

SCENE REALISM: THE SYNERGY OF DATA BASE TECHNOLOGY AND CIG HARDWARE

Howard Wilkerson
Thomas C. Brown

Evans & Sutherland Computer Corporation
Simulation Division
Salt Lake City, Utah

ABSTRACT

Combat readiness training for pilots in the next decade requires realistic computer-image-generator (CIG) scenes that are correlatable to large real-world gaming areas. Major advances in hardware texture, memory, and polygon capacities by themselves are insufficient to sustain the desired increases in training effectiveness. Innovative data base technologies can leverage improvements in hardware capacity, architecture, and algorithms.

This paper describes how this synergy of software and hardware is achieved in current CIG technology. The availability of many thousands of texture maps makes possible the use of high-resolution, remotely sensed image data to create Geo-Specific texture over the entire gaming area. This textured terrain has much higher fidelity because it contains small-scale feature elements that correspond to real-world objects. Geo-Specific texture reduces the number of required ground polygons, which then permits increased density of three-dimensional polygonal features. The powerful data base compression technique of instancing, previously somewhat restricted in application mainly to terrain polygons and cultural objects, is now extended (with an enhanced hardware architecture) to the application of Geo-Specific texture with smooth shading.

INTRODUCTION

High-performance computer image generators (CIGs) for flight simulation have undergone considerable evolution during the 1980s in their ability to provide realistic, out-the-window, visual images with complex scene content. The use of photo-derived texture, especially Geo-Specific texture, has matured so that the terrain is recognizable and can be correlated to real-world areas.

Image generators and their data bases need to satisfy a broad range of requirements for use in flight and vehicle simulation. In addition to providing optical flow, closure rates, and scene density, CIG data bases must represent the real world with as much fidelity as possible while simultaneously covering as large a gaming area as possible.

Recent Enhancements to CT Image Generators

The challenges above are met with the following enhancements to the CT family of image generators:

- Many thousands of separate texture maps are now available for use in the data base.
- The on-line (semiconductor memory) and off-line (disk) storage capacity for data bases has been quadrupled.
- Hardware data base compression techniques have been improved.

- Faster algorithms have been incorporated into both the hardware and the microcode.

Concurrent with hardware development, new data base methodologies have been developed. In particular, satellite imagery and aerial photographs have been used to generate Geo-Specific texture for the terrain surface.

Texture on the CT Hardware

A texture map is a two-dimensional array of data elements used to modify the intensity, color, or transparency of a polygon. Each texture element in the array is called a texel, which is the analog of a pixel in raster displays.⁽⁶⁾

The data within the texture map may be designed so that the map is self-repeating; the applied texture is a two-dimensional, spatial, periodic function. A self-repeating map may be either photo-derived⁽¹⁾ or algorithmically generated. An example of the use of a photo-derived, self-repeating map is for the ocean surface. Self-repeating maps cover large regions with minimal on-line and off-line storage requirements.

A Geo-Specific texture map, however, is discrete (non-repeating) and one map is allocated to a specific geographic section of terrain. When maps are juxtaposed, their shared boundary is not evident. The availability of thousands of texture maps in the image generator allows an entire data base, and not just a high-resolution inset, to be carpeted with Geo-Specific texture.

CORRELATION TO THE REAL WORLD

Correlation Criteria

Geo-Specific texture on the terrain model is one criterion for creating a data base that corresponds to the real world. Two other criteria include: 1) geometric terrain fidelity and 2) three-dimensional feature representation.

The terrain polygons must provide correct elevations, and the shape (especially the ridges on the horizon) must be accurate. The degree of real-world correlation for the terrain increases as polygon edge lengths decrease and polygon counts increase.

Small-scale terrain elements, such as gullies or the intricate patterns created by vegetation, are modeled with Geo-Specific texture. A metric for how well the actual world is captured is texel size. Texel-size limiting factors include the source-data resolution and/or the on-line texture storage capability.

Large-scale polygonal items, such as roads, major buildings, towers, and coastlines, are recognizable in the data base and on maps and Joint Operations Graphics (JOG) charts. However, merely recognizing and matching items in the data base to a map or chart may be inadequate to meet the needs of training for combat readiness. Although every specific object within the data base may correspond to an item on a map or a chart, the reverse has not always been true. The advent of Geo-Specific texture (1) opened up the possibility of using texture to display much more real-world scene detail in a data base than would otherwise be possible using polygons. Now almost every item on a JOG chart or large scale (1:24000) topographical map may be represented in the data base.

As the ability to produce detailed data bases that closely match the actual portion of the Earth being modeled increases, improved source material for texture maps, terrain elevations, and feature placement is needed. In addition, this source information must be processed rapidly and automatically if large data bases are to be built without incurring prohibitive man-hour costs.

Previous Work with Digital Cartographic Data

In the past, digital cartographic data has seen widespread use for efficient production of terrain models and for feature placement.⁽⁷⁾ This data is available from a number of sources, and is based upon the Defense Mapping Agency (DMA) Digital Terrain Elevation Data (DTED) and Digital Feature Analysis Data (DFAD). Level 2 DTED consists of elevations in a grid at three-arc-second increments of latitude and longitude, which are equivalent to 100-meter intervals on the Earth's surface at the equator.⁽⁴⁾ DFAD provides locations of point, linear, and areal features such as lakes, roads, coastlines, buildings, and bridges within the DTED grid, with a point-to-point accuracy of 26 meters.⁽³⁾

However, digital cartographic data may not always be adequate. The 100-meter resolution is low, and can only provide a coarse terrain model. Fine-grained topographic features such as gullies cannot be represented in a polygonal terrain model derived from DTED, even though the image generator would be able to display them if they were available.

DFAD is also limited in the resolution of features captured. A river less than a certain width would not be in DFAD, but it may be visible from an aircraft in actual flight operations and may provide important navigational cues. Only a subset of the real-world features is present, and the feature code may not adequately describe the object. For example, an areal feature designated as forest does not indicate the subtle, slope-dependent variations in vegetation density that occur in mountainous terrain. The maps and DFAD may not be current because man-made cultural features change more rapidly than maps can be updated and published.

DTED and DFAD may not be available for the region desired for the image-generator data base, and obtaining digital data for a new area may require excessively long lead times. In particular, data for regions outside the United States may be difficult or impossible to obtain. Commercially available digital, cartographic data with higher resolution (based on 1:24000 scale quadrangles) is becoming available, but it may be years before all the maps can be digitized.

While DTED and DFAD will continue to be of use, additional source data for the generation of terrain models and three-dimensional features is needed to fully exploit the increasing memory and display capacity of state-of-the-art image generators. This new source material must have high spatial resolution, must be available for worldwide areas, and must lend itself to procedural processing for the large data bases typically used in flight simulation. Satellite imagery and digitized aerial photography are two sources that meet these requirements.

The Use of Remotely Sensed Data

An example of the kind of remotely sensed data that is commercially available in digital format is satellite imagery from the French SPOT (Satellite Pour l'Observation de la Terre) satellite. The data is available on a worldwide basis without restriction, and planned future SPOT satellites assure long-term data availability well into the next century. The 60-kilometer x 60-kilometer image area allows a large CIG data base to be covered with a reasonable number of images. The ground resolution is up to 10 meters for panchromatic images or 20 meters for multispectral images. A rapid revisit capability allows most parts of the Earth to be captured with short lead times if an image is not available from the archives.

Raw satellite data needs to undergo several levels of processing before it is useful for building visual data bases. The data must be radiometrically corrected to equalize each sensor's output, and geometrically corrected to account for satellite and earth motion

due to oblique image acquisition. Ground control points taken from topographical maps can then be used to rectify the image data to a standard map projection.

Stereo pairs of images may be used to extract digital elevations in a grid with up to 10-meter resolution in cartesian *x*, *y*, and *z* coordinates. This allows many more sample elevation points per unit area than are available with DMA data. The result is a digital terrain model that is registered with the panchromatic image.

Satellite imagery lends itself to procedural processing in data base creation. Since the imagery is in digital format and is usually rectified to a known map projection, neighboring or overlapping images may be registered and mosaicked together. These images may then be converted to Geo-Specific texture maps that form a seamless mosaic covering the entire data base.

In addition to satellite imagery, aerial photographs with sub-meter resolution are available from the U.S. Department of Agriculture for most parts of the United States. These photographs may be used to supplement satellite imagery in areas of the data base where higher detail is required.

APPLICATION OF REMOTELY SENSED DATA IN VISUAL DATA BASES

The image data produced from satellite images and aerial photographs may be used as source data to generate Geo-Specific texture maps for visual data bases. The elevation data array obtained from stereo pairs may be used to generate a polygonal terrain model.

Because they are derived from satellite imagery and aerial photographs, the Geo-Specific texture maps give the terrain a much more natural and realistic appearance than if the terrain had texture derived from other sources. The content of the texture map closely matches the character of the real-world region being modeled. The vegetational variations and shading of natural areas such as forest and desert are captured, as are the street and building patterns of urban areas. Every real-world feature larger than the source data resolution is present in the Geo-Specific texture map.

This provides footprints for exact feature placement of desired three-dimensional objects in the data base.

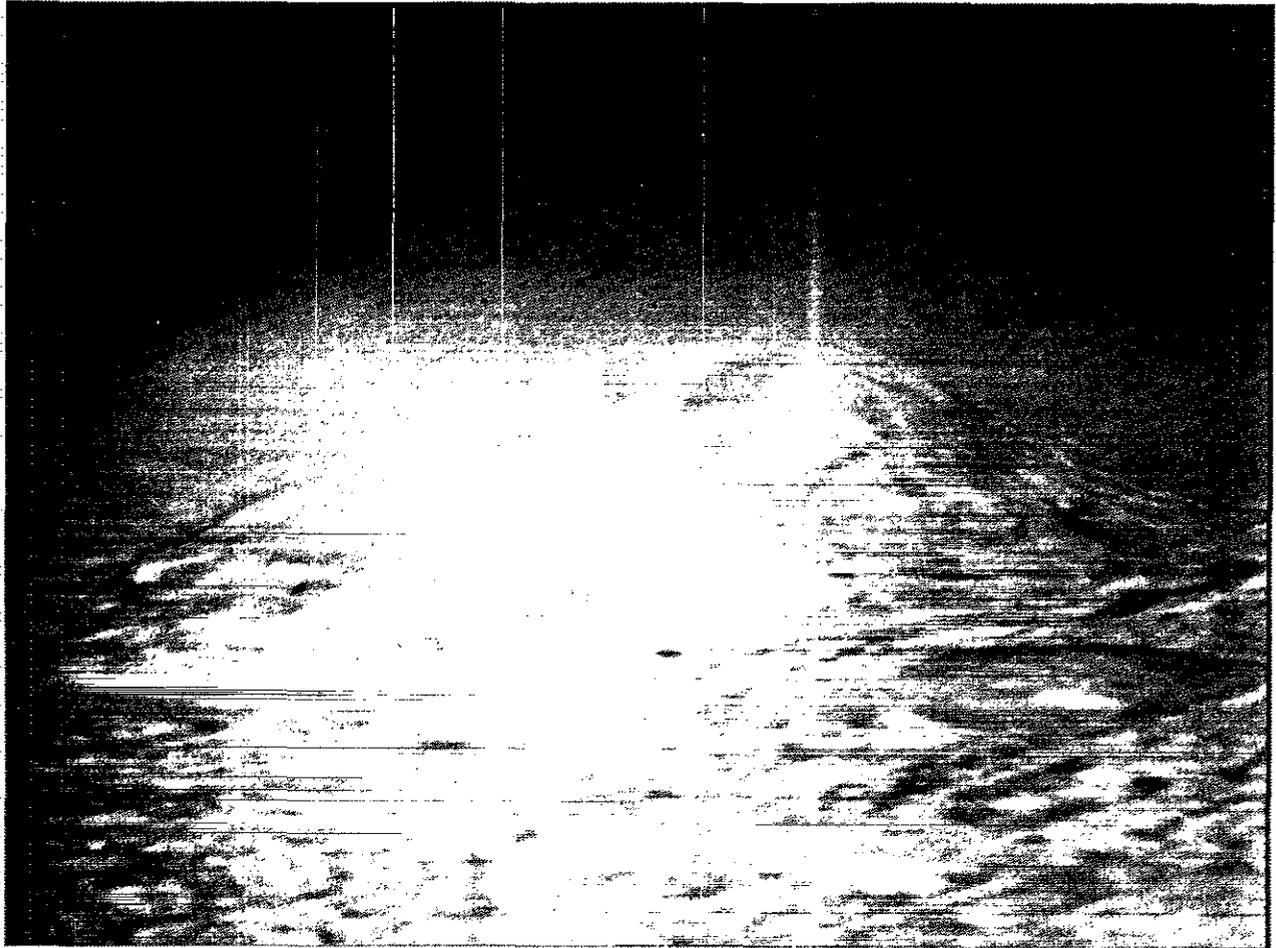
The three-dimensional appearance of the surface is also enhanced because the texture map contains shading information based upon the slope orientation with respect to the sun. Undulations in the terrain, too small to be represented with polygons, appear three-dimensionally because shading provides the eye with important three-dimensional cues.⁽⁵⁾ Depending upon the flight regime, this three-dimensional effect may be used to reduce the number of terrain polygons without introducing any scene degradation. The polygons formerly allocated to terrain may then be used to increase the three-dimensional feature density.

At some minimum altitude above ground level, the texture derived from satellite imagery becomes less useful as the limiting resolution of the source imagery is approached and individual texels become apparent.

If higher-resolution source data is available, as from digitized aerial photographs, this data may be used to create additional higher-resolution, Geo-Specific texture maps. The high-resolution and low-resolution maps must be spatially and contrast-blended to prevent discontinuities. These maps have the most appropriate application as high-resolution inserts into parts of the data base where a higher degree of faithfulness to the real world is required, such as in target areas or in specific flight corridors.

Geo-Specific texture maps may also be enhanced with photo-derived, self-repeating maps. The photo-repeating map is chosen to represent the local ground motif, such as a sagebrush pattern over a desert area. As the altitude decreases, the satellite-derived texture over a desert region transitions properly into a higher-frequency texture pattern. This strategy is especially effective in natural areas, because the texture behaves properly down to the ground level.

Figures 1 through 3 demonstrate the efficacy of using satellite imagery to generate Geo-Specific texture for a visual data base. The texture and elevation data was derived from a stereo pair of panchromatic satellite images supplied by the SPOT Image Corporation. The imagery was acquired by the satellite in mid July of 1987.

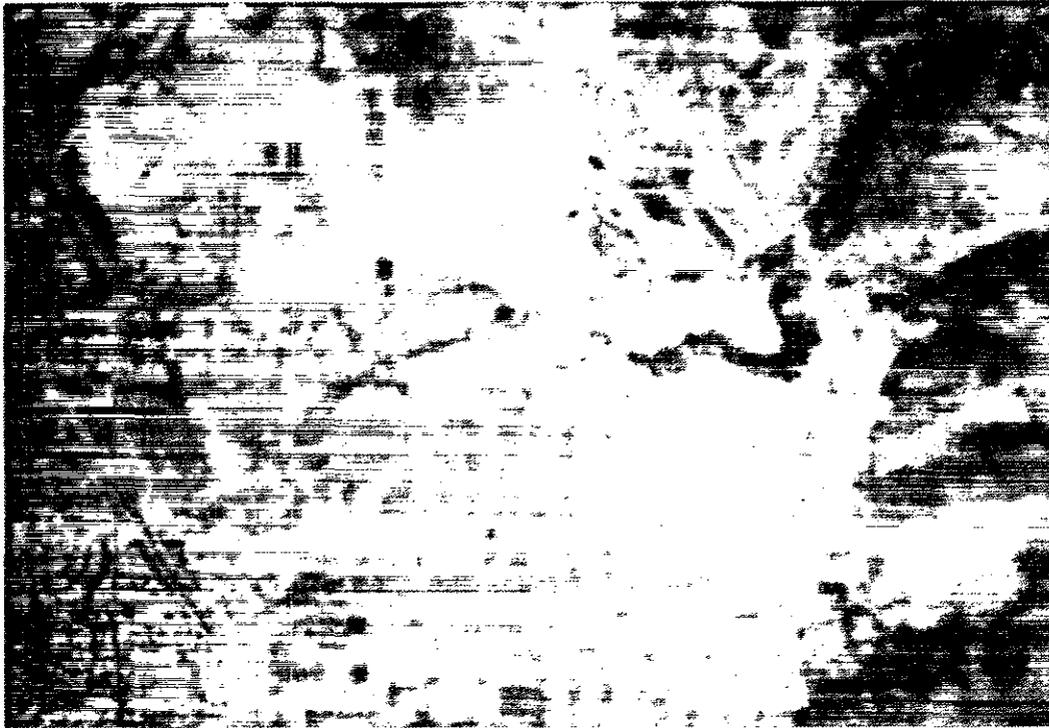


PHOTOGRAPHED DIRECTLY FROM A REAL-TIME CT6 SYSTEM

EVANS & SUTHERLAND

Figure 1

Figure 1 shows the juxtaposition of a high-density, urban area with neighboring mountains in a visual data base. The altitude is 2500 feet above ground level. This was photographed from a real-time image displayed on an Evans & Sutherland image generator. (Source Imagery ©CNES 1988 SPOT Image Corp.)



PHOTOGRAPHED DIRECTLY FROM A REAL-TIME CT&S SYSTEM

EVANS & SUTHERLAND

Figure 2

Figure 2 is an orthogonal view of the same geographic area shown in Figure 1. Note the high-resolution inset in the lower left-hand corner. (Source imagery ©CNES 1988 SPOT Image Corp.)

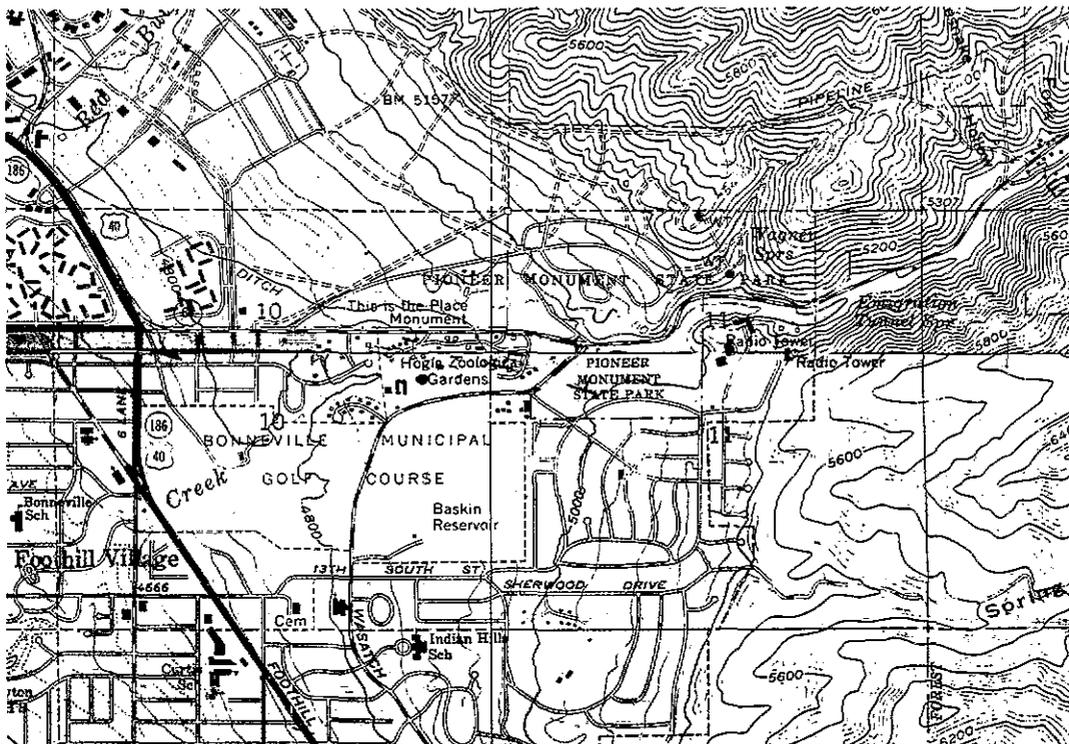


Figure 3

Figure 3 shows two juxtaposed, United States Geological Survey (USGS) topographical maps (Sugarhouse and Fort Douglas, 7.5-minute quadrangles) of the same geographic area pictured in the real-time image shown in Figure 2. Notice that nearly all of the cultural details on this map are also present in the real-time image. However, some details displayed in Figure 2 are not on this map because the map has not been updated recently.

ENHANCED INSTANCING

A direct relationship exists between off-line disk storage capacity and the geographic extent of a data base gaming area. For a data base with homogeneous scene density, the data base area increases in proportion to the quantity of data that can be stored. On the CT image generator, quadrupling the off-line storage enables much larger gaming areas to be modeled.

With instancing and the use of basis sets, a many-fold increase in data base area can be achieved. Hardware instancing and basis sets are effective data base compression methods that have been used for a number of years.(7)

The basic concept of hardware instancing is that one copy of an item, such as a tree, is stored in the data base and the CIG then replicates it many times in the displayed scene.(2) Since data decompression occurs in the CIG at run time, the displayed visual scene may have many more virtual polygons than are stored on disk or in the on-line, image-generator environmental memory. The hardware algorithms to implement instancing have been made significantly faster. This results in a greater on-screen density of instanced features and terrain polygons.

Of particular interest is the use of instancing to model terrain. A terrain surface can be decomposed into a set of triangles, many of which are geometrically identical, but geographically unique. The geometric data for each triangle must be created and stored in the data base only once, and it can then be instanced throughout the terrain skin wherever the geography requires it. The leverage of instancing arises because a relatively small set of polygons can be used to model large regions of terrain.

Each Geo-Specific texture map is unique in its geographic context. Therefore, an instanced terrain polygon must have knowledge of its final destination on the terrain skin if it is to be used with Geo-Specific texture. The CT hardware architecture allows Geo-Specific texture to be applied to instanced terrain polygons with no effect on data base gaming area or terrain-polygon storage requirements.

TEXTURE LEVEL OF DETAIL

A different form of data base compression is used with texture maps. With the thousands of Geo-Specific texture maps that are required for a high-fidelity data base, on-line texture storage requirements would be enormous and quite expensive. In the CT hardware, only those maps needed for the current scene are stored on-line. The maps are then dynamically updated in real time from the off-line disk memory. A hierarchical level-of-detail structure is used in the Geo-Specific texture maps, so that the highest-resolution texture maps, with the smallest texels and the greatest on-line storage requirements, are displayed closest to the eyepoint. The lower-resolution texture maps are used for that part of the image further from the eyepoint.

SUMMARY

Evans & Sutherland has developed an enhanced architecture for its CT family of visual image generators used in flight and land-vehicle simulation. Faster hardware algorithms and increased on-line memory allow higher scene densities because more objects can be stored and processed during each field. Increases in off-line storage, including texture storage, were designed to allow much larger geographic areas to be modeled and to allow extensive use of Geo-Specific texture throughout the data base.

Satellite imagery and aerial photographs are used as source data for generating the terrain surface texture. Stereo pairs of aerial photographs and satellite imagery are used to extract elevations that provide a more accurate terrain model than was previously available with DMA DTED. The data base content is now much more texture-intensive. Geo-Specific texture lends the visual scene a much more real-world appearance.

Instancing has been extended to permit instanced terrain polygons to be used in combination with Geo-Specific texture. Hence, no additional cost in polygon storage capacity is incurred when Geo-Specific texture is used. This allows large-area data bases to be constructed with no loss of terrain fidelity or scene density and with reasonable disk storage requirements.

REFERENCES

1. Clark, L. Charles and Thomas C. Brown. "Photographic Texture and CIG: Modeling Strategies for Production Data Bases," *Proceedings from the 9th IITSC*, November 1987.
2. Costenbader, Jan. "CIG Data Bases in an Instance: Bits and Pieces," *Proceedings of the Image III Conference*, May 30 - June 1, 1984.
3. Defense Mapping Agency. *Product Specifications for Digital Feature Analysis Data (DFAD) Data Base Level 1 and Level 2*, Second Edition, 1986.
4. Defense Mapping Agency. *Product Specifications for Digital Landmass System (DLMS) Data Base*, First Edition, 1977.
5. Livingstone, Margaret S. "Art, Illusion, and the Visual System," *Scientific American*, January 1988, pp. 78-85.
6. Robinson, John and Steve Zimmerman. "Exploiting Texture in an Integrated Training Environment," *Proceedings of the 7th IITEC*, November 1985.
7. Wales, Cary E. and Michael Cosman. "DMA and CIG: A Shotgun Wedding," *Proceedings of the 5th IITEC*, November 1983.

ACKNOWLEDGMENTS

The concepts presented in this paper are the result of an evolution of ideas. Many individuals at Evans & Sutherland have contributed to this paper, and their assistance is gratefully acknowledged.

ABOUT THE AUTHORS

Howard Wilkerson is a Software Engineer for Evans & Sutherland's Simulation Division where he is currently assigned to the CT Exploitation group. He received a BS in Physics from Frostburg State College. He has been with Evans & Sutherland for five years, the last three of which have been in CT Data Base Engineering.

Thomas C. Brown is a Software Engineer for Evans & Sutherland's Simulation Division and is currently assigned to the CT Exploitation group. He received a BS in Computer Science from the University of Utah with an emphasis on medical imagery. His work with medical imagery and computer-assisted surgery has been published. Tom has devoted the past ten years to the simulation product line at Evans & Sutherland and has been a major contributor to the development of photo-derived texture for the CT Data Base Engineering group.