EXTENDING THE TERRAIN COMMON DATA MODEL TO TRAINING SIMULATIONS ON LOW-COST VISUAL SYSTEMS

Guy A. Schiavone and Brian Goldiez

Institute for Simulation and Training
University of Central Florida

Abstract - Most current training simulation systems employ custom generated, ad hoc approaches to building terrain models. Recently, the DARPA STOW Worldwide Terrain Database generation effort, the Joint Simulation System (JSIMS) Synthetic Natural Environment (SNE), and the Joint Warfighting Simulation (JWARS) SNE have cooperatively defined a Terrain Common Data Model (TCDM) for low and medium resolution simulations. The U.S. Army STRICOM Synthetic Natural Environment Science and Technology Objective (SNE STO) seeks to extend the TCDM to other important classes of simulations, thus improving overall interoperability between networked simulation, as well as defining the basis for terrain data production requirements, terrain integration constraints and expectations, SEDRIS transmittal contents, and runtime terrain data use. At IST, our efforts are concentrated on extending the TCDM to low-cost visual systems. Our approach is to develop a detailed specification of terrain database requirements based on tasks commonly performed using low-cost visual systems. In this paper, we chart a course for accomplishing this goal, including describing the development of a structured database of design elements, the definition of a multi-systems view to include multiple resolutions and sensors, a plan for being consistent with initiatives by others working on the TCDM, and our plans for future work in this area.

Guy A. Schiavone has over 10 years of teaching and research experience working in the areas of computer modeling and simulation, numerical analysis, image processing and electromagnetics. He has numerous published papers and conference presentations in the areas of electromagnetics, terrain database analysis and visualization, and distributed computing for modeling and simulation. Dr. Schiavone received the Ph.D. in Engineering Science from Dartmouth College in 1994, and the BEEE from Youngstown State University in 1989.

Mr Brian Goldiez is the Deputy Director of the Institute for Simulation and Training. Mr. Goldiez has a wide range of management and technical experience in modeling and simulation. He has in excess of 20 years experience as an engineer and analyst of various simulations and training devices. Mr. Goldiez also has over 10 years of increasing responsibility in management of high technology research and development groups. His background encompasses all phases of product development from concept research and development to delivery and support. His service includes lecturing, writing, and serving on several simulation organizations. Mr. Goldiez recently contributed to a book on simulation. He is actively involved in several simulation groups and conferences. He consults regularly to industry and government. Mr. Goldiez s principle areas of technical expertise are in systems engineering, modeling physical entities, testing, and graphics systems.

EXTENDING THE TERRAIN COMMON DATA MODEL TO TRAINING SIMULATIONS ON LOW-COST VISUAL SYSTEMS

Guy A. Schiavone and Brian Goldiez

Institute for Simulation and Training
University of Central Florida

BACKGROUND

A great deal of research effort has been expended in investigating and mitigating interoperability problems related to representing the Synthetic Natural Environment in distributed simulation (e.g. DIS, HLA). The effort has been focused on source material, databases, and data formatting. First SIF and now SEDRIS have been developed as a means for providing a common framework for representing the SNE. However, fully representing the SNE in a run-time form suitable for particular combinations of hardware and software applications requires a process of data fusion, manipulation and modeling known as terrain database generation process Traditionally, terrain database generation systems have been built around special-purpose processes and software, with no thought given to interoperability, modularity, or extendibility with anything other than the intended host simulator. This lack of attention to interoperability at the level of terrain database generation unnecessarily complicates the task of devising frameworks for promoting terrain data exchange and reuse. It has become clear that, to promote the cost-saving potential of widespread terrain data reuse, what is required is a data model applicable throughout the terrain database generation process that unambiguously identifies objects and properties of objects that model the synthetic natural environment (SNE).

TERRAIN DATABASE GENERATION

The objective of terrain database generation is to integrate and quickly process into usable form source data such as high resolution digital terrain elevation data, basic feature data such as roads, rivers and vegetation, vector data for 3D modeling of buildings, attribution data such as material codes and naming taxonomies, and corresponding high resolution imagery for photo-realism. Products of the terrain database generation process may include polygonal data suitable for visual simulation, and a correlated CGF terrain database. These products provide the foundation to support a wide range of military training objectives.

Although it is not generally recognized, terrain database generation is often one of the most expensive and time-consuming tasks in the development of human-in-the-loop visual simulations. Much of the expense incurred in terrain database generation is due to the lack of automation throughout the data pipeline, creating the need for extensive "hand-modeling" and artistic license.

Figure 1 shows a terrain generation process that is in some respects typical, but is also very individualized to match the tasks at hand and the software tools available. This particular process was used to generate a terrain database of an urban area for visual simulation using the SVS system, as well as a correlated CGF database that included SAIC s Multiple Elevation Structure (MES) enhancements. In a typical terrain generation pipeline there are several points in the process that require extensive human intervention. The process of extracting from remote sensing imagery the required database primitives, such as geospecific textures, geometry, and attributes (e.g. surface material types), may be semi-automated using modern image processing techniques but almost always requires a large degree of intervention in the current state of the art. Although automated extraction of buildings and other 3D features from imagery has been an active area of open research for several years, this task also still requires extensive interactivity at present. Construction of detailed 3D models is a process that is well-known as being labor intensive, requiring a great deal of domain-specific knowledge and artistry. Another set of hand-modeling (or at best, semi-automated) operations generally referred to as cut and fill operations are applied in the process of integrating into the terrain areal and linear features such as roads, rivers and lakes. These operations refer to local alterations and repolygonizations to the terrain surface necessary to prevent erroneous conditions such as rivers flowing uphill, lake surfaces higher than the surrounding terrain, and roads with incorrect grades, improper banks, or overly-steep shoulders. Cut and fill operations are usually necessary in the process of integrating 3D models into the terrain, as well.

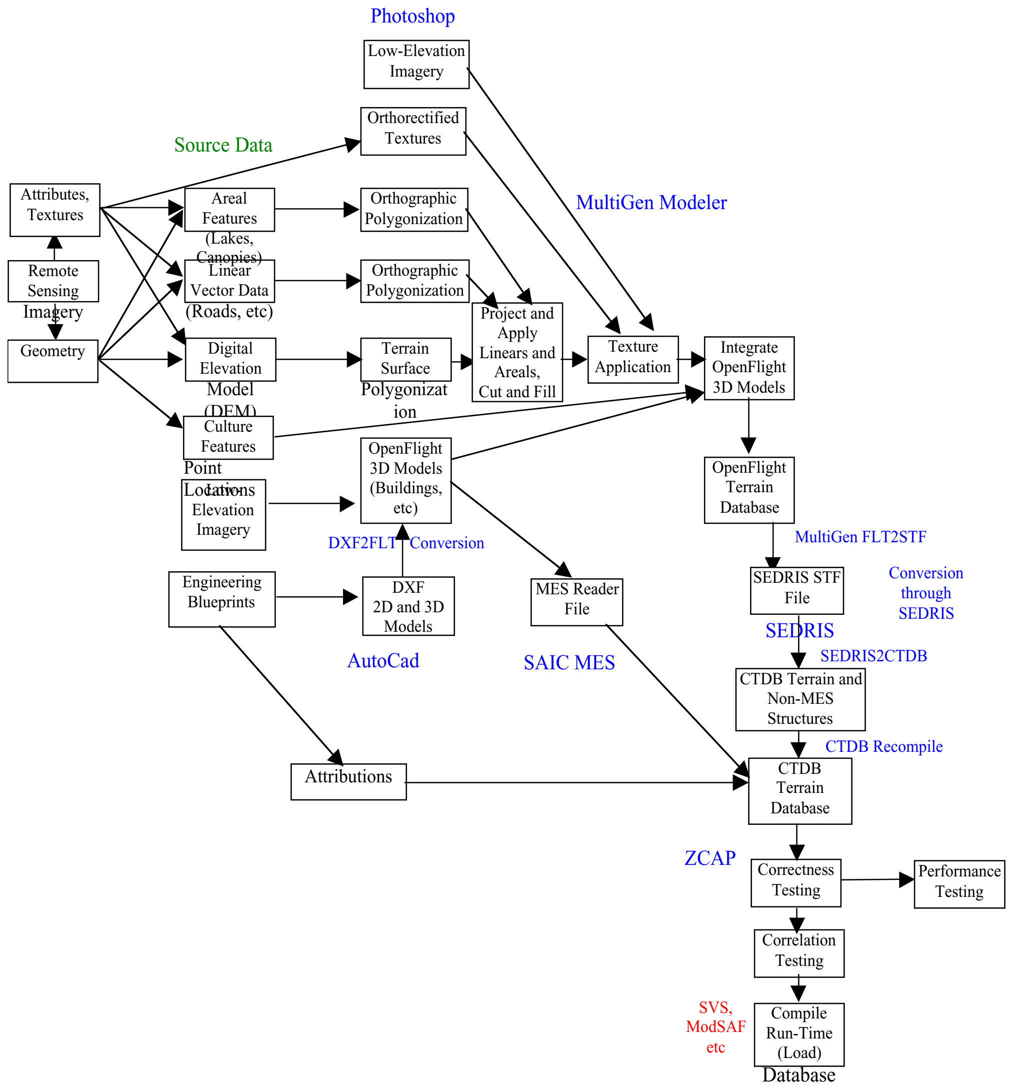

Figure 1. Example of a terrain database generation process to produce a correlated visual and CGF MOUT terrain database including Multiple Elevation Structures..

One of the most time-consuming and expensive steps in terrain database generation comes at the point of acceptance testing, where the end product is tested for correctness and performance. A successful acceptance testing procedure for visual terrain database inspection must assess both the correctness of the database and the suitability of the database for the target application and hardware. Terrain database correctness refers to a lack of gross errors in the database topological structure, with geometric errors within specified tolerances. Terrain database suitability for a given target application and hardware platform is established by meeting the performance requirements of the simulation. Visual terrain databases are specified and developed based on the intended simulation application and targeted computer image generator (CIG). The application drives requirements like minimum frame rate, levels of detail, fidelity, geometric complexity, textures, and the dynamics of the point of view. The need to interoperate with other simulation applications in a distributed simulation exercise imposes additional requirements. The usual approach in the selection of CIG hardware is to determine the degree of complexity to be present in the simulated environment (SE) and specify enough hardware to render that environment at the desired frame rate. The frame rate is also dependent on the application, for example, ground vehicle simulation applications require less dynamic control on the point of view than aircraft simulators. Although performance prediction tools that can be applied at appropriate points in the terrain database generation pipeline are generally non-existent, veteran database modelers have developed the experiences to intuitively predict polygon budgets and database characteristics, so that at run time the desired frame rate is often obtained without excessive iteration. This knowledge on fixing and distributing polygonal density and texture sizes, and other geometry and rendering optimization procedures, is crucial for preventing excessive iteration and attendant expenses. Rejection of the product at the point of acceptance testing most often incurs the substantial penalty of having to iterate far back into the database generation pipeline, due to the high degree of dependency between terrain database elements. For example, poor performance due to excessive polygonal load may require a return to the point of terrain polygonization, which in turn requires a repetition of all subsequent steps including expensive cut and fill and other hand-modeling operations.

GRAPHICS API S

API s for graphics systems have normally been proprietary to the developing company. Open GL broke from that trend when Silicon Graphics released their proprietary graphics API to the public. The API was generally at too low a level to support sophisticated data base developers who required rapid throughput to meet clients needs. Higher level API s, in particular SGI s Performer, grew to support the needs of these sponsors. However, Open GL did not fully support the needs of the M&S community so developers created versions to suit their own product architectures.

In parallel with Open GL were a proliferation of vendor unique API s such as Lockheed Martin s TARGET, Evans & Sutherland s EaSIEST, and Quantum 3D s GLIDE. These API s are tailored to the needs of M&S developers and tuned to their companies hardware.

Microsoft s Direct 3D is a promising offering where Microsoft offers a proprietary API oriented to the Windows platforms for game developers. Many of the features in Direct 3D have applicability to M&S. The status of Direct 3D, though, continues to evolve with new versions and with respect to its residing on Windows.

We were recently hopeful that a joint SGI and Microsoft effort called Fahrenheit would address cross platform and cross operating system issues. Little was released publicly on Fahrenheit other than it would meld the best of Direct 3D and Open GL. However, at the IMAGE Conference in July, 1999 it was disclosed that Fahrenheit would be a proprietary API. At Siggraph 99 SGI disclosed that they abandoning Fahrenheit.

One can conclude from the above discussion that issues with API s remain unsettled. This situation adversely affects optimal reuse of SEDRIS transmittals as well as inhibits the incorporation of emerging data sources into image generators. The authors of this paper believe that the issues noted above indicate an optimal time to combine several related initiatives to achieve improved terrain data interchange and reuse through a focused research program. The suggested strategy is to accommodate the variety in data base approaches being developed by others by conducting analyses and research into new image generator structures and tools that better accommodate the variety in data base sources and user needs. Our goal is to reduce the customizing of data bases for real time rendering and to provide tools and techniques developers can use to support the development of SNE s. A simple illustration of the image generation process that uses SNE data bases, shown in Figure 2, will serve as an orienting point for the research topics described in the following section.

LOW-COST VISUAL SYSTEMS

Low cost visual systems (LCVS) are defined in this paper as Intel-based platforms with hardware acceleration for graphics processing. Until recently, LCVS were restricted to platforms running Microsoft Windows operating systems, but the recent increased availability of image generation hardware drivers for the Linux operating system has made Linux-based LCVS a viable option. An LCVS can be found commercially available today for between $2000 and $20,000. Performance of LCVS is increasing rapidly, as documented by Goldiez and Rogers, 1999. Recent advances move polygon processing from the host processor to the graphics card. Future advances will embed the graphics pipeline on a chip that will either be on a separate card or included on the mother board. Third parties are adding features such as data base paging, gen lock, level of detail control, and fixed frame rates to LCVS making them into real time image generators suitable for modeling and simulation.

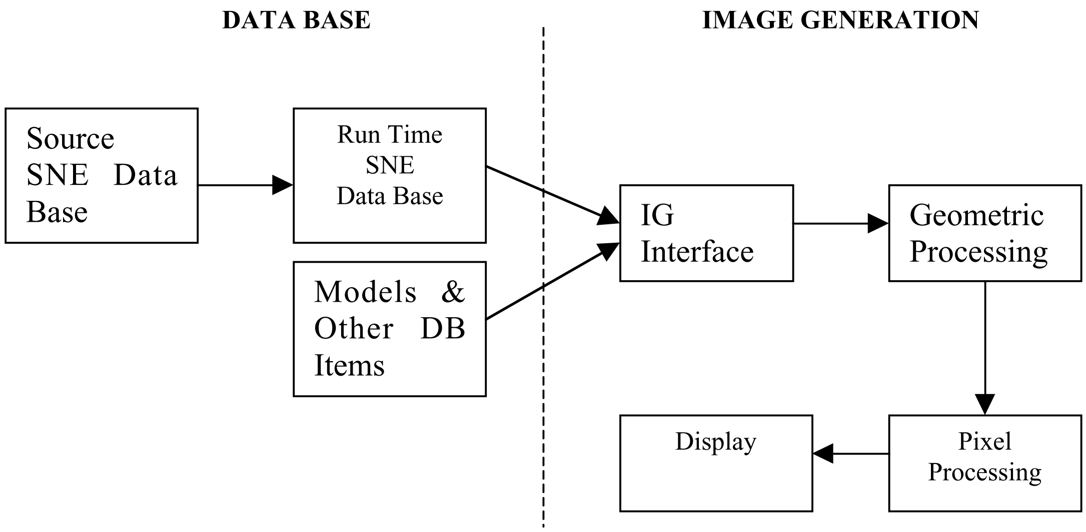

Figure 2. Image Generation for the Synthetic Natural Environment

LCVS are clearly the future for modeling and simulation. Advances being made by commercial companies in processing power are useful to and being harnessed by the modeling and simulation community. A new set of companies have emerged that are adapting commercial computer graphics systems to the needs of the modeling and simulation community. In addition, the commercial graphics community has welcomed the modeling and simulation community s participation in advancing the state of the art in graphics technology. Likewise modeling and simulation is doing its part in increasing LCVS through efforts such as IST s Beowulf Cluster research (Schiavone, Tracy and Palanniapan, 2000).

As with most high performance computers, LCVS generally have a narrow range of high performance. In particular, the software drivers that handle graphics data bases and allocate computing resources to their processing are very sensitive to the data base organization. If one backs up from the driver, it is easy to see that data base format and organization play a strong role in the ultimate efficiency of processing the data. For example, most LCVS have special purpose interfaces for handling data that operate on the data more efficiently than so-called open interfaces such as OpenGL. Performance under open systems, such as OpenGL typically suffer in comparison to the special purpose handler. Steps are being made by the OpenGL community to adapt OpenGL to the special needs of computer graphics companies. For example, the OpenGL Architecture Review Board (ARB) is an industry standards organization that is putting updates into OpenGL to improve its efficiency. LCVS companies participate heavily in the OpenGL ARB. Linkages between groups such as the OpenGL ARB and the Terrain Common Data Model or SEDRIS would provide an important linkage between the format of data and its handling and consumption by LCVS.

TERRAIN COMMON DATA MODEL

In order to improve the reusability and interchange of terrain and synthetic environment data that has been value-added for modeling and simulation, efforts are underway to define a common data model that applies throughout the terrain database generation process. An ongoing standardization effort being undertaken by NIMA to define a similar data model, called the USIGS Conceptual Data Model (UCDM), that is applicable to source data and geographic information systems (Fauquet, 1999). The most promising effort that applies to terrain databases designed for Modeling and Simulation is the extension of the Terrain Common Data Model.

The Terrain Common Data Model (TCDM) is a beneficial mechanism for cataloguing terrain and cultural attributes of data bases for ultimate processing to create computer models of the synthetic natural environment (Birkel, Miller and Root, 1999). The TCDM is based on the SEDRIS EDCS, which is currently being developed as an international standard through the International Standards Organization (Richbourg,, Miller, Foley and Schiavone, 2000). The TCDM, though, was organized around the WARSIM constructive model that deals with aggregated units and is not interactive with human trainees. This is an important distinction because humans require a richness in the synthetic environment for reasoning that is fundamentally different from current computer generated models of human activity. It is important to note that the TCDM is not a physical format, rather, it is a conceptual model for defining environmental data requirements. As such, the TCDM has potential for extension to include concepts that are important in low-cost visual systems. For example, additional attribution could be needed for real time image generators using the TCDM due to the needs of the human consumers of that data. Some steps have already been taken in this direction, such as the proposed extension of the TCDM to general virtual simulations described by Janette, Miller, Nordstrom and Birkel, 2000. The authors of this paper believe that the TCDM can accommodate the needs of visualization in addition to those of constructive models if data content is added and linkages established between communities.

Visual systems, in general, and LCVS, in particular, require means to have TCDM attributes ported to data base generation systems for their ultimate use by the graphics engine. The dilemma in porting TCDM arises in several ways for LCVS. First, are issues in having knowledge of and accessing specific data items from a general structure such as TCDM (the problem of, which house do I select from the TCDM library?). Second, is the issue related to data format, possible need for conversion to a LCVS specific format, and the possible loss of content of the TCDM attribute (the problem of, the red house that needs paint . Third, is data content for specific TCDM elements (e.g., polygon count, texture type, etc.) and impacts of data content on target LCVS (the problem of the red house that needs paint and has red shutters). These issues lead rise to the need for the TCDM to have the needed data included as part of the TCDM framework and for development of so-called data extractors that can extract TCDM elements and format these elements for a target LCVS. In addition, integrated terrain sets, such as towns or basic terrain, need reference metadata, so users know the degree of richness and pedigree in the data they will be using.

There are issues related to the cost effectiveness of data extractors that must be investigated further. Issues include stability of the TCDM structure, completeness of TCDM data elements, the amount of data in the TCDM library, and most importantly, the utility a company would gain from extracting TCDM elements (that are somewhat generic) versus developing them or using existing libraries the vendor may own, and other applications where a TCDM extractor could be useful.

There is a good benefit to TCDM type structures. Vendors that must model or operate on a synthetic natural environment could have common data structures that describe the environment (in our case terrain) and its attributes and link these descriptions to spatial models of the synthetic natural environment. Attributes of terrain and culture should also be described in a lexicon understood and accepted by a wide group of users and developers of synthetic natural environments. The lexicon understanding should extend outside of DoD for maximum coverage and to support economies of scale.

An example helps illustrate the benefit just noted. Consider the requirement to create a synthetic natural environment of a geo-specific region. Terrain elevation data is obtained from a variety of sources in the form of digital elevation points on a grid. The gridded environment is attributed with TCDM culture based on spatial coordinates supplied by a cultural map. Most image generators, including LCVS, must then polygonize the digital elevation model based on what are normally proprietary skinning algorithms. Some culture must then be moved (normally manually) to align with new spatial models of the terrain skin. Attribute characteristics of the terrain and culture can provide valuable clues if alignment is necessary. For example, rivers must flow at the minimum point between adjacent polygons, while trees can tolerate less stringent requirements. A single, consistent mapping of terrain and attribution on a spatial and logical basis would be beneficial. But who is going to create, validate, and maintain such a mapping? This question must be answered if the TCDM is going to extend beyond its WARSIM orientation.

An automated process of bringing data into LCVS from TCDM, whether for attribution or actual rendering through a system specific library is needed. . The authors are currently working with Evans and Sutherland Computer Corporation on a program that will eventually automate the validation process of cultural alignment. But the proliferation and changing architectures of LCVS make the development of a single automated process problematic for using TCDM unless the DoD community gets better engaged, participates, and supports interactions with the commercial sector, such as the OpenGL ARB. LCVS and many other computer graphics systems are driven by commercial games and work station needs, not by M&S forces. One might think that the picture is pretty bleak, but two solutions are offered. The first solution is the DoD nurturing of activities that will nurture TCDM formatters for LCV s. The Army (AMSO, STRICOM, TEC, and others) are nurturing such activity through their activities on SEDRIS and the SNE STO IPT. The second solution is the Army s support of activities to bring M&S and commercial computer graphics into better alignment through participation in commercial graphics activities, including research, prototyping, commercial standards development, and testing techniques.

Progress is being made along both solution fronts. The first solution meets today s needs for TCDM content and products using it, namely WARSIM. The SNE STO IPT has membership that includes, among others, the WARSIM contractor team and its support elements. Other members of the IPT support the needs of the non-WARSIM community and include, for example, linkage with programs sponsored by the NSF that are modeling specific rain forests as well as STRICOM sponsored programs involving LCVS. The second solution involves a longer term commitment to participating in activities that are supportive of the wider computer graphics community. One of these programs involves a new computing architecture for M&S, including graphics engines, while the other involves interacting with the community that is updating OpenGL.

Updates to OpenGL are being made to support new LCVS architectures, such as set top boxes which are actually small scale image generators. As stated previously, there is a very close coupling between real time image generation and the drivers that condition data for use by the image generator hardware. OpenGL deals with primitives that do not directly map to TCDM. TCDM must be shown to be useful to the commercial graphics community for companies to develop links between TCDM and OpenGL. In addition, changes are being made to OpenGL that could dramatically increase its efficiency for real time image generation (e.g., texture compression). It would be very beneficial for LCVS and image generators if the TCDM stored or identified textures in these compressed formats, thereby eliminating the need for special conversion software. Other changes are being considered for OpenGL include accommodation for streaming video, either separately or within a polygonal boundary, and volumetric processing. The point here is that consistent and regular interaction between the TCDM activity and that of the OpenGL development community is necessary to ensure that TCDM is not obsolete by the time it is released.

TCDM should include attribution supporting multi-spectral rendering, by including attributes such as material codes, moisture content, and ambient temperature. Actual algorithms for multi-spectral imaging of specific attributes would be beneficial as a TCDM feature. Including multi-spectral algorithms with an attribute alleviate the need of the data base modeler to guess at the algorithm and opens the door for interface software to use the algorithms, as is.

CONCLUSION

Traditionally, the terrain database generation process has been highly individualized to meet the needs of particular combinations of hardware platforms and software applications, with little or no accommodation for interoperability and eventual exchange and reuse of the resulting terrain database. Efforts underway to extend the Terrain Common Data Model to a wide variety of virtual and constructive simulations should prove beneficial for ameliorating the problems of data reuse and realizing the cost-saving potential of data exchange mechanisms such as SEDRIS. In this paper, we have outlined a plan to extend the TCDM to low-cost visual systems, appropriate to both military and civilian modeling and simulation applications.

REFERENCES

Birkel, Paul A., Miller, Dale D., and Eric Root, The Terrain Common Data Model (TCDM): A Joint Effort of JSIMS, STOW, and JWARS , Proceedings of the 8 th Conference on Computer Generated Forces, Orlando, FL, 1999.

Fauquet, Ronald L., Use of Data Models and Standards in a Distributed Heterogeneous Environment , Proceedings of the Third IEEE META-DATA Conference, April 6-7, 1999, National Institutes of Health, Bethesda, Maryland, Paper #12, (Proceedings are published electronically at http://computer.org/proceedings/meta/1999/)

Goldiez, Brian, and Rogers, Woodard. Real-Time Visual Simulation on PCs , IEEE GG&A, January/February, 1999

Janette, Annette, Miller, Dale D., Nordstrom, John, and Birkel, Paul A., Extending the JSIMS Terrain Common Data Model to Support Virtual Simulations , Spring 2000 Simulation Interoperability Workshop, Orlando, FL, Paper # 00S-SIW-112, pp. 667-678.

Richbourg, Robert F., Miller, Dale D., Foley, Paul G., and Schiavone, Guy A., Standardizing the Codification of Environmental Objects: The Environmental Data Coding Specification , Fall 2000 Simulation Interoperability Workshop, Orlando, FL, Paper # 00F-SIW-071

Schiavone, Guy, Tracy, Judd, and Palaniappan, Ravishankar, Preliminary Investigations into Distributed Computing Applications on a Beowulf Clusters , Proceedings Fourth IEEE International Workshop on Distributed Simulation and Real-Time Applications, CA, August 24-26, 2000, pp.13-17.