Visualisation of topographic data using video animation

J-P Muller, T Day, J Kolbusz, M Dalton, S Richards, J C Pearson.

Department of Photogrammetry and Surveying
University College London
Gower Street
London WC1E 6BT, U.K.

ARPANET/BITNET: jpmuller@cs.ucl.ac.uk

ISPRS INTERCOMMISSION III/IV

Abstract

A wide range of software tools have been developed for visualising three-dimensional range data of topographic surfaces. Digital Elevation Models (DEMs), derived from photogrammetric measurements and from digitised contour maps are here used for illustration. To increase our understanding of the information content in DEMs in isolation or in combination with remotely-sensed images and to facilitate the quality assesment of automated terrain extraction techniques a number of tools have been developed for displaying perspective views of time sequences of landscapes with a variety of different surface shading models.

We consider a variety of shading models from the texture mapping of a satellite image over a DEM to attempts to represent different surface reflectances. The texture mapping problem integrates both single band (monochromatic) and multispectral (false colour of original or colour-space transformed) satellite images. Artificial shading models from simple Lambertian shading to the ray-tracing of terrains, using Bi-directional Reflectance Distribution Functions are presented.

1 Introduction

Perspective views of the surfaces of three-dimensional data are a significant aid to an appreciation of the spatial nature of natural phenomena. Examples in remote sensing include physical interpretation of satellite image brightness values; understanding of geological structures and analysis of hydrological networks and catchment areas. Examples in aerial-based topographic mapping include understanding the detailed nature of forestation and planning issues involved in building and highway construction.

The paper map has for centuries been the primary means of visual communication of spatial data and its properties, advantages and drawbacks have been debated for almost as long. However, apart from flight simulator projects (see, for example,[Schachter, 1983],[Zyda et al., 1988]) and military route planning ([Garvey, 1987]) little research into the potential of dynamic perspective views of landscape has been made. It is one of the primary objectives of our research programmes to explore the issues involved in the visualisation of a wide variety of scientific and engineering data including photogrammetric measurements of industrial parts in CAD systems ([Muller and Anthony, 1987]) and of anthropometric data ([Kolbusz, 1988]).

Visualisation of scientific data has received a great deal of attention in the last year as a result of a US National Science Foundation study (see [McCormick et al., 1987], [Winkler et al., 1987], [Frenkel, 1988]) which presented a detailed list of scientific, technical and socio–political arguments in support of its crucial role in interpreting the results of numerical simulation experiments; understanding the massive amounts of digital data now being generated in scientific experiments and in visual communication of ideas and products to wider audiences.

In this paper, however, we focus our attention on the visualisation of small-scale Digital Elevation Models (DEMs) using different surface shading techniques. It should be noted that the previewing and video animation techniques described here are quite general enough to encompass most problem domains, the main difference being the representation of surface geometry in the different application areas.

Static perspective views of landscape from aircraft (usually refered to as oblique photography, see [Brew and Neyland, 1980]) have been used extensively since the Second World War for both photo-interpretation and reconnaisance mapping. However, apart from [Gelberg and Stephenson, 1987] and some unpublished work by [Quam, 1987] and the Animation Group at the Jet Propulsion Laboratory [Wolff, 1987] little if any use has been made of perspective viewing of terrain.

The breakthrough in the spread of visualisation techniques is coming about because of two technological developments : graphics workstations and special-purpose graphics rendering hardware, such as VLSI matrix manipulation chips (so-called "geometry engines", see [Zyda et al., 1988]). These enable visualisations to be accomplished for wireframe representations of small DEMs (without hidden line removal) at video refresh rates and for simple shaded representations of DEMs at rates of approximately 1 frame per minute. Increasing RAM memories (up to 256MB) and more sophisticated "geometry" and "rendering engines" suggest that within a few years, the types of visualisations presented here will be done at video refresh-rates and we suggest will be commonplace on engineering workstations.

Three dimensional data is considered here in terms of its bounding surface rather than its existence as a volume. To produce visual representations of such bounding surfaces, we consider this as a problem of rendering (i.e. simulating the interaction of light with a 3D surface). A number of techniques have been developed in computer graphics over the last 20 years (see, for example, [Rogers, 1985]) for the rendering of 3D data and the most relevant of these is discussed in section 4. Section 2 gives a definition of the objectives of our visualisation work whilst section 3 discusses the various approaches which have been adopted for the mathematical representations of DEMs. Section 5 describes our techniques for image display and previewing and how we define our view-path for our animations as well as what equipment we currectly employ. Finally, section 6 suggests ways in which the work can be extended in the future. All figures shown in this paper are produced using the SPOT-PEPS data-set described in more detail in [Day and Muller, 1988].

2 Objectives

In recent years, there has been an explosion in the amount of data collected about the Earth's surface from satellite imagery. However, work on analysing this imagery has tended to concentrate on two-dimensional image processing of its multispectral content whilst ignoring the increasing amounts of structural (e.g.cultural) and terrain-related information which have come from narrower spectral resolution and/or higher spatial resolution.

Several authors ([Woodham, 1985],[Muller, 1988a],[Muller, 1988b]) have recently suggested an alternative approach using a computational vision model of remote sensing. At the centre of this new approach is the attempt to model image formation through computer graphical techniques to try to understand how the patterns visible in satellite images originate and hence devise sensible data processing strategies for their future automated analysis. These studies have also included the automated extraction of DEMs from SPOT satellite images (see [Day and Muller, 1988] and [Muller et al., 1988]) which form one key input to image formation models, particularly in areas for which the appropriate scale of map contour data is unavailable and the cost of digitisation may be unacceptable (see [Dowman and Muller, 1986]).

To test our image formation models (for representing surface reflectance and/or sky radiance) as well as assess the geomorphological veracity (see, [Muller and Saksono, 1988]) of our automatically-derived DEMs, we have a number of alternative strategies which can be grouped as follows :

- render a synthetic satellite image(s) from the same viewpoint as the satellite with a simple pinhole camera;
- render a synthetic satellite image(s) from the same viewpoint as the satellite using an exact representation of the satellite optical system and sensor characteristics (see [O'Neill and Dowman, 1988]);
- render a synthetic satellite image(s) from a perspective viewpoint;
- render the actual satellite image(s) from a perspective viewpoint;
- repeat items listed above using false colour representations.

3 Surface representation

The input to a DEM can come from a variety of different sources. To date we have investigated the following (see examples in Figure 1) :

a Extract from SPOT image, for comparison;

Figure 1: SPOT image extract, and intensity range images of area from various sources

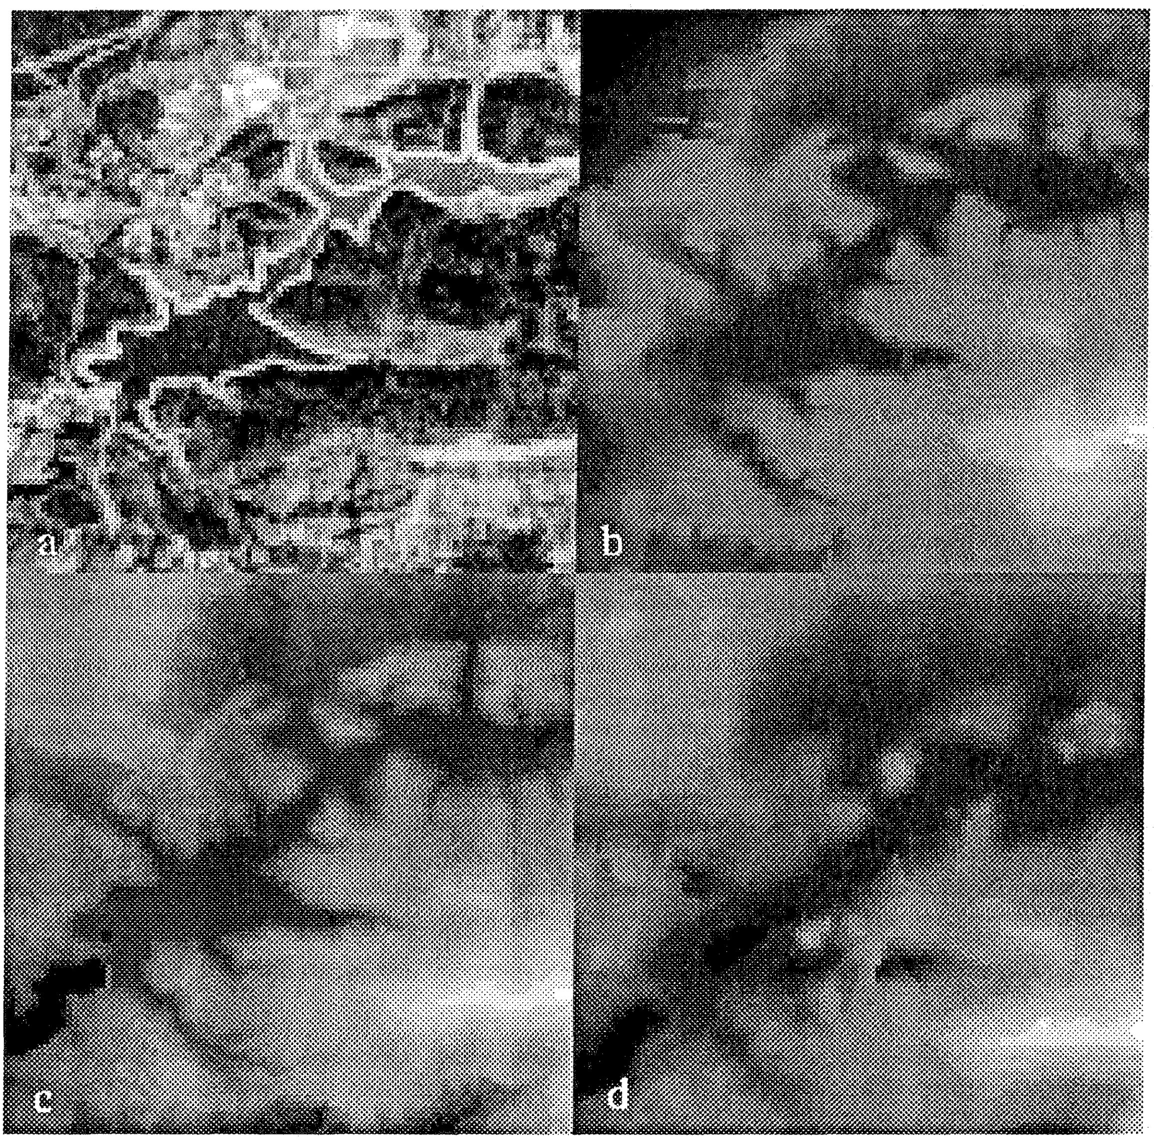

b Interpolated manually digitised contours;

c Manual photogrammetric measurements taken from aerial photographs;

d Automatically stereo-matched SPOT images interpolated to the same grid intervals.

Both the digitised contours and the stereo matched data are in the form of a collection of spot elevations which do not have a data-structure amenable to easy handling or rapid rendering. Delauney triangulation ([McCullagh and Ross, 1980]) is therefore employed to construct a Triangulated Irregular Network (TIN) (using Laser-Scan Laboratories' MATRIX software) from which either irregular triangle data-structures can be extracted or a regular grid can be interpolated.

To build a regular triangular data-structure from a gridded DEM we divide each square (or rectangular) DEM cell into two triangles. Instead of splitting each cell in the same direction and introducing directionality, we choose at random between the two possibilities. (see example in Figure 2)

Figure 2: Wireframe image of low resolution DEM

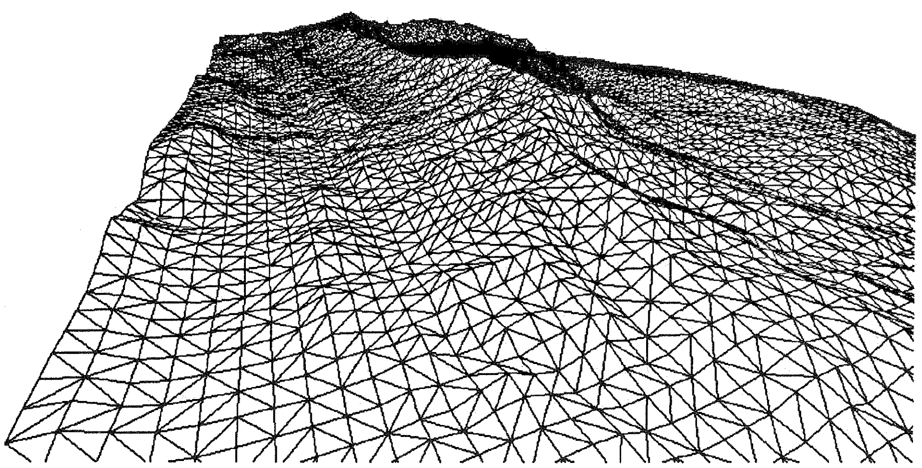

For DEM data, this surface representation scheme is likely to remain the most efficient until suitable alternatives can be found for storing compressed versions of elevations related to the appropriate sampling densities for the terrain type (see discussion in [Muller and Saksono, 1988]).

There have been a number of suggestions for different sampling schemes for smooth surfaces derived from range data (see [Bhanu and Ho, 1987]) which are more suitable for computer manipulation. Experiments with surface fitting methods using Bernstein-Bezier curved surface patches (see [Kolbusz, 1988]) compacts much detail to a given surface tolerance. Our use in this work is restricted to the interpolation of DEM data. This is useful for constructing high 'resolution' versions of DEMs from lower resolution ones to minimise spatial undersampling (the so-called 'aliasing' problem).

4 Rendering

Rendering software takes a surface representation model structure and produces images, taking into account the absolute position and relative orientation of the objects, light sources and the viewer.

Rendering software must solve a number of problems with DEM data including the removal of hidden surfaces (see review in [Sutherland et al., 1974]) and the approximate simulation of light interactions with matter (see review in [Magnenat-Thalmann and Thalmann, 1987]). These computationally expensive tasks are usually best performed on graphics workstations with specialized peripherals (see[England, 1988]).

Increasing levels of complexity can be used to approximate light interaction with matter (so-called shading models). Together with this complexity comes greater realism.

At the simplest level, we can assume that every facet of the surface has a constant colour. This approximation can be refined by taking into account some measure of how much the reflected light from a surface varies as a function of illumination and viewing geometry.

These simple shading models include Lambertian ([Woodham, 1985]) and a number of analytical approximations to the relative proportions of diffuse and specular components of scattering which have been named after their inventors :[Gourard, 1971],[Phong, 1975],[Blinn, 1977] and [Cook and Torrance, 1981].

To model light interaction between surfaces and with matter or gases in the atmosphere, more advanced models such as ray-tracing, radiosity and Monte-Carlo ray tracing must be considered (see references in [Magnenat-Thalmann and Thalmann, 1987]). These allow us to consider shadows, reflections, distributed light sources, colour bleeding and the dynamic sampling of non-stationary objects.

In the following sections, rendering is treated at two different levels of complexity: fast rendering of satellite images using texture mapping and DEMs using simple Lambertian shading and very much slower rendering of DEMs using ray-tracing to incorporate diffuse and specular surface reflectance models and sky radiance.

4.1 Fast rendering of satellite images using texture mapping

Single-band and false colour composites of planimetrically-geocoded satellite images can be used as input to the rendering of DEMs where the grid-points or Delauney triangles are used as vertices.

Each triangle can be assigned a particular colour (at its centroid) derived from the image. Since each triangle is represented by a single colour, we are effectively reducing the image resolution down to the DEM resolution. This can be overcome by interpolating a higher resolution DEM or by increasing the number of triangles (by recursive subdivision, see [Fournier et al., 1982]) before sampling colours from the image. In either case, an image of higher resolution than the DEM should not be sampled due to the possibility of 'aliasing' (e.g. intermittent sampling of a bright feature such as a road). Instead the image resolution should be reduced by area-averaging to around the DEM resolution. Figure 3 is an example of a static perspective view of a SPOT panchromatic image shown alongside a black-and-white version of a false-colour composite image derived from LANDSAT-TM.

To resolve hidden surfaces, a number of approaches have been tried. In general, three dimensional surfaces composed of, for example, intersecting opaque and translucent facets require a much more complex rendering schema than that of a DEM (see, [Sutherland et al., 1974]). After all, a DEM already contains sorted position information, and for our application, surfaces are known to be opaque and non intersecting.

4.2 Resolving hidden surfaces with depth buffers

Depth buffering is a simple method for rendering objects, particularly with surfaces composed of planar polygons and triangles into an image. For each image pixel we store the colour, and the range to the object visible in the pixel. Initially the image is blank and the depth of all pixels is set to the largest representable value. Rendering primitives are then scan converted to pixels, with the pixels only rendered on the framebuffer if the range ('z-depth') is less than the previous value.

The renderer for our triangle based system uses the depth-buffer provided on a special-purpose graphics accelerator board attached to one of our Sun workstations. This device includes hardware to perform 3D transformations using homogeneous co-ordinates ([Newman and Sproull, 1973]). However, this device is limited in the accuracy to which it can store range (16 bits), so objects at similar depths projected to the same image pixels may not be rendered in the correct order. Because homogeneous co-ordinates are used, the depth value stored in the buffer is actually a nonlinear increasing function of the true depth, with the effect of increasing depth resolution in the foreground, but lowering resolution in the background. In video animated sequences, this can cause distant objects to appear to 'crawl'.

To solve the problem of aliasing artefacts, the A-buffer (standing for anti-aliased, area averaged, accumulation buffer, see [Carpenter, 1984]) has been developed. This is a general hidden surface technique that can resolve visibility among a completely arbitary collection of opaque, transparent and intersecting objects, to a resolution many times greater than the traditional z-buffer.

Rendering operations using this system is identical to that of the Z-buffer: modelling primitives are diced to flat polygons which are then thrown at the A-buffer in an arbitary order. Every polygon is resolved after first being clipped to pixel extent. The visibility resolver, using a box filter, maintains a list in z of sub-pixel sizes. When all objects are resolved, we do a weighted summation of any sub-pixel values.

The A-buffer is useful in many areas, such as correctly resolving very small polygons. Its generality may be more than is needed for simple surfaces such as terrain, but for complex self-enclosed 3D objects, it would appear to be ideal.

Figure 3: 30m DEM overlayed with SPOT panchromatic (top) and LANDSAT-TM (bottom)

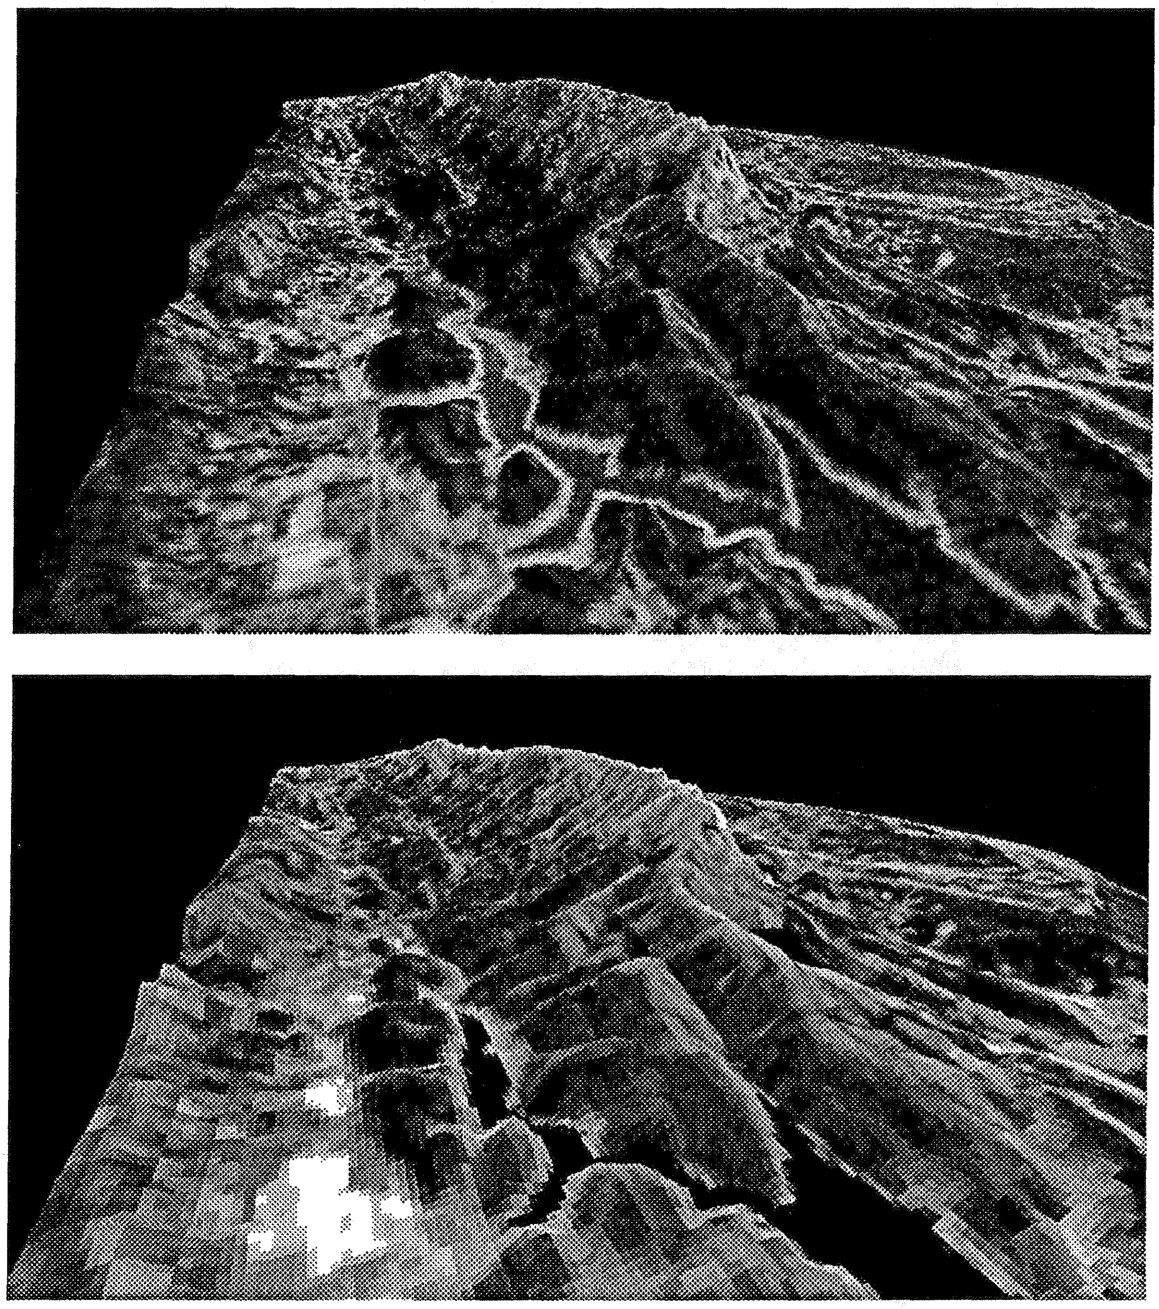

4.3 Ray-tracing

A numerical modelling scheme has been developed in which physical processes can be simulated. These involve the production of an image on a detector from the scattering of light from a terrain. At the present time, no attempts have been made to include atmospheric effects.

The model was constructed using a computer simulation technique known as Monte Carlo sampling (see [Kajiya, 1986]). This allows a finite number of rays to represent a huge number of actual photon(wave) interaction events with matter by ensuring that optimum sampling is done of the critical parameters (e.g.ray direction, time of emission, wavelength and polarisation).

The critical path for each ray starts at the detector where rays are fired through the principal point into the environment. This is done because we know that rays emitted from the detector will contribute to the image whereas most of the rays emitted from the light source will not. After intersection with the terrain, the rays are absorbed and then re-emitted until they either hit another surface or a source of illumination such as the sky or the Sun.

Different surface reflectance models can be represented using the Bi-directional Reflectance Distribution Function (BRDF) (see [Nicodemus, 1977]) which is a measure of the ratio of exitance radiation to irradiance for all solid angles of incidence and exitance for all wavelengths.

Monte Carlo sampling uses probability functions to bias the sampling to those parameters which most contribute to the measurement. For example, to reduce sampling errors at the detector, a Poisson distribution is used to jitter randomly the direction of the emitted rays within the sample area ([Mitchell, 1987]). An analagous technique is used to reduce the sampling of rays with wavelengths where the detector has a poor response and to bias the re-emission of rays from a surface into the direction of the principal plane of illumination.

An example of the application of this technique is shown in Figure 4 which demonstrates the potential of this technique to produce realistic visualisations of terrain including the effects of sky radiance.

Figure 4: Ray traced image of Montagne Sainte Victoire at 8am on 12th May

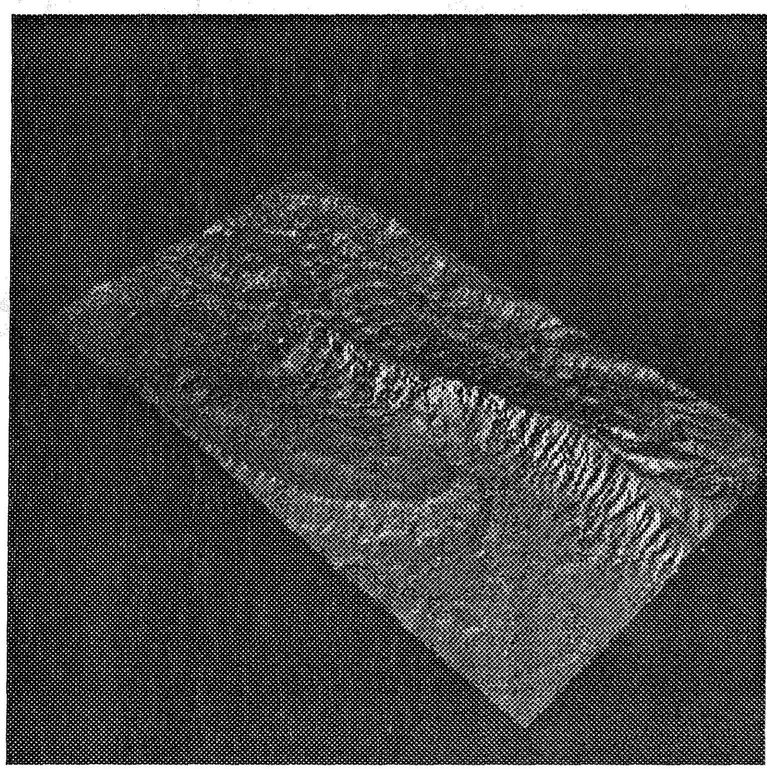

5 Display

There are a number of animation techniques which can be used to display rendered perspective views of terrain.

These display techniques include :

1. Single camera viewpoint;
2. Stereo camera viewpoint;
3. Multiple camera viewpoint (e.g. ahead, from above, at the side);
4. Fixed focus of interest for camera(s);
5. Movable focus of interest for camera(s);
6. Interactive specification of camera(s) flightpath;
7. Pre-defined specification of camera(s) flightpath (e.g. view from a pre-determined flightpath)

In addition to describing the camera viewpoints (i.e. absolute orientation parameters) we can also specify the velocity and acceleration of the camera and the duration of each rendered sequence. From this, the number of frames required for rendering any one piece of animation can be calculated and the animation can be left as a batch task. Finally a description of our equipment for the production of video animations will be made.

5.1 Multiple viewpoints

Single views of landscape can easily be generated using any of the rendering techniques described in the last section. However, their use is limited as the investigation of any single area requires repititive video animations to be made. A more efficient scheme is to render a number of views concurrently on different workstations (or with transputers, see [Muller et al., 1988]) using the same DEM database, where each collection of views at a single instant of time either shows a stereo view of each scene or multiple views from a number of, possibly, orthogonal viewpoints. This also offers the possibility of focusing interest on particular areas of the landscape by fixing one viewpoint whilst moving the others and varying the focal length of our virtual cameras. These multiple viewpoints can easily be displayed using the NeWS/X11 [Gosling et al., 1988] window management system.

Stereo views (see example in Figure 5) have been developed for the analysis and editing of blunders in automatically generated DEMs. They can be either currently be viewed using a mirror stereoscope (which limits it to a single observer) or using a variety of anaglyptic (Red/Green) techniques. Future planned developments include 3-D digitising pens or gloves for editing and/or camera viewpoint manipulation and Liquid Crystal Shutter screens placed in front of CRT monitors or video projectors (see, for example, [Fergason, 1983]) which will permit 24-bit colour images to be displayed stereoscopically.

5.2 Animation

Techniques have been developed to enable the easy interactive specification of the imaginary flight path of our virtual camera(s) around the landscape. Since it is not possible yet to visualise our landscape in real time, we preview the 'flight' with a wireframe model of the terrain, and use this specification to drive the renderers.

Figure 6 shows a typical screen with this flight specification software. Notice in this figure that there are a number of distinct functional units including

- (at top left) a control panel where camera parameters can be specified and changed;
- (at bottom left) the flight-path can be drawn superimposed on an intensity-range image of the DEM;
- (at top right) the vertical profile corresponding to the planimetric path defined in the bottom left can be specified and changed;
- (at bottom right) a wireframe without hidden line removal which can be previewed at video rates.

An optical mouse attached to the workstation allows us to draw a planimetric flight path over a view of the DEM (or satellite image or both). To control viewing orientation, two sub-windows within the control panel allow us interactively to specify camera direction. The top window contains a small slider which represents the position in the sequence that the currently displayed frame is and this can be modified by the left button. The right button contains a menu for the options of locking the displayed frame to the spline control points. Also in this menu is the option to (re-)play the sequence. Figure 7 shows an illustration of several views generated from this interactive specification.

Figure 5: Stereo pair (SPOT overlay)

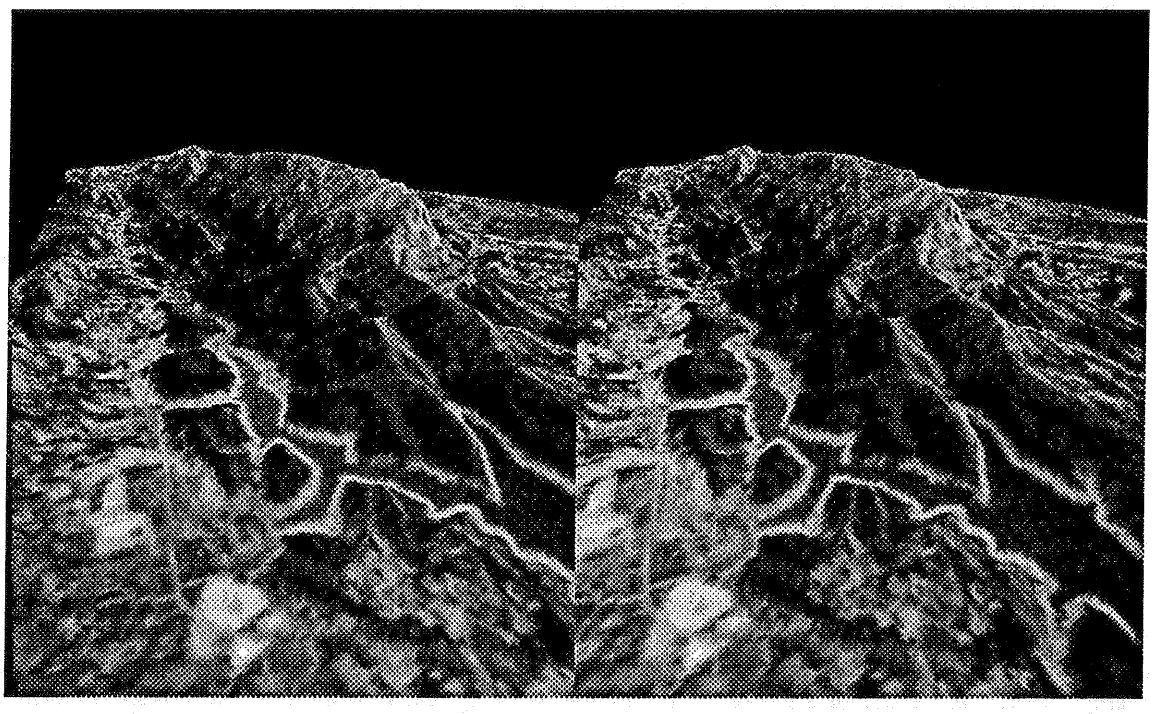

Figure 6: Flight path specification

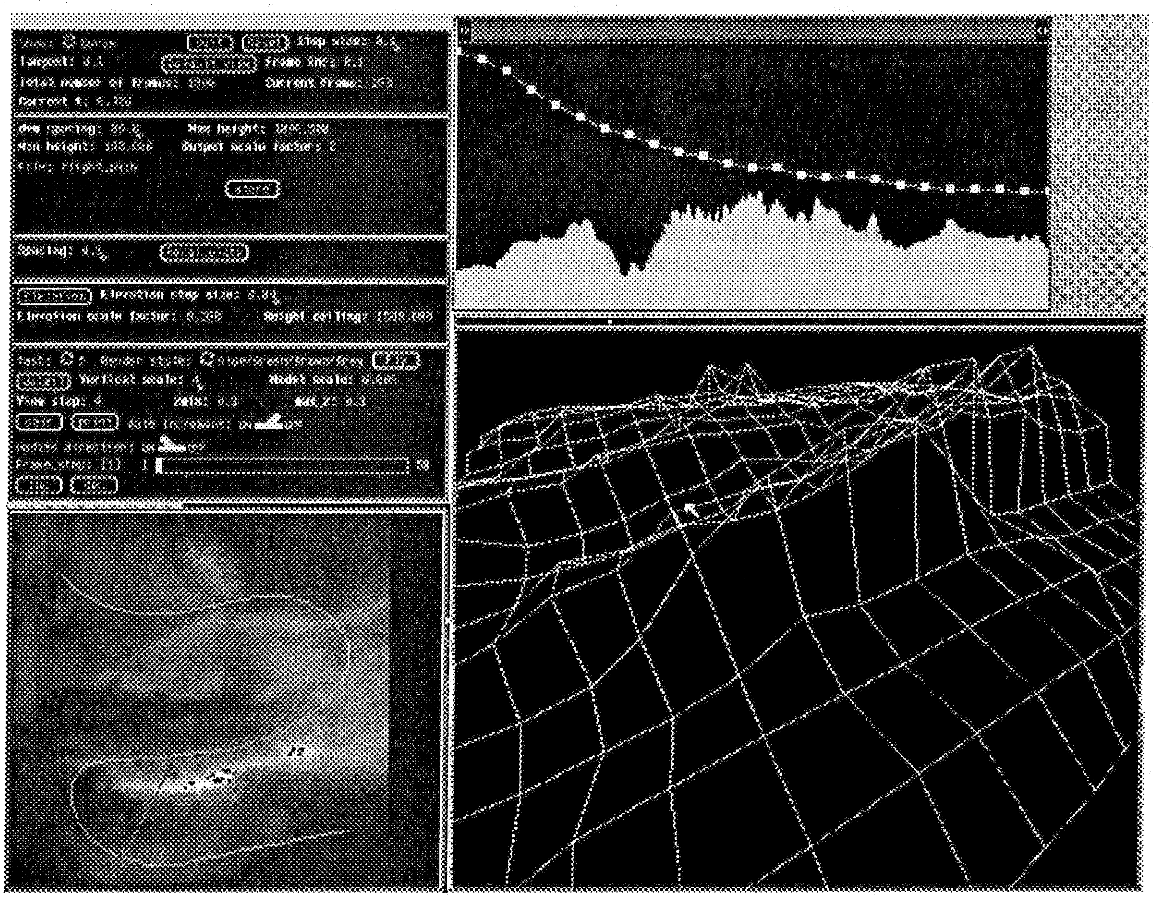

Figure 7: Views from various points along flightpath

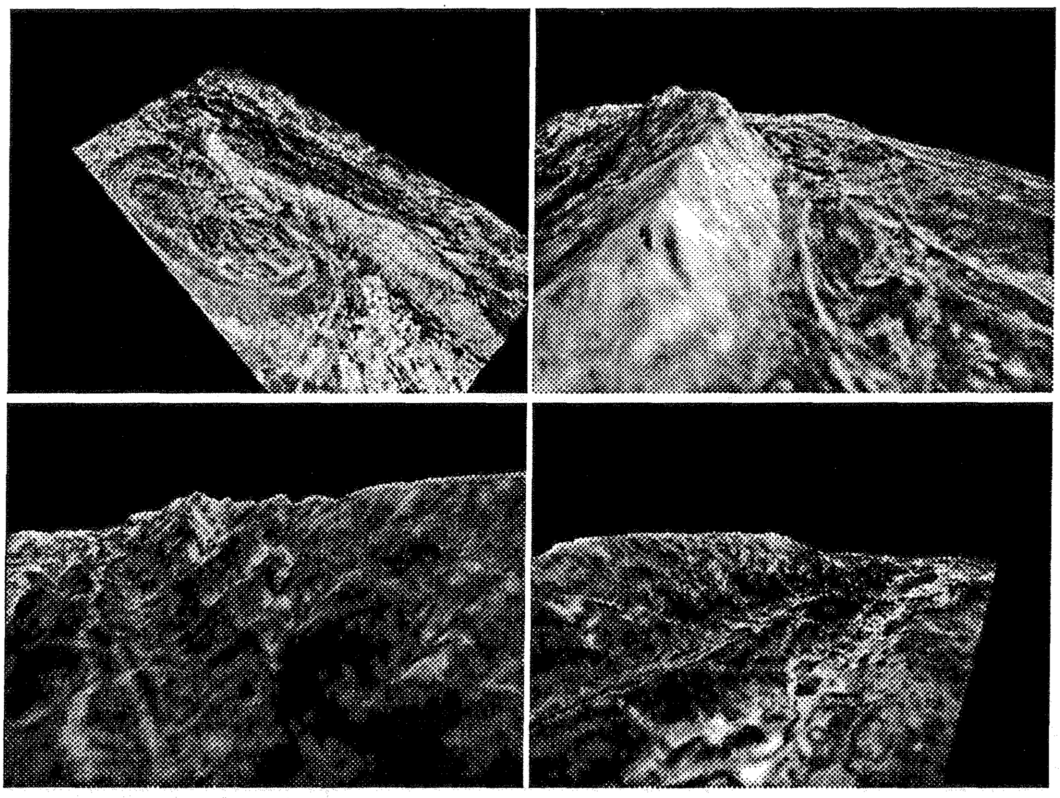

In addition to interactive specification, views can be generated for a pre-specified flightpath. This can be used to re-visit a site after modifications have been made to a stereo matching algorithm or rendering technique. In the future this will be linked into our digital mapping system (LITES2 from Laser-Scan Laboratories) to enable "flights" to be made from ground-level along a road and use digitised aerial or ground-based photographs to render the views.

5.3 Hardware

Figure 8 shows a schematic diagram of the hardware being used to produce video animations. The primary means of data transfer is digital and uses our co-axial Ethernet cabling. Data can be transferred from our rendering engine (the GP accelerator) on our Sun-3/160 to the 24-bit Primagraphics CCIR/RS-170 framestore on our Sun-3/260. Data can also be created on the VaxstationII/GPX using the LITES2 or MATRIX packages and transferred using software protocol conversions to our Sun-3/260. Finally data from the two digitising CCD cameras in our Kern DSR-11 can be transferred from the VMS/GPX to our Unix/Sun via the same route.

The Primagraphics hardware contains a device driver which enables the Sun window management system to be used and enables any Sun user on the network to generate and display frames. Finally, the EOS animation controller allows single-frames to be written from the Primagraphics framstore to standard half-inch U-matic videotape from which VHS dubbings can be made (CCIR/PAL only).

Timings for renderings of CCIR 768x575x8-bit frames are shown in Table 1. The transfer from one workstation to another is a negligible fraction of the total time.

Table 1: Rendering times

Equipment	Shading	No. triangles	Speed	
Sun 3/160	Lambertian shading	191730	320	triangles/sec
3/160 + GP1	SPOT image overlay	426560	4080	triangles/sec
Sun 3/260	Ray-traced (32 rays/pixel)	191730	1.5	pixels/sec

Figure 8: Visualisation hardware and data sources

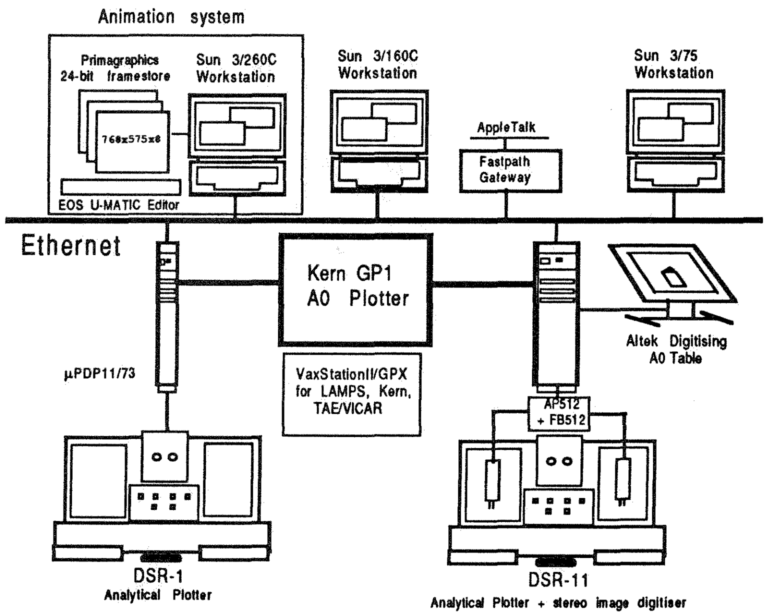

6 Discussion and conclusions

Video animated sequences of topographic data open up new possibilities for communicating geographical spatial information which is currently encrypted in paper maps or digital databases. It frees the viewer from the tyranny of two-dimensional degradations of our three-dimensional world and can provide significant aids to specialist photo-interpreters who may be interested in only small areas of a satellite image or only one aspect of the information contained within a satellite image.

Future work will include the incorporation of larger databases (we have already produced a number of animations of the whole planet using a 5' DEM database from NOAA) and complete stereo-matched SPOT scenes using scanline techniques based on [Robertson, 1987]. Experiments will shortly begin on using a new generation of graphics accelerator (see [England, 1988]) to speed up rendering times and extensions to arrays of transputer elements (see [Muller et al., 1988]) are currently being designed. It is our firm belief that within a few years commercial systems will be available for doing animations at the pace of the analyst and this will become a standard technique for the analysis of satellite data and the output of image understanding systems.

7 Acknowledgements

Small sections of the satellite image renderer were written under an SERC grant concerned with the development of "real-time 2.5D Vision Systems" (see [Muller, 1988a]). The ray-tracer was developed under a research contract with BP Petroleum Development Limited. We would like to thank Lyn Quam (SRI), Jeffrey Hall (JPL) and Alex Pentland (MIT) for useful discussions and M. P Foin, IGN for providing the SPOT data, underflight photographs and control points through the aegis of the SPOT-PEPS campaign.

References

[Bhanu and Ho, 1987] B Bhanu and K Ho. CAD-based 3D object representation for robot vision. *IEEE Computer*, 22–36, August 1987.

[Blinn, 1977] J F Blinn. Models of light reflection for computer synthe-iszed pictures. *Computer Graphics (Proc. SIGGRAPH 77)*, 11(2):192–198, 1977.

[Brew and Neyland, 1980] A N Brew and H M Neyland. Aerial photography. In *Manual of Photogrammetry*, pages 280–283, ASPRS, 1980.

[Carpenter, 1984] L Carpenter. The A-buffer, an antialised hidden surface method. *Computer Graphics (Proc.SIGGRAPH '84)*, 18(3), 1984.

[Cook and Torrance, 1981] R L Cook and K E Torrance. A reflectance model for computer graphics. *ACM Trans. on Graphics*, 1(1):7–24, 1981.

[Day and Muller, 1988] T Day and J-P Muller. Quality assessment of digital elevation models produced by automatic stereo matchers from SPOT image pairs. *IAPRS*, 27-A3, 1988.

[Dowman and Muller, 1986] I J Dowman and J-P Muller. Real-time photogrammetric input versus digitised maps: accuracy, timeliness and cost. *Proc. Auto-Carto London*, 1:583–543, 1986.

[England, 1988] N England. Application acceleration: development of the TACC-1. *SunTechnology*, 34–41, Winter 1988.

[Fergason, 1983] J L Fergason. *Liquid Crystal Display with Improved Angle of View and Response Times.* Technical Report 4,385,806, US Patent, May 1983.

[Fournier et al., 1982] A Fournier, D Fussell, and L Carpenter. Computer rendering of stochastic models. *Comm. ACM*, 25(6):371–384, 1982.

[Frenkel, 1988] K A Frenkel. The art and science of visualising data. *Comm. of the ACM*, 31(2):111–121, 1988.

[Garvey, 1987] T D Garvey. Evidential reasoning for geographic evaluation for helicopter route planning. *IEEE Trans. on Geoscience and Remote Sensing*, GE-25(3):294–304, 1987.

[Gelberg and Stephenson, 1987] L M Gelberg and T P Stephenson. Supercomputing and graphics in the earth and planetary sciences. *IEEE CG & A*, 26–33, July 1987.

[Gosling et al., 1988] J Gosling, T Hoeber, and D Rosenthal. Programming with news. *SunTechnology*, Winter:54–59, 1988.

[Gourard, 1971] H Gourard. Continuous shading of curved surfaces. *IEEE Trans. Computers*, C-20(6):623–629, 1971.

[Kajiya, 1986] J Kajiya. The rendering equation. *ACM SIGGRAPH86, vol.20(4)*, 1986.

[Kolbusz, 1988] J Kolbusz. *Visualisation of 3D Scientific Data.* Bsc Project Report, Dept. of Computer Science, University College London, March 1988.

[Magnenat-Thalmann and Thalmann, 1987] N Magnenat-Thalmann and D Thalmann. An indexed bibliography on image synthesis. *IEEE CG & A*, 27–38, August 1987.

[McCormick et al., 1987] B H McCormick, T A DeFanti, and M D Brown. Visualisation in scientific computing. *Computer Graphics*, 21(6):6–26, 1987.

[McCullagh and Ross, 1980] M J McCullagh and C G Ross. Delauney triangulation of a random data set for isarithmic mapping. *Cartographic Journal*, 17(2):93–99, 1980.

[Mitchell, 1987] D. Mitchell. Generating antialiased images at low sampling densities. *Computer Graphics (Proc. SIGGRAPH87)*, 21(4), 1987.

[Muller, 1988a] J P Muller. Key issues in image understanding in remote sensing. *Phil. Trans. R. Soc. Lond. A*, 324:381–395, 1988.

[Muller, 1988b] J-P Muller. Image understanding system components in remote sensing. In *Proc. ICPR, Beijing, China*, October 1988.

[Muller and Anthony, 1987] J-P Muller and A Anthony. Synergistic ranging systems for remote inspection of industrial objects. In *Proceedings of "2nd Industrial and Engineering Survey Conference", London*, 1987.

[Muller and Saksono, 1988] J-P Muller and T Saksono. Fractal properties of terrain. *IAPRS*, 27-A3, 1988.

[Muller et al., 1988] J-P Muller, K A Collins, G P Otto, and J B G Roberts. Stereo matching using transputer arrays. In *Proceedings of the XVIth International Congress of ISPRS, Kyoto, Japan, IAPRS 27-A3*, 1988.

[Newman and Sproull, 1973] W M Newman and R F Sproull. *Principles of Interactive Computer Graphics.* McGraw-Hill, New York, 1973.

[Nicodemus, 1977] Nicodemus. *Geometrical Considerations and Nomenclature for Reflectance.* NBS Monograph 160, US. Dept. of Commerce, 1977.

[O'Neill and Dowman, 1988] M A O'Neill and I J Dowman. The generation of epipolar synthetic stereo mates for SPOT images using a DEM. In *Proceedings of the XVIth International Congress of ISPRS, Kyoto, Japan, IAPRS 27-A2*, 1988.

[Phong, 1975] B-T Phong. Illumination for computer-generated pictures. *Comm. of the ACM*, 18(6):311–317, 1975.

[Quam, 1987] L Quam. private communication, 1987.

[Robertson, 1987] P K Robertson. Fast perspective views of images using one-dimensional operations. *IEEE CG & A*, 47–56, February 1987.

[Rogers, 1985] D F Rogers. *Procedural Elements for Computer Graphics.* McGraw-Hill, New York, 1985.

[Schachter, 1983] B Schachter, editor. *Computer Image Generation.* John Wiley & Sons, New York, 1983.

[Sutherland et al., 1974] I E Sutherland, R F Sproull, and R A Schumacker. A characterisation of ten hidden surface algorithms. *ACM Computing Surveys*, 6(1):1–55, 1974.

[Winkler et al., 1987] K-H A Winkler, J W Chalmers, S W Hodson, P R Woodward, and N J Zabusky. A numerical laboratory. *Physics Today*, 28–37, October 1987.

[Wolff, 1987] R Wolff. Panel on "the physical simulation and visual representation of natural phenomena". *Computer Graphics (SIGGRAPH 87)*, 21(4):337–338, 1987.

[Woodham, 1985] R J Woodham. A computational vision approach to remote sensing. In *Proc. IEEE Computer Vision & Pattern Recognition Conf.*, pages 2–11, 1985.

[Zyda et al., 1988] M J Zyda, R B McGhee, R S Ross, D B Smith, and D G Streyle. Flight simulators for under $100,000. *IEEE CG & A*, 19–27, January 1988.