Technical Books.
- RECKER, R.J., D.W. GEORGE & D.P. GREENBERG 1990. "Acceleration Techniques for Progressive Refinement Radiosity", *Computer Graphics*, 1990 Symposium on Interactive 3D Graphics Proceedings, 24(2, March): 59-66.
- REEVES, B. 1989. "SIGGRAPH '89 Panel Session, Bloopers, Outtakes, & Horror Stories of SIGGRAPH Films", *Computer Graphics*, 23(5, December): 235-239.
- RUSHMIER, H.E. & K.E. TORRANCE 1987. "The Zonal Method for Calculating Light Intensities in the Presence of a Participating Media", *Computer Graphics, SIGGRAPH '87 Proceedings*, 21(4, August): 303-310.
- SILLION, F. & C. PUECH 1989. "A General Two-Pass Method Integrating Specular and Diffuse Reflection", *Computer Graphics, SIGGRAPH '89 Proceedings*, 23(3, July): 335-344.
- WALLACE J.R., K.A. ELMQUIST & E.A. HAINES 1989. "A Ray Tracing Algorithm for Progressive Radiosity", *Computer Graphics, SIGGRAPH '89 Proceedings*, 23(3, July): 315-324.
- WHITTED, T. 1980. "An Improved Illumination Model for Shaded Display", *Communications of the ACM* 32(6, June): 343-349.
- WILLIAMS, L. 1990. "3D Paint", *Computer Graphics*, 24(2, March): 225-233.
- ZHU, Y, Q. PENG & Y. LIANG 1988. PERIS: "A Programming Environment for Realistic Image Synthesis", *Computers and Graphics*, 12(3/4): 299-308.