

Visualizing Terrain and Navigation Data

David Coombs¹, Billibon Yoshimi², Tsung-Ming Tsai, Ernest Kent³

National Institute of Standards and Technology, Intelligent Systems Division
100 Bureau Dr Stop 8230, Gaithersburg MD 20899-8230 USA
tsung-ming.tsai@nist.gov

Abstract

Terrain and navigation data have been collected from sensors mounted on an autonomous offroad vehicle. Data include range images from scanning laser range imagers, position and orientation estimated from GPS and inertial devices, and color images registered with the laser range data. These data have been visualized as static models with interactive displays and as 3D movies (sequences of 3D range images) displayed nominally from the sensor's viewpoint but also with some interactive control of the vantage point. Display technologies have ranged from VRML (Virtual Reality Modeling Language) alone to immersive stereo displays. Design choices in the display of the data and their effect on the visual impact are discussed. The ability to rapidly and interactively visualize these data has enabled critical insights into the data that would otherwise have been difficult to achieve.

Keywords: 3D data visualization, terrain visualization, range data visualization, vehicle navigation data visualization, data sequence visualization.

1 Introduction

One of the challenges of developing offroad autonomous driving systems is understanding the data provided by the sensors that the vehicle uses to

perceive its surroundings. First, it can be difficult to verify that the sensors are functioning correctly and that interfaces are working properly. The logistics of collecting data make the process very expensive. Therefore it is invaluable to be able to quickly verify the integrity of the data so experiments and collections can be repeated if necessary while the personnel and equipment are available.



Figure 1: The NIST HMMWV is actuated and instrumented with a laser range scanner, inertial sensors and GPS.

Second, the data themselves can be difficult to interpret. It can be helpful to get an intuitive understanding of the properties of the data in order to develop algorithms that perform well. For instance, the two scanning LADAR (LAser Detection and Ranging) range imagers that have been used produce range images of $64 \times 128 = 8192$ pixels and $32 \times 180 = 5760$ pixels. Traces of navigation data from GPS (Global Positioning System) and INS (Inertial Navigation System) that extend over several minutes produce large volumes of data that are difficult to interpret. However, three-dimensional visual displays of these

¹now with CYTOMETRICS Inc., 12 Penns Way, New Castle, DE 19720.

²now with IBM T.J. Watson Research, P.O. Box 218, Yorktown Heights, NY 10598.

³now with Sidereal Designs, Inc., 18920 Festival Drive, Boyds, MD 20841.

inherently 3D data make the data immediately accessible, particularly when viewed in stereo or with motion.

The NIST HMMWV (Highly Mobile Multipurpose Wheeled Vehicle⁴, Figure 1) drives autonomously offroad, avoiding obstacles that are detected by the vehicle's sensors. The vehicle drives at speeds up to 35 km/h (10 m/s, 20 mph) on rolling grass-covered meadows where the obstacles are large trees and shrubs. The vehicle is commanded to follow a route given by a sequence of GPS coordinates located a few hundred meters apart. As the vehicle drives, it repeatedly plans an obstacle-free path to follow the commanded route in real time using data sensed while the vehicle proceeds.

The present approach has been in use for several years at NIST. It has been adopted by the Office of the Secretary of the Defense Demo III UGV (Unmanned Ground Vehicle) program and was demonstrated at the Demo III A (Alpha), and Demo III B (Bravo) field trials.

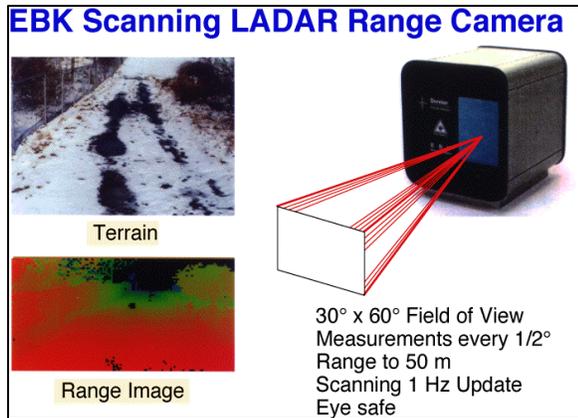


Figure 2: The EBK LADAR range camera scans over its field of view in about 0.5 s, collecting 8192 range measurements that can be treated as a range image or as a collection of 3D points.

The ability to visualize the data from the range imagers and navigation sensors has been instrumental in developing this capability. It was possible not only to quickly verify that collected

⁴Certain commercial equipment, instruments, or materials are identified in this paper to foster understanding. Such identification does not imply recommendation or endorsement by the National Institute of Standards and Technology, nor does it imply that the materials or equipment identified are necessarily the best available for the purpose.

data were valid, but also to interactively explore the data collected under varying conditions in order to clearly understand the phenomenology of the sensors.

2 Range Image Movies

The range imagers that have been used are time-of-flight scanning LADAR devices. The EBK range camera is shown in Figure 2. It scans its $30^\circ \times 60^\circ$ FOV (field of view) in about 0.5 s and retrace time makes the frame rate about 1 Hz.

The GDRS range camera is conceptually similar but its construction is considerably different, and it scans its $20^\circ \times 90^\circ$ FOV at 15 Hz v. the EBK's 1 Hz.

The range data are interpreted as illustrated in Figure 3 to detect obstacles that the vehicle must avoid. Neighboring range values in each column of range data are examined for patterns that indicate positive and negative obstacles (e.g., rocks and holes, respectively).

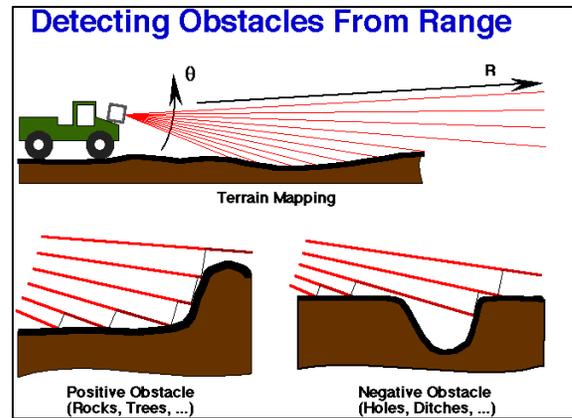


Figure 3: Obstacles are detected in range data by sweeping upward along each column of pixels in the range data in search of range discontinuity. Positive obstacles present clusters of similar range. Negative obstacles produce range gaps.

One of the most natural presentations of range image sequences is a 3D movie. EBK images are presented as sequences of images in which each pixel is represented by a tile in space (Figure 4). The tiles are oriented normal to the optical axis of the range camera. The tiles are sized to slightly overlap one another when viewed from the camera's origin. In order for each tile to subtend approximately 0.5° of visual angle, the size of each tile scales with its range from the camera. The tiles could be any appropriate shape for the sensor and the data. In this case, we used squares.



Figure 4: Movies of range images are presented. Each range image pixel is represented as a tile. The tiles are sized to slightly overlap when viewed from the range camera's origin. The viewer's vantage point is above and right of the camera's origin, revealing gaps between ground tiles and a "shadow" cast by a shrub in the left side of the image. The color of each tile is a function of its range and a small random factor to distinguish it from its neighbors. The display device is an ImmersaDesk that presents binocular stereo displays by back-projecting on a drafting table style screen. The viewer's glasses are tracked, and the display is adjusted for the wearer's viewpoint.

Each tile is colored as a function of its range from the camera, over the device's range of 0 m to 50 m. Initially each tile was colored only as a function of its range, but when overlapping tiles were near each other in 3D space, their colors were so similar that binocular fusion of the scene was difficult and the powerful immediacy of the 3D image was lost. In order to distinguish each tile from its neighbors, each tile's color is dithered by a small random factor. Alternative methods for distinguishing the tiles could use distinct patterns texture-mapped onto each tile, but dithering the color of each tile is faster to render. Another approach would be reducing tile size to avoid overlap, but the dithered color offers the ability to see a range image without gaps when the viewer's head is positioned in the LADAR camera's location. If color imagery registered with the range data is available, the tiles can be colored according to the color image. Even the slight color variations that occur in real imagery are likely to suffice to distinguish the tiles enough to enable the viewer to correctly fuse the scene. In addition, the natural color and shading cues can have powerful effects on the viewers' perception of the scene.

The display device is an ImmersaDesk (IDesk) [4], which presents binocular stereo displays by back-projecting on a drafting table style screen. (The ImmersaDesk was derived from the earlier CAVE [3].) The stereo system uses field sequential stereo, presenting the even field to one eye, and the odd field to the other. The frame rate is 96 Hz, so each eye experiences an effective video rate of 48 Hz. The viewer wears StereoGraphics LCD-shuttered glasses. There is a wand with a 2D strain gauge thumb knob and three buttons. An earlier version we used had three mouse buttons. The position and orientation of the glasses and wand are tracked independently by a magnetic SpacePad tracker with one tracking sensor attached to each device. The wand's pose, buttons, and knob can be used to navigate through a scene. The display is adjusted for the viewpoint of the person wearing the tracked glasses. Because only one viewer is tracked, any other viewpoint is distorted. Further, when the privileged glasses move, the display changes, which can be disconcerting for viewers who are not moving the same way as the tracked glasses.

The ImmersaDesk has a significant advantage over desktop systems for visualizing and exploring data. The IDesk's relatively large work volume and viewpoint-adjusted display enable viewers to examine data by walking around it and peering into it from different vantage points. The natural kinesthetic experience of exploring data in this fashion seems more immediately natural than navigating and manipulating a data set with a mouse at a desktop.

To enable the viewer to more effectively explore the data, the model is scaled and translated so it occupies the space just in front of the screen. The viewer's glasses are tracked and the scene is rendered to appear correctly only from the vantage of the wearer of the glasses. By scaling the model and translating it out in front of the screen, relatively small movements of the viewer's head are amplified into effectively large navigation with respect to the range camera's origin.

With a sufficiently powerful machine and graphics engine, these data can be displayed at rates faster than real time. For example, the EBK LADAR completes its scan and dwell cycle collecting a 64×128 frame of data every 0.9 s. If the visualization code displays a sequence of EBK frames at 1.1 frames/s, the display is at real time for the sensor. The ability to display faster than real time presents the opportunity to present the data in ways other than simply frame by frame. For example, range points can be added to the displayed frame in the order they are scanned by the camera, and the same number of the oldest pixels can be deleted. This presents the data approximately as

they were scanned, but it requires knowledge of the motion of the camera. Another good use of the time is to interpret the data to merge surfaces and objects.



(a)



(b)

Figure 5: Images of registered color and range. The model is scaled and translated to present it in the space just in front of the screen. This enables the viewer to easily move about even within the terrain to explore the structure of the screen. The viewer's glasses are tracked and the model is rendered to appear correctly from that viewpoint alone.

3 Color Range Images

When color intensity images are registered with range data, 3D models can be displayed on the ImmersaDesk with real colors, as shown in Figure 5. In this display, the range data are presented as a surface comprised by triangular patches (two triangles for each four adjacent range points). The color is applied in a per-vertex style, with the color of each patch a function of the three colors of its corners. The color imagery is higher resolution than the LADAR data. The color for each vertex is the color of the color pixel nearest each range pixel.

Given ample computational and rendering power, texture mapping would give the benefit of high color image resolution presented upon the lower resolution range framework. Even if texture mapping is ruled out, results could be improved by cleverly selecting the color for each vertex. For instance, consider a relatively small object with a distinctly different color than its background, e.g., a beige sapling tube standing in a grassy field. Choosing the distinct color of the beige tube in a neighborhood of green grass would enhance the viewing of the sapling tube in the data.

The model shown is scaled and translated to occupy the space in front of the screen so the viewer can easily look around inside the model (Figure 5 (b)). This enables an immediate grasp of the scene structure. The model can be presented with its true dimensions. However, given the scale of the scene (ranging from 0 m to 50 m), gross navigation is needed to reveal the structure of the full-scale scene. Using a scaled model amplifies head and body movements so they achieve the same effect as gross navigation in a full-scale scene. Most viewers with little experience found it much easier to use head and body movements rather than wand-based navigation.

Tile presentations of the color range images (of the style used for range image movies) have been used, but the triangulated surface seems more natural than tiles. Similarly, triangulated surfaces can be used to present range image movies, but tiles seem to offer better immediate insight to the scene structure when no color images are available.

An obvious development to take is presenting movies of color range data. The technical difficulties center on the volume of video data that must be managed either live or off-line.

4 GPS and INS Navigation Traces

The NIST HMMWV is equipped with GPS and INS navigation instruments. GPS offers high resolution differential mode as well as ordinary GPS mode. GPS relies on satellite visibility so even ordinary GPS is lost at times, e.g., when trees block the signal. Differential GPS relies on the visibility of a base station in addition to the satellites. Resolution is improved in differential mode by removing correlated errors in the signal seen by both the stationary base station and the mobile GPS unit.

Data have been collected in the field at NIST and elsewhere. 3D model displays of the navigation data enable examining the data interactively as well as quickly validating the collected data.

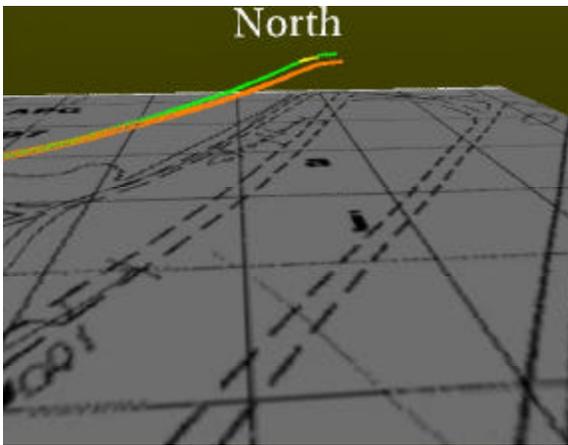


Figure 6 GPS and INS navigation data displayed over a site map. INS data are plotted in orange. Differential GPS data are shown in green and non-differential GPS data in yellow.

GPS and INS navigation data are displayed over a map of the site (Figure 6). Displaying the data over the site map provide immediate visual validation of the data. INS data are plotted as orange. GPS data are shown in green (differential) or yellow (non-differential). The yellow GPS segment shows where the GPS signal fell back from the higher-resolution differential GPS to ordinary GPS when the differential signal dropped out.

Each GPS or INS point is displayed as a simple icon for fast rendering. Crosses and open squares are fast to render, and open icons don't occlude or overlap one another much. The speed of the vehicle can be judged even at a glance from the spacing of the icons along the vehicle's path. Each icon is also of known dimension, 2 m by 2 m. As an option to

increase the rendering frame rate, the viewer can elect to display only a fraction of all the points.

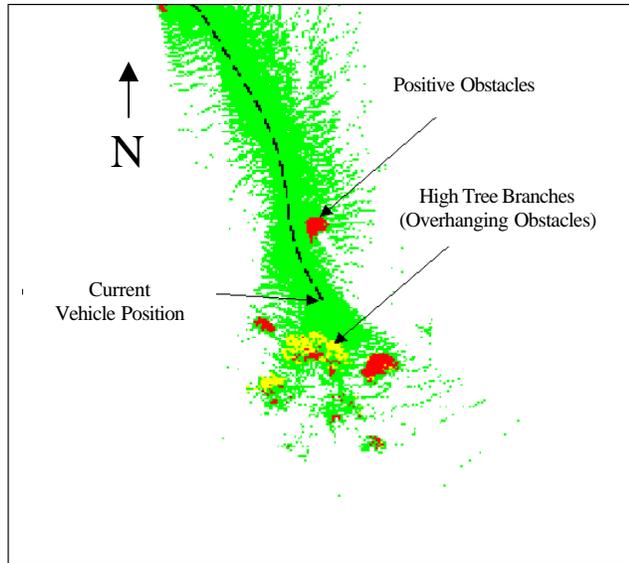
For some locations, a site map is available, and the navigation data are displayed above a flat deck with the site map texture-mapped onto it. This helps place the data in the context of a familiar map and permits rapid subjective verification that there are no major errors in the data. If terrain elevation data were available as well, the site map could be texture-mapped onto it instead of flat deck. This would enable verifying elevation of the data. The computational cost of rendering a large elevation map (e.g., as ElevationGrid) is high, and added benefit of the 3D relief for each application on the subject terrain must be considered versus the baseline effectiveness of a flat site map underlying the data.

5 Elevation Maps

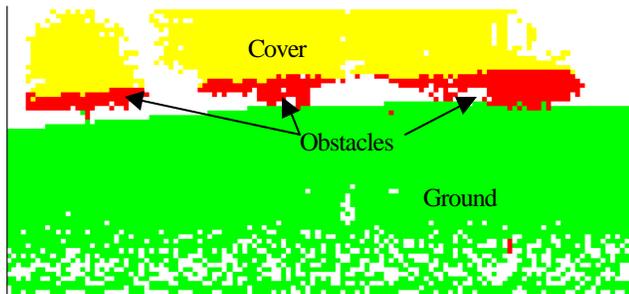
2D displays have been used that show obstacle detection and classification of cover under which the vehicle might be concealed from aerial view. Figure 7 (b) shows such an image from the EBK LADAR range camera. Obstacle detection and terrain analysis are based on data from the time-of-flight scanning LADAR range camera [1]. The range data and vehicle position and attitude are integrated in a North-oriented reference frame (Figure 7 (a)).

There is a basic assumption that the terrain is smooth except for discrete obstacles. These techniques perform well on benign terrain [2]. This approach enables the vehicle to travel over rolling meadows at speeds up to 35 km/h (10 m/s, 20 mph) while avoiding obstacles that are within the vehicle's current sensing capabilities, such as large trees and shrubs. The vehicle has no *a priori* knowledge of the obstacles.

Traversing more challenging, convoluted terrain will require analysis of the terrain to estimate the speed at which each patch of terrain can be traversed by the vehicle. Traversable speed limits are also a function of the direction of travel across the patch. To understand the terrain as sensed and modeled, we have visualized elevation maps. 2D displays don't easily reveal the terrain. 3D interactive displays of the terrain present the structure immediately, and the ability to interact with the terrain data enables the viewer to explore the terrain in ways that are difficult to accomplish with the 2D displays we have used.



(a) Local Scrolling Map



(b) Obstacle Image

Figure 7: A north-oriented obstacle map (a) integrates data over time. A classified LADAR image (b) shows ground on which the vehicle can drive, obstacles to avoid, and concealment (labeled cover) under which the vehicle could be concealed from aerial view.

In order to visualize the terrain over which path plans are generated [5], we display the elevation model the vehicle will use to plan its path (Figure 8). This local map covers $40\text{ m} \times 40\text{ m}$ with each cell $0.2\text{ m} \times 0.2\text{ m}$. A scale vehicle ($3\text{ m} \times 1\text{ m} \times 1.5\text{ m}$) is rendered in the center of the map as a reference point. The model is described in VRML (Virtual Reality Modeling Language) [6]. Depending on the VRML browser used, the viewer can easily navigate the scene to explore the data from different viewpoints.

The elevation model is described as a VRML ElevationGrid with a model of the vehicle placed on the terrain according to the vehicle's navigation data (from GPS and INS). The vehicle is placed at the elevation reported in the elevation model in the map's center and the vehicle is oriented according the vehicle pose information in the model base.

The map can be annotated with color displays of the classifications of the terrain, e.g., positive and negative obstacles, concealment, etc. To visualize the model accumulating over time, the model can be animated showing successive elevation models as the display steps through the data. The viewpoint can be fixed to the map coordinate frame and the vehicle will appear to pitch, roll, and turn in the center of the map as the map scrolls under the vehicle. Alternatively, the viewpoint could be attached to the vehicle's heading so the vehicle will appear to pitch and roll over the terrain that scrolls and rotates under the vehicle.

Perhaps also revealing would be displaying the next larger elevation model, which extends $500\text{ m} \times 500\text{ m}$ at 2 m horizontal postings. This map is fixed in the world coordinate frame, so the vehicle would appear to travel over a fixed terrain model. This map is scrolled less frequently, so it is a more stable reference frame than the small, high-resolution 40 m map. This larger map would illuminate larger terrain features beyond 20 m from the vehicle.

6 Conclusions

A variety of visualizations of terrain and navigation data have been explored. Their utility has been greatest in rapidly validating data collected in the field and in interactively exploring the nature of the data collected by various sensors and processed to varying degrees in order to understand their phenomenology. Both of these tasks have been greatly facilitated by these visualization tools and techniques for examining such inherently 3D data in a natural way.

Acknowledgments

Development of an intelligent perception and vehicle navigation control system is a jointly-sponsored effort that leverages technology from the NIST Intelligent Machines Initiative program and receives financial support from the Office of the Secretary of the Defense Demo III UGV program and the AUTONAV (Autonomous Navigation) research project agreement between the German Ministry of Defense and the United States Department of Defense. These and other visualization efforts using the ImmersaDesk have received financial support from NIST's SIMA (Systems Integrations for Manufacturing Applications) Program.

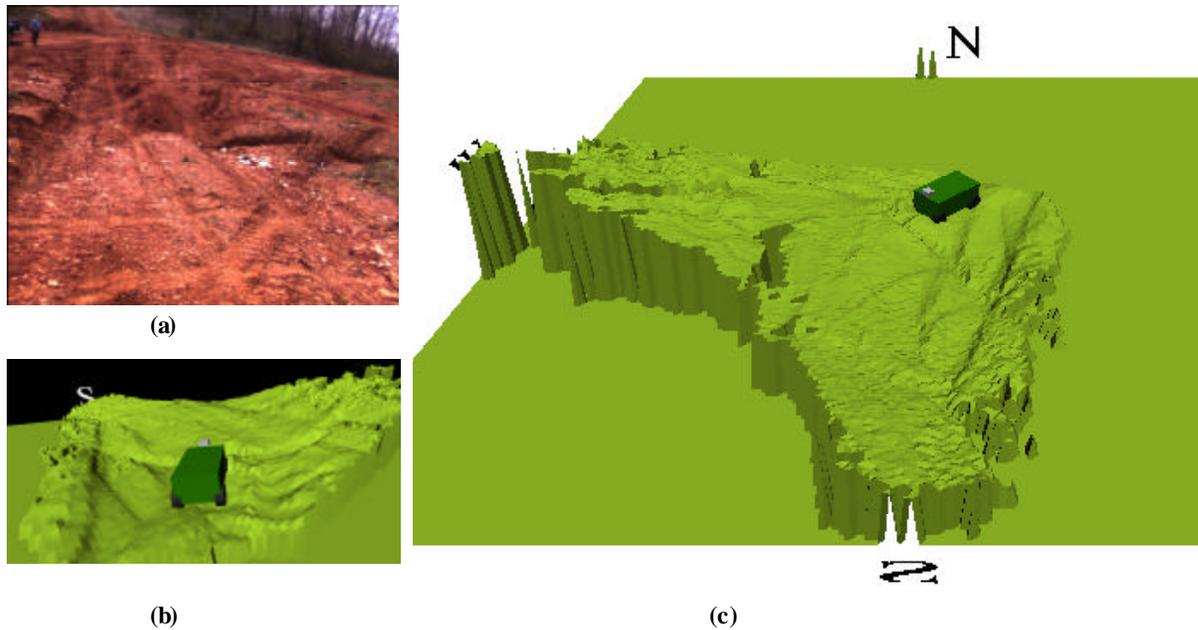


Figure 8: Elevation map of a gulch at Fort Knox Kentucky. This terrain model is accumulated from a sequence of Schwartz LADAR range images and vehicle navigation data. A color image of the gulch is shown in (a). The terrain model seen from a similar viewpoint is shown in (b). An overview of the accumulated terrain is shown in (c). The map covers 40 m \times 40 m with 0.2 m horizontal pistons. The vehicle is centered in the map. The map has been clipped to enlarge the portion containing elevation data. The vehicle model is 3 m long by 1 m tall by 1.5 m wide.

This work has benefited from innumerable discussions with Karl Murphy, Steven Legowik, Alberto Lacaze, Tsai-Hong, Stephen Balakirsky, Maris Juberts, Jim Albus, John Evans, Marilyn Abrams, Tommy Chang, and other Demo III colleagues. Thanks are due to Tsai-Hong Hong and Marilyn Abrams for the color imagery registered with range data, to Karl Murphy for Figure 2 and Figure 3, and to Tsai-Hong Hong for Figure 7.

References

- [1] Chang, T., Hong, T.-H., Legowik, S., Abrams, M.N. "Concealment and Obstacle Detection for Autonomous Driving," *Proceedings of the Robotics & Applications 1999 Conference*, Santa Barbara, CA, October 28-30, 1999.
- [2] Coombs, D., Murphy, K., Lacaze, A., Legowik, S. "Driving Autonomously Offroad up to 35 km/h," *Proceedings of IV 2000 the IEEE Intelligent Vehicles Symposium*, Dearborn, Michigan, October 4-5, 2000.
- [3] Cruz-Neira, C., Sandin, D.J., DeFanti, T.A., Kenyon, R.V., and Hart, J.C., "The CAVE: Audio Visual Experience Automatic Virtual Environment," *Communications of the ACM*, Vol. 35, No. 6, June 1992, pp. 65-72.
- [4] Czernuszenko, M., Pape, D., Sandin, D., DeFanti, T., Dawe, G. L., and Brown, M. D., "The ImmersaDesk and Infinity Wall Projection-Based Virtual Reality Displays," *Computer Graphics*, Vol. 31 Number 2, May 1997 pp 46-49.
- [5] Lacaze, A., Moscovitz, Y., DeClariss, N., Murphy, K. "Path Planning for Autonomous Vehicles Driving Over Rough Terrain," *Proceedings of the ISIC/CIRA/ISAS 1998 Conference*, Gaithersburg, MD, September 14-17, 1998.
- [6] Abernathy, M., and Shaw, S., "Integrating Geographic Information in VRML Models", *Proceedings of VRML Symposium 1998*, Monterey, CA, February 16-19, 1998.