

Enveloping Users and Computers in a Collaborative 3D Augmented Reality

Andreas Butz*, Tobias Höllerer, Steven Feiner, Blair MacIntyre†, Clifford Beshers

Department of Computer Science

Columbia University

New York, NY 10027

butz@cs.uni-sb.de, {htobias,feiner,beshers}@cs.columbia.edu, blair@cc.gatech.edu

Abstract

We present *EMMIE* (Environment Management for Multi-user Information Environments), a prototype experimental user interface to a collaborative augmented environment. Users share a 3D virtual space and manipulate virtual objects that represent information to be discussed. We refer to *EMMIE* as a hybrid user interface because it combines a variety of different technologies and techniques, including virtual elements such as 3D widgets, and physical objects such as tracked displays and input devices. See-through head-worn displays overlay the virtual environment on the physical environment, visualizing the pervasive “virtual ether” within which all interaction occurs. Our prototype includes additional 2D and 3D displays, ranging from palm-sized to wall-sized, allowing the most appropriate one to be used for any task. Objects can be moved among displays (including across dimensionalities) through drag & drop.

In analogy to 2D window managers, we describe a prototype implementation of a shared 3D environment manager that is distributed across displays, machines, and operating systems. We also discuss two methods we are exploring for handling information privacy in such an environment.

1. Introduction

In the early 1990s, Weiser coined the term *ubiquitous computing* to describe a world in which large numbers of computing devices were woven into the fabric of our daily life [40]. These devices include not only displays (ranging from palm-sized to wall-sized), but also an assortment of embedded computers that add computational behavior to physical objects and places that would not otherwise have them (such as doors or desks). Because these computers can

be networked together, they add a (mostly) invisible virtual layer to the physical reality surrounding us.

In contrast to the proliferation of computing devices in such an environment, *augmented reality* (AR) [4] typically focuses on the use of personal displays (such as see-through head-worn displays) to enhance a user’s senses by overlaying a directly perceptible virtual layer on the physical world. Because information is presented on a small number of displays, computation usually takes place on the few relatively powerful machines driving those displays. This contrasts with the ubiquitous computing paradigm, which is typically widely distributed and decentralized.

AR interfaces can enhance a ubiquitous computing environment by allowing certain parts of its hidden virtual layer to be visualized, as well as displaying personal information in a way that guarantees it remains private and customizable for each user. However, one important drawback of pure AR interfaces is that their interface elements are drawn from purely virtual environments, such as 3D widgets and 3D interaction metaphors, and thus remain within the virtual realm. Such interfaces can be hard to deal with, partially because the affordances offered by more concrete interfaces are absent. We believe that AR systems can profit from the use of physical objects and the interaction techniques they afford [24]. Therefore, by integrating elements of ubiquitous computing with AR, we can leverage the ubiquitous displays to allow users to manipulate information in a concrete way when appropriate.

In this paper, we present the design of an experimental hybrid user interface for collaboration that combines AR, conventional 2D GUIs, and elements of ubiquitous computing. We use the term *hybrid user interface* to refer to the synergistic use of a combination of user interface technologies [14]. In the hybrid user interface we describe here, see-through head-worn displays are used in conjunction with other displays and devices, ranging from hand-held to desktop to wall-sized. Our goal is to create an environment in which information displayed on the 3D AR and conventional 2D displays complements each other, and can be eas-

*Current affiliation: Fachbereich Informatik, Universität des Saarlandes, 66041 Saarbrücken, Germany

†Current affiliation: College of Computing, Georgia Institute of Technology, Atlanta, GA 30332

ily moved between the various displays.

1.1. Design Approach

Our prototype uses AR as an encompassing 3D multimedia “ether” that envelops all users, displays, and devices. This not only allows interaction and display to take place in a common, shared space, but also visualizes interactions among the physical devices that populate the space. Since users may wish some virtual material to be private, we have explored ways of allowing them to determine and inquire whether selected items can be viewed by others.

We address the often conflicting needs that collaborating users have to focus on each other and on the computer-based tasks they are performing, by allowing both to occupy the same space. Since users increasingly enter meetings carrying their own laptops or personal digital assistants (PDAs), and many tasks benefit from or require information that may reside on these personal machines, we make them an intrinsic part of the interaction. Because different tasks and interaction styles benefit from the use of different displays and devices, we have attempted to create a unified architecture that supports a wide range of hardware. And, we have tried to do this within a dynamic collaborative structure in which users and their computers can freely join and leave the group.

The system that we are developing is a working prototype. Like much research that uses experimental devices, our goal is not to suggest a current practical alternative to existing mature technologies, but rather to explore now, using commercial hardware pushed to or past its limits, directions that will become feasible later when the needed technologies reach maturity. Thus, the see-through head-worn displays we use are relatively low-resolution, heavy, odd-looking, and insufficiently transmissive; the 3D trackers suffer from limited range and noise-induced jitter; and adding another computer to the network requires the familiar tedium of setting parameters and connecting cables. However, we remain confident that these current impediments will be overcome by ongoing research and development efforts that address them; for example, see-through head-worn displays that look and feel much like conventional eyeglasses [35], accurate wide-range motion tracking [18, 17], and standards for short-range wireless ad hoc data and voice networking [11]. Therefore, our testbed provides a way to explore future user interaction paradigms that will become increasingly relevant as new hardware breaks down these technological barriers.

1.2. Environment Management

In analogy to the window manager of a 2D GUI, we use the term *environment manager* to describe a component that organizes and manages 2D and 3D information in a heterogeneous world of many virtual objects, many displays, and



Figure 1. A meeting scenario using EMMIE. Two users wear tracked see-through head-worn displays, one of whom has also brought in his own laptop. The laptop and a stylus-based display propped up on the desk are also tracked, using sensors attached to their backs. All users can see a wall-sized projection display. The triangular source for one tracker is mounted at the left of the table; additional ceiling-mounted trackers are not visible here.

many users. Traditional window managers handle a relatively small number of windows on a single display (possibly spread across multiple screens) for a single user. In contrast, an environment manager must address the more complex task of managing a global 3D space with a combination of virtual and real objects, and a heterogeneous set of computers, displays and devices, shared by multiple interacting users.

We have named our prototype collaborative hybrid user interface, EMMIE (Environment Management for Multi-user Information Environments). EMMIE’s rudimentary environment manager supports a dynamically changing mix of displays and devices, allows information to be passed between 2D and 3D devices, and provides mechanisms for handling privacy in multi-user environments and services such as searching.

2. A Collaboration Scenario

We developed EMMIE to experiment with supporting collaboration among participants in a meeting. The participants share a 3D physical space, for example by sitting around a table, as shown in Figure 1. This shared space contains computers of different kinds, such as workstations and PCs installed in the meeting room, as well as laptops and




Figure 2. An EMMIE user manipulates a 3D model with an optically tracked hand-held pointer. Other virtual objects shown include simple iconic 3D representations of applications (e.g., a “movie projector” that can play videos) and data objects (e.g., “slides” that represent still images).

PDAs the participants have brought with them. These computers provide displays ranging from wall-sized to palm-sized, and various interaction devices, such as keyboards, mice, touch pads, and pens. Each of the workstations, PCs, laptops, and PDAs runs its own unmodified operating system and 2D GUI.

In addition to the physical space, participants also share a 3D virtual space that is overlaid on the physical space by means of AR technology, in our case tracked, see-through, head-worn, stereo displays. As shown in Figure 2¹, the virtual space contains graphical objects that visually appear to be located in the physical space.

The mapping between the 3D physical and virtual spaces is achieved by tracking relevant physical objects, such as computers, displays, input devices, and participants, using a variety of commercial tracking techniques (infrared, ultrasonic, and magnetic). The 3D position and orientation of these physical objects is used to control the behavior of virtual objects in both the 3D and 2D environments.

In EMMIE, most of the objects in the 3D virtual space are 3D icons. They represent information, such as text,

graphics, sound or animation, much like the icons in a conventional 2D GUI. For example, Figure 2 includes simple iconic 3D representations of applications (e.g., a “movie projector” that can play videos) and data objects (e.g., “slides” that represent still images). In a straightforward adaptation of 2D GUIs, dragging data objects to application objects allows them to be processed. Other kinds of virtual objects include 3D widgets, such as menus or sliders and 3D models (e.g., the model of our lab that the user in Figure 2 is holding in his hand).

3. Related Work

Our research relates to and incorporates elements from work in different areas: AR, virtual worlds, ubiquitous computing, and CSCW.

Our notion of a hybrid user interface is closely related to Rekimoto’s explorations of *multi-computer direct manipulation interfaces* [29, 30, 32]. Like him, we are interested in user interfaces that make it easier to work in a heterogeneous computing environment employing different devices and displays. We go beyond the scenario that Rekimoto describes in that our users can share a global 3D AR space with environment management facilities that support privacy through the use of see-through head-worn displays.

i-LAND [20] is an integrated cooperative work environment with specifically designed *roomware* components (electronically enhanced walls, tables, and chairs) that can share digital information via a physical transportation mechanism using passive objects similar to the mediaBlocks Ullmer et al. propose [38]. EMMIE, on the other hand, provides information management facilities in a global AR space, linking different devices the user is already familiar to (their PDAs, laptops, or workstations) into the global space and to each other, supplying virtual intermediate representations for information exchange.

There is current research at Xerox PARC that focuses on augmenting the physical world seamlessly and invisibly with electronic tags to connect physical objects with the computing environment, essentially forming a “calm” augmented environment [15, 39]. As increasing numbers of physical objects are linked to the world of computation, automated management will become increasingly important. We believe that these systems could benefit from a customized visualization of the electronic layer and from an environment management component.

UNC’s “Office of the Future” [28] provides a vision of how today’s low-resolution AR tracking and display technologies, such as those used in EMMIE, could ultimately be replaced with a combined projection and tracking system to better support a multi-user collaborative environment. The PIT project at UNC [27] presents a two-person, two-screen, stereo display workspace for collaborative study of a 3D

¹All images in this paper that show overlaid graphics (Figures 2–6) were shot directly through an optical, see-through, head-worn display mounted on a fiberglass dummy head. The head’s right eye socket contains a miniature NTSC video camera. This makes it possible to produce images that provide a better approximation to what a user actually sees through the display than do screen dumps. Fuzziness and artifacts in the images are caused by the low resolution of the camera and the lower resolution of the head-worn display, compounded by the interlaced recording and digitization process.

model. PIT shares some overall goals with ours (shared 3D graphics space and access to common devices). In contrast to EMMIE, it is currently targeted to a specific two-person collaboration task (protein fitting); uses a fixed set of displays, each of which has a specific purpose; and does not support general information exchange mechanisms among the displays.

Open Shared Workspace [23] is based on the premise that continuity with existing individual work environments is a key issue in CSCW. Users of our environment also bring in their own tools, such as laptop computers, and can work with the desktop environments with which they are familiar. However, we differ substantially from this and other CSCW work in that instead of relying on video conferencing (or, for that matter, virtual 3D and multimedia worlds [9, 10, 34]), we view a 3D AR environment as an embracing shared virtual space, incorporating, instead of replacing, existing UI technologies. With EMMIE’s AR environment we are trying to achieve *seamlessness* [22] between different computer platforms, display devices of different sizes and dimensionalities, and among different (local or remote) users.

Researchers at the University of Washington [6] and at the Technische Universität Wien [37] have developed AR Interfaces for CSCW that use see-through head-worn displays. While this work shows the potential value of AR for collaboration, we go beyond the pure deployment of AR for visualizing 3D objects or representing teleconferencing avatars to include environment management.

Since Fitzmaurice’s pioneering work on the *Chameleon* tracked hand-held display [16], several researchers have employed similar displays as “see-through” lenses to overlay computer generated imagery on the real world [3, 31]. We use this technique as one of many valuable tools for collaboration.

Finally, EMMIE builds on our own previous work combining 2D and 3D information displays, in which we embedded a physical 2D display in a virtual 3D information space [14], overlaid conventional 2D windows on the 3D world using a see-through head-worn display [12], and developed a wearable outdoor hybrid user interface that combined a tracked see-through head-worn display with a hand-held pen-based computer [13].

4. EMMIE’s Hybrid User Interface

4.1. Interaction with virtual objects

In EMMIE, virtual objects are manipulated with 3D pointing devices that combine a tracker target and two buttons to control a 3D arrow. We use both the hand-held version shown in Figure 2, and one in the form of a ring worn on the index finger, which allows thumb access to the two buttons. An object is highlighted when the projection of

the arrow’s tip intersects the object’s projection in the view-plane (of the user’s dominant eye, in the case of our head-worn stereo displays). A user can pick up a highlighted object by pressing the first button, causing the arrow to turn into a hand. The object can be moved until the button is released, which drops the object at the pointing device’s current location. This variation of the techniques discussed in [7] allows easy access to remote objects.

Certain virtual objects represent applications or tools embedded in the virtual space, such as image viewers or sound players. Dropping a virtual object of the appropriate type onto a tool opens the object (e.g., plays a sound file in the head-worn display’s earphones or displays an image on the virtual projection screen of an image viewer). Pressing the second button in empty space creates a pie menu [21] around the pointer, from which one of a set of tools can be selected and instanced. Pressing the second button over a highlighted data object immediately creates the appropriate tool and opens the object with it.

4.2. Interaction with physical objects

The physical objects that EMMIE manages are the computers present in the physical environment and their input devices and tracked displays. There are two ways of looking at these computers within the EMMIE framework. On one hand, they can be seen as self-contained systems with their own operating system, user interface and software. For example, a conventional laptop can be a perfectly adequate tool for displaying and manipulating text and it can be used this way within EMMIE.

On the other hand, we can look at the same computers as the sum of the interaction devices and displays they provide: keyboard, mouse, pen, screen, and speakers. For example, in addition to their normal use for displaying data, tracked displays facilitate an additional kind of interaction, since their position and orientation can influence what they display. This additional mode is used for some of the hybrid interaction techniques we have developed.

4.3. Hybrid interaction

By *hybrid interaction*, we mean those forms of interaction that cut across different devices, modalities, and dimensionalities [14, 33, 29, 30]. For example, to use a physical wall-sized display to display an object in the 3D virtual space, we have to provide a way to move data back and forth between the 2D desktop of the wall sized display’s workstation and the 3D virtual space surrounding us.

In EMMIE, this transition between spaces is done by simple drag & drop mechanisms. The desktop of each workstation known to EMMIE provides a special icon representing the virtual space. By dragging any regular file onto this icon, a corresponding virtual object (3D icon) is




Figure 3. Drag & drop of virtual objects. A virtual object, displayed as an iconic “slide,” is picked up using a 3D pointing device that controls an iconic hand (left image). The object is dragged to a laptop, whose spherical bounding volume highlights, and is dropped onto the laptop (center image). The object then appears on the laptop’s screen (right image).




Figure 4. The same display tablet can serve as a physical magic lens and magic mirror. In the left image, a user without a head-worn display is looking through the magic lens at a 3D CAD object. In the right image, a user with a head-worn display has just dragged a virtual image slide in front of the magic mirror for further inspection.

created in the virtual space above the workstation’s display. This 3D virtual object can now be manipulated with EMMIE’s tools. It can be shared with another user by handing it over, or it can be dropped onto any workstation managed by EMMIE (see Figure 3), which makes the corresponding data available on the desktop of that workstation and starts up the application associated with its data type.

The effect of these mechanisms is similar to the *pick-and-drop* technique presented in [29, 30] with an important difference: there is a visible and useful representation for the data in the virtual environment while it is being moved between the physical machines. Augmented surfaces [32], developed at the same time as EMMIE, provide such a visible representation on a shared projected wall and desk. However, the presence of this representation raises a variety of privacy issues, which we discuss later, that cannot easily be addressed with publicly viewable surfaces.

Another form of hybrid interaction is the use of a tracked

display (in 3D physical space) for displaying virtual objects (in the overlaid 3D virtual space). Borrowing the terminology of [5], we have used a tracked flat-panel display to implement both a physical magic lens (inspired by [16]) and a physical magic mirror, which show the 3D virtual objects that can be seen through or reflected in the display, respectively (see Figure 4). The physical magic lens and magic mirror open a portal from the real world to the virtual world for those EMMIE users who are not wearing head-worn displays and who otherwise cannot see the 3D virtual objects. The lens and mirror also provide additional interaction techniques for all users; for example, allowing otherwise invisible properties of an object, such as its privacy status, to be inspected and modified, as described in the following section.

Note that the tracked flat-panel display is embedded within the visual field of a user wearing a see-through head-worn display. Based on the flat panel’s size and pixel count, and its position relative to the user and virtual objects seen using it, the flat panel can provide a selective high-resolution view of any portion of the 3D virtual space.

To experiment with hybrid interaction techniques, we implemented a simple 3D search function for virtual objects. A tracked display (acting simultaneously as a magic mirror or lens) presents the user with a set of sliders and buttons through which a subset of the objects in the environment can be specified by criteria such as their data type or size. A bundle of 3D leader lines in the virtual space connects the tracked display to the objects that meet the specified criteria, as shown in Figure 5. Since the leader lines are virtual objects, they are visible in the see-through head-worn displays as well as in the magic mirror. Readjusting the search criteria causes the set of leader lines to change interactively, implementing a dynamic query facility [2] embedded in the 3D world.




Figure 5. A simple interactive search mechanism creates 3D leader lines emanating from the tracked display to objects satisfying the query.

4.4. Privacy management

Privacy is an important issue whenever data is represented in a shared environment. On small hand-held and desk-top displays, privacy is typically maintained by shielding and orienting these devices as needed, but these techniques can't easily be applied to larger, fixed displays that are mounted for group visibility. Since an EMMIE user may want some data to be private and other data public, we need to provide a way to modify the privacy of a virtual object in the shared ether. Furthermore, we would like an object's privacy status to be either constantly visible or easily accessible to its owner. The challenge in an AR environment is to achieve this without also being visually annoying or outright obstructive of other interactions.

Szalavari, Eckstein, and Gervautz's work on multi-user games in augmented reality [36] supports the display of private information. However, it does not provide the user with explicit control over privacy beyond what is implied by game conventions; for example, the identities of board game pieces are always kept private to their owners until they are used in play. Earlier research by Agrawala et al. [1] on a two-user, time-multiplexed projection display used similar application-specific, static privacy policies. In our previous work [8], we considered the conceptual design of two techniques for allowing users to view and control explicitly the privacy status of objects in a collaborative teleimmersive environment: *privacy lamps* and *vampire mirrors*. Here we present prototype implementations of both within EMMIE and discuss how they satisfy the above criteria.

Privacy lamps (see Figure 6) are cone-shaped virtual



Figure 6. A user changes an object's privacy status by moving a privacy lamp above it.

light sources that emit privacy as colored (currently red) light, that distinguishes it from the ambient lighting conditions. Any objects in the environment that lie substantially within the light cone of a privacy lamp will be marked private. These objects will also be rendered as if lit by the lamp, providing direct visual feedback to the user about their privacy state. Our privacy lamps float, typically facing downward onto the world. The higher the lamp, the larger the area of the light cone that intersects with any plane below it, and hence the more objects that can be made private with one interaction.

Privacy lamps satisfy our design criteria nicely. Both the lamps and their lighting effects are always visible, so users can tell privacy state at a glance. The lamps themselves do not obscure other interactions, because they float above the normal workspace. Changing the lighting attributes of objects adds no clutter to the scene, and, because it mimics a common physical phenomenon, is easy to interpret visually. Finally, the lamps make it easy to find all private objects simply by following their beams.

Vampire mirrors, implemented using a tracked, stylus-sensitive, LCD panel (see Figure 7), act as magic mirrors in the virtual environment, reflecting a user's virtual objects. Public objects are reflected fully, while private objects are either invisible or optionally displayed as a ghost image. By placing a vampire mirror at the back of the workspace, oriented so only that user can see it, a user can review the privacy state of all objects quickly: only public objects will appear bright and full in the mirror. To change an object's privacy state, the user touches its image on the vampire mirror.

As with the privacy lamps, the vampire mirrors give us a means of viewing and modifying the privacy state without



Figure 7. A user changes an object's privacy status on the vampire mirror.

cluttering the scene. Interpreting the mirror is easy, if one considers that it shows the owner what other users can see, making it immediately obvious whether an object is public or not. Because the mirror is placed behind objects, it does not obscure or impede any other interaction with those objects. The LCD panel can also be used as a privacy lens, allowing the user to look through the display to view and modify the privacy status of objects on its other side.

5. Implementation

EMMIE is implemented on top of Coterie, our testbed for exploratory research in augmented environments [25]. Coterie is implemented in Modula-3, and runs on a variety of platforms (including many versions of UNIX, and Windows NT/95). Coterie provides the programmer with an environment in which it is easy to rapidly prototype and experiment with distributed interactive applications. To that end, it supports an object-based distributed shared memory programming model, allowing the programmer to implement distributed applications much as they would implement multi-threaded applications in a single process. Communication is done through shared objects, which may exist at one site and be accessed remotely from all others, or be replicated across multiple sites. Replication is required to support the highly interactive applications we develop, as data that is needed to refresh a display many times per second must be local to the process refreshing the display.

Coterie presents this model to the programmer via both compiled (Modula-3) and interpreted (Repo) languages, and includes libraries for 3D graphics and tracker control. By allowing programmers to prototype distributed programs in an interpreted language, Coterie greatly speeds the

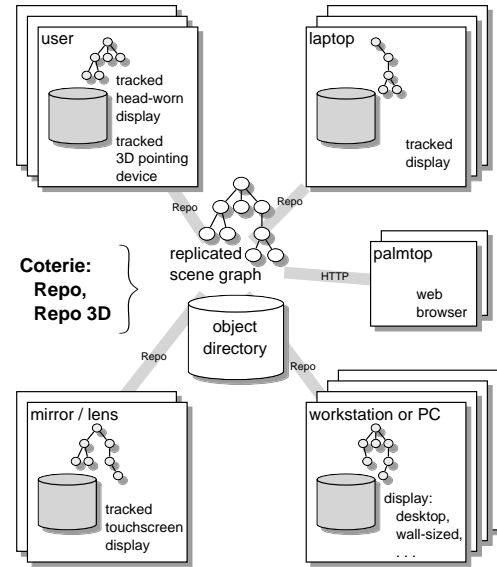


Figure 8. EMMIE architecture. Using Coterie's shared object facilities, an object directory and a scene graph are replicated in all interested processes.

development process.

EMMIE takes significant advantage of two components of Coterie: the Repo interpreted language, and the Repo-3D distributed 3D graphics library [26]. EMMIE is distributed over several machines. Its primary structure is a simple replicated object directory implemented in Repo, similar to the one described conceptually in [25]. Each EMMIE process has a copy of this replicated directory. When any process adds or removes an object, all copies of the replicated directory are automatically updated through Coterie's replicated object infrastructure. Thus, from the standpoint of the application programmer, all processes are equal, with no centralized master process required to coordinate the application. (Transparent to the application programmer, Coterie ensures that changes to all copies of a replicated object are serialized by a sequencer process designated for that object.)

The object directory is replicated in each process to ensure fast access to objects when needed for real-time graphics. The items in the object directory are well-defined object structures that contain all the information needed to manipulate them in any of the processes. One of each object's components is a Repo-3D scene graph that defines the appearance of the object. This scene graph is constructed of a hierarchy of Repo-3D choice groups, each of which allows the processes in EMMIE to choose among a variety of different local appearances (e.g., highlighted or not, in a mirror or on a head-worn display), as well as to control

the global appearance (e.g., publicly visible or private to one process). Because this single well-defined object hierarchy is replicated in all processes that import the object, the clients can be defined in a straightforward manner, and we can experiment with various interaction techniques and object representations simply and cleanly.

Figure 8 shows EMMIE’s architecture. Some users wear Virtual i.O see-through head-worn displays with hear-through earphones. Each head-worn display is connected to a 3D-hardware accelerated PC or workstation that also controls its user’s 3D pointing device. The 3D position of each head-worn display and pointing device is tracked with an Origin Instruments DynaSight infrared LED tracker and each head-worn display’s orientation is tracked with a built-in magnetometer and inclinometer. The magic mirror and lens are implemented on a Wacom PL-300 SVGA LCD panel with pen-input facilities, driven by a PC. The panel is tracked by a Logitech 6DOF ultrasonic tracker.

Other workstations and laptops join the environment by running a background thread that implements EMMIE’s drag & drop functionality, allowing them to be fully integrated in the environment. While we assume that workstation displays stay in fixed positions, laptop displays are tracked with Logitech 6DOF trackers. Hand-held devices, such as 3Com Palm organizers, are included by running a web browser on them and sending them HTML over a PPP link from a Coterie process on another machine. All processes of the distributed system share access to the same database of virtual objects, discussed above.

6. Conclusions and Future Work

We have presented EMMIE, a prototype hybrid user interface to an augmented information environment. EMMIE supports collaborating users by providing and coordinating virtual, physical and hybrid interaction techniques for different parts of the environment. By applying and merging these techniques, each in the situation where it is most appropriate, a hybrid user interface is created whose potential is much greater than that of the sum of its parts. While the relative immaturity of the head-worn display and tracking technologies prevents our current EMMIE prototype from being a viable system for regular use, it has been a useful testbed for experimenting with the design of hybrid user interfaces that will address future improvements in these technologies.

While EMMIE’s architecture was frozen in October 1998 with the departure of the first author, we have started to continue work on it. We began by upgrading to a new set of more accurate trackers (InterSense IS600 Mark II, a hybrid ultrasonic and inertial 6DOF system) and higher-resolution head-worn displays (Sony LDI-D100B SVGA resolution Glasstron). We are now integrating EMMIE with

our outdoor augmented reality system [13, 19] to allow indoor and outdoor users to communicate. Using a 3D model of the environment, indoor users create virtual objects and highlight real objects for outdoor users to see, and maintain histories of outdoor users’ activities; in turn, outdoor users point out interesting objects and events for indoor users to view.

Another area of research we are exploring is the development of dynamic, context-sensitive techniques that can be added to EMMIE’s environment manager. These include analogues of techniques used in window managers to help users position objects when they are created, and keep them organized. For example, when a new object is created, we would like the environment manager to place it in a reasonable initial location, such as an unoccupied location near the focus of the user’s attention. Some window managers, such as X11’s *twm*, provide such assistance by positioning new windows on unoccupied parts of the screen when possible. However, the definition of “unoccupied” is more complicated in an augmented environment than in a desktop interface, because real objects, as well as virtual ones, must be considered. For example, we may not want to place a new virtual object on top of some physical object, such as a telephone, unless both are meaningfully associated.

Other techniques that we are designing address the dynamic nature of augmented environments. Consider a conventional window manager’s single-user desktop: it is generally static, with activity occurring only inside windows whose position and size are rarely modified. In contrast, an environment manager must contend with the dynamic nature of a collaborative augmented environment in which one user has no control over other users and the virtual and physical objects they manipulate. For example, if several virtual objects are suspended above a table, and two users of the system wish to talk to one another, they may not want the objects to block their view of each other. Rather than forcing the users to move themselves or the objects to maintain a clear line of sight, we would instead like the environment manager to move the virtual objects to accommodate the users. To accomplish this, the system could make use of knowledge about which objects are currently important to the users, and which are not. For example, the system could be instructed (or infer) that in some situations it is important for a user to be able to see certain other users, displays, and objects. The system would then be responsible for ensuring that the virtual objects do not violate these visibility constraints, even when users, displays, and physical objects move.

Acknowledgements

This research was supported by a 1997–98 stipend to Andreas Butz from the German Academic Exchange Service

(DAAD), ONR Contracts N00014-97-1-0838 and N00014-99-1-0249, the Advanced Network & Services National Tele-Immersion Initiative, and gifts from Intel, Microsoft, and Mitsubishi.

References

- [1] M. Agrawala, A. Beers, B. Fröhlich, and P. Hanrahan. The two-user responsive workbench: Support for collaboration through individual views of a shared space. In *Proc. ACM SIGGRAPH '97*, pages 327–332, Los Angeles, CA, August 3–8 1997. ACM Press.
- [2] C. Ahlberg, C. Williamson, and B. Shneiderman. Dynamic queries for information exploration: An implementation and evaluation. In *Proc. ACM CHI '92*, pages 619–626. ACM Press, 1992.
- [3] D. Amselem. A window on shared virtual environments. *Presence*, 4(2):130–145, 1995.
- [4] R. T. Azuma. A survey of augmented reality. *Presence*, 6(4):355–385, Aug. 1997.
- [5] E. A. Bier, M. C. Stone, K. Fishkin, W. Buxton, and T. Baudel. A taxonomy of see-through tools. In *Proc. ACM CHI '94*, pages 358–364. ACM Press, 1994.
- [6] M. Billinghurst, J. Bowskill, M. Jessop, and J. Morphet. A wearable spatial conferencing space. In *Proc. 2nd Int. Symp. on Wearable Computers*, pages 76–83, 1998.
- [7] D. A. Bowman and L. F. Hodges. An evaluation of techniques for grabbing and manipulating remote objects in immersive virtual environments. In *Proc. 1997 Symp. on Interactive 3D Graphics*, pages 35–38, Providence, RI, 1997.
- [8] A. Butz, C. Beshers, and S. Feiner. Of vampire mirrors and privacy lamps: Privacy management in multi-user augmented environments. In *Proc. ACM UIST '98*, pages 171–172, November 2–4 1998.
- [9] C. Carlsson and O. Hagsand. DIVE—A platform for multi-user virtual environments. *Computers and Graphics*, 17(6):663–669, November–December 1993.
- [10] E. Churchill and D. Snowden. Collaborative virtual environments: An introductory review of issues and systems. *Virtual Reality*, 3(1):3–15, 1998.
- [11] Ericsson, IBM, Intel, Nokia, and Toshiba. Bluetooth mobile wireless initiative. <http://www.bluetooth.com>, 1998.
- [12] S. Feiner, B. MacIntyre, M. Haupt, and E. Solomon. Windows on the world: 2D windows for 3D augmented reality. In *Proc. ACM UIST '93*, pages 145–155, 1993.
- [13] S. Feiner, B. MacIntyre, T. Höllerer, and A. Webster. A touring machine: Prototyping 3D mobile augmented reality systems for exploring the urban environment. In *Proc. ISWC '97 (Int. Symp. on Wearable Computers)*, Cambridge, MA, October 13–14 1997.
- [14] S. Feiner and A. Shamash. Hybrid user interfaces: Breeding virtually bigger interfaces for physically smaller computers. In *Proc. ACM UIST '91*, pages 9–17. ACM Press, 1991.
- [15] K. P. Fishkin, T. P. Moran, and B. L. Harrison. Embodied user interfaces: Towards invisible user interfaces. In *Proc. EHCI '98*, Heraklion, Greece, 1998.
- [16] G. W. Fitzmaurice. Situated information spaces and spatially aware palmtop computers. *CACM*, 36(7):38–49, July 1993.
- [17] E. Foxlin, M. Harrington, and G. Pfeifer. Constellation: A wide-range wireless motion-tracking system for augmented reality and virtual set applications. In *Proc. ACM SIGGRAPH '98*, pages 372–378, July 19–24 1998.
- [18] S. Gottschalk and J. Hughes. Autocalibration for virtual environments tracking hardware. In *Proc. ACM SIGGRAPH '93*, pages 65–72, Anaheim, August 1–6 1993.
- [19] T. Höllerer, S. Feiner, and J. Pavlik. Situated documentaries: Embedding multimedia presentations in the real world. In *Proc. ISWC '99 (Int. Symp. on Wearable Computers)*, San Francisco, CA, October 18–19 1999.
- [20] T. Holmer, L. Lacour, and N. Streitz. i-LAND: An interactive landscape for creativity and innovation. In *Proc. Conf. on Computer-Supported Cooperative Work (ACM CSCW'98)*, Videos, page 423, 1998.
- [21] D. Hopkins. Directional selection is easy as pie menus! *login: The Usenix Association Newsletter*, 12(5), Sept. 1987.
- [22] H. Ishii, M. Kobayashi, and K. Arita. Iterative design of seamless collaboration media. *CACM*, 37(8):83–97, Aug. 1994.
- [23] H. Ishii and N. Miyake. Toward an open shared workspace: Computer and video fusion approach of Teamworkstation. *CACM*, 34(12):37–50, December 1991.

- [24] B. MacIntyre and S. Feiner. Future multimedia interfaces. *Multimedia systems*, 1996(4):250–268, 1996.
- [25] B. MacIntyre and S. Feiner. Language-level support for exploratory programming of distributed virtual environments. In *Proc. ACM UIST '96*, pages 83–94, Seattle, WA, November 6–8 1996.
- [26] B. MacIntyre and S. Feiner. A distributed 3D graphics library. In *Computer Graphics (Proc. ACM SIGGRAPH '98)*, Annual Conference Series, pages 361–370, Orlando, FL, July 19–24 1998.
- [27] The PIT: Protein interactive theater. URL: <http://www.cs.unc.edu/Research/graphics/GRIP/PIT.html>, 1998.
- [28] R. Raskar, G. Welch, M. Cutts, A. Lake, L. Stesin, and H. Fuchs. The office of the future: A unified approach to image-based modeling and spatially immersive displays. In *Proc. ACM SIGGRAPH '98*, pages 179–188, 1998.
- [29] J. Rekimoto. Pick-and-drop: A direct manipulation technique for multiple computer environments. In *Proc. ACM UIST '97*, pages 31–39. ACM Press, 1997.
- [30] J. Rekimoto. A multiple device approach for supporting whiteboard-based interactions. In *Proc. ACM CHI '98*, pages 344–351. ACM Press, 1998.
- [31] J. Rekimoto and K. Nagao. The world through the computer: Computer augmented interaction with real world environments. In *Proc. ACM UIST '95*, pages 29–36, 1995.
- [32] J. Rekimoto and M. Saitoh. Augmented surfaces: A spatially continuous work space for hybrid computing environments. In *Proc. ACM CHI '99*, Pittsburgh, PA, May 15–20 1999. ACM Press.
- [33] S. Robertson, C. Wharton, C. Ashworth, and M. Franzke. Dual device user interface design: PDAs and interactive television. In *Proc. ACM CHI '96*, pages 79–86. ACM Press, 1996.
- [34] D. Seligmann and J. Edmark. Automatically generated 3D virtual environments for multimedia communication. In *Proc. Fifth Int. Conf. in Central Europe on Comp. Graphics and Visualization (WSCG97)*, February 10–14 1997.
- [35] M. Spitzer, N. Rensing, R. McClelland, and P. Aquilino. Eyeglass-based systems for wearable computing. In *Proc. First Int. Symp. on Wearable Computers*, pages 48–51, Cambridge, MA, October 13–14 1997.
- [36] Z. Szalavari, E. Eckstein, and M. Gervautz. Collaborative gaming in augmented reality. In *Proc. VRST '98*, pages 195–204, Taipei, Taiwan, November 1998.
- [37] Z. Szalavari, D. Schmalstieg, A. Fuhrmann, and M. Gervautz. "Studierstube": An environment for collaboration in augmented reality. *Virtual Reality*, 3(1):37–48, 1998.
- [38] B. Ullmer, H. Ishii, and D. Glas. mediaBlocks: Physical containers, transports, and controls for online media. In M. Cohen, editor, *Proc. ACM SIGGRAPH '98*, Annual Conference Series, pages 379–386. ACM SIGGRAPH, Addison Wesley, July 1998.
- [39] R. Want, K. Fishkin, A. Gujar, and B. Harrison. Bridging physical and virtual worlds with electronic tags. In *Proc. ACM CHI '99*, Pittsburgh, PA, May 15–20 1999. ACM Press.
- [40] M. Weiser. The computer for the 21st century. *Scientific American*, 3(265):94–104, 1991.