Digital Earth: Building the New World

Yvan G. Leclerc, Martin Reddy, Lee Iverson, and Nat Bletter

SRI International
333 Ravenswood Avenue
Menlo Park, CA 94025, USA

Abstract: Digital Earth is a massive, distributed framework for managing and visualising georeferenced information over the Internet. Imagine being able to walk down the streets of Paris, seeing the sights and hearing the sounds around you, without leaving your home. You are able to converse with a friend who is also online, and the two of you decide to fly to Tahiti together. As you travel around the globe, you stop momentarily at places of interest to view three-dimensional (3D) representations of local structures, learn about representative works of art, browse historical information for the area, or view real-time video and audio feeds of local events. This vision presents virtual reality with many substantial challenges, for example: rendering complex georeferenced data for the entire planet, managing and displaying embedded multimedia data, visualizing massive scenes in real-time, and developing useful navigation interfaces for interacting with these scenes.

Keywords: 3D geography, georeferenced information, multiuser, terabytes, terrain visualisation, virtual heritage, virtual environments, visual simulation, VRML.

1. Introduction

The Digital Earth project at SRI (http://www.ai.sri.com/digital-earth) aims to build the infrastructure and visualisation tools for a massive, open, distributed, interactive, multi-resolution, 3D representation of the earth, in which multifarious information can be embedded that refers to specific points on that earth. No system has been previously attempted on the scale that we propose here. Digital Earth is a compelling multimedia experience that presents virtual reality with a number of significant challenges. Before proceeding, we will expand upon our opening description to clarify its meaning and implications.

- By the term massive, we mean quantities of data that could not be stored on a single disk or maintained by a single organisation. Most existing terrain visualization or flight simulator applications work by pre-loading all data into main memory on start up. However we will be dealing with aggregate data sizes in the order of many terabytes and therefore cannot afford this luxury.
- SRI's Digital Earth project is an open venture in that we intend to use open standards to build the Digital Earth, as well as contribute freely-available data and

tools to the general public. We feel that this is an important issue in today's marketplace.

- Considering the massive quantities of information that we envisage, it is necessary to distribute the Digital Earth contents over many servers around the world and to allow multiple organisations to contribute and manage their own data. This work will take advantage of emerging network technologies such as the Next Generation Internet or Internet2.
- The Digital Earth will be interactive in that users can actively control their destination and flight path around the planet. They can also select types of information that they wish to browse and view multimedia content such as movies, images, and text. We also desire collaborative capabilities so that multiple remote users can communicate with each other, follow the same paths, or perform teacher-student learning activities.
- In order to interact with massive quantities of distributed data in real-time, we need to employ hierarchical level of detail (LOD) techniques. A single computer could not possibly handle all of the data in the Digital Earth if it were transmitted as a flat structure. For example, colour imagery for the entire world at 1 m resolution would require over 1 petabyte (10^{15} bytes) of memory. If we had elevation data for the world at 30 m resolution, this would produce a geometric model of over 500 billion (5×10^{11}) polygons.
- Finally, the Digital Earth involves many different types of georeferenced data, not just terrain elevation and imagery. For example, we could insert virtual representations of buildings at their correct geospatial location, visualise weather phenomenon, or annotate locations with place names.

This vision sprung from a speech that the Al Gore, Vice President of the United States, made to the California Science Center in 1998 entitled, "The Digital Earth". In this speech, Gore challenges the scientific community to build a *"multi-resolution, three-dimensional representation of the planet, into which we can embed vast quantities of geo-referenced data"* (Gore, 1998). Gore gives the following example of the utility of such a facility:

> *"Imagine, for example, a young child going to a Digital Earth exhibit at a local museum. After donning a head-mounted display, she sees Earth as it appears from space. Using a data glove, she zooms in, using higher and higher levels of resolution, to see continents, then regions, countries, cities, and finally individual houses, trees, and other natural and mad-made objects. [...] She is not limited to moving through space, but can also travel through time. After taking a virtual field-trip to Paris to visit the Louvre, she moves backward in time to learn about French history, perusing digitized maps overlaid on the surface of the Digital Earth, newsreel footage, oral history, newspapers and other primary sources."*

In this paper, we present a technical infrastructure that will enable the Digital Earth to be built and maintained. We also introduce solutions for interacting with and visualizing these data sources using virtual reality technologies. A prime focus of this work is to enable users with a standard PC configuration to view the Digital Earth over the web using standard solutions such as VRML97 (Virtual Reality Modeling Language); but also to allow for more complex configurations such as a CAVE or Immersadesk interface via custom browsing software.

2. Infrastructure

The Web today already allows access to trillions of bytes of data of a wide variety, including text, maps, images, 3D models, music, and video. However, there is currently no way to find all available data about a given geographic location or area because the only organizing principles that span all data are either keywords/keyimages (used by search engines) or hyperlinks. In other words, data on the Web today is typically not georeferenced in any way.

Also, the sheer volume of geographically indexable data about our world, including satellite and aerial imagery, 3D models of terrain and buildings, and maps, is so large that no single organization can possibly produce, store, or even index all of the data. What is needed, therefore, is an organizing principle, or infrastructure, that can be used to georeference any type of data from around the world without having to modify any of the data itself. This infrastructure must be able to:

1. Store the terabytes of data that represents a baseline set of data covering the real world.
2. Allow very fast indexing of petabytes of data of all kinds by geographic area and various metadata attributes.
3. Make it simple for producers to either add their data, or a pointer to their data, within the infrastructure.

Our solution is based on a hierarchy of Web servers that would be distributed around the world. Each server in this hierarchical organization has a DNS (Domain Name System) name that represents a given geographic area of the earth, called a cell. The server is responsible for all data contained within this area and having detail within a limited range. It is also responsible for providing DNS service to those servers corresponding to smaller cells within its own cell, and these servers are responsible in turn for correspondingly more detailed information about their smaller cells.

For example, instead of a domain name like www.ai.sri.com, we might have an address that defines a geographical location using a hierarchal format such as minutes.degrees.tendegrees.geo. An example of this would be 37e47n.1e5n.10e20n.geo. The first part of this address defines a 1 minute x 1 minute grid of the earth, the second part defines a 1 degree x 1 degree grid of the earth, and the third part defines a 10 degree x 10 degree grid of the earth.

Each of these geographic scales is managed by a (potentially) separate server. By limiting the area and range of details in this way, each server becomes responsible for only a tiny fraction of all of the data available around the globe. Furthermore, by using DNS names that represent the geographic service area of the server, clients can immediately determine which server to query without the need for a search engine or global name server.

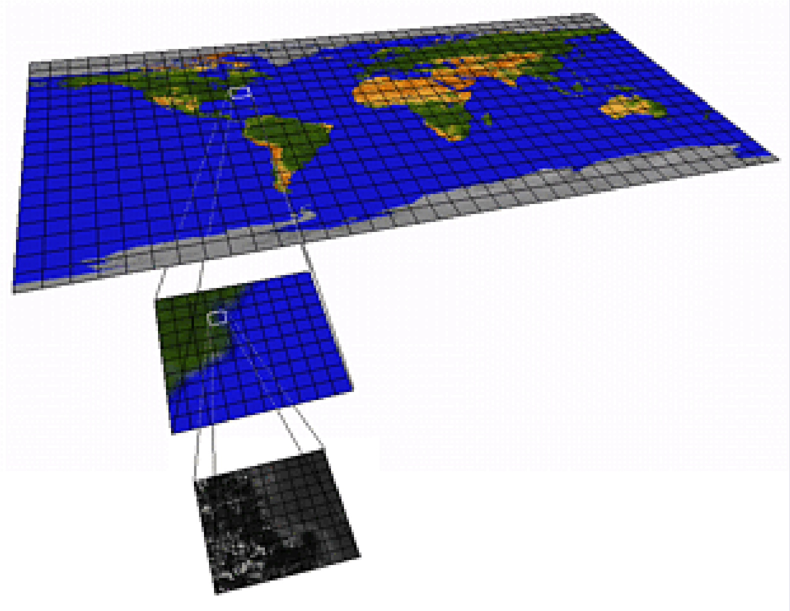

Figure 1. An illustration of the use of DNS to index geographic data hierarchically. Each cell in the top layer is a 10 degree x 10 degree region of the planet with a unique DNS address. The lower layer represents the cell containing Washington, D.C. This in turn contains a grid of uniquely addressable 1 degree x 1 degree cells, which in turn contain a grid of 1 minute x 1 minute cells.

3. 3D Representation

The Digital Earth will consist of huge volumes of distributed data. We require that the user be able to interact with these, experiencing minimal download times and high frame rates. This requires implementing level of detail techniques to manage the streaming and display of terrain data.

Multi-resolution terrain is best represented using a hierarchical structure such as a tiled, pyramid scheme (Falby et al., 1993; Hitchner and McGreevy, 1993; Leclerc and Lau, 1995). This involves progressively downsampling an image or elevation bitmap to produce the multi-resolution pyramid. Each level of this pyramid is then segmented into a grid of equally sized rectangular tiles, for example, 128 x 128 pixels. A tile at one level of the pyramid will therefore map onto four tiles on the immediately higher-resolution level, that is, the tiles at the higher-resolution level cover half the geographical area of the former. Using such a representation, we can progressively transmit higher resolution data around the area of interest while other regions remain in low resolution (see Figure 2, below). The use of tiling also allows us to fetch and display only those sections of the Digital Earth that are visible from a certain vantage point.

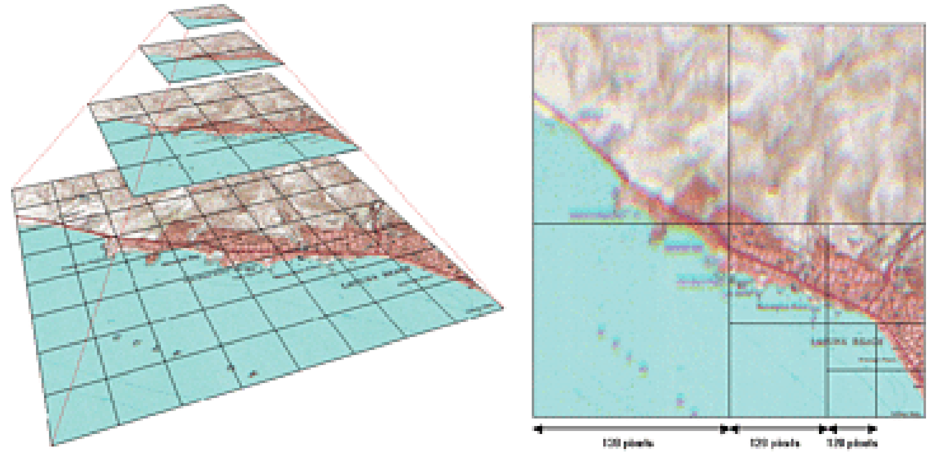

Figure 2. (a) A tiled, pyramid representation for terrain, and (b) Illustrating the use of this structure to produce viewpoint-dependent, multi-resolution imagery.

One of the principal goals of this work is to allow multiple types of georeferenced data to be embedded into the global infrastructure. For example, we could have 100 km resolution data for the entire globe but recursively insert higher-resolution datasets for smaller regions of interest, for example, a 1 km resolution dataset for the conterminous United States and a 1 m resolution dataset for Yosemite Valley, CA. Then we could insert 3-D models for buildings in Yosemite Valley, some of which could contain hyperlinks to various multimedia presentations about the National Park.

4. Coordinate Systems

Most computer graphics system use a Cartesian coordinate system to model all objects in 3-D space (e.g. OpenGL, 1997; ISO/IEC, 1997). In terms of georeferencing, this coordinate system is most similar to a geocentric coordinate system, where all locations are specified in units of meters as an (x, y, z) offset from the center of the planet. However, most georeferenced data are provided in some geodetic or projective coordinate system.

A geodetic (or geographic) coordinate system is related to the ellipsoid used to model the earth, for example the latitude-longitude system. A projective coordinate system employs a projection of the ellipsoid onto some simple surface such as a cone or a cylinder, for example the Lambert Conformal Conic (LCC) or the Universal Transverse Mercator (UTM) projections, among many others (Synder, 1987). Each of these coordinate systems was designed for slightly different applications and offers particular advantages and restrictions. For example, some projections can only represent small-scale regions, others are conformal (same scale in every direction), and others can be equal area (projected area corresponds to the earth's physical area over the entire projection). Figure 3 provides an illustration of a number of contemporary coordinate systems.

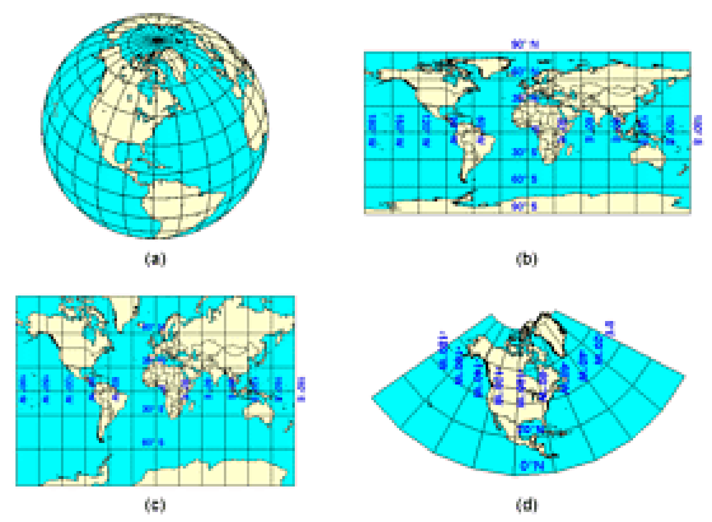

Figure 3: Examples of geodetic and projective coordinate systems. (a) Orthographic projection, used for perspective views of the Earth, Moon, and other planets, (b) Latitude/Longitude graticule, used to locate points on the Earth's surface via a grid of meridians and parallels, (c) Mercator projection, used for navigation or maps of equatorial regions, (d) Lambert Conformal Conic, used by USGS for topographic maps. (Images adapted from Dana, 1998. Reproduced with kind permission.)

We use the SEDRIS (Synthetic Environment Data Repository and Interchange Specification: http://www.sedris.org/) Coordinate System API in order to perform conversions between various geographic coordinate systems. We have also converted part of this library to Java (the GeoTransform package: http://www.ai.sri.com/geotransform) so that we can perform these conversions on the fly inside a VRML97 scene. Using the GeoTransform package we have developed new VRML nodes to provide support for geographic data in VRML97 and contributed this work to the GeoVRML Working Group, an official working group of the Web3D Consortium. All of the above software are freely available under open source like licensing schemes.

There is a further problem to face when dealing with geographic coordinate systems: that of precision. Most graphics systems only support rendering using single-precision floating point numbers, and VRML97 can only store floating point numbers in single-precision (ISO/IEC, 1997). IEEE single-precision values are represented using 32-bits with a 23-bit mantissa (IEEE, 1985), thus providing around 6 digits of precision (22^{3} = 8.39 x 10^{6}). Given that the equatorial diameter of the earth is 12,756,274 m under the WGS 84 ellipsoid (Synder 1987), then we will only be able to represent geocentric values down to the order of 10s of meters. In order to get round this problem, we define an absolute georeferenced location for each terrain dataset or feature, and an implicit local coordinate frame against which the geometry is referenced. Internally, the absolute location is stored as double-precision geocentric coordinates. Each vertex of a model is transformed into a double-precision geocentric coordinates and the single-precision difference of these two values is used to render the vertex.

5. Visualisation

We intend the Digital Earth to be browsable using commercial off-the-shelf software. In order to do this, we rely on the VRML97 format to describe the terrain data so that the user will be able to interact with the data using a standard VRML97 plug-in for their web browser. We also intend to enhance an existing terrain visualisation system, called TerraVision, in order to efficiently read these VRML representations. A general user can therefore use the VRML browser interface to easily gain access to the Digital Earth, while a more serious user could install TerraVision in order to use an application that has been specifically tuned to the data and which offers various virtual reality interfaces.

5.1. The TerraVision System

TerraVision is a real-time, distributed terrain visualization system that has been developed over several years at SRI International (see Figure 4). It was originally developed as part of the DARPA-funded Multi-dimensional Applications Gigabit Internet Consortium (MAGIC) and Battlefield Awareness and Data Dissemination (BADD) projects and has been specifically designed to browse massive terrain and other data distributed over a fast network (Leclerc and Lau, 1995, Reddy et al., 1999). It incorporates features such as an active map display, 2-D pan and zoom display, 3-D flythroughs, time of day and fog selection, incorporation of georeferenced models such as buildings or roads, and support for virtual reality devices such as head-mounted displays and the CAVE.

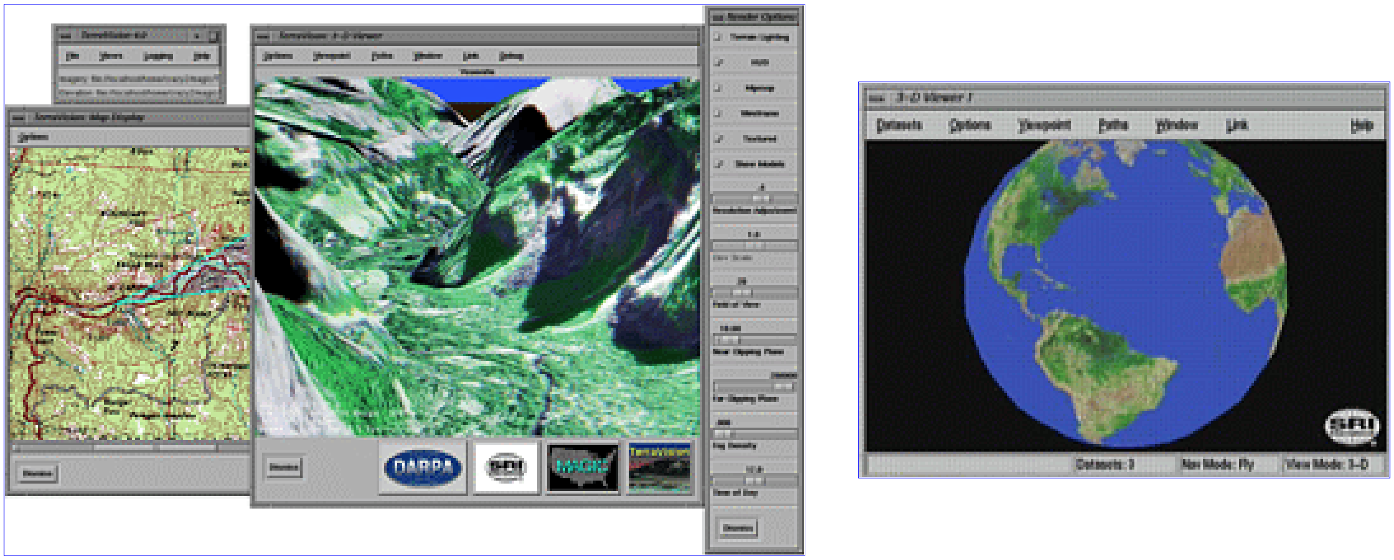

Figure 4. Left: screenshot of the TerraVision terrain visualization system showing the 3-D viewer and the co-registered Map viewer. Right: an MPEG movie showing a TerraVision fly-through from the globe to Menlo Park, CA. You can download the TerraVision system for free (SGI platform only) from: http://www.ai.sri.com/TerraVision

5.2. A VRML Browser

An important goal of our work is to enable open solutions. One facet of this goal is the adoption of various open standards, such as VRML, to give a wide cross-section of users access to the content. By employing VRML as the file format to represent the multi-resolution structure of the Digital Earth, we allow users to interact with it using standard off-the-shelf VRML browser software. VRML browsers are produced by several companies and are available for a range of platforms. These are often provided for free as plug-ins for Internet browsers such as Netscape Communicator (NC) or Microsoft's Internet Explorer (IE). IE5 was shipped with a pre-installed VRML plug-in.

By default, a VRML browser will display a 3-D scene, perform any key-frame animations that are specified, and allow the user to interact with the scene by using a mouse. Certain objects can be defined as hyperlinks so that when the user clicks over them, an action is performed such as loading a new VRML scene or displaying an HTML page. It is possible to extend the base functionality of a scene by embedding Java code directly into objects to define their behavior, or to control the VRML browser from an external Java applet running in the Internet browser. These features enable us to encapsulate much of the Digital Earth functionality into a standard VRML application. For example, we will be able to navigate around a multi-resolution, 3-D representation of the globe; embed multiple terrain datasets as well as other features such as buildings, roads, and textual annotations; and click over features to display other multimedia objects. However, it is likely that certain capabilities will not be available in a standard VRML browser, or that they will be available at a lower performance level. For example, TerraVision will offer more sophisticated and optimised tile management techniques, perform stitching between tiles of different resolutions, implement terrain-specific navigation models, and so on.

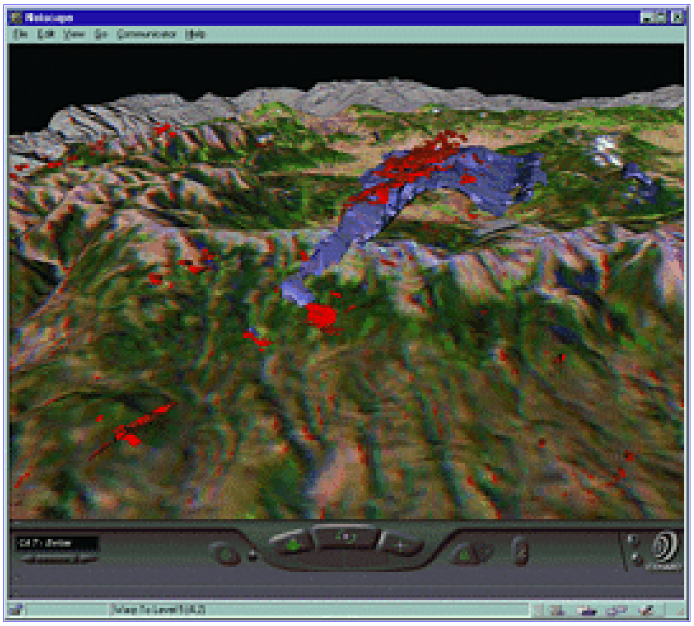

Figure 5. An example showing a VRML browser display a terrain model of the Colorado Rockies and showing the embedding of a georeferenced object representing a clear air turbulence isosurface. (Weather model provided by National Center for Atmospheric Research, Boulder, CO.) You can browse this model yourself from the SRI VRML Terrain Dataset web page at: http://www.ai.sri.com/VRMLSets

6. Navigation

The issue of efficient navigation of our global structure is a crucial one. Without a means to interact with the data in a timely and appropriate manner, our work will be of limited practical value. There are a number of dimensions to the issue of navigating large, planetary models. We will discuss some of the most important issues here.

1. **Terrain Following:** The earth is round, not flat. As we navigate over the earth's surface we should therefore expect to follow a curved flight path. We therefore desire a navigation method that will maintain a particular height above the surface of the earth. In order to do this we need to know the gravitational up vector for a particular region of terrain so that we can calculate height above the terrain and also orient the user correctly. This up vector could be found by using the 3D normal to the plane which is tangent to the reference ellipsoid at the region in question.

2. **Altitude-based Velocity:** The velocity at which the user can navigate should be dependent upon their height above terrain. For example, when flying through a valley at a height of 100 m above the terrain, a velocity of 100 m/s could be considered relatively fast. However, when viewing the entire globe from space at an altitude of 20,000 km, zooming in at the same speed would be painfully slow. We therefore scale the velocity of the user's navigation in an attempt to achieve a constant pixel flow across the screen. We have found that a simple linear relationship, assuming height above the earth's ellipsoid, gives satisfactory results.

3. **Active Maps:** When flying low over an area of terrain, it is often difficult to maintain a context of your position in the world. We therefore employ overhead and map displays to provide this context. We can project the user's 3-D geocentric location onto the map display so that the user can easily ascertain their location in the world. Additionally, we can allow the user to click over the map and then move the viewpoint directly to that location.

4. **Multi-Modal Navigation:** Mouse gestures are often insufficient to provide good navigation in a complex environment. TerraVision will provide the ability to navigate using a number of inter-related modalities such as pen-based gestures and spoken language. This capability will be built using SRI's Open Agent Architecture (http://www.ai.sri.com/~oaa) and prior experience in building such an interface to a 2D map. For example, the user could say something like "Take me to Dundee". The system would then automatically descend to this city. Or the user

could ask "What's the name of this river?" while sketching it on the terrain, and so on.

7. Conclusions

We have described the Digital Earth concept: the dream of building a massive, distributed model of the earth where users can fly from outer space down to street-level scales and interact with large amounts of georeferenced multimedia data. More than just a pipe dream, we have shown how we may build a scalable framework to support this vision by introducing a geographically based DNS naming convention; we have presented a number of problematic issues and provided solutions to these, for example modeling massive, distributed multi-resolution terrain, supporting geographic coordinate systems, dealing with double-precision data, navigating around a large globe structure; and finally, we are developing solutions to allow users around the world to interact with the Digital Earth in an efficient and seamless manner.

References

Dana, P. (1998), "Geographer's Craft Project Web Page", Department of Geography, University of Texas at Austin. Available on-line at http://www.utexas.edu/depts/grg/gcraft/notes/mapproj/mapproj.html.

Falby, J. S, Zyda, M. J., Pratt, D. R. and Mackey, R. L. (1993). "NPSNET: Hierarchical Data Structures for Real-Time Three-Dimensional Visual Simulation". Computer and Graphics, 17(1):65-69. Pergamon Press, UK.

Gore, A. (1998). "The Digital Earth: Understanding Our Planet in the 21st Century". Speech delivered at the California Science Center (CSC), Los Angeles, CA, 31 January. Available on-line at http://www.ai.sri.com/digital-earth/gore.html

Hitchner, L. E. and McGreevy, M. W. (1993). "Methods for User-Based Reduction of Model Complexity for Virtual Planetary Exploration". In Proceedings of the SPIE: the International Society for Optical Engineering, vol 1913, pp. 622-36.

IEEE 754-1985 (1985). "IEEE Standard for Binary Floating-Point Arithmetic". IEEE Standards Publication. http://standards.ieee.org/.

ISO/IEC 14772-1:1997 (1997). "The Virtual Reality Modeling Language". December 1997. Available on-line at: http://www.web3d.org/Specifications.

Leclerc, Y. G. and Lau, S. Q. (1995). "TerraVision: A Terrain Visualization System". Technical Note 540, Artificial Intelligence Center, SRI International, 333 Ravenswood Avenue, Menlo Park, CA 94025. Available on-line at http://www.ai.sri.com/pubs/technotes/aic-tn-1994:540/document.ps.gz

OpenGL Architecture Review Board (1997). "OpenGL Reference Manual: the Official Reference Document to OpenGL, Version 1.1" (Second Edition). Editors: Kempf, R. and Frazier, C. Addison Wesley, MA. http://www.sgi.com/software/opengl/guides.html.

Reddy, M., Leclerc, Y. G., Iverson, L. and Bletter, N. (1999). "TerraVision II: Visualizing Massive Terrain Databases in VRML". IEEE Computer Graphics and Applications, Special Issue on VRML, 19(2): 30-38. Available on-line at http://www.ai.sri.com/~reddy/pubs/CGAVRML/

Snyder, J. P. (1987). "Map Projections - A Working Manual". U.S. Geological Survey

Professional Paper 1395, US Government Printing Office, Washington, DC.

For information contact: Yvan G. Leclerc / leclerc@ai.sri.com