Strike Visualization in Stereo on the Virtual Workbench

Ranjeev Mittu and Jeffrey K. Uhlmann
e-mail: {mittu or uhlmann}@ait.nrl.navy.mil
Decision Support and Force Level Planning section, code 5584
Information Technology Division
Naval Research Laboratory
Washington, D.C. 20375-5320
Tel: 202.404.8716
Fax: 202.404.7887

Paper Categories: Human Computer Interfaces, Decision Support Systems.

1 Introduction

The ability to view strike plans, such as routes for aircraft, radar detection envelopes and terrain is important to strike planners. A 3D view of the scenario provides the planner with an understanding of the important interactions and spatial relationships between the objects in the environment. Although current strike mission planning/ visualization systems provide an acceptable interface for viewing certain aspects of the strike environment, there are two areas that still need addressing in order to enhance mission planning and previewing. The first is the lack of a true 3D representation of the scenario. Current systems do work with 3D information such as terrain, routes, etc. However, this 3D information is displayed on a monitor that is 2D in nature. What is needed is a better approach to visualizing the important spatial relationships on a 2D monitor. The second item which is deficient is a poor visual representation of Radar Terrain Masking (RTM). The problem of representing RTM is of great importance to strike planners. When terrain obscures the view of a radar in a certain direction, we say that the terrain has masked the view of the radar. The ability to accurately identify RTM gives the planner a visual representation of where the "holes" are in the radars coverage. These "holes" represent safe areas through which strike aircraft may avoid detection. Additionally, if the strike planner has to modify the routes, then the ability to visualize RTM is very important.

NRL has developed a Decision Support System (DSS), hereafter referred to as the STrike Optimized

Mission Planning Module, or STOMPM [3]. This decision aid serves as a testbed for research and development of asset routing algorithms, where the assets may be strike aircraft or missiles. These routing algorithms take into consideration the positions of enemy radars, intended targets, and the features of the terrain. We have investigated various techniques within this DSS which could improve current strike planning/visualization systems, particularly with regard to more realistic 3D displays and a better visual model for RTM.

With regard to better 3D displays, we have incorporated stereo graphics within the STOMPM interface. This stereo interface is further enhanced by the use of a virtual workbench. The virtual workbench consists of a translucent table top containing a mirror placed at a 45 degree angle beneath the table top. Video is projected onto the mirror, which reflects the image onto the translucent table top. Since this video contains 3D stereo information, the images contained in the video appear as floating above/below the table top. This approach has led to a more realistic display of the strike environment. With regard to RTM, we have developed a line-of-sight radar model. In this model, a cone, with its vertex placed at the location of a radar, sweeps continuously in animated sequence from minimum to maximum elevation. RTM effects are easily visualized using this model, due to the fact that the portions of the cone blocked by the terrain at each elevation angle, are removed from the cone itself. This provides a strong visual cue as to what is visible to the radar at each elevation angle and direction.

Section 2 will cover the 3D line-of-sight model for RTM, and discuss its advantages over current methods for visualizing RTM. Section 3 will describe the stereoscopic interface that has been implemented on the virtual workbench within the DSS to enhance the visual capabilities of the system. We will also detail some of the limitations associated with the use of stereo. A brief summary, with conclusions, is presented in section 4.

2 Visualizing RTM Effectively in 3D

In the strike application, the deficiencies with common methods for visualizing radar detection regions were already well known. The most common representation of detection regions is as hemispherical domes defined by the maximum range of the radar (similar to Figure 1, but with non-transparent domes). These domes have portions cut out due to terrain masking of the radar, but these cut regions are usually occluded by non-masked regions. In other words, the most important features – the masked regions defining safe corridors for strike aircraft – cannot easily be seen. Using semi-transparent (fog/haze) graphical representations are often used so that the masked regions can be at least partially seen through the detection regions. In practice, unfortunately, the result is less than satisfactory (see figure 1).

Moving to a 3D display seems to be a good solution to the RTM problem because it promises to allow the user to quickly find a point-of-view from which the various masked regions can be easily

seen and studied. Experiments with detection domes reveal, however, that while the masked regions can be seen well from various angles, it is still not possible to study the regions from the direction of a proposed strike route because they are usually occluded by non-masked parts. What is needed is a way to strip away the occluding portions so that the masked region along the route can be seen clearly. Sophisticated interaction techniques using sensor gloves might provide an interface for the user to manipulate the portions of the dome that are displayed (or not displayed), but this would probably result in a tedious exercise for the user to retain precisely the desired regions of interest. A faster and more flexible visualization tool is needed.

Recognizing that the user needs to see different slices of the detection dome independently in order to make sense of the complicated interactions between terrain and the radar's region of coverage, some method is needed to generate and display the relevant slices of the dome automatically. It is of course impossible to produce a general purpose strategy for identifying what slices are "most relevant" to a specific application, so some unintelligent strategy is required. The most flexible approach along these lines is to use the fourth dimension, time, to iteratively cycle over some fixed decomposition of the radar dome. One could, for example, iteratively display different planar slices of the dome, one after another. This type of display produces a succession of flat surfaces. This way of decomposing the dome, however, does not seem particularly useful because the planar slices do not have any physical relevance to the geometry involved in RTM.

The geometry of a line-of-sight radar inherently defines a polar coordinate system with its origin at the position of the radar. In such a polar coordinate frame the important dimensions are range, azimuth, and elevation. By displaying two of the dimensions and iterating over the third, an animated display can be generated that (a) allows the user to see different portions of the region in isolation and (b) gives the user intuition about the overall 3D structures of the detection regions. The question that must be answered is which, if any, of the choices of coordinates to animate provides adequate visualization of the masked regions.

Iterating on range yields an animated sequence in which the detection region starts as a point and then expands to the full radar dome (figure 1). Unfortunately, this display provides little information beyond what is provided by the full radar dome alone. Iterating on azimuth produces a wedge defined by the minimum and maximum elevation and detection range of the sensor at each discretized azimuth value (Figure 2). The wedge sweeps around the sensor and clearly traces over the terrain masked regions. One of the principal features of the wedge display is that when combined with a static plot of a proposed strike route, a viewer can see as the wedge sweeps perpendicular to a leg of the route and determine whether the route goes through a masked region or is within the detection region of the radar. The only drawback of the the wedge display is that it does not provide the viewer with a strong sense of the 3D extents of the masked regions.

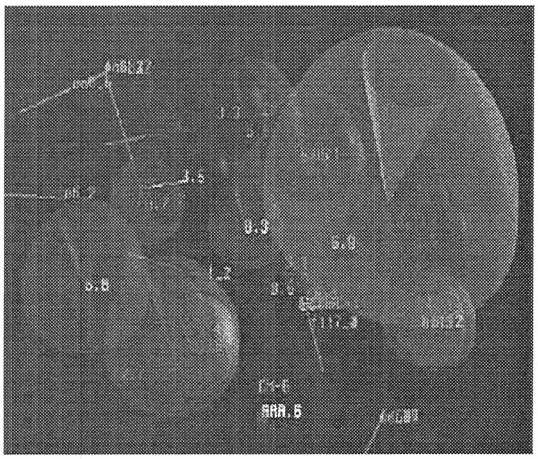

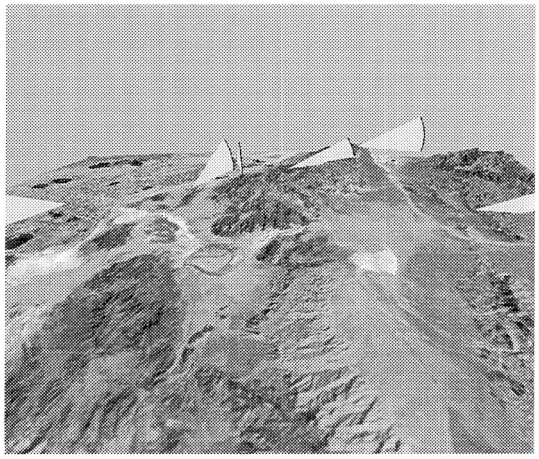

Figure 1 (left): Semi-Transparent Radar Dome; Figure 2 (right): Wedge Implementation of RTM

Iterating on elevation generates a cone that continuously sweeps from minimum to maximum to minimum elevation (Figures 4-6). This cone appears to melt over the terrain as it sweeps down from maximum elevation onto the terrain features and, in so doing, gives a strong sense of where the masking effects are most pronounced. In strike applications, the cone display provides an ideal visualization tool for determining when a proposed route is within the detection region of the radar. Specifically, by watching as the cone sweeps up and down it is possible to see exactly where and how it intersects a static plot of the route.

In summary, by displaying individual slices of a radar detection dome in isolation, the viewer is able to discern critical information that would otherwise be occluded if all of the radar detection information were displayed at once. Furthermore, by animating a sequence of slices the viewer is able to mentally construct a complete picture of the 3D structure of the radar dome while at the same time is able to scrutinize individual slices in the animated sequence. There does not appear to be any comparably effective visualization strategy that will work on a 2D display.

3 Stereoscopic visualization on the Virtual Workbench

Once strike plans have been generated for a given scenario, displaying the plans and the scenario itself in a form in which the important spatial relationships are easily seen has been an important issue in strike visualization. In some instances, one can get sufficient knowledge of the scene by viewing 3D information without the use of stereo. However, extracting detailed information and relationships in the data sometimes is made easier by the use of stereo. Consider the simple case of viewing the scenario (consisting of radar cones, terrain, and aircraft routes) from an orthogonal viewing position on the computer monitor [1]. It may be difficult to precisely obtain altitude information because altitude is in the same direction as a vector perpendicular to the computer screen, and hence there

[1] by "orthogonal viewing position", we mean an eye position which is looking in the -z direction onto the scenario.

are no depth cues for the viewer. This situation is alleviated when the same scene is viewed in stereo. The use of stereo makes visualizing depth easier. In fact, the use of stereo allows the user to view details of the strike environment from a relatively orthogonal viewing position, without having to set the viewpoint to some other perspective.

The advantage of the use of stereo on the virtual workbench is clear. Viewing the scenario from an orthogonal eye position gives the most overall information because very little is hidden to the viewer, whereas in some other projection, the effects of z-buffering [4] may hide important parts of the terrain, routes, etc. This is one of the reasons orthogonally projected views are used extensively in mission planning systems. However, the orthogonal viewing position has limitations in non-stereo mode. Stereo allows one to visualize details of the terrain and routes from an orthogonal viewing position, not having to view the terrain/routes with respect to the horizon to obtain important visual information.

We have modified the STOMPM interface to incorporate stereo graphics [1] [2]. Stereo graphics works in the following way: produce two images of the scene, one in the top half of the display and one in the bottom half of the display. When the monitor's refresh rate is doubled to 120HZ, only half of the 60HZ display is used for each refresh. One of these images is for the right eye (the one that was drawn to the bottom of the display) and one for the left (the one that was drawn to the top of the display). The user can then wear Liquid Crystal Display (LCD) shutter glasses to view the image in stereo. The shutter glasses work by showing the left eye the image intended for the left eye at the 60HZ refresh rate, similarly for the right eye. However, each eye sees the image meant for it in alternating sequence. The shutters in the glasses open and close in synchronization with the monitors refresh rate (the synchronization signal is sent to the glasses from an emitter). The overall effect to the user is a view which more closely resembles 3D – the stereo image can be either projected in front of, or behind, the computer screen, by adjusting a certain parameter in the software.

Now the question arises whether to view the orthogonally displayed image (mentioned in the second paragraph of this section) on the computer monitor or the workbench? It may appear that one is just as appropriate as the other. However, the ability to interact with the scenario (i.e., clicking near entities to get information about them, repositioning entities, modifying routes, etc) is most easily accomplished when it is displayed on a table top environment, such as the workbench. This is because users naturally perceive altitude in the same direction as a vector which is perpendicular to the earth. On a computer monitor, altitude would be in the same direction as a vector perpendicular to the computer screen. This is a little awkward to work with, especially when trying to adjust the routes in the z direction.

We have enhanced the stereo interface further by displaying this stereo image on a virtual workbench. The workbench was developed by the Virtual Reality (VR) group of the Advanced Information

Technology (AIT) branch at NRL. The schematic for the virtual workbench can be seen in Figure 3.

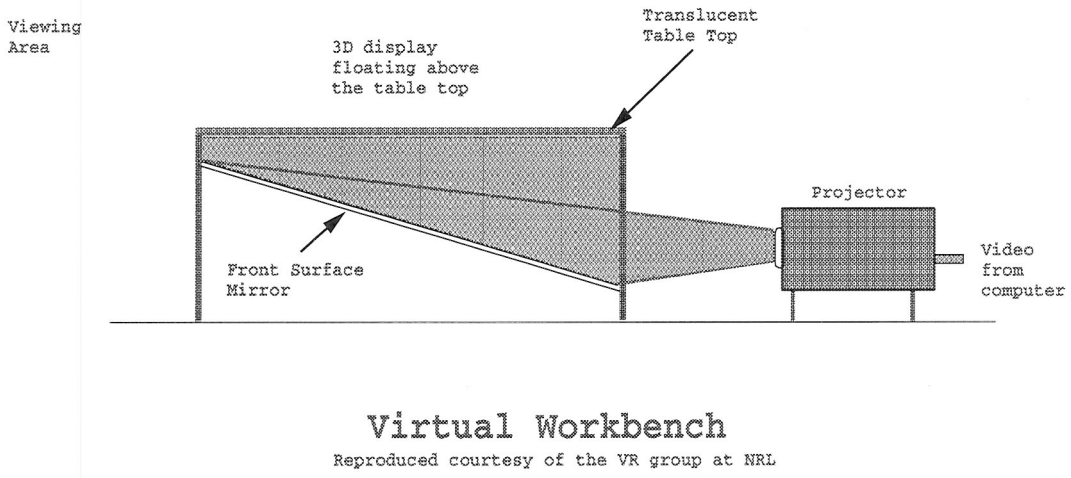

Figure 3: Schematic diagram of Virtual Workbench.

Video from the computer is sent to the projector, which projects the image onto a mirror. The mirror reflects this image onto a translucent table top. Two pairs of emitters are placed at the back two corners of the table top. The use of two pairs of emitters provides a stronger synchronization signal for the shutter glasses. When displaying an image on the workbench, the user can control whether the image appears to float above, or just below, the table top. This is controlled by adjusting a certain parameter in the software (just as the image was made to project as either inside or outside the computer screen).

We are also working to develop techniques for interacting with the scene on the workbench. The AIT VR group has integrated a sensor glove within the workbench environment. The user wears this glove when interacting with the environment on the workbench. By touching two fingers together wearing the glove, an EVENT is registered. The application program can then use this event to manipulate the scene in some fashion. This sensor glove technology can be used, for example, to click near a radar, routes, or targets to get specific information about them. The ability to manipulate the scene is also important (i.e., adjusting the routes, changing the position of radars and targets, etc). Again, the sensor glove can be used in a very natural way to move and edit objects in the scenario region.

One question that will need to be addressed is how to best display, on the workbench, textual attributes such as positions of radars, targets, ranges of radars, etc, when the user clicks near them. Experimentation with various types of techniques, (such as whether to display "free" text at the

point where the user touches the sensor glove fingers together, or to show this text on some type of "status" board at some location on the workbench) will have to be done to determine the best interface for textual information. In any approach that is taken, the text, or any "status" board containing text, will have to be drawn as part of the scene in stereo as well. One cannot just open up a separate display window for textual information because the window will not be projected either on top of or underneath the table top (it will appear to lie at the interface and can disturb the stereo effect).

There are several items worth mentioning about the use of stereo. As was already mentioned, the stereo image can be made to float above or just below the virtual workbench table top, by adjusting a certain parameter in the software. When one sets the parameter such that the image is floating above the table top, the interaction of the stereo image with the clipping planes determines how well the stereo effect is preserved. Zooming or panning effects can disturb the stereo effect because this can cause the stereo image to be clipped by the clipping planes. This is an artifact of stereo, and does not depend on whether the stereo image is projected on the computer screen or the workbench. It is important to have the entire scene visible when one wishes to project the image outside of the computer screen, or on top of the workbench. However, by having the entire scene visible on the workbench (or computer screen), it may be impossible to view the important details associated with the scenario. In many instances, it is desirable to project the image underneath the table top, or inside the computer screen. What we've noticed is that in our particular application, when the maximum height of the viewable terrain is underneath the table top (even after panning or zooming), the clipping planes have no adverse effect on the stereo image. In this regard, it may be easier to zoom or pan, thus enabling more details associated with the scenario to be seen.

3.1 Examples

The scenarios in Figures 4-6 are placed at the Naval Air Station (NAS) Fallon test range located near Fallon, Nevada. It should be mentioned that only the Digital Terrain Elevation Data (DTED), overlayed with Landsat imagery is used to represent the strike environment – these are not actual strike scenarios. The scenarios consist of cones (representing the radars' 360 degree view) sweeping from their minimum to maximum detection elevations. Areas cut out from each cone reveal the effects of terrain on the detection region of the radar. These "terrain masked" regions represent holes through which strike aircraft may avoid detection. The animation of the cones sweeping continuously from minimum to maximum elevation of detection represents a novel visualization tool for identifying holes in a radar's coverage. Figure 7 shows a suppression route (which destroys one of the radar sites). This leaves a safe corridor for the attack aircraft to pass through in reaching their desired targets.

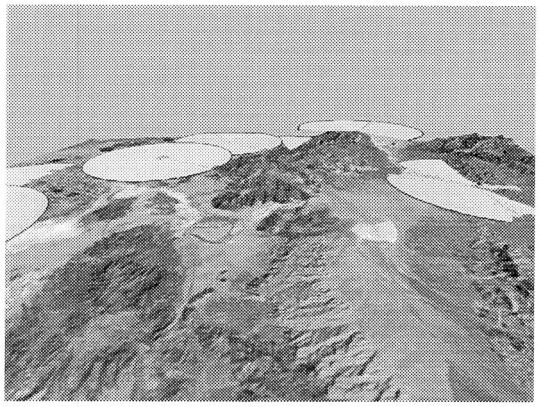

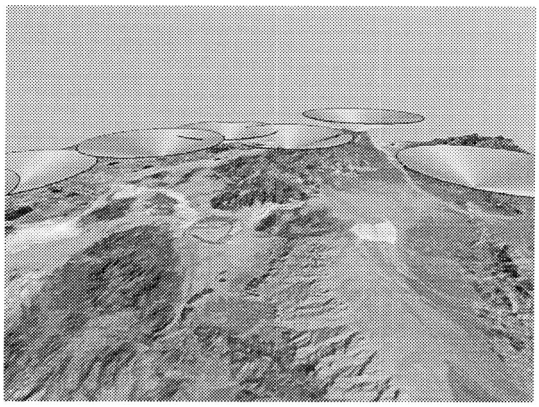

Figure 4 (left): Cones at minimum elevation; Figure 5 (right) Cones at median elevation (both figures in non-stereo mode.)

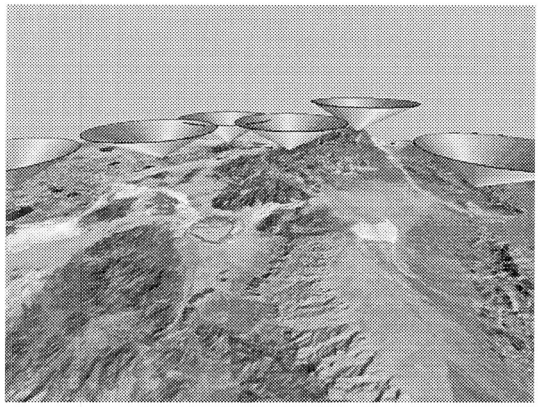

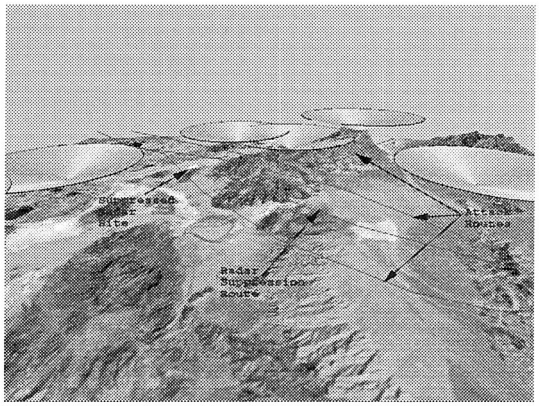

Figure 6 (left): Cones at maximum elevation; Figure 7 (right) Suppression and Attack routes (both figures in non-stereo mode.)

We have also built a simulation capability in which the routes evolve over time to their destinations. This provides the user with important information, such as where the routes are with respect to each other at specified times. The ability to simulate the routes, with the animation of the cones moving from minimum, median, to maximum detection elevation in stereo on the virtual workbench provides the user with a detailed account of the spatial relationships in the scenario, particularly RTM.

4 Summary and Conclusions

Current planning systems are non-stereo and generally use displays that are set to an orthogonal viewing position, but previewing is done in either an orthogonal, or other, position (however, z-buffering effects may hide some aspects of the scenario when the display is not in orthogonal viewing mode). Although planning systems use orthogonal viewing positions, there are certain limitations that still apply.

Without stereo graphics, 3D planning has certain obvious limitations. As an example, consider the case of trying to adjust routes in 3D space while in orthogonal viewing mode. In this mode, there are no depth cues to adjust by. In current systems, one has to manually enter the latitude and longitude for the routes in 2D (either by clicking in the scenario region or typing the latitude, longitude values), but then one has to adjust the altitude via a separate window (adjustment also may include moving any waypoint on the route in any given direction). This becomes a time consuming task. In non-orthogonal viewing mode, one may be able to adjust the routes, but if the routes are hidden behind objects such as terrain, they may not be seen from a given non-orthogonal mode and may require rotations for proper viewing, etc. Rotations are a cumbersome calculation and if the number needed becomes prohibitively large (which might well occur if the terrain is extremely mountainous, the number of planners is large, or one has to adjust/rotate/adjust/rotate in succession), valuable planning time may be wasted. Previewing the scenario in orthogonal mode also does not provide much in the way of substantial depth cues. However, previewing in non-orthogonal mode may be fine since very little user interaction is occuring with the scenario during previewing. Generally, a few good viewpoints will be enough to see what is going on in the scenario, therefore reducing the cost of rotations.

With the use of stereo on the workbench (in addition to the use of a sensor glove for interaction purposes), one can effectively plan and preview while the system is in orthogonal viewing mode (moreso than if the stereo image were projected on a computer screen), thus reducing the cumbersome calculations needed for scene rotations in non-orthogonal viewing mode. If the planner insists on using non-orthogonal mode for previewing, in stereo, then certain conditions should be met in order for the stereo effect to work properly. In fact, the first two conditions should be met whenever using stereo:

1. the clipping planes do not adversely affect the scenario when it is projected above the table top or
2. the scene is projected underneath the table top.
3. either 1 or 2 being coupled with the effects of z-buffering, or cost of the scene rotations, being tolerated (applies also in non-stereo mode).

Planning in non-orthogonal viewing mode in stereo may suffer the same drawbacks as in the case of non-stereo mode (i.e., cost of rotations). It should be clear that certain parts of the planning process (like adjusting the routes) are best done quickly and efficiently in stereo using an orthogonal viewing position (and it should also be clear why most systems stay away from non-orthogonal views for planning). For previewing, the use of a non-orthogonal previewing position in non-stereo mode may be sufficient, but does not provide the planner with the best depth cues as obtained via stereo. The use of a [non]orthogonal viewing position coupled with stereo mode provides better depth cues

than the non-stereo, non-orthogonal previewing position.

We have enhanced the capability of our system by using stereographics on a virtual workbench. This method of visualization, along with the line-of-sight model for RTM, has greatly improved the ability to plan and preview by enhancing the ability to see the spatial relationships in a strike scenario. We feel that this is a better approach than standard VR techniques for strike visualization because it still allows multiple planners to interact with each other, as they preview or modify the scenario (because the shutter glasses are transparent). In virtual reality displays, the ability to interact with other planners would be lost mainly because a head-mounted display must be worn by the user, thus restricting his ability to interact with the other planners.

The stereo interface coupled with the workbench has also been used for displaying and interacting with visualization objects generated using the results from a numerical simulation of an axially excited free square jet [5], visualizing the airwake and exhaust gas trajectories over the DDG51 destroyer [5] [6], and medical visualization [7] It is also being investigated for visualizing commercial aircraft routes for the Federal Aviation Administration (FAA). We feel that the stereo interface or stereo interface coupled with the workbench has the capability to enhance the user interface of any decision support system (particularly when it comes to visualizing the spatial relationships between the data), whether it be in the military or business enterprise.

References

[1] Tessman, Thant., "Perspectives on Stereo", IRIS Universe, Summer 1989 issue.

[2] Akka, Robert., "Writing Stereoscopic 3D Graphics Software using Silicon Graphics GL", StereoGraphics Corporation Internal report, April 17, 1991.

[3] Mittu, Ranjeev., "A Prototype System for the Evaluation of Interdependent Routing Algorithms for Military Aircraft", *11th Annual Decision Aids Conference*, Naval Postgraduate School, Monterey, CA. *Precision Strike Technology Symposium*, Johns Hopkins University Applied Physics Laboratory, Columbia, MD, 1994.

[4] IRIS 4D Series Graphics Library Programming Guide. Silicon Graphics Computer Systems, 5/90.

[5] Virtual Workbench - A Non-Immersive Virtual Environment for Visualizing and Interacting with 3D Objects for Scientific Visualization. Upul Obeysekare, et al.

[6] Landsberg A., J. Boris, W. Sandberg, and T. Young, "Analysis of hte NonLinear Coupling Effects of a Helicopter Downwash with an Unsteady Airwake", *33rd Aerospace Sciences Meeting*, Reno, NV, January 9-12, 1995.

[7] http://www.ait.nrl.navy.mil/vrlab/pages/research.projects/