VISUALIZATION TECHNIQUES TO SUPPORT MONITORING AND CONTROL OF AUTONOMOUS PLATFORMS

Matthew Plumlee, Rick Komerska, Roland Arsenault, Colin Ware, and Larry Mayer
Data Visualization Research Lab
Center for Coastal and Ocean Mapping
University of New Hampshire
Durham, New Hampshire
{mdp| komerska | roland}@ccom.unh.edu
{colinw | larry.mayer}@cisunix.unh.edu

Steven G. Chappell
Autonomous Undersea Systems Institute
Lee, New Hampshire
chappell@ausi.org

ABSTRACT

We have developed a suite of computer graphics visualization and interaction techniques that have direct application in aiding a human operator in mission planning for, monitoring of, and control of multiple autonomous and semi-autonomous entities, including AUVs. The suite is built around the techniques of center-of-workspace interaction (CWI), subordinate 3D windows (zoomports), and frame of reference operations (FoR-ops). CWI provides a natural and efficient interface for navigating a dynamic 3D virtual environment, while zoomports provide a way to monitor several areas of interest without losing the broader context. Frame of reference operations make it possible to monitor groups of AUVs for specific monitoring tasks. These techniques have been integrated into a package called GeoZui3D (Geographical Zooming User Interface) that supports real-time data acquisition, integration, and display. In this paper, we describe several applications of our techniques to mission planning, real-time monitoring, and vehicle control. We use the CADCON simulation toolset, which presents a virtual undersea environment populated with multiple AUVs performing a simple mission, to provide a context in which to demonstrate our techniques in these applications.

INTRODUCTION

The current state of AUV operational capability dictates that, at least for the next 5-10 years, most AUVs will operate in a semi-autonomous mode, with monitoring and control update frequencies ranging from near real-time to several days (Blidberg 2001). In an autonomous oceanographic sampling network scenario, where data delivered from multiple sampling platforms feeds large ocean computer models, forecasts that could shape revised AUV mission profiles can be expected to take place over the course of hours to days (Curtin, Bellingham et al. 1993). Human operators, using specialized visualization tools for task planning, progress monitoring, and control of individual vehicles and fleets of AUVs, will close the loop in the control of these systems.

Given the requirements for semi-autonomous monitoring and control of AUVs for the near future, the techniques we describe lend themselves well to the AUV robotics domain. In the interim, these techniques may also be useful for designing subsequent approaches towards true autonomy. Sheridan (1992) describes five generic functions involved in the supervisory control of a robotic system: (1) task planning, (2) programming new instructions, (3) monitoring the robot's progress, (4) intervening in the robot's control when required, and (5) learning from experience. We have applied our techniques across several of these functional areas. For task planning, we demonstrate how our interaction techniques aid in the exploration of baseline data and in running through missions via simulation. For monitoring, we demonstrate how our techniques can be used for quality control and assurance, status monitoring, and detection of hazardous conditions. Post-collection analysis benefits from our techniques in ways similar to task planning and monitoring. Finally, we demonstrate some ways in which these techniques might enhance an operator interface to control a single vehicle or groups of vehicles.

The first technique we describe is center-of-workspace interaction (CWI). This technique allocates the central portion of a 3D workspace to be the focal region, to which objects under investigation are brought. If the object happens to be moving, this technique "locks-on" to the object's frame of reference, so that further investigation can be done with respect to the object rather than the surrounding environment. CWI also provides a natural place to situate often-used tools and object-specific tools.

The second technique we describe is the integrated use of subordinate 3D windows (zoomports), each of which uses CWI. Zoomports appear similar to the “blow-up” maps seen in National Geographic magazine, but are taken into the interactive three-dimensional realm. Like the “blow-up” maps, each zoomport is visually linked to an overview window through a tether line and a box (which we call a zoomport proxy). These tethers and proxies help a human operator determine the absolute locations (with respect to the surrounding context) and relative locations of zoomports and their current points of interest.

The third technique we describe is the creation of composite frames of reference, which we call *frame of reference operations* (FoR-ops). While there are many different FoR-ops possible, we describe two of particular interest to our discussion. The first is the *overview* operation, which tracks a group of objects (including AUVs and possibly stationary entities such as buoys or waypoints). The second is the *closest-proximity* operation, which monitors a group of objects and triggers alerts when the closest two objects in the group are within a predetermined distance. Zoomports can be attached to FoR-ops in a number of different ways, each of which presents trade-offs for different tasks.

CENTER-OF-WORKSPACE INTERACTION

Many people are familiar with two metaphors for navigating a 3D virtual environment: anyone who has played a flight simulator game has experienced the flying metaphor; anyone who has played Doom or Quake has experienced the walking metaphor. Other metaphors like eye-in-hand (Ware and Osbourne 1990), zooming (Bederson and Hollan 1994) or point of interest navigation (Mackinlay, Card et al. 1990) are less familiar, but they still share the same basis: all of these metaphors rely on moving the viewpoint. Not just moving the viewpoint, in fact, but also *thinking* of traversing a virtual scene fundamentally in terms of moving the viewpoint.

When interacting with the real 3D world, however, humans generally prefer to keep focus objects within arms’ length. We commonly construct workspaces in which objects and tools are brought to a central working region: a kitchen stove, an office desk, or a technician’s workbench. The utility of these real-world environments lies in the tools and resources made readily available within them, such as the heat and cooking implements available at a kitchen stove. When possible, objects are generally brought to these environments to be operated on or transformed, though some objects are too bulky. In a virtual environment, no object is too large as long as one can control the scale of the world relative to the workspace. Thus it makes sense to give up the common notion of the centrality of the viewpoint in a 3D environment in favor of a center of focus within arms’ reach.

We can define a center of workspace more formally as a frame of reference (FoR) with its origin not far behind the center of a display window. We also define the *central region* of the workspace to be the largest cube centered at the center of the workspace that fits in the viewable workspace (the cube abuts the display monitor). The center of the workspace FoR has six degrees of freedom (DoF): 3-DoF for the position of the reference point, 2-DoF for orientation (azimuth and elevation), and 1-DoF for scale. The viewpoint is fixed with respect to this FoR and the user is given control of the FoR with respect to the virtual world. This means the same operations can be used to look around the virtual world as to spin an object at the center of the workspace.

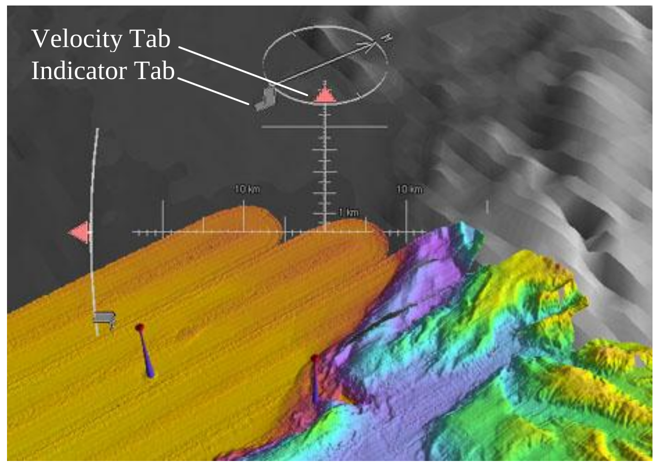

Figure 1: CWI in the context of GeoZui3D. The center of the workspace is at the intersection of the embedded central crosshairs. The azimuth and elevation control widgets are anchored near the crosshairs. Scale indications are dynamically updated along the length of the central crosshairs.

The CWI technique (as well as the techniques we describe later) has been integrated into a software package called GeoZui3D (Ware, Plumlee et al. 2001). GeoZui3D runs under Windows and Linux operating systems and supports the simultaneous rendering of geo-referenced objects from landmasses, to tracks of underwater vehicles, to temperature information in the water volume. GeoZui3D makes the center-of-workspace FoR explicit as a cross to which control widgets are attached as illustrated in Figure 1. Selecting and dragging either the triangular, pink velocity tabs or the rectangular, gray indicator tabs with the left mouse button affect rotation in the workspace FoR. Azimuth is free to vary in all 360°, but elevation is constrained between –30° (looking up) and 90° (plan view). There is no mechanism to induce roll, and azimuth rotation is always applied before elevation (except to the widgets). This maps well to human degrees of freedom: when we move around in the world we typically rotate our heads about a vertical axis (azimuth) and tilt our heads forward to look down (elevation), but we do not often tilt our heads sideways (roll). This also means that the up direction is always in the plane perpendicular to the screen X-axis, and that the horizon is always horizontal when it appears in the window.

Scaling and translational motion in GeoZui3D combine features of point of interest navigation described by Mackinlay (1990) with the NestedVision3D center-of-workspace navigation described by Parker (1998). Using the middle mouse button, a user "clicks on" a point of interest on any object in the scene to quickly translate it to the center of the workspace. Releasing the middle mouse button halts the translation immediately. Scaling is performed about the center of the workspace by keeping the middle button depressed and moving the mouse forward (to zoom in) or backward (to zoom out). If the point of interest that a user clicks on is a moving target, such as an AUV, the center of workspace "locks on" to the AUVs frame of reference. This makes it trivial for a user to "follow" a moving object, and look around according to that object's vantage point.

The CWI concept brings a number of important benefits. Its biggest contribution is the simplicity and efficiency it brings to virtual environment exploration during the planning and analysis stages. However, it also solves several interface problems that permeate throughout the planning, monitoring/control, and analysis stages of a given mission. It solves some of the problems with stereoscopic displays; it guides design interface decisions, such as where to place tools such as rulers and directional indicators; it can be used to support rendering that maximizes both interactivity and resolution; and it opens up new possibilities in creating and visualizing relationships between multiple 3D views.

MULTIPLE WORKSPACES: ZOOMPORTS

The CWI concept facilitates the implementation of a multiple workspace environment. We call each workspace in GeoZui3D a *zoomport*. Each zoomport is a window with its own center-of-workspace FoR, and the same set of widgets and interaction techniques we have already described. In general we conceive of zoomports as a way of showing a selected area in greater detail, although the view can also be scaled to give an overview, or otherwise transformed to provide an alternate view. There are two important structures behind the implementation of zoomports, both partially inspired by the interlinked 2D windows described by Ware and Lewis (1995). The first structure is logical, and the second is visual.

The logical structure of zoomports is hierarchical. The hierarchy is rooted at an ever-present zoomport we call the *main window* (illustrated in Figure 1). Additional zoomports are created as *subwindows* of the main window, which can in turn parent their own zoomport subwindows. The user creates a zoomport by selecting the zoomport tool from a menu (or using a hotkey) and clicking on a point on any surface in the desired parent window. The new zoomport's FoR is created with its location at the selected surface point, its orientation matching its parent's, and its scale set to make the new focus appear roughly twice as large as it appears in the parent window.

The visual structure of zoomports reflects the hierarchy while aiding the user in fusing the information from multiple zoomports. The visual structure consists of three components as illustrated in Figure 2:

1. The zoomport itself, with its window management border helping to distinguish it from other zoomports.
2. A *zoomport proxy*: a representation of the zoomport's FoR displayed in the window of its parent.
3. *Zoomport tethers*: lines with one endpoint at the center-of-workspace proxy in the parent view, and the other at one of two near corners of the zoomport.

The zoomport proxy design represents the viewpoint with a box, resembling a spotlight or a camera. Two semi-transparent triangles representing the vertical and horizontal viewing angles emanate from the box to suggest the shape of the viewing frustum. A line is dropped from the box through the surface below, providing a better grounding of its position and thereby indicating zoomport orientation. A semitransparent post joins a cross that represents the zoomport's center of workspace. The post is anchored at the center of the cross by a small square, and has a marker at the top. All three parts of the post are scale-independent, providing a way to mark the presence of a zoomport whose scale is too small for the rest of the proxy to be seen. At such a scale, no further effort is made to indicate orientation or scale.

The tethers perform a number of important interface functions. Tethers make it easier to locate a zoomport's proxy, and better integrate information between views (Yamaashi, Cooperstock et al. 1996). Tethers act to disambiguate the relationships between multiple zoomports and their proxies in a common parent window.

The user can click on a zoomport proxy and drag it along visible surfaces in the parent window, causing the FoR of the associated zoomport to update accordingly. This provides an easy way for a zoomport's focus to be changed from another context: all the user needs to do is place the zoomport proxy on the new focus object.

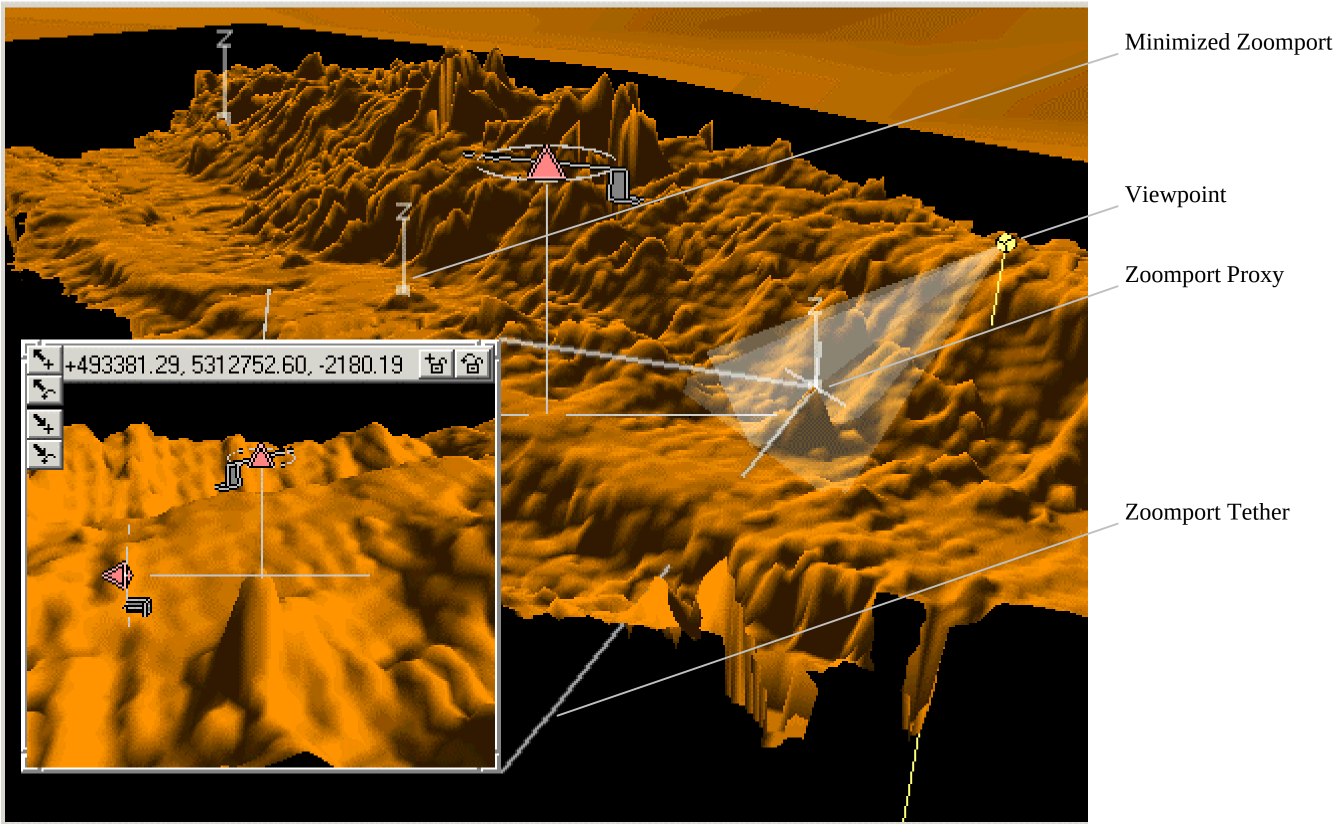

Figure 2: A zoomport, with proxy and tethers linking its center of workspace to its parent's view. Two minimized zoomports appear back and to the left; the base of each post is at the center of workspace for the respective zoomport.

The use of multiple zoomports with CWI works well without the distortions intrinsic to other techniques such as "fisheye" views (Furnas 1986; Sarkar and Brown 1994; Lamping, Rao et al. 1995). A small feature can be seen at one scale in the context of the surrounding environment at another scale. With three zoomports, two features at the same scale but separated by a large distance can be viewed simultaneously in the context of the area between them. Several AUVs can be monitored independently by assigning one zoomport to each. The power of zoomports to aid in planning simulations, monitoring, and control is further enhanced by assigning zoomports to the results of frame of reference operations.

FRAME OF REFERENCE OPERATIONS

In GeoZui3D, we consider that each object, such as a vehicle or waypoint, has its own FoR. This FoR could be at a central point such as the center of mass, or at a focus point in front of a camera location, or at arm's length of a robotic manipulator arm. This formulation allows us to create composite objects that aggregate the information from several FoR's into a single FoR to which zoomports can be attached. We call the aggregation process a *frame of reference operation* (FoR-op). We now describe two examples of FoR-ops. The first is the *overview* FoR-op, and the second is the *closest-proximity* FoR-op.

Figure 3 shows a 2-dimensional representation of the idea behind the overview FoR-op. A cube is created that encompasses the FoR's of the vehicles under investigation, generated from the minimum and maximum *x*, *y*, and *z* coordinates among the five FoR's. The center of the cube becomes a perfect location to attach a zoomport's center of workspace, and the cube itself maps perfectly to the central region of a zoomport's workspace. This makes it possible to attach a zoomport to a group of vehicles and follow them no matter where they go, or how far apart they stray from each other.

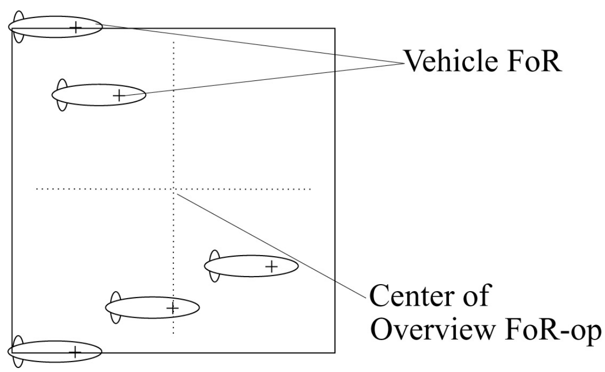

Figure 3: 2-dimensional representation of the overview FoR-op, applied to five AUVs.

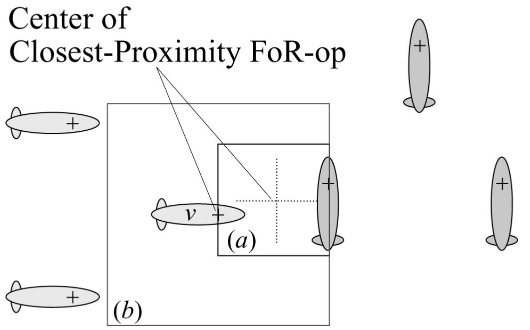

Figure 4: 2-dimensional representation of the closest-proximity FoR-op, applied to six AUVs. The smaller square (a) shows how the FoR-op maps to a zoomport when the zoomport's center of workspace is attached to the FoR-op itself. The larger square (b) shows how the FoR-op maps to a zoomport when the zoomport's center of workspace is attached to vehicle *v*, with its scale set by the FoR-op.

Figure 4 depicts a 2-dimensional representation of the idea behind the closest-proximity FoR-op. A cube is created that encompasses the FoR's of the two closest vehicles of the group under investigation, similar to the cube generated for the overview FoR. The center of the cube becomes a good location to attach a zoomport's center of workspace, such that the central region of the workspace contains the vehicles that are closest together. This makes it possible to pop up a zoomport to alert an operator when two vehicles get too close to each other. If vehicle control is supported, however, it may be better to attach a zoomport to the FoR of one of the vehicles, so that one of the vehicles could be commanded to avert the hazard. For such a purpose, the closest-proximity FoR-op can adapt to a new center, and can rescale appropriately (Figure 4(b)).

FoR-ops can be performed across vehicles, as well as other geo-referenced objects, and can be made dependent on any system variable of interest. Such vehicle-related variables might include energy level, current mission task, organizational role, or estimated time to next task transition. This provides an operator with a powerful tool for monitoring sets of vehicles and objects based on user-defined criteria.

APPLICATION TO AUV DOMAIN

All three of the interaction and display techniques we've described (CWI, zoomports, and FoR-ops), can be used together to address issues within the AUV domain. We now describe three specific application areas and how these techniques can be utilized in each. We start with a task-planning application, continue with a real-time monitoring application, and finish with a discussion on operational vehicle control. The techniques that we describe in each of these application areas can also be applied to post-process data, scenario analyses, and simulation.

To aid us in our exploration of these issues, we have developed a CORBA-based interface that allows us to link GeoZui3D as a visualization client to the CADCON environment server. CADCON (Chappell, Komerska et al. 1999) is a simulation harness expressly designed to provide a shared environment for multiple distributed interacting vehicle and instrument platforms. Its primary purpose is to provide a testbed to explore research issues associated with higher-level interactions among multiple heterogeneous participants. CADCON provides a relatively simplistic but complete simulation of an underwater environment including simulations of various AUV platform types.

Task Planning

Task planning is employed during the initial mission scenario definition process. Vehicle commands are created by the operator and scripted into mission profiles. These command sequences would ideally then be executed in a simulated virtual environment at faster than real-time speeds. This would allow for relatively rapid assessment, refinement and verification of the sequences before they are actually programmed onto the platform(s). Later, during the operational phase of the mission, data collected from monitored vehicles and other sources can be used to help the operator re-plan aspects of the mission when there is a need to alter the original vehicle plan or respond to new circumstances. Because task planning relies closely on aspects of vehicle monitoring and control, many of the techniques described below also apply in the subsequent AUV domain application sections.

The visualization and interaction techniques discussed in this paper can be employed to aid in this iterative planning/re-planning effort. For example, zoomports can be attached to follow the state of one or more vehicles. CWI can be used to view the moving vehicle from many useful vantage points during the simulated mission. CWI makes it very easy to switch back and forth between vehicles, or between a vehicle and a static feature. During the simulation, closest-proximity FoR-ops can be employed to pop up and follow vehicles that come within a given radius of other vehicles or objects in the environment, providing a kind of collision detection tracking display. Alternately, a zoomport can be attached to the FoR-op to optimally encompass vehicles or objects that meet any user-defined criteria. An example would be the tracking of vehicles whose battery charge is running low.

Figure 5 further demonstrates these ideas. In this figure, the overview of a previously mapped harbor terrain appears in the main window. We assume that a set of tasks has been developed and appropriate commands have been downloaded to the vehicles. Simulation tools are then used to model the expected behavior of the vehicles under these new commands, and visualization techniques are used to aid in evaluating their effectiveness. For example, a closest-proximity FoR-op tracks an AUV in the process of docking at an energy/data transfer station, shown in the upper left

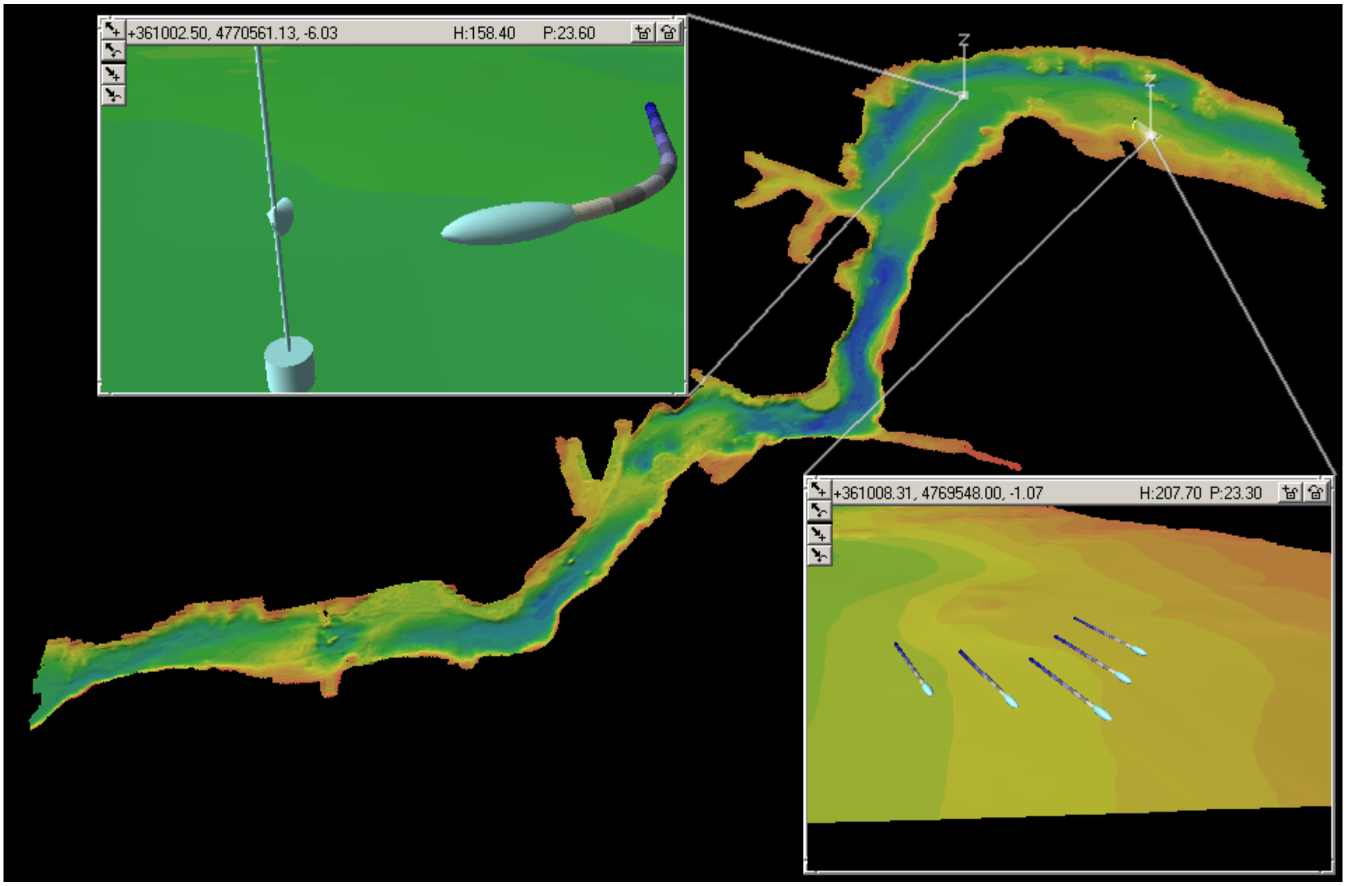

Figure 5: Zoomports used to monitor various vehicles' progress during task planning simulation and at-sea operations.

zoomport, adjusting its scale automatically as the vehicle docks. In another zoomport, AUVs identified as part of a group executing a "search" behavior are tracked in a single window. By being attached to the overview FoR-op, the zoomport window dynamically changes scale and position to encompass or exclude vehicles as these user-defined criteria are met over time. Based on analyses supplemented by these views, the operator can then decide whether her commands are in need of revision or are ready to be downloaded to the actual vehicles for execution.

Real-Time Monitoring

Even with perfect planning, human intelligence is required during task execution to deal with unforeseen events. For example, given a survey mission, these events could be geographical in nature, caused by poor knowledge of the shape of the ocean bottom; they could be mechanical, creating gaps in the data or unexpected behavior in the vehicles themselves; or they could come from less-than-perfect planning, potentially creating hazardous conditions for which a vehicle's intelligence is unprepared.

Visualization of primary data (such as bathymetry) and control metadata (such as vehicle position), integrated in a 3D display, can aid in monitoring a vehicle's progress. We have recently done this, using the CORBA-based interface to bring data into GeoZui3D from a hull-mounted sonar unit. We are also actively pursuing its use with telemetered data from a Hugin AUV in conjunction with C&C Technologies, Inc. Figure 6 illustrates a view of data being collected from the Hugin AUV, with a zoomport tied its position.

GeoZui3D's real-time data integration combines with its interaction techniques to create new monitoring capabilities. For instance, it is possible to interactively explore the new data using CWI to discover probable outliers, and to find places where another pass may be required to collect more data. Likewise, photographic or backscatter imagery can be viewed in-situ with respect to the bathymetry to ensure that the data agree. Zoomports can be used in conjunction with FoR-ops to monitor the progress and status of AUVs, in ways similar to those described in the task planning discussion.

Teleoperation & Supervisory Control

All of the uses we've described for our interaction techniques so far are monitoring uses that currently have supporting implementations in GeoZui3D. The interaction ideas could be taken a step further, however, into the realm of single or group vehicle control. The idea is to facilitate the assignment of commands to vehicles, both in real-time control intervention situations as well as the building of mission scripts that use these commands. Care must be taken to insure that there is a proper balance struck between the visual clutter of too many tools and status indicators, and the requirement of effective monitoring and control of the vehicle.

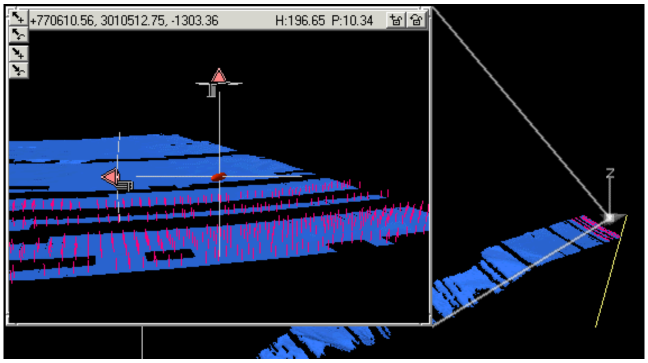

Figure 6: GeoZui3D display of binned sonar data collected in real-time from a Simrad EM2000 mounted in a Hugin AUV, operated by C&C Technologies in the Gulf of Mexico (Arsenault 2001).

For an ROV or an AUV with real-time communications capability, these techniques could be used as the basis for a teleoperative control interface. For each vehicle, vehicle-specific widgets could be anchored near the center of the workspace when a workspace is centered on that vehicle. Examples include widgets for thruster directional control, or for manipulator-arm movement and grasping. In a multiple vehicle environment, centering on another vehicle in the display would reset the control widgets as appropriate to that vehicle's capabilities and status.

The CWI and zoomport techniques could be used in a supervisory control capacity, assuming an AUV that supports supervisory control. In this case, the command interface might consist of a zoomport with child windows for information such as status and mission scripts. The status window would provide information related to the last poll of vehicle status, such as speed and heading. The mission script window would provide a view of the scripted command file(s) for that vehicle, in addition to the set of vehicle commands. Typical commands might include "Transit", "Survey", and "Dock" and would be based on the capabilities of that particular vehicle. The zoomport itself would provide a vehicle-centered reference frame through which the operator could view the surrounding environment. These ideas can be extended to groups of vehicles: commands displayed to the operator would reflect the higher-level behaviors of the group as a whole, and would presumably allow the operator to command the group as a unit. Figure 7 shows a hypothetical zoomport attached to a vehicle, including its current mission profile and status information. Implicit in these cases is the requirement that the monitoring and command software utilizing these interaction techniques be aware of the supported vehicle command capabilities and status.

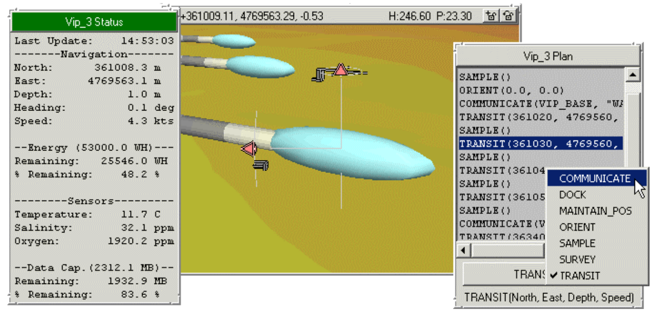

Figure 7: Hypothetical zoomport showing vehicle status, mission profile and allowed command set.

SUPPORTING INTERACTION TECHNOLOGIES

Many of the benefits that our techniques bring to planning, execution, and analysis of tasks are enhanced through the addition of supporting visualization techniques. We describe briefly some of these techniques below, many of which are currently implemented in GeoZui3D.

During the post-process analysis phase, just as important as making a discovery from the data is presenting that discovery to others. To facilitate such dynamic presentations, we have developed a scripting interface for GeoZui3D based on Tcl (Welch 1997) that allows zoomports and objects in the scene to be created, moved around, and destroyed at program run-time. The interface also supports the attachment of zoomports to moving objects and groups of objects. The scripting interface extends CWI by giving each FoR a name and textual commands that are analogous to the interactions available with the mouse. In addition to enabling stored scripts to be played to reenact previous event sequences, the scripting interface provides a simple hook into GeoZui3D for real-time mission monitoring applications.

We have found it useful to provide the operator with an interactive lighting control to visually detect (or to hide) certain targets or systematic aberrations in the data. We have also found it useful to employ an interactive mechanism for exaggerating terrain height, allowing an operator to make subtle features more easily seen. Finally, we have found it useful to provide a way of draping 2D images (such as sidescan sonar, error surfaces, or generic grids) over 3D surfaces, allowing an operator to fuse this more abstract 2D spatial data with the underlying topography.

CONCLUSION

We have described three novel visualization-based interaction concepts: center-of-workspace interaction, 3D zoomports, and frame of reference operations, and described how they might be applied to support AUV mission planning and operations. We have incorporated these techniques into a software tool called GeoZui3D, which has been successfully used to support bathymetric mapping projects at the Center for Coastal and Ocean Mapping at the University of New Hampshire. In addition, we have begun to explore the use of these techniques in the AUV domain, through our work with real-time monitoring of the Hugin AUV and simulated multiple vehicle monitoring and control in CADCON. Our hope is that these techniques will not only make the process of interacting with AUVs more intuitive and less prone to error, but will better enable researchers to advance the state of the art in the development of AUV systems.

REFERENCES

Arsenault, R. (2001). Real-Time Display of Hydrographic Data for QA and QC. U.S. Hydrographic Conference.

Bederson, B. B. and J. D. Hollan (1994). Pad++: A Zooming Graphical Interface for Exploring Alternate Interface Physics. UIST '94, ACM Press.

Blidberg, D. R. (2001). Personal communication. R. J. Komerska.

Chappell, S. G., R. J. Komerska, et al. (1999). Cooperative AUV Development Concept (CADCON) - an Environment for High-Level Multiple AUV Simulation. Eleventh International Symposium on Unmanned Untethered Submersible Technology, Durham, NH. See also http://www.ausi.org/cadcon/simulation.html

Curtin, T. B., J. G. Bellingham, et al. (1993). "Autonomous Oceanographic Sampling Networks." Oceanography **6**(3): 86-94.

Furnas, G. W. (1986). Generalized Fisheye Views. CHI '86, ACM Press.

Lamping, J., R. Rao, et al. (1995). A Focus+Context Technique Based on Hyperbolic Geometry for Visualizing Large Hierarchies. CHI '95, ACM Press.

Mackinlay, J. D., S. K. Card, et al. (1990). Rapid Controlled Movement Through a Virtual 3D Workspace. SIGGRAPH '90, ACM SIGGRAPH.

Parker, G., G. Franck, et al. (1998). "Visualization of Large Nested Graphs in 3D: Navigation and Interaction." Journal of Visual Languages and Computing **9**: 299-317.

Sarkar, M. and M. H. Brown (1994). "Graphical Fisheye Views." Communications of the ACM **47**(12): 73-84.

Sheridan, T. B. (1992). Telerobotics, Automation, and Human Supervisory Control. Cambridge, MA, MIT Press.

Ware, C. and M. Lewis (1995). The DragMag Image Magnifier. CHI '95.

Ware, C. and S. Osbourne (1990). "Explorations and Virtual Camera Control in Virtual Three-Dimensional Environments." Computer Graphics **24**(2): 175-183.

Ware, C., M. Plumlee, et al. (2001). GeoZui3D: Data Fusion for Interpreting Oceanographic Data. Oceans 2001 (In Press) See also http://www.ccom.unh.edu/vislab/GeoZui3D

Welch, B. B. (1997). Practical Programming in Tcl and Tk. Upper Saddle River, New Jersey, Prentice-Hall.

Yamaashi, K., J. R. Cooperstock, et al. (1996). Beating the Limitations of Camera-Monitor Mediated Telepresence with Extra Eyes. CHI '96, ACM Press.