

**Nicholas R. Hedley**  
nix@u.washington.edu

**Mark Billingham**  
**Lori Postner**  
**Richard May**

Human Interface Technology  
Laboratory  
University of Washington  
Box 352-142  
Seattle, WA 98195, USA

**Hirokazu Kato**

kato@sys.im.hiroshima-cu.ac.jp  
Faculty of Information Sciences  
Hiroshima City University  
3-4-1 Ozuka-Higashi, Asaminami-ku  
Hiroshima 731-3194, Japan

# Explorations in the Use of Augmented Reality for Geographic Visualization

## Abstract

In this paper, we describe two explorations in the use of hybrid user interfaces for collaborative geographic data visualization. Our first interface combines three technologies: augmented reality (AR), immersive virtual reality (VR), and computer vision-based hand and object tracking. Wearing a lightweight display with an attached camera, users can look at a real map and see three-dimensional virtual terrain models overlaid on the map. From this AR interface, they can fly in and experience the model immersively, or use free hand gestures or physical markers to change the data representation. Building on this work, our second interface explores alternative interface techniques, including a zoomable user interface, paddle interactions, and pen annotations. We describe the system hardware and software and the implications for GIS and spatial science applications.

## 1 Introduction

As computers become more ubiquitous, they are increasingly being used to support remote collaboration. Adding a camera and microphone to a computer converts it into a powerful tool for remote teleconferencing, and software such as Microsoft's NetMeeting allows users to share applications and effortlessly exchange data.

Although computer-supported remote work is common, fewer interfaces support face-to-face collaboration. This is particularly true for viewing and interacting with three-dimensional data sets and models. In this setting, it is particularly important for the interface to provide intuitive viewing and manipulation of spatial models while at the same time supporting the normal conversational cues of a face-to-face meeting.

Augmented reality (AR) is one promising technology for face-to-face collaboration on spatial tasks. AR superimposes three-dimensional virtual images on the real world so users can see each other, the surrounding real environment, and the virtual models in their midst. Our previous research has found that users can collaborate more effectively in an AR setting than on the same task in an immersive VR interface (Billingham, Weghorst, & Furness, 1998) and that AR interfaces can enhance the cues already present in face-to-face collaboration (Billingham, Bowskill, Jessop, & Morphet, 1998).

One of the greatest benefits of AR interfaces is that they can be integrated into the existing workplace and combined with other, more traditional interface technologies. For example, the EMMIE system is a hybrid user interface

that synergistically merges information in an AR headset with data shown on monitor and projection displays (Butz, Hollerer, Feiner, MacIntyre, & Beshers, 1999). Within EMMIE, users can seamlessly move virtual objects from being overlaid on the real world to being placed on a desktop monitor. Despite this, there are few examples of interfaces that combine AR techniques with other technologies.

We report on hybrid interfaces that combine technologies such as augmented reality, immersive virtual reality, and computer vision methods for natural user input. Our prototype interfaces are designed to allow multiple users to view and interact with geospatial data. However, the techniques we describe could be applied to many different application domains. It is our hope that this work will motivate others to consider how AR technologies can be used in hybrid user interfaces for their own applications.

In the remainder of this work, we first describe the application area in more depth and then each of the technologies used in our systems and the prototype hybrid interfaces developed around these technologies. We briefly present user experiences with the system and contrast our work with those of other researchers in the field. Finally, we outline our directions for future work and lessons we have learned from developing these interfaces.

## 2 Application Domain: Collaborative GIS

Visualization of geospatial data is one of the most important application areas for virtual reality and computer graphics technology. Geographic information systems (GIS) facilitate powerful spatial analytical capabilities in academic, private, government, and corporate arenas. These capabilities are largely a result of the analytical metaphors (such as a map overlay) and functionality facilitated by GIS software. Interface research during the past twenty years have brought us technologies that may fundamentally change the nature of how we interact with and understand spatial data. Advanced environments and tools render the interface far more transparent than ever before. Researchers have already


shown the value of immersive virtual reality for data visualization (Dede, Salzman, Loftin, & Sprague, 1999), (Erickson, 1993) and science education (Lin & Loftin, 1998). These new technologies need to be investigated to evaluate their potential for integration into meaningful and valuable spatial applications. Geographers have been developing 3-D GIS visualizations and moving independently of computer science and human-computer interaction (HCI) toward virtual environments for some time (Kraak, 1994; Hedley & Campbell, 1998; MacEachren et al., 1999; MacEachren & Kraak, 2001; Slocum et al., 2001). Convergences of computer science, HCI, and geographic visualization have resulted in compelling collaborative 3-D “geo-visualization” interfaces that represent a significant advance in GIS interfaces (Hedley, Drew, Arfin, & Lee, 1999). Initial efforts have informed researchers of methodological approaches to evaluating the effectiveness of collaborative GIS settings (Nyerges, Moore, Montejano, & Compton, 1998).

Advanced interface technologies, environments, and associated modes of interaction are significant because of the way in which they may change the way in which we acquire spatial knowledge. The adoption of such visualization interfaces has practical and theoretical implications. Cartographers and spatial scientists need to reconsider the nature of human-visualization relationships as new visualization technologies emerge. This is especially important as these emergent technologies employ different mechanisms by which people interact with visual spatial information and acquire knowledge. This in turn requires us to evaluate the aspects of interfaces, visualization content, and the individual users that influence the formation of spatial mental models.

Linked to these altered relationships are significant questions that affect the meaningful use and application of such tools. Do new visualization techniques match traditional forms of spatial visualization training, or do they engage different faculties of individuals? Does use of an AR 3-D map (Figure 1) improve performance in spatial decision-making over traditional 2-D maps, or 2-D and 3-D GIS visualizations (figure 2)? Are the interactions different? What relative influence do 3-D spatial model content and interface type (AR, desktop,



**Figure 1.** AR geographic visualization.



**Figure 2.** Desktop 3-D GIS geographic visualization output.

map) have over successful acquisition of spatial knowledge and collaborative spatial visualization activities? Do these factors affect the meaning of geographic visualizations and the information they present? Several parallel projects have begun to explore such questions, comparing AR versus desktop 3-D spatial visualization use by individuals (Hedley, 2001a, 2001b) and groups. These are intended to complement the project reported in this paper, as we develop a suite of AR interfaces and validation experiments, and conduct usability testing in applied settings.

From observing how GIS researchers use existing tools, is it clear that any interface developed to support them must satisfy the following requirements:

- the ability to represent multi-attribute data
- support for multiple views into the same data set, such as ego- and exocentric viewpoints
- the use of real tools and natural interaction metaphors
- support for a shared workspace

It would be difficult for any single interface technology to satisfy all these requirements. For example, AR interfaces can provide an exocentric view of a data set and public and private information views, but it is difficult for them to provide the egocentric views that are

common in an immersive VR interface. For this reason, we believe that combining AR technology with other technologies in a hybrid user interface can more naturally support the GIS process. In the remainder of this paper, we describe two hybrid user interfaces for GIS visualization: AR PRISM and its successor, the GI2VIZ interface.

### 3 The AR PRISM Interface

The AR PRISM interface is based on the underlying idea that multiple users gathered around a real map should be able to see and manipulate three-dimensional virtual geographic information superimposed on the map. This manipulation should be based on physical markers and tangible interaction techniques. To achieve this, AR PRISM uses a combination of three different interface technologies:

- computer vision-based AR tracking
- natural gesture-based interaction
- transitional interfaces

### 3.1 Computer Vision-Based AR Tracking

One of the most difficult parts of developing an AR application is precisely calculating the user's viewpoint in real time so that the virtual images are exactly aligned with real-world objects. In our work, we use ARToolKit, a free software library that uses computer vision techniques to calculate camera position and orientation relative to physical markers (Kato, 2001). Once the real camera position is known, the virtual camera viewpoint can be set to the same position and computer graphics can be drawn that appear to exactly overlay the video image of the real marker. Rekimoto's CyberCode software (Rekimoto & Ayatsuka, 2000) also uses computer vision techniques to provide a similar capability.

Users typically wear a lightweight head-mounted display (HMD) with a small video camera attached. The HMD and camera are connected to a computer that performs image processing on the video input and composes computer graphics onto the video image for display in the HMD. When users look at a tracking marker, they see virtual images overlaid on the marker in a video see-through augmented reality. This ability to overlay virtual models on real physical objects naturally lends itself to physically based interaction metaphors, such as virtual interactions based on real object proximity or placement.

### 3.2 Natural Gesture Interaction

A second interface technique is wireless natural gesture interaction. This allows users to interact with data sets using natural hand gestures or physical object motion. To accomplish this, we use a projection table and an overhead camera. The surface of the table is made of opaque glass, and beneath the table are several lamps that produce a large amount of infrared (IR) light. Above the table is a fixed camera with a lens filter that filters out everything but IR light (figure 3).

When users stand around the table and place their hands or other objects on the table surface, these create dark patches in the IR image seen by the overhead camera (figure 4). This image is passed to a computer (con-




Figure 3. Rear-projection display table.

nected to the camera) that uses background subtraction and edge detection routines to find the outlines of hands and objects on the table (figure 5 and 6). The shape and position of the hands and other objects' position can then be used to change application parameters or as input channels. This software can track as many hands and objects as there are on the table, enabling several users to gather about the table and interact with the data set. For a complete description of how this software works, please see the Hi-Space Web site (May 2001).

### 3.3 Transitional Interfaces

The AR PRISM interface uses techniques developed in the MagicBook project (Billinghurst, Kato, & Poupyrev, 2001) to allow users to smoothly move between physical reality, augmented reality, and immersive virtual reality. Given that different degrees of immersion may be useful for different visualization tasks, we believe that a transitional interface has the flexibility to meet the changing needs of users throughout the GIS collaborative process.



Figure 4. Input IR image.



Figure 5. Subtracted image.




Figure 6. Edge detection.




Figure 7. Avatar in an exocentric AR view.

In the AR PRISM interface, if several users are gathered about the real map and are looking at the same tracking marker, they will see an AR image from their own exocentric viewpoint. When a user finds an AR model interesting, he or she can fly into the virtual world and experience it immersively, freely moving through the world and looking about it from an egocentric viewpoint. Users immersed in the virtual world will see each other represented as virtual avatars. More interestingly, if one user is viewing the computer graphics in an AR mode while another is in immersive VR, the AR user will see an exocentric view of the avatar of the immersive user in the virtual world (figure 7).

Thus, the AR PRISM interface supports collaboration on three levels:

- *Level 1 as a physical object:* Users can view and interact with the physical objects in a face-to-face setting.
- *Level 2 as an AR object:* Users with AR displays can see virtual objects appearing on the physical objects.
- *Level 3 as an immersive virtual space:* Users can fly into the virtual space together and see each other represented as virtual avatars in the VR scene.

### 3.4 The Hybrid Interface

In the AR PRISM hybrid interface, users stand around the projection table with a real map spread out





**Figure 8.** Using the AR PRISM interface.

on top of it and several physical tracking markers (figure 8). They each wear or hold a lightweight HMD with a small video camera attached. Each camera is connected to an SGI O2 computer, and the video output from the O2 is fed back into the headset, enabling the user to experience video see-through AR. A separate overhead camera connected to a Pentium PC with a Matrox video capture board is used to track hand and object position on the table. The SGI O2's and Pentium computers are networked to exchange tracking information. The network bandwidth between the machines is very small, and so the number of users could scale to as many as dozens around the same table.

There are two different display options. Users can either wear a pair of Olympus EyeTrek displays (figure 8), or they can hold a handle with a Sony Glasstron mounted on it (figure 9). Both displays are full color with NTSC resolution and are biocular rather than stereoscopic. The EyeTrek has the advantage that the user's hands are free while wearing it, whereas the Glasstron is easy to remove from the face and pass among several users. The Glasstron configuration is also designed to support transitions between the AR and VR scenes. To do this, the Glasstron handle has an InterTrax 30 inertial tracker mounted on it and a pressure pad and switch. The inertial tracker provides head orientation information when the user is immersed in the



**Figure 9.** The Sony Glasstron mount

virtual scene. Flicking the switch causes the user to fly in or out of the virtual scene, and pressing on the pressure pad moves them through the immersive world in the direction they're looking.

### 3.5 AR PRISM Software

The AR PRISM software is composed of two different components: a hand/object tracker and several AR graphical clients (one for each user). The hand/object tracker is run as a simple server application, broadcasting output to any of the AR graphical clients connected to it. The output is simply in the form: `<EventType> <xPosition> <yPosition>`, where `EventType` is an integer specifying if the position information is related to a hand or object position,

and `xPosition` and `yPosition` are the normalized coordinates of the tip of the hand or marker centerpoint on the table. So, if a user placed a marker object in the center of the table, the following message would be sent to all the connected clients: “2 0.5 0.5”

Each of the SGI O2s runs the AR PRISM graphical client. This software integrates the ARToolKit tracking library with a VRML97 parser to enable any VRML content to be viewed on a physical marker. The VRML97 parser is based on the open source lib-VRML97 library with modifications made to allow ARToolKit to set camera viewpoint information.

The actual content shown in the displays depends on the messages sent from the overhead camera and PC server. For example, if a physical marker is placed in the middle of the map then the  $(x,y)$  coordinates of the marker (in this case,  $(0.5,0.5)$ ) are used in a look-up table to load the VRML model corresponding to the 3-D virtual terrain model for that portion of the real map. In this way, multiple VRML models can be associated with each physical marker.

### 3.6 User Interactions

The AR PRISM interface uses a “sensordata” metaphor for user interactions. That is, a physical marker (the terrain marker) represents the portion of the data set being viewed, and other markers represent “sensors” that modify the way this data is represented. For every grid location in the real map, three-dimensional topographic information was available, and for some locations there was also soil, hydrology, and well information.

As the user moves his or her hand across the surface of the real map, a graphic overlay of longitude and latitude information of the hand location is shown in the user’s headset display. Icons representing data available at that point are also visible. If the user wants to view some of the geological data set, he or she places a terrain marker on the map (figure 10). Once the user has removed his/her hand, the hand/object tracking system determines that a marker has been placed on the map and sends the coordinates to the AR client software. The AR software then recognizes the marker and looks



Figure 10. User placing a terrain marker.





Figure 11. VRML model of terrain.

up its map location to retrieve the correct VRML model and render it over the marker card (figure 11).

If the user wants to see a new portion of the virtual terrain, the user simply moves the marker to a new map location, and the system recognizes the new location and swaps VRML models. (See figure 12.) The AR scene and real map are aligned so that features seen on the map can be seen in the three-dimensional virtual model overlaying the same map location. So, in figure 12, the coastline in the virtual model is aligned with the coastline on the real map. Once the terrain model has



**Figure 12.** Updated VRML model.

been loaded, there is no need to keep the marker on the map. The user can look at the visualization from any angle and at any distance (figure 13). The card can be picked up so that the user can hold it in his/her hand, rotate it, and see more detail. Each of the users about the table can see the terrain model from their own viewpoint, and the model can even be picked up and passed around the users.

Many of the map grid locations have content other than just topography. If the icons in the display indicates that hydrology, soil, or well information is available, the user is able to see this by using the appropriate sensor card. Each type of data sensor has one physical marker, and placing this marker next to the terrain marker causes the VRML model to change to the appropriate data set. For example, placing the soil sensor next to the terrain marker changes the texture on the terrain model to show soil features (figure 14a). This swapping is based on the physical proximity of the data and sensor cards.

Whereas viewing the soil information changes the model texture, entirely new models can also be loaded. For example, when the user places the well sensor card beside the terrain model, the model becomes a volumetric display of subsurface features including vertical wells (figure 14b).

Once a terrain model has been loaded, the interactions are all performed using the AR client software, so the user does not have to place both cards on the map

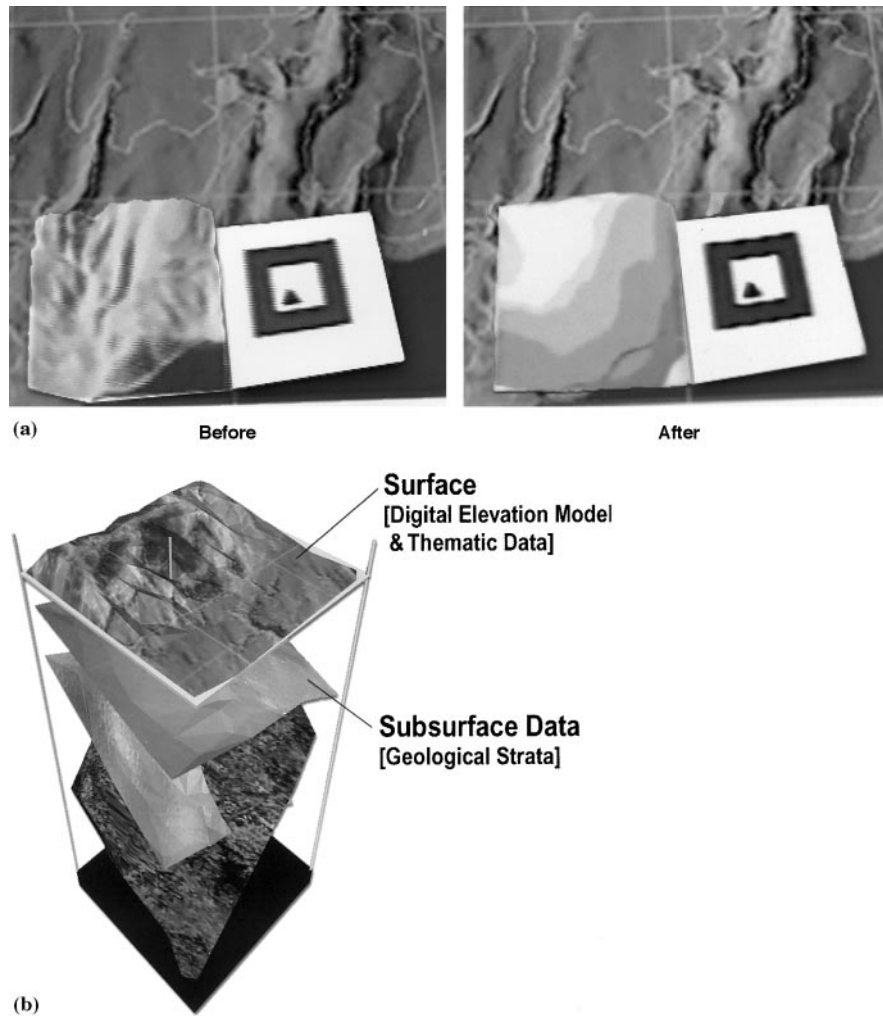


**Figure 13.** User getting close to terrain model.

for this data swapping to take place. The views into the data set are also independent, so, if one user wants to see the hydrology information of a particular region, he or she can pick up the terrain card, place the hydrology sensor card next to it, and see the hydrology model loaded. However, if another user does not see those two cards together, his or her AR client will not load the hydrology data set, and the original terrain model will continue to be displayed.

Users can also fly into a scene and experience it immersively. Once they have selected a region of the map and loaded the AR scene, if they flip the switch on the Glasstron handle, the camera viewpoint smoothly animates from being an AR exocentric view to an immersive VR egocentric view (figure 15). The ARToolKit tracking is now turned off, and the inertial tracker is used to allow the user to freely look around in any direction. Pushing the pressure pad flies them in the direction they're looking, removing the need for them to change their physical position in the real world. As described previously, users who are not immersed in the VR scene can see a miniature virtual avatar of the VR user, or they can fly into the virtual scene themselves to see each other represented as life-sized avatars.





**Figure 14.** (a) Loading soil data.


**Figure 14.** (b) Subsurface geological model.

### 3.7 User Experiences

We are in the process of conducting formal user studies, although almost one hundred people have tried the AR PRISM software informally. These users have generally found the software very easy to use. Associating a virtual data set with a physical object means that this data is as easy to manipulate as the real object. They also like the ability to see other users in the real world at the same time as seeing the AR model, and felt that this supports natural face-to-face collaboration. Generally, users preferred the Sony Glasstron display mount to the

hands-free Olympic EyeTrek configuration. This was not only because it was easy to share views with others, but also because it reduced their feeling of separation from the real world and reduced incidences of nausea.

However, there is room for improvement. For instance, almost all the users that fly into immersive VR scenes ask how they can move backwards in the scene. When we developed the interface, we assumed that we should mimic the real world (in which people rarely walk backwards or in the direction opposite to which they are facing). However, in most first-person com-




**Figure 15.** Transition from AR to VR.

puter games and VR environments, this is common, and users often ask for it. Users also needed some reminding not to cover tracking pattern and sometimes had to be shown how to orient the physical markers for best tracking results. This implies that we may want to spend more time on developing a physical marker with a form factor that makes it difficult to obscure the tracking pattern.

### 3.8 AR PRISM Content

Although this interface could be used for interacting with any geological data sets, the content we have used is from the Big Beef Creek area in the Puget Sound region of Washington, USA (figure 16). The Big Beef Creek field station is run by the University of Washington's School of Fisheries and is used to investigate salmon propagation (among other things). The School of Fisheries has detailed topological data of the region as well as data about soil types, hydrology, and environmental features.

To use this content, a 2-D base map of the area was divided into 35 grid cells (seven rows and five columns) (figure 16), and VRML models were generated for each of the cells. These models were generated from USGS digital elevation model (DEM) data, using the IRIS Explorer data visualization application.



**Figure 16.** Location of Big Beef Creek relative to Puget Sound (looking north).

## 4 Recent Work: The GI2VIZ Interface

Although the AR PRISM interface provides an intuitive way to collaboratively view geospatial data, it has a number of limitations. Firstly, there is limited support for interaction with the data set; users cannot annotate the data or place their own virtual models into the AR scene. The application also has no support for viewing 2-D imagery. In many GIS applications, 2-D images such as satellite photos can be as important as the 3-D terrain models. A drawback with the ARTool-Kit tracking technique is that it uses a single tracking



**Figure 17.** *GI2VIZ interface.*

marker, meaning that virtual models can disappear if this marker is even partially obscured. Finally, the use of a real map prevents the zooming and panning techniques that are possible with a projected map image.

To address some of these limitations and to extend the AR PRISM interface, we are working on its successor, the GI2VIZ interface. The GI2VIZ project is a hybrid interface that has many of the same characteristics of the AR PRISM interface. In addition, it supports the use of projected 2-D imagery, allows annotation of both 2-D and 3-D content, allows manipulation of the virtual models and addition of content to the virtual scenes, and supports a new, more robust, multiple-marker AR tracking technique. In this section, we briefly describe these enhancements.

Like AR PRISM, the GI2VIZ interface is based around a projection table. However, in this case, 2-D map imagery is projected onto the surface of the table (figure 17). The glass table also has a MIMIO input device mounted on it which uses ultrasonic and infrared signals to track pen input on a flat surface. The MIMIO pen can be used for normal mouse input in any Windows PC application. The application run on the projection table supports loading and viewing of 2-D map images. It also has a zoomable user interface based on the Pad++ work (Bederson et al., 1996). This allows the user to pan around the map images and zoom into any arbitrary resolution using the MIMIO pen. Some of



**Figure 18.** *Virtual model overlay.*

the 2-D images also have square ARToolKit tracking markers embedded in them (figure 18). When the user looks at these marks through the headset display, he or she sees virtual 3-D terrain imagery overlaid on the projected 2-D image. In this way, the GI2VIZ interface supports simultaneous viewing of 2-D and 3-D images.

There is also a connection between the 2-D map images and the 3-D content. The user can use the MIMIO pen to annotate the 2-D projected imagery. As they draw lines on the 2-D images, virtual lines appear on the associated 3-D VRML models. These drop lines follow the contours of the model and provide an easy way for users to mark regions of significance in the virtual data set. (See figure 19.)

The ARToolKit library has also been enhanced to support additional tangible interaction techniques. The basic ARToolKit library was extended to support camera pose determination from sets of tracking markers. This increases the robustness of the AR tracking and means that one or more of the markers can be occluded and the camera pose still found. Figure 18 shows a virtual model overlaying a set of projected marks. This enhancement also enables the use of a wider range of tangible interaction techniques. One of the first we have been exploring is a real paddle that can be used to manipulate virtual objects. The paddle is marked with a tracking marker that enables it to be tracked relative to the set of markers that are overlaid with a virtual terrain model. The position and orientation of the paddle can



**Figure 19.** VRML terrain model with annotations (black markings).

be found, so a number of tangible interaction techniques can be used based on paddle tilt, motion, and proximity to virtual objects. Users can pick up virtual models by placing the paddle over the models and then dropping them onto the virtual terrain data set. (See figure 20.) Once placed on the terrain, simple collision detection is used to allow the user to push the models around with the paddle. In this way, users can quickly add objects to the geospatial AR scenes. This paddle interaction technique is based on earlier work presented at ISAR 2000 (Kato, Billinghamurst, Poupyrev, Imamoto, & Tachibana, 2000) but modified to work within the GI2VIZ hybrid interface.

## 5 Related Work

A number of researchers have developed interfaces that we can compare our work to, although none have presented a hybrid interface with the same combination of technologies.

There are many examples of augmented tabletop interfaces. Some, such as the perceptive workbench work



**Figure 20.** Paddle interaction with the AR scene.

of Leibe et al. (2000), use infrared sources to aid in hand tracking, whereas others, such as the pioneering work of Krueger's VIDEODESK (Krueger, 1983), use visible-light cameras. Projection tables such as the responsive workbench (Kruger et al., 1995) are capable of showing stereoscopic virtual images and typically use magnetically tracked input devices to interact with these images. Other systems, such as Ullmer's metaDESK (Ullmer & Ishii, 1997) support the use of real objects as interaction devices. Our work is related to these efforts, but different in a number of key ways:

- The use of AR displays allows it to support an arbitrary number of face-to-face collaborators.
- In the AR PRISM interface, the table surface is covered by a physical map and is used just for interaction, not projected display.
- The physical props are not only interaction devices but are also used as containers for virtual data.

The use of physical objects as input devices builds on the work of Hiroshi Ishii's Tangible Media group (<http://tangible.media.mit.edu>). The many projects developed by this group are based around the notion of the tangible user interface (TUI) in which real-world objects are used as computer input and output devices (Ishii & Ullmer, 1997). As Fitzmaurice points out, tangible interfaces are extremely intuitive to use because



physical object manipulations are mapped one-to-one to virtual object operations, and they follow a space-multi-plexed input design (Fitzmaurice & Buxton, 1997).

Although intuitive to use, information display with tangible interfaces can be a challenge. For example, in the Triangles work (Gorbet, Orth, & Ishii, 1998), physical triangles are assembled to tell stories, but the visual representations of the stories are shown on a separate monitor distinct from the physical interface. Presentation and manipulation of 3-D virtual objects on projection surfaces is difficult (Fjeld, Voorhorst, Bichsel, Lauche, & Rauterberg, 1999), particularly when trying to support multiple users each with independent viewpoints. Most importantly, because the information display is limited to a projection surface, users are not able to pick virtual images off the surface and manipulate them in 3-D space as they would a real object.

The AR PRISM interface overcomes these problems by combining the enhanced display possibilities of augmented reality with the intuitive manipulation of tangible user interfaces. We call this combination *tangible augmented reality*. Earlier AR interfaces used more-traditional input devices, such as a handheld mouse or tablet (Rekimoto & Ayatsuka, 1998; Hollerer & Feiner, 1999) or magnetic trackers to allow people to view and edit AR content (Kiyokawa, Takemura, & Yokoya, 1999). In many cases, these interfaces used existing 2-D and 3-D interface metaphors from desktop or immersive virtual environments. As can be seen from our work, tangible AR interfaces are anchored in the real world, and AR content manipulation is based on interaction between real-world objects. This means that tangible AR interfaces are as intuitive to use as it is to manipulate real objects.

## 6 Conclusions and Future Work

In this paper, we have described how we have combined different interface technologies to create two hybrid interfaces for geospatial data visualization. Combining AR and VR techniques, we have created an interface in which users can seamlessly move between the real and immersive virtual worlds. Combining computer

vision hand and object tracking with AR techniques, we have produced an intuitive means for loading data sets into the AR environment. Adding support for pen-based input and 2-D image viewing, we have enabled the user to make annotations on the 3-D virtual models.

In the future, we will explore more interesting interaction techniques that go beyond model loading. For example, if we have an AR model of a virtual column of data appearing above a card, then we should be able to move a second card up and down this model to view different aspects of this data set using a variation on the “magic lens” technique (Bier, Stone, Pier, Buxton, & DeRose, 1993). Other interaction techniques to explore include being able to show layers of information (for example, weather on top of terrain), and the addition of speech input for true multimodal commands.

Finally, we will develop techniques for feeding live data into the interface. So, for example, as weather patterns change, the weather texture maps will be updated in near real time. This will also enable this interface to be used for emergency management and other command-and-control type applications wherein it is critical to know where real vehicles and people are in remote environments at all times.

## Acknowledgments

The GI2VIZ portion of this work was supported and monitored by the Advanced Research and Development Activity (ARDA) Center, and by the National Imagery and Mapping Agency (NIMA) under contract number NMA 201-01-C-0006. The AR PRISM portion of this work was partially supported by National Science Foundation Grant REC-9873620, *Learning in Virtual Environments*. We would like to thank Mark Stoermer for developing some of the models used in the AR PRISM interface, Randy Stiles for the content used in the GI2VIZ interface, and Richard Home for the use of sample models developed using 3DEM terrain visualization software.

## References

- Bederson, B., Hollan, J., Perlin, K., Meyer, J., Bacon, D., & Furnas, G. (1996) Pad++: A zoomable graphical sketchpad

- for exploring alternate interface physics. *Journal of Visual Languages and Computing*, 7(3), 3–31.
- Bier, E., Stone, M., Pier, K., Buxton, W., & DeRose, T. (1993). Toolglass and magic lenses: The see through interface. *Proceedings of SIGGRAPH 93*, 73–80.
- Billinghurst, M., Weghorst, S., & Furness, T. (1998). Shared space: An augmented reality approach for computer supported cooperative work. *Virtual Reality*, 3(1), 25–36.
- Billinghurst, M., Bowskill, J., Jessop, M., & Morphett, J. (1998). A wearable spatial conferencing space. *Proceedings of the Second International Symposium on Wearable Computing (ISWC '98)*, 76–83.
- Billinghurst, M., Kato, H., & Poupyrev, I. (2001). The MagicBook: A transitional AR interface. *Computers and Graphics*, 25, 745–753.
- Butz, A., Hollerer, T., Feiner, S., MacIntyre, B., & Beshers, C. (1999). Enveloping users and computers in a collaborative 3D augmented reality. *Proceedings of the 2nd International Workshop on Augmented Reality (IWAR 99)*, 35–44.
- Dede, C., Salzman, M., Loftin, B., & Sprague, D. (1999). Multisensory immersion as a modeling environment for learning complex scientific concepts. In W. Feurzeig and N. Roberts (Eds.), *Computer modeling and simulation in science education* (pp. 282–319). New York: Springer-Verlag.
- Erickson, T. (1993). Artificial realities as data visualization environments: Problems and prospects. In Alan Wexelblat (Ed.), *Virtual reality: applications and explorations* (pp. 3–22). Boston: Academic Publishers.
- Fitzmaurice, G., & Buxton, W. (1997). An empirical evaluation of graspable user interfaces: Towards specialized, space-multiplexed input. *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI'97)*, 43–50.
- Fjeld, M., Voorhorst, F., Bichsel, M., Lauche, K., & Rauterberg, M. H. K. (1999). Exploring brick-based navigation and composition in an augmented reality. *Proceedings of HUC' 99*, 102–116.
- Gorbet, M., Orth, M., & Ishii, H. (1998). Triangles: Tangible interface for manipulation and exploration of digital information topography. *Proceedings of CHI 98*, 49–56.
- Hedley, N. R., Drew, C. H., Arfin, E., & Lee, A. (1999). Hagerstrand revisited: Interactive space-time visualizations of complex spatial data. *Informatica*, 23(4), 155–168.
- Hedley, N., & Campbell, B. (1998). *Distributed Collaboration Project Final Project Report* (Technical Report R-99-3). Seattle: Human Interface Technology Laboratory.
- Hedley, N. R. (2001a). Exploring the dimensions of spatial mental models formed in augmented and virtual environments. Thomas F. Saarinen Award Paper, Association of American Geographers' Annual Meeting (New York City, March 2001). Paper available from author upon request.
- . (2001b). Virtual and augmented reality interfaces: Empirical findings and implications for spatial visualization. *Proceedings of the 20th International Cartographic Conference, Beijing, China. August 6–10, 2001*. International Cartographic Association in conjunction with the Chinese Society of Geodesy, Photogrammetry and Cartography: Vol. 4, 2606–2613.
- Hollerer, T., & Feiner, S. (1999). Exploring MARS: Developing indoor and outdoor user interfaces to a mobile augmented reality system. *Computers and Graphics*, 23(6), 779–785.
- Ishii, H., & Ullmer, B. (1997). Tangible bits: Towards seamless interfaces between people, bits and atoms. *Proceedings of CHI 97*, 234–241.
- Kato, H. (2001). See ARToolKit website ([www.hitl.washington.edu/artoolkit/](http://www.hitl.washington.edu/artoolkit/)).
- Kato, H., Billinghurst, M., Poupyrev, I., Imamoto, K., & Tachibana, K. (2000). Virtual object manipulation on a table-top AR environment. *Proceedings of the International Symposium on Augmented Reality (ISAR 2000)*, 111–119.
- Kiyokawa, K., Takemura, H., & Yokoya, N. (1999). A collaboration supporting technique by integrating a shared virtual reality and a shared augmented reality. *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics (SMC '99)*, VI, 48–53.
- Kraak, M-J. (1994). Interactive modeling environment for three-dimensional maps: Functionality and interface issues. In A. M. MacEachren and D. R. F. Taylor (Eds.), *Visualization in modern cartography*, 269–285. Oxford: Pergamon Press.
- Krueger, M. (1983). *Artificial reality*. Reading, MA: Addison-Wesley.
- Kruger, W., Bohn, C., Frohlich, B., Schuth, H., Strauss, W., & Wesche, G. (1995). The responsive workbench: A virtual work environment. *IEEE Computer*, 28(7), 42–48.
- Leibe, B., Starner, T., Ribarsky, W., Wartell, Z., Krum, D., Weeks, J., Singletary, B., & Hodges, L. (2000). Toward spontaneous interaction with the perceptive workbench. *IEEE Computer Graphics and Applications*, 20(6), 54–65.
- Lin, C., & Loftin, R. B. (1998). Application of virtual reality in the interpretation of geoscience data. *Proceedings of the ACM Symposium on Virtual Reality Software and Technology (VRST 1998)*, 187–194.
- MacEachren, A. M., Edsall, R., Haug, D., Baxter, R., Otto, G.,

- Masters, R., Fuhrmann, S., & Qian, L. (1999). Virtual environments for geographic visualization: Potential and challenges. *Proceedings of the ACM Workshop on New Paradigms in Information Visualization and Manipulation*, 35–40.
- MacEachren, A. M., & Kraak, M.-J., (Eds.). (2001). Research challenges in geovisualization. *Cartography and Geographic Information Science* (special issue on geovisualization in collaboration with the ICA Commission on Visualization and Virtual Environments), 28(1).
- May, R. (2001). The Hi-Space website: <http://www.hitl.washington.edu/hispace.html>
- Nyerges, T. L., Moore, T. J., Montejano, R., & Compton, M. (1998). Interaction coding systems for studying the use of groupware. *Journal of Human-Computer Interaction*, 13(2), 127–165.
- Rekimoto, J., & Ayatsuka, Y. (1998). Augment-able reality: Situated communication through physical and digital spaces. *Proceedings of the International Symposium on Wearable Computers (ISWC'98)*, 68–75.
- . (2000). CyberCode: Designing augmented reality environments with visual tags. *Proceedings of Designing Augmented Reality Environments (DARE 2000)*.
- Slocum, T. A., Blok, C., Jiang, B., Koussoulakou, A., Montello, D. A., Fuhrmann, S., & N. R. Hedley. (2001). Cognitive and usability issues in geovisualization. *Cartography and Geographic Information Science*, 28(1), 61–75.
- Ullmer, B., & Ishii, H., (1997). The metaDESK: Models and prototypes for tangible user interfaces. *Proceedings of the ACM Symposium on User Interface Software and Technology (UIST 97)*, 223–232.