How Virtual is Your Reality?

By Clive "Max" Maxfield

Our very own Clive "Max" Maxfield reports from the land of virtual reality at the I/ITSEC show.

It's rare that you see something you believe could make a significant impact on the world, but this happened to me at the beginning of December, 1998.

One of the more interesting conferences and trade shows of the year is *I/ITSEC (Interservice/Industry Training, Simulation, and Education Conference)*, which is held in Orlando, Florida, USA. This conference is unusual in that it is held immediately after Thanksgiving weekend, but it was explained to me that this gives the American Generals an excuse to bring their families to Walt Disney World, so that's OK.

The reason I/ITSEC is so interesting is that it offers a showcase for the state-of-the-art in computer-based visual simulation and virtual reality. Governments spend huge amounts of money on training, and in today's restricted economy they want to get the "biggest bang for their buck". For example, by the time you take into account wear and tear, fuel, cost of munitions, and environmental impact, it can cost as much as $100,000 US Dollars every time you start up a tank (and this doesn't take account of the danger to personnel that is inherent in live-fire exercises)! In fact the National Training Systems Association estimates that US government investment in modeling and simulation (M&S) is currently running at about $3 billion US Dollars per year, which equates to some honking big and mega-cool flight and tank simulators, let me tell you!

MODELING THE WORLD

The company that almost blew my socks off and takes my personal "most amazing idea" award is Geometrix Inc, San Jose, CA (**www.geometrixinc.com**). The first product from this small (12 person) company is 3Scan™, which provides a way of quickly creating models for use in 3D visualization, simulation, and animation environments. The idea is that you put the object to be scanned on a computer-controlled turntable and scan it with a single video camera Fig.1).

The camera I saw them using was a $120 "cheap-and-cheerful" version – you could get significantly better results with a more expensive unit. Once you've instructed the computer how many images you want to use, it automatically rotates the turntable and captures the appropriate data; for example, 20 images = 18 degrees between images.

The key points here are "Standard Video Camera" (i.e. no lasers or anything like that), and "Just one camera", which means that you don't require multiple cameras for stereoscopic viewing to calculate depth.

Once the scanning process is finished, the 3Scan program examines the first frame to locate and identify key features in the image (the algorithm that

Fig.1. A computer-controlled turntable and a Video camera are used to scan an object.

Copyright © 1998 Wimborne Publishing Ltd and Maxfield & Montrose Interactive Inc

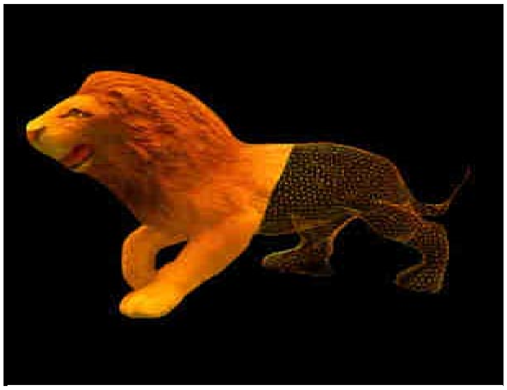

Fig2. The result of the scanning process is a very tasty 3D model.

does this is pretty sophisticated in its own right). The program then determines how far each of these features moves from frame to frame. Also some features become obscured and others become visible as the object rotates, so the program adds new features and discards old ones as it proceeds.

The program uses parallax to determine the relative positions of each of these key features (which ones are closer and which ones are further away). (In this context parallax is the apparent displacement of each of the features from one frame to another.)

The end result is that 3Scan processes the raw video data and generates a very tasty 3D model with appropriate textures (surface characteristics) (Fig 2.)

These models are essentially a mesh of polygons (triangles) with textures applied to them. This means that the models wouldn't really be suitable for individual animation, but they would make great general-purpose models for inclusion in 3D scenes.

But then I saw …

I thought that 3Scan was pretty cool, but then the guys from Geometrix introduced me to SoftScene™, which is where things start to get really amazing. With SoftScene, the clever rascals have extended the 3Scan concept to allow someone with a single hand-held camera to say walk around a building filming it.

As for 3Scan, SoftScene uses advanced image processing algorithms to detect and extract prominent visual features in the scene (Fig.3). In this image, green cross-hairs indicate features that have been detected in the current frame.

Fig.3. SoftScene uses advanced image processing algorithims to detect and extract prominent visual features.

SoftScene tracks the prominent features from frame to frame, re-invoking the feature extraction algorithm whenever additional features are required. In Fig.4, cyan lines are used to indicate the *image displacement vectors* of the selected features (the green crosshairs indicate where the features are in the current frame). SoftScene then employs advanced *structure-from-motion* algorithms based on analyzing motion parallax to spatially locate the tracked features.

One very clever aspect of all of this is that the person holding the camera obviously isn't going to walk around the building following an exactly circular path. So one of the things the SoftScene program has to do is to calculate the path of the camera in order to compensate for the operator moving closer to and father from the building. The image in Fig.5 shows the fully reconstructed camera path along with a "point cloud" representing the final set of tracked features in 3D space. The green and red camera icons represent the first and last frames, respectively (camera icons are shown 12 frames apart in this image).

Fig.4. cyan lines are used to indicate the image displacement vectors of the selected features

Copyright © 1998 Wimborne Publishing Ltd and Maxfield & Montrose Interactive Inc

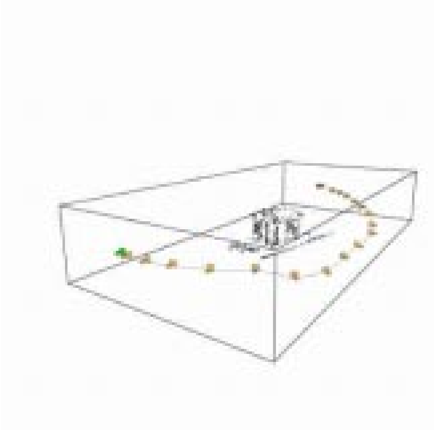

Fig.5 shows the fully reconstructed camera path

In fact some people (like the creators of Jurassic Park, for example) are really interested in knowing the path of the camera for the purposes of combining live video with animated characters. Using the SoftScene process, they could just give the program a video clip and it could tell them the path of their camera (this is essentially a product in its own right).

But returning to our building, to complete the process SoftScene uses advanced topology and computational geometry algorithms to construct a polygonal mesh from the data contained in the 3D point cloud (Fig.6). Textures are then automatically extracted from the video imagery and applied to the surfaces of the model, resulting in the model shown in Fig.7.

We should note that this was an early beta of the SoftScene program – Geometrix are planning on formally unveiling the full product at the Siggraph conference in mid-1999. Also, Geometrix have recently been awarded $950,000 in development funding to attach a GPS receiver to the camera and to record GPS data on the audio track of the video, which will allow them to generate extremely accurate models.

Tremendously exciting

I think that the SoftScene technology is tremendously exciting. It seems to me that there are almost limitless applications in all walks of life

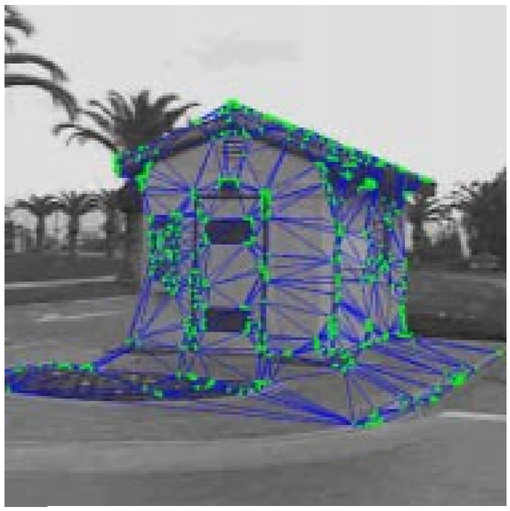

Fig.6. A polygon mesh is generated from data contained in the 3D point cloud.

that would benefit from the ability to create models using this technique. For example, imagine a police officer at the scene of an accident walking round with a video camera filming everything. Later, that video could be converted into an accurate 3D model that you could "walk through". This would make it much easier to determine whether people could see what they said they saw from the positions they were in, and so forth.

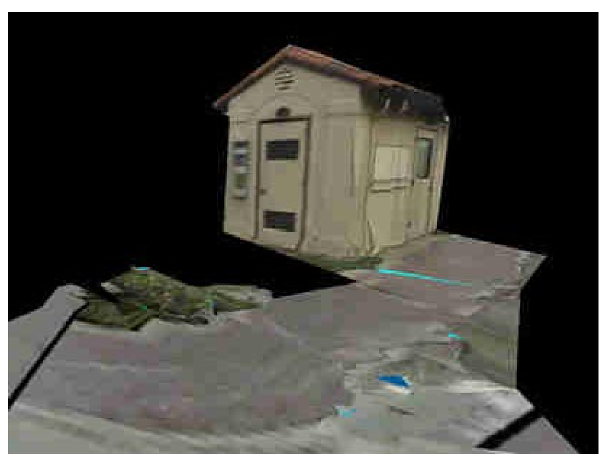

Fig.7. Final 3D model

Another potentially interesting market is television studios, which are increasingly making use of "virtual sets". These computer-generated sets offer many advantages, such as the fact that you don't have to set them up, take them down, or store them, and also that you can quickly and easily change things like colors and textures. The problem is that creating 3D computer models for these virtual sets can be extremely time-consuming – SoftScene technology would make creating the models for these virtual sets MUCH easier.

Geometrix emphasize that their technology is not currently targeted towards the commercial (home) market, but I can also see many potential applications here. For example, imagine filming the inside of your house with your video camera and using SoftScene to create an accurate 3D model of your home. This would allow you to experiment with alternative decors in a virtual environment – the possibilities really are endless!

Copyright © 1998 Wimborne Publishing Ltd and Maxfield & Montrose Interactive Inc

EYE EYE …

Several companies were demonstrating eye-tracking systems, including SMI (SensoMotoric Instruments from Germany, **www.smi.com**) and ASL (Applied Science Laboratories, Bedford, MA).

The most sophisticated version is attached to a head-mounted assembly, with sensors that track the position of each eye (and the size of your pupils and suchlike).

This data can then be displayed on a computer screen superimposed on an image of what you were looking at the time. This lets you model the path the observ-ers' eyes took, combined with other data such as how long the eyes rested at each point.

Applications for this sort of technology range from flight simulators (for example, see what instruments the pilot looked at and in what order during a simulated emergency) to web marketing (to see how users read web pages – what they look at and what they miss or ignore).

Although the head-mount unit was a little obtrusive, it did have the advantage that it was position-sensitive, which allows the system to compensate for movements of the head. Also, this would be far less obtrusive if it were mounted directly into a fighter pilot's helmet, for example.

Of course applications of the web-marketing ilk don't require this level of sophistication, and simpler units are available that can be mounted on your desk or on top of the computer screen and track your eye motions from that perspective.

"I SEE NO SHIPS…"

Another technology that scored high on my "cool-meter" was the concept of virtual binoculars, which were being displayed by companies such as Virtual Research Systems Inc, Santa Clara, CA (**www.virtualresearch.com**) and n.Vision, McLean, VA (**www.nvis.com**).

These little rascals look and feel like a standard pair of binoculars – for example, you can adjust the focus of each eye, widen or narrow the spacing between the eyes, and so forth. The difference being that the images seen by each eye are generated by a computer connected to the binoculars and displayed on small LCD screens inside the binoculars.

One clever aspect of this is that the binoculars track their position in 3D space, so as you turn your head (and body) the scene changes appropriately. For example, when I looked through them, I could see a plane coming towards me and passing overhead, and by moving my body I could track the plane, zoom in on it, and so forth.

This sort of device can be very useful for such tasks as training forward observers to recognize planes and vehicles in a battlefield situation without the cost of actually fielding the real units.

IF AT FIRST …

Whenever I hear the term "simulators" I think of flight simulators (e.g. planes and helicopters) or vehicle trainers (e.g. tanks and suchlike), so I was somewhat surprised when I meandered into the booth of Jason Associates Corporation, Draper, UT (**www.jason-sed.com**) to find a smorgasbord of crane simulators.

The idea is that you sit in a seat from a real crane (well, a seat made by the same company that makes the seats for the real cranes) with real controls facing a large screen showing the load on the end of the hoist, and then you practice, and practice, and …..

Of particular interest was a dock-side simulator that modeled the effects of the swell of the sea, the motion of the boat, the effects of the wind on the load, and so forth (in the more sophisticated units the seat has motion and the controls have tactile feedback).

It obviously makes a lot of sense to practice this sort of thing in a virtual world before you are let loose on a real crane.

Jason Associates started off modeling cranes for military applications, but now they're moving into the commercial space. As the representative I was talking to said: "Have you any idea how many cranes there are in the world?" I must admit that I hadn't previously spent a lot of time pondering this poser, but it certainly gives one pause for thought.

VISIONDOME

If you place your hands to either side of your head, and restrict yourself to about a 60° field of view, the resulting image is approximately the same as you see on a computer screen. Some applications, however, require a greater field of view to more fully immerse you into the simulation.

One very nice solution demonstrated at I/ITSEC was the VisionDome® from Alternate Realities Corporation, Durham, NC (**www.virtual-reality.com**)

Copyright © 1998 Wimborne Publishing Ltd and Maxfield & Montrose Interactive Inc

(Fig.8). These little beauties range from 16 feet wide by 10 feet high (for 2 to 5 viewers) up to models that are 26 feet in diameter and can accommodate over 45 viewers.

The nice thing about the VisionDome is that it provides a fully immersive multi-user virtual reality environment without requiring goggles, glasses, helmets, or other restrictive devices. This makes the VisionDome an ideal choice for multi-user simulations, training, design, engineering, and so forth. For example, the virtual crane simulations discussed above would be even more realistic if they were to use a VisionDome as their display device.

As I was to discover, there are three very clever portions to this system:

- The dome itself
- The projector/lens
- The software

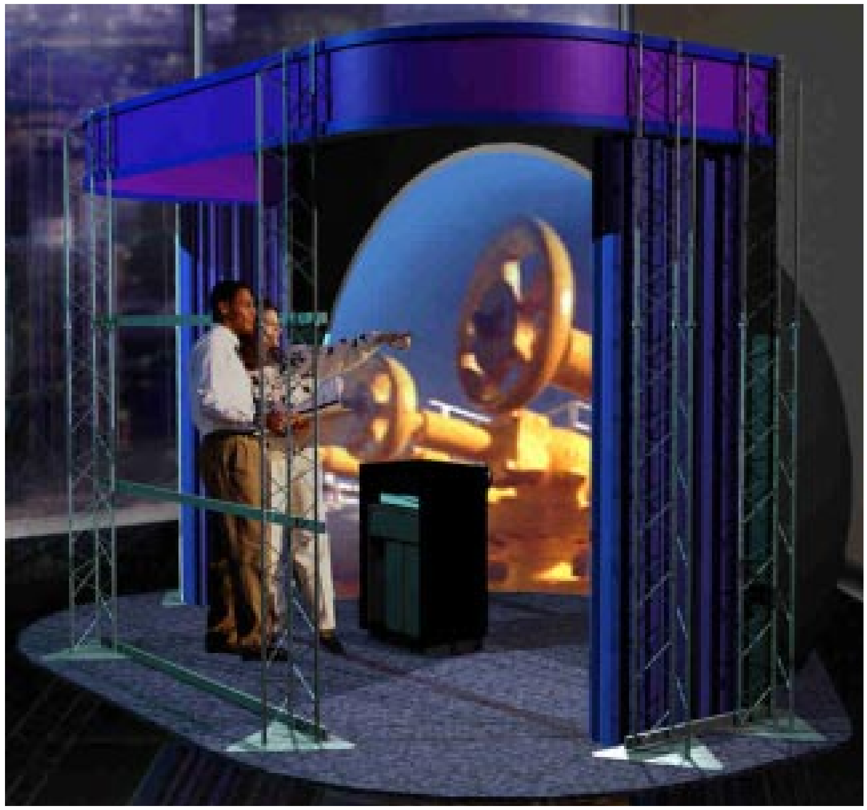

Fig.8. The VisionDome from Alternative Reality Corp

The dome

The reason I say the dome was clever, is that the way it was constructed made it extremely portable and easy to erect. First of all you raise an aluminum frame, then you clad the outside of the frame with a thick black covering and the inside with the white lining material. The cunning part of all this is that if you were to try to attach the white lining to the struts of the frame, you would end up with distortions on the surface of the dome. Instead, the white lining is only attached around the edges of the frame, then a small pump is used to create a slight negative pressure by sucking out the air from between the black cladding (on the outside of the frame) and the white cladding (on the inside of the frame). This results in the inner lining being stretched taught and forming a perfect hemisphere.

The projector & lens

The next interesting point was the projection system. Instead of having red, green, and blue “guns”, the VisionDome employs a single-unit projector, which means that there’s no problems with alignment. But the really clever part of the projection unit was the lens, which somehow manages to project a perfectly focused image across a 180° field of view.

Furthermore, this lens has an infinite focal length. This means that you can hold a piece of paper a few inches from the lens and see that portion of the image in focus. Then as you move the paper further and further away the image stays in focus without any correlation problems. I don’t have a clue how this works, but the end result is that the image was in perfect focus at every point inside the dome.

The software

Last but not least is the software. In order to appreciate what this does, you first have to understand a few basic concepts about 3D graphics. First of all, a 3D graphics application performs a process called “tesselation”, in which all of the objects in the 3D world are converted into collections of simple polygons (usually

Copyright © 1998 Wimborne Publishing Ltd and Maxfield & Montrose Interactive Inc

triangles). Second, the vast majority of today's high-performance 3D graphics applications use a programming interface called OpenGL®, which allows the application to pass requests to graphics accelerator boards in a well-defined format.

The problem is that OpenGL was conceived with flat-screen displays in mind. In the case of the VisionDome, if you want to display a straight line, then that line has to be curved so that it appears straight on the curved screen (if you see what I mean). In order to handle this, Alternate Realities' software takes the OpenGL calls and re-tesselates everything on the fly (for each frame) so that it can correct for the curvature of the screen.

The overall effect is one of total immersion – I now know what it feels like to fly through a 3D molecule, and that's not something you expect to hear yourself saying too often!

3D GRAPHICS

If you want to know more about 3D graphics, including the meaning behind terms like tesselation, OpenGL, textures, and suchlike, then we would recommend the easy-to-understand introductory book *3D Graphics Supercomputing on Windows NT*, as described below.

Graphics Book: Special Offer!

As computer generated images become more complex and beautiful, the technology behind them has become more involved and harder to understand. Choosing the best graphics system for your applications and needs can therefore be a daunting task. Fortunately help is at hand, because the computer graphics technology leader on Windows NT – Intergraph Computer Systems – has written a really great book that explains computer graphics technology in an interesting and understandable way.

The book commences by introducing 2D bitmapped and vector graphics, and then moves on to discuss 3D graphics concepts. These introductions are easy to understand and set the scene for what is to come. As you will discover, the main portion of the book is occupied by an extensive glossary, which covers everything from the different schemes for representing color and translucency, to lighting, shading, texturing, and rendering.

In addition to the easy to understand text, the book is jam-packed with superb color graphics that fully illustrate the important points behind each topic. The cover price of this 147 page book is $49.95 US Dollars, but readers of *EPE Online* can purchase this extremely useful reference for only $19.95 (plus shipping & handling).

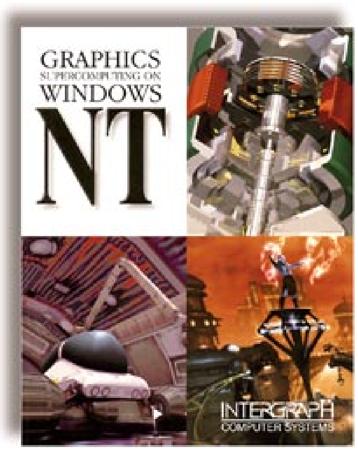

You can download a sample chapter of this book (the chapter on 3D Graphics) by bouncing over to the *EPE Online* Web site at **www.epemag.com** and strolling into the *Library.* If you like what you see, you can purchase the full printed book directly from the *EPE Online Store*.

Go to next section

Copyright © 1998 Wimborne Publishing Ltd and Maxfield & Montrose Interactive Inc