

Computer-aided art direction

Paul Richens

Martin Centre for Architectural and Urban Studies, University of Cambridge,
Department of Architecture

Presented at the Cinema and Architecture Conference, Cambridge 1995.

Published in Penz, F. and M. Thompson (eds) Cinema and Architecture: Melies, Mallet-Stevens, Multimedia. British Film Institute London 1997

Introduction

Computer-aided design for architects began to be possible about 25 years ago; in the last five it has become commonplace. In a few years time, it will be as ubiquitous as the word-processor is today. But the construction industry as a whole is not a sophisticated user of Information Technology. The state of affairs in the film industry is quite different; computer graphics of the utmost sophistication plays an increasingly important part in the *production* of film, but very little in their *design*. Is it possible that an opportunity is being missed?

This paper is written from the standpoint of a designer of architectural CAD software, and seeks to explore the extent to which computer graphics techniques, which have proved useful in architecture, could be used in Production Design, and the simulation of Cinematography. Several experiments, using commercial software of the sort that architects find useful, have lead to the realisation that much is possible, but that a full realisation of the benefits would require software specially adapted to the task. The bulk of this paper describes the nature of this adaptation; it is in the nature of a preliminary specification for software for *Computer-aided art direction*. The initial

investigation centred on the needs of film students; it has since broadened to look at the needs of established practitioners.

Objectives

The main use of CAD in architecture, though the least interesting, is for drafting the huge quantities of final drawings that tell the builder what he is to build. Film designers also produce large format “blueprints” to describe sets to the construction crew. To an architect they are strange, not only for being invariably dimensioned in feet and inches, but because they show only the inner, visible surface of the set, and give almost no information about its construction or structural support. This of course, reflects the way in which responsibilities are divided between the designer (what it looks like) and the rigger (how it stands up), which is quite different from that found in architecture (at least in Britain, and at the present day). So the volume and complexity of construction drawings is less, and standard packages for drafting could be used without difficulty.

Much more interesting is the use of preliminary drawings, renderings and models in the process of arriving at a design. In architecture, extensive use is made of three-dimensional computer models and rendering systems for presentation to clients. Use of a computer for conceptual design is rather unusual; when it does happen the software employed is likely to be lightweight, easy to use and fast to respond – it is more likely to be running on a Macintosh than a high-end workstation. A number of engineering disciplines rely extensively on computers - such as structural design, artificial lighting, heating and air-conditioning.

In the film world, we can expect geometric modelling to be helpful in set design, computer rendering to be useful for lighting design and cinematography, and animation for simulating tracking shots and special effects. To the professional designer (who we can assume knows how to do his job without a computer),

the most useful result will probably be the assistance that computer simulation gives him in *negotiating* a design with the director and other members of the production team. The student will find a much wider range of benefits - the simulator will teach him or her about set design, lighting and photography, in much the same way that a well-made desktop publishing program teaches its users about page grids, fonts, point sizes and leading - all the esoterica of the sub-editors craft (The realisation that a software application could be constructed as a teaching machine - to teach both its own behaviour and some area of craft skill – is one of the profound insights that underlies the Apple Macintosh, and the modern Graphical User Interface.)

The primary aids to negotiating a film design at present are drawings, scale models and story-boards. A computer simulation would act like a scale model, but with the added advantage of being able to simulate lighting and the camera's eye. Discussions can centre on the simulation, the model can be adapted as they proceed, story-boards can be produced as a sequence of stills from the simulation. Indeed, a computer based story-board could go further by mixing stills with animation clips, for example to show the effect of a moving camera. In education, it is likely that such an elaborated story-board could be used to develop narrative skills.

Modelling

Very sophisticated geometrical modelling techniques are used for industrial design, and indeed in film animation. Architectural modellers tend to be simpler, and would be quite adequate for set design. Modellers (such as ModelShop or UpFront), which are popular in architectural schools and tend to emphasise simplicity and speed over geometrical precision, show the way to go. The model is composed of “primitive” shapes such as blocks, spheres and cylinders, combined with more elaborate extrusions and solids of revolution. An elementary spline-surface element is desirable for more organic shapes. These basic elements can be produced in any shape and size, and “grouped” to

make composite objects which can be repeated, scaled and oriented as needed.




Figure 1. Primitive shapes in Modelshop.

Sets frequently include “floaters” pieces which can be removed temporarily to gain access for lights or camera. Obviously, the geometrical model should have similar flexibility. It may also enable checks to be made that such apertures are big enough, for example to allow a crane to operate.

Two things which occur frequently in film sets (but have hardly been seen in architecture since the ending of the baroque) need special treatment. Both are to do with representing objects in the distance. In an urban exterior, parts of the scene in the foreground (containing the action) will be built full size. Background buildings will be painted on a cloth. In between, the set may contain scale models of objects in the middle distance, or more elaborate perspective pieces

built with some degree of foreshortening as well.

A modeller adapted for set design should be able to render a 3D model of the background onto a surface, and then use the image as a backcloth. It should also be able to calculate perspective pieces from the full scale geometry and a given camera position. It will also be able to simulate the strange distortions that result if the camera moves from the centre of projection, or if an object or actor approaches too closely.




Figure 2. Computer generated backcloth behind a computer generated set.

Texturing

When a set is built on a stage, the first task is to construct the bold forms; this corresponds to the geometrical modelling described above. The second is to

apply surface texture, which may be by applying vacuum-formed plastic sheet, by rendering with cement or plaster, or by attaching fibrous plaster casts. Sometimes, even, unsimulated real materials are used. This process, which gives relief and texture to the surfaces, has its simulacrum in the computer-graphics technique known as “bump-mapping”. An image is drawn in shades of grey to represent the various degrees of relief, and attached to a surface. When the scene is rendered, the relief is used to modify the local orientation of the surface, and hence the way that it reflects light.




Figure 3. Bump mapping.

The third stage of preparing a set is to paint it. The corresponding computer-graphics technique is called (rather misleadingly) texture-mapping. A full-coloured image is mapped to the surface, and affects its local colour when it is rendered. The texture map can be generated in a computer painting program, but is more commonly scanned from a photograph. In set design it would

frequently be obtained by scanning a painting, or a sample of the desired finish. Other computer-graphics techniques which would be of use are reflectance mapping (which deals with varying degrees of shininess), transparency mapping (eg for dirty glass) and luminosity mapping (eg for placing lighted windows on a distant building).




Figure 4. Texture mapping for a variety of materials.

These techniques are not found in the simple architectural modellers mentioned earlier. Texture mapping is carried out very speedily by special hardware of the sort found in expensive graphics workstations; it is likely to be available in PCs in the near future. The other techniques are found in the more elaborate but slower rendering programs. They need to be foreground techniques in any system aimed at set design, where the way in which surfaces react to light is of

prime importance.

Lighting

The simulation (and teaching) of lighting is the most challenging, but potentially the most rewarding area for research. Conventional computer rendering programs only begin to acknowledge the complexities of film lighting, and even the most advanced research software would be severely stretched.

Light sources used in film studios include tungsten, tungsten/halogen, metal/halide and carbon arc. The distinctions matter, because they affect the colour temperature of the light. The power of sources is typically 1 or 2 kW, but may commonly vary from 200w to 15kW.

Sources are used in a range of different types of fittings. Open face fittings have a reflector and source in a container, but no lens. They may be adjustable to some extent, and there are a variety of shapes. Sometimes up to a dozen open face fittings are combined to form a modular array. These provide a diffuse light from a large-area source. Softlites and skylights use reflectors or diffusers to intercept the beam, and again provide large-area soft lighting.

The most commonly used closed fitting contains a reflector and a fresnel lens. It provides a widely-adjustable beam with a soft edge, but coming from a fairly small source, so shadows are defined. Focus spots, on the other hand, have precise optics and an aperture plate capable of giving a precisely shaped beam from practically a point source.

It is interesting to compare this diversity of fittings, with the types offered by conventional computer rendering programs. These typically offer a choice of three, a remote source (sun), a point source (naked bulb) and a spot. This has an adjustable beam angle, and perhaps a variable rate of cut-off. However, altering

the focus does not affect the brightness of the beam, as happens with a real fresnel spot. There is no simulation of large-area sources at all, nor any notion of colour temperature. What is needed is a catalogue of industry-standard fittings, each characterised by precise photometric data over the whole range of focus settings.

Fittings may be modified in many ways. The beam may be restricted by barn-doors or snoots. The colour can be modified by filters, the intensity and spread by diffusers or scrims. Dimmers (which are ubiquitous in theatre lighting) are used with restraint, as they alter the colour temperature.

Film lighting may be extensively modified after it leaves the lantern by reflectors (matt or metallic), diffusers (eg of glass fibre, tracing paper or gauze), and gobos, which are devices used to cast shadows, or to keep light off places where it is not wanted. Gauzes may also be used, partially fogged by side lighting, to give an effect of distance to the background.

Control of intensity is occasionally by dimmers, but more often by choice of fitting, spreading the beam, fitting scrim, or simply by moving it further from the subject. The required level of light may be expressed photometrically (400 foot-candles), or perhaps in terms of the aperture that will be used on the camera (eg T5.6, with some implied assumptions about the shutter speed, film stock, and processing options).

Conventional computer-graphics does not provide a very adequate simulation of all this. The normal techniques are not photometrically accurate, and fail to deal with large-area sources, soft-edged shadows, diffusers and reflectors. Most importantly, they deal only with light received directly from the source, and not with light reflected from other surfaces in the set. In film and theatre lighting it is essential to take into account at least the first inter-reflection between surfaces (this is also a severe problem in architectural rendering of daylit interiors, where

five or more levels of inter-reflection may be significant). Two techniques to solve this problem are currently at the research stage - radiosity and stochastic ray-tracing. Both are slow to calculate, but a variant of the first (progressive radiosity) can give an approximate rendering quickly, and then improve it over time by taking more inter-reflections into account. Curiously, most radiosity renderers deal only with large-area diffuse sources, though there is no insuperable reason for this. There are known difficulties in combining radiosity with bump-mapping, which need to be overcome.

The ideal lighting simulator for film design will use industry-standard fittings, and produce a visual simulation showing shadows, modelling, colour and the balance of fill to key. It should be possible to introduce a simulated photometer, in order to measure the calculated light levels (incident or reflected). If, as is likely, the most accurate simulation is slow, there should be a range of faster approximations. Ideally it should be possible to position, focus and filter the lanterns, while seeing the effect immediately on the computer screen. Rapid feedback is invaluable to the practitioner, in enabling a design to be realised quickly, and to the student, in developing an intuition about how things behave.

Another mode of operation would be to use idealised lanterns, which can be positioned and focused to get the angles right, and then coloured and dimmed to balance the lighting, again with immediate feedback of the result. When the lighting has been fully composed, the idealised lights would be replaced by real lanterns of appropriate wattage, so as to achieve the same effect without the aid of dimmers.




Figure 4. Direct manipulation of lights, with a simulated dimmer board.

Photography

One problem with photometrically accurate lighting simulation is that the range of brightness in the scene, as calculated, will vary as widely as it does in the real world, and will not be representable on the computer screen which has a dynamic range limited to about 1:30. Some technique is necessary to reduce the range, and it would most usefully be a simulation of the photographic parameters such as exposure, film stock and processing.

Camera simulation includes positioning and orienting the camera, and picking the correct aspect ratio. The lens should (as with lanterns) be chosen from a catalogue of industry-standard types, which gives their focal length, and aperture range, both photometric (T) and optical (f). Zoom lenses need the range

of focal lengths, and possibly details of how the apertures change over the range.

It should be possible to set the aperture from a simulated scene, perhaps by using a simulated 18% grey test card. Alternatively the aperture can be set arbitrarily, and the resulting photographic image simulated. This simulation might include the effect of film stock and processing options on sensitivity, latitude, gamma and colour rendering. It would be valuable also to simulate the effect of focusing, and of changing aperture on the depth of field, by selectively blurring parts of the image.

The camera simulation might also take account of filters of various types, such as neutral density, graduated, diffusing and colour correction. The last implies a detailed treatment of colour temperature, and would be important if location filming was to be taken into account.

Animation

As well as simulating a static shot, it is desirable to be able to render a sequence as a short animation. The most obvious case is where the camera moves on a track or crane, but there are others, such as the simulation of zooms, focus pulling, changeable lighting, and moving properties or vehicles. Lights may move with the camera, or be attached to other objects (practical lights).





Figure 5. Setting up a crane shot animation, with (below) the start and end frames of the sequence.



The simulation of movement requires laying down a track (generally a curve in three dimensions), specifying how the camera orientation changes as it moves along the track, and controlling its speed, especially at the start and end of a movement.

As digital effects become more and more prominent in film-making, the art director will increasingly need to try things out at the design stage. So the simulations might well extend to include matting, chroma key, computer controlled model photography, and indeed computer animation itself.




Figure 6. Simulation of special effects photography.

Conclusions

The initial impetus for this investigation came from the National Film and Television School, which trains people in most of the disciplines involved in film-making (except acting) in a very practical way, by making films. However the cost in time and materials of doing this is considerable, and tends to limit the experiences that can be offered to students. A student production designer does not often get the opportunity to design and build for more than one production, but with computer simulation could easily explore half a dozen. Similarly a

lighting-cameraman could learn a great deal about the behaviour of light, lenses and film-stock without actually shooting anything.

Our experiences with student productions indicate that it is quite possible, using normal architectural software, to get a good approximation to the visual “feel” of a set. A simple model is useful to explore camera angles, and quickly shows up problems (like the set is not tall enough) and opportunities (angles that had not been thought of). Simulating lighting turned out to need a good deal of trickery, because the normal computer lighting model is far from physical reality. Photographic simulation we did not attempt. But the overall impression was that though it is possible to do these things, it is far from easy. Though the majority of the ambitions described in this paper have been realised piecemeal in commercial software, or at least in the research laboratory, they have not ever been put together in a way that the student of film could use.

We have kept in the background the possibility that the same software might be used to simulate theatre or television design, and possibly other events such as exhibitions and rock shows. Stage sets are designed for a more limited set of viewing positions, but the geometry of the stage and the auditorium becomes an important part of the problem. Stage lighting uses hundreds of lanterns in fixed positions, but with elaborate computer-controlled dimmers, which enable the whole production to be lit. Film set-ups use fewer lights, other means of control, and may be reconstituted for every scene. Television studios seem to hold an intermediate position.

It is likely that software aimed at the educational objectives would also be of use to practitioners. The most interesting area of use is in the negotiations which have to occur in any kind of collaborative design. So we see computer simulation being used in negotiations between the set designer and director, and then between cinematographer and director, or perhaps all three. It should help to promote common understanding, and avoid expensive mistakes. We have noticed that often far more of a set is built than is ever filmed, which could

perhaps be avoided by simulating the shooting script. Similarly, the commissioning of special effects or computer animation would be helped if the art director can present a sketch simulation of what is required.

Acknowledgements

The initial impetus for this investigation came from the staff of the National Film and Television School, particularly Maureen Thomas, and the Director, Henning Camre. Diana Charnley and Sharon Springel have been most helpful in providing information about film-making practice. I am deeply grateful to the production designers who have been most generous with their time in discussing particular projects - Stuart Craig (*Mary Ryan*), and Christopher Hobbs (*Cold Lazarus*).