



Visualization System

We have described how we wish to organize and store georeferenced data. However, data are useless without an effective interface to communicate their meaning to the user. We therefore propose the development of a revolutionary application that enables users to browse and interact with the Digital Earth in a novel manner. We intend to radically enhance SRI International's terrain visualization system, TerraVision, to provide this next-generation capability. However, it is also our intention that these data be accessible to a wide range of users---although perhaps with less functionality---using standard and freely available commercial off-the-shelf (COTS) software on personal computers.

The TerraVision System

[TerraVision](#) is a real-time, distributed terrain visualization system that has been developed over several years at SRI International (see [Figure 3](#)). 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 (VR) devices such as head-mounted displays (HMDs) and the CAVE.




Figure 3 Screenshot of the TerraVision terrain visualization system showing the 3-D viewer with embedded building models, and the co-registered Map viewer.

TerraVision offers the following powerful capabilities that make it a suitable springboard to develop a Digital Earth application.

Network Awareness

TerraVision was designed from the onset to browse large quantities of data distributed over a wide-area network. As such, TerraVision is able to gracefully handle situations where high-resolution data for an area have not been received over the network yet. It will simply fall back to the highest-resolution data that are available and continue to let the user interact with these data. As the higher-resolution data are streamed into the application, the display is progressively updated.

Real-Time Performance

TerraVision employs various sophisticated optimizations to deliver maximum performance to the user. These include level-of-detail management to optimize the bandwidth of the network and the graphics pipeline, caching techniques to store recently used data so that they are not continually retransmitted over the network, high-performance visibility culling algorithms that are specialized to the hierarchical nature of the data, and prediction algorithms that attempt to pre-load data for regions where the user is heading so that these data are instantly available.

Dataset Scalability

Many contemporary terrain visualization packages work by first loading all terrain data into main memory first. These systems are therefore limited by the size of available memory. In contrast, TerraVision requires only a small subset of the total dataset to be in memory at any one time: that portion of the terrain that is visible to the user, and at the appropriate level of detail. As a result, TerraVision can browse arbitrarily large terrain datasets and has already been demonstrated with datasets in the order of tens of gigabytes.

Cross-platform Capability

In the past, TerraVision has been deployed largely on SGI graphics workstations such as the O2 and Octane. However, we have engineered TerraVision to be easily portable to other platforms and we have recently performed a port to Microsoft's Windows NT platform.

These features are basic requirements for a Digital Earth application. However, we propose to build on these foundations to create an application with advanced capabilities that redefines the state of the art. The Digital Earth is a potentially enormous undertaking and so we are required to build an application that is uniquely enabled to manage vast volumes of data distributed over advanced NGI networks.

In addition to engineering TerraVision to support the Digital Earth framework, we propose to make profound advances to TerraVision's capabilities. These will encompass the following features.

Collaboration

We believe that a vital capability is to allow numerous, remotely located users to interact with each other while navigating the Digital Earth. This would allow collaboration between colleagues and friends, as well as enhancing mission planning operations and communication with workers in the field. This collaboration should involve the ability to communicate with each other in real time; to see some representation of each other in the Digital Earth for visual feedback; and some ability to work together, such as with a virtual whiteboard.

One way in which this could be achieved is through the use of the [Open Agent Architecture](#), developed at SRI International. The OAA is a framework for integrating a community of heterogeneous software agents in a distributed environment through a central facilitator which coordinates agents capabilities and message passing.

Intelligent Network Agents

A useful facility would be if the user could be automatically informed about certain events that occur. For example, if new data have been recently added for a region of interest or if a colleague or friend has just gone online. This might also be achieved through the use of the OAA. Clients could register interests with the OAA's facilitator, and the facilitator would inform them whenever any of these events occurred.

Multimedia Support

The Digital Earth should be a rich medium for storing many types of information, for example, QuickTimeVR movies, RealAudio sounds bites, HTML texts, and VRML animations. The Digital Earth application therefore needs to be able to understand all of these formats and be able to present these to the user in a suitable fashion. It is likely that this will be done by using several COTS helper applications, in much the same way that current Internet browsers use external applications to support file formats that are not understood by the browser itself.

Novel Navigation Schemes

Standard navigation techniques such as flying over terrain or zooming into small regions will be insufficient for a Digital Earth application. TerraVision already supports a variety of input devices such as the mouse, HMDs, and the CAVE. However, more sophisticated interaction techniques are required to ease the user's ability to navigate effectively around a large, multi-resolution, global network of information. We propose to investigate the integration of the following novel navigation schemes to maximize a user's efficiency in the Digital Earth:

Speech Recognition

Through the OAA, we have the capability to introduce speech recognition into TerraVision so that users could enter commands verbally, thus enabling a hands-free and intuitive interaction with the system.

Place Name Queries

It is likely users know where they want to go in the Digital Earth. They should therefore be able to directly specify this to the system. For example, if a user requested, "Take me to Edinburgh," then the system should automatically descend to this city, perhaps asking the user for further clarification if there exist multiple places with the same name. This would require a database of place names along with their latitude/longitude extents for TerraVision to cross-reference. Such a database might be potentially massive and should be queryable over the network so that the entire database does not need to be stored locally.

Co-registered Maps

A user may also be interested in investigating a particular region, but does not know, or does not wish to specify, a particular place name. An intuitive way to do this is to present a 2D map of the world with which the user can point and click to home into the region of interest. This interface is useful because it is generally easier to locate a place on the earth by using a map rather than trying to navigate through three dimensions.

Selective Views

The Digital Earth could contain a diverse selection of many different types of information. Presenting the user with all of these at one time would be obfuscating and complex. Instead, the user should be able to customize a view of the Digital Earth by specifying the types of information to be made aware of. A simple example might be to display only airport buildings and suppress all other kinds of buildings.

Off-the-shelf Browser Software

An important goal of this proposal is to enable open solutions. We have no desire to restrict the utility of the Digital Earth to a single application or operating system. Instead, we have proposed a framework that gives a wide cross-section of users access to the content. This is done through the adoption of various open standards, such as VRML.

By employing VRML as the file format to represent the multi-resolution structure of the Digital Earth, we allow for the possibility of users interacting 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). In fact, Windows 98 was shipped with a pre-installed VRML plug-in for IE4, and it is hoped that in the future, VRML support will continue to be integrated directly with Internet browser software. This would mean that users would not need to download any supplemental software to view the Digital Earth. A user could just direct an Internet browser to the appropriate location and instantly begin accessing its facilities.

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 currently offers the following advantages over a standard VRML browser. (N.B. it is feasible that some of the following could be implemented for a standard VRML browser through the use of various Java scripts embedded in the scene, or running externally to the browser.)

Performance

TerraVision is a high-performance, multi-threaded application that has been designed with the sole purpose of rendering large geographic databases in real time. As such, it can employ more efficient, optimized solutions to various generic real-time graphics operations, for example, visibility culling is performed using a fast quad-tree search of the multi-resolution hierarchy.

Interaction Techniques

Many of the novel interaction techniques that we have suggested for TerraVision will not be available in a standard VRML browser. For example, the use of speech recognition, intelligent agents, or collaboration will not be available by default. VRML does provide a good suite of interaction techniques, however, so useful interaction will still be possible.

Seamless Terrain

Any tiled, multi-resolution representation will suffer from the problem of tearing. That is, adjacent tiles of different resolution do not share all the same vertices, and so holes can appear in the terrain along tile boundaries. In TerraVision, we employ specialized techniques to stitch these holes so that we display a continuous landform. Also, TerraVision can employ the more accurate criterion of projected screen size to decide when to reduce terrain detail, whereas a VRML browser performs level of detail based upon the user's presence or absence inside a predefined volume.

Advanced Tile Management

TerraVision employs tile management techniques to improve interactivity, for example, it attempts to predict where the user will be in the near future by a simple extrapolation of the current flight path and then pre-fetches the tiles for that region so that they will be immediately available for rendering. TerraVision also maintains a local cache of tiles so that it does not need to reload and parse data for regions of the terrain that have been recently browsed.

Despite these points, it is clear that VRML introduces an attractive scalability feature to our proposal. If the resources are available, then a user can use TerraVision running on a fast graphics workstation to quickly and intuitively navigate around the Digital Earth. Alternatively, these same data can be accessed from a laptop

machine with a standard VRML browser. This level of capability may be of particular use to military personnel performing mission planning and battle damage assessment, or to emergency teams coordinating a natural disaster relief effort. To summarize, the support of VRML adds the following benefits to this proposal:

- Encourages the use of commercial off-the-shelf software.
- Enables the Digital Earth to be automatically browsable with numerous third-party applications, thus not tying it to a single application.
- Uses VRML browsers that are freely and widely available, so a large number of users will have access to the features of the Digital Earth.
- Supports multiple delivery platforms, including lower-end platforms such as Macs and PCs, so that the Digital Earth is not reserved for people who have access to high end graphics workstations.
- Introduces scalability, in terms of performance as well as available features, hardware support, and so forth.

Sample Digital Earth Applications

The existence of massive quantities of georeferenced data distributed over a Next Generation Internet, coupled with an application to navigate these data appropriately, could have profound commercial, educational, and societal impact. For example, the following application areas could benefit from the existence of the Digital Earth.

Education

Students could easily perform investigations about the history, culture, and people of an area by simply flying to that region of the world. They could perform virtual field courses by visiting remote sites that are represented in the Digital Earth. In addition, given the ability to communicate with other online users, students could collaborate with each other on projects, or groups of students could be directed by a teacher.

Virtual Tourism

Tourists could use the Digital Earth to plan a holiday or experience the attractions at some remote location. This could be used by travel companies to advertise a particular site, or to provide local information for tourists who are visiting an area. We might also envisage a portable computer that can link to the Digital Earth, coupled with a GPS, to provide a portable map system.

Crisis Management

The availability of a high-resolution, 3-D terrain representation could provide crisis management teams with powerful tools to oversee, plan, and coordinate disaster relief efforts, for example, emergency teams fighting a forest fire or organizing hurricane relief efforts, or environmental workers evaluating a flood or other time-critical conditions.

Military Mission Planning

The Digital Earth framework could be used to greatly enhance DoD operations and its warfighting capability. For example, access to a real-time 3-D representation of a target site could enable greater mission planning and assessment flexibility. This could be coupled with a collaborative capability where commanders can make and communicate annotations to 2-D maps or 3-D terrains to highlight strategies, targets, or enemy capabilities.

Virtual Real Estate

It is likely that portions of data in the Digital Earth will be available for purchase or lease, just as is the case on the Internet today. This establishes an environment for commercial entities to provide data and services for a fee. It is possible that business may also grow out of the acquisition and control of certain regions of the Digital Earth.

Global Circulation Models (GCMs)

The Digital Earth data indexing method proposed here can provide enormous advantages for the dissemination of GCM results and enable the distribution of GCM's computation. The global change community uses GCMs to analyze what changes have occurred and will likely occur. The five primary models in use (i.e., GISS, NCAR, GDFL, OSU, and UKMO) all use a similar data representation of a three-dimensional grid of widely spaced points. A current grand challenge in climate studies is to substantially increase the resolution of these models and to incorporate regional analyses. Currently, most models have grid cells of 2.5 degrees or 5 degrees square. A doubling of resolution requires more than an order of magnitude increase in the required computing. Using the `geo` domain addressing system these models will be able to distribute their data access and storage. Further, this referencing scheme will aid the dissemination of these data to diverse users.

Previous: [3D Representation](#)

[SRI's Digital Earth Project](#)
[SRI's Digital Earth Proposal](#)

digital-earth@ai.sri.com -- Fri Apr 16 11:21:17 1999