CUCS-005-99

An Experimental Hybrid User Interface for Collaboration

Andreas Butz[1], Tobias Höllerer, Clifford Beshers, Steven Feiner
Dept. of Computer Science
Columbia University
butz@cs.uni-sb.de,
{htobias,beshers,feiner}@cs.columbia.edu

Blair MacIntyre
College of Computing
Georgia Institute of Technology
blair@cc.gatech.edu

Abstract. We present EMMIE (Environment Management for Multi-user Information Environments), an experimental user interface to a collaborative augmented environment. Users share a 3D virtual space and manipulate virtual objects representing information to be discussed. This approach not only allows for cooperation in a shared physical space, but also addresses tele-collaboration in physically separate but virtually shared spaces. 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. Our research 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] Current affiliation: Fachbereich Informatik, Universität des Saarlandes, Germany

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 [41]. These devices include displays (ranging from palm sized to wall sized), but also include 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 all of these computers can be networked together, they add a (mostly) invisible virtual layer to the physical reality surrounding us in our daily lives.

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 and headphones) to enhance a user's senses by overlaying a directly perceptible virtual layer on the physical world. Because information is displayed 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 the hidden virtual layer of a ubiquitous computing environment to be visualized, as well as displaying personal information in a way that guarantees it remains private and can be customized for each user. However, one important drawback of pure AR interfaces arises because the 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. As we suggest in [27], AR systems can profit from the use of physical objects and the interaction techniques they afford. 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 CSCW 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 [18]. In the interface described 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 easily moved between the various displays.

Design Approach

Our prototype uses AR as an encompassing 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. 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, directions that will become feasible later when the needed technologies reach maturity. Thus, the 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 much like conventional eyeglasses [37], accurate wide-range motion tracking [22, 21], and standards for mobile wireless data and voice networking [15]. 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.

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 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.

In this paper, we introduce EMMIE (Environment Management for Multi-user Information Environments), a prototype environment manager. EMMIE supports a

Figure 1: A meeting situation 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. 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.

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.

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 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 2[2], the virtual space

[2] All images in this paper that show overlaid graphics (Figures 2, 3, 5, and 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 correspond to what a user actually sees through the display. 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 digitization process.

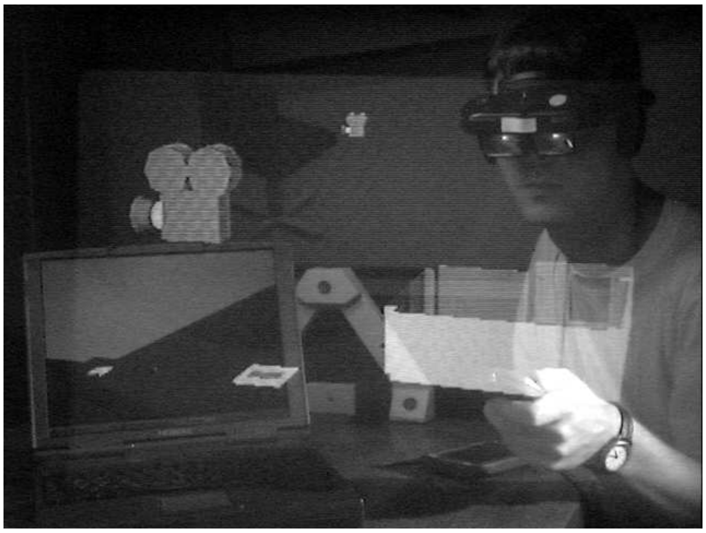

Figure 2: A user manipulates a 3D model with an optically tracked hand-held pointer. Other virtual objects shown include simple iconic 3D representations of applications (e.g., a “movie projector” that can play videos) and data objects (e.g., “slides” that represent still images).

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).

An alternative scenario we kept in mind while designing our system, is the set of *telecubicles* described in [31, 10]. In this environment, each user sits at a desk in the corner of a room where the two adjacent walls and the desk surface are stereo projection displays creating a CAVE-like [13] immersive experience. A set of remote telecubicles can be assembled electronically into a large virtual room. A user’s local cubicle contains both physical and virtual objects, while the physical and virtual objects in the remote cubicles, as well as the remote users, appear locally only as rendered models. Since rendering can be selectively omitted, physical objects and even users can be kept hidden from other participants. One of our rationales in the design of EMMIE was to simulate such a telecubicle environment and to provide a

UI design testbed for it. The techniques presented in this work are thus intended for both local and remote collaboration.

Previous and 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* [32, 33]. 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 AR space with environment management facilities through their see-through head-worn displays.

i-LAND [23] is an integrated work environment supporting cooperative work 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 [39]. 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 [19, 40]. 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 an environment management system.

UNC's "Office of the Future" [31] 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.

We recently learned about the PIT project at UNC [30], which presents a two-person two-screen stereo display workspace for collaborative study of a 3D model. Their system 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 [26] 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 [11, 12, 36]), we view a 3D AR environment as an embracing shared virtual space, incorporating, instead of replacing, existing UI technologies. In fact, with EMMIE's AR environment we are trying to achieve *seamlessness* [25] between different computer platforms, display devices of different sizes and dimensionalities, and among different (local or remote) users.

Researchers at the University of Washington [7] and at the Technische Universität Wien [38] have proposed AR Interfaces for CSCW. Both groups use see-through head-worn displays to overlay computer graphics on the real world. The University of Washington group also performed experiments to evaluate user performance in an augmented versus a totally immersed setting and in an body-stabilized versus head-stabilized setting [6, 8]. 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 [20], several researchers have employed similar displays as lenses or "see-through" devices to overlay computer generated imagery on the real world [2, 34]. We use this technique in the broader context of a hybrid user interface environment management system, recognizing it 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 [18], overlaid conventional 2D windows on the 3D world using a see-through head-worn display [16], and developed a wearable outdoor hybrid user interface that combined a tracked see-through head-worn display with a hand-held pen-based computer [17]. A number of other researchers have worked on embedding 2D windows in 3D environments, mainly for making information accessible from within virtual worlds [14, 3]. EMMIE extends this research by using see-through displays to integrate multiple, heterogeneous displays into one unified, collaborative information space.

Interaction with virtual objects

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 viewplane (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 [9] 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 back 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 [24] 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.

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 can be used for some of the hybrid interaction techniques we have developed.

Hybrid interaction

By *hybrid interaction*, we mean those forms of interaction that cut across different devices, modalities, and dimensionalities [18, 35, 32, 33]. For example, to use a workstation in the physical environment and the wall-sized display connected to it 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 workstation and the 3D virtual space surrounding us.

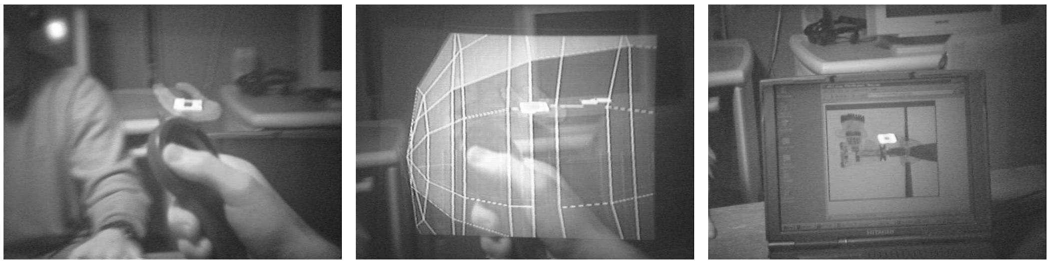

Figure 3: Drag & drop of virtual objects. A virtual object is picked up using a 3D pointing device (left image), dragged to a laptop, whose spherical bounding volume highlights, and dropped onto it (center image). The object then appears on the laptop's screen (right image).

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 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 [32, 33] 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. (The presence of this representation also raises a variety of privacy issues, which we discuss later.)

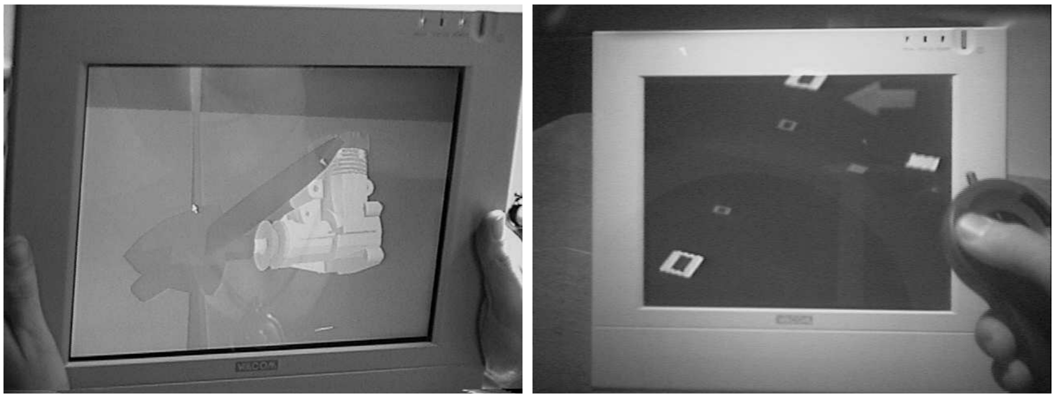

Figure 4: The same display tablet can serve as physical magic lens and magic mirror. In the left picture a user not wearing a HMD is looking through the magic lens at a 3D CAD object. In the right picture a user wearing a HMD has just dragged a virtual image slide in front of the mirror for further inspection.

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 [20]) 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 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 [1] embedded in the 3D world.

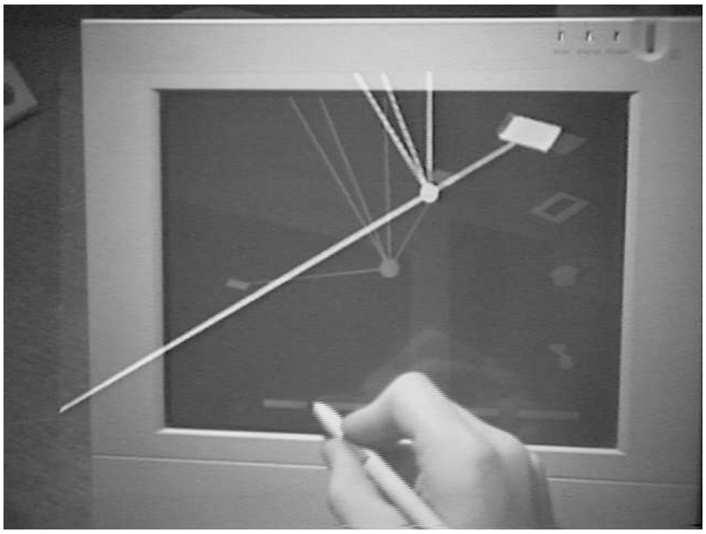

Figure 5: A simple interactive search mechanism creates 3D leader lines emanating from the tracked display to objects satisfying the query.

Privacy management

Privacy is an important issue whenever data is represented in a shared environment. Since an EMMIE user may want some data to be private and other data public, we need to provide a way to view and modify the privacy of the 3D virtual objects. Furthermore, because the best security measure is eternal vigilance, we want the privacy information to be either constantly visible or highly accessible. The challenge in an AR environment is to achieve this without also being visually annoying or outright obstructive of other interactions. In [10] we considered the conceptual design of two user interfaces for controlling privacy in a collaborative teleimmersive environment: *privacy lamps* and *vampire mirrors*. Here we present prototype

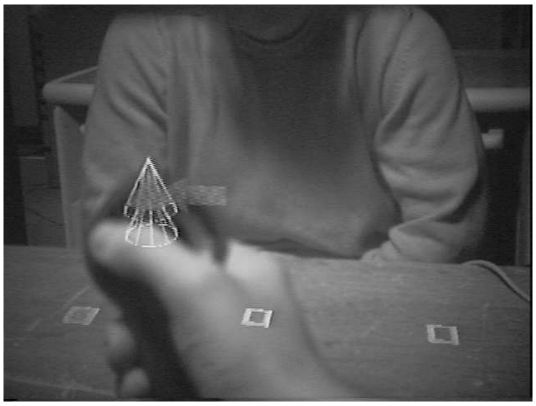

Figure 6: A user changes the privacy status of an object by moving a privacy lamp above it

Figure 7: A user changes the privacy status of an object on the vampire mirror

implementations of both within EMMIE and discuss how they satisfy the above criteria.

Privacy lamps (see Figure 6) are cone-shaped virtual light sources that emit privacy as colored light, typically red, 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. Typically these lamps float, 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 (see Figure 7) act as magic mirrors in the virtual environment, reflecting a user's virtual objects, except that they reflect only public objects fully. Private objects are either invisible or optionally displayed as a ghost image. By placing a vampire mirror at the back of the workspace, 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 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.

In our original conception [10], the vampire mirror was a virtual object. In EMMIE, we use a tracked, touch-sensitive LCD panel as a vampire mirror. Users can place it on the desktop in a convenient spot to get a high-resolution view of the privacy state, and can touch the actual display to change the privacy state, giving them passive haptic feedback.

In addition, this LCD panel can be used as a privacy lens, allowing the user to look through the display at objects to determine their privacy state. As with the vampire mirror, private objects are invisible or appear as ghost images, and public objects are rendered normally. Again, the owner can toggle an object's privacy state by touching its image on the lens. An advantage of the lens is that it allows users to work with a physical object close to their bodies, making the interaction real and comfortable.

Implementation

EMMIE is implemented on top of Coterie, our testbed for exploratory research in augmented environments [28]. 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 the compiled (Modula-3) and interpreted (Repo) languages, as well as various libraries for such things as 3D graphics and tracker control. By allowing programmers to prototype distributed programs in an interpreted language, Coterie greatly speeds the development process.

EMMIE takes significant advantage of two components of Coterie: Repo, the interpreted language, and Repo-3D, the distributed 3D graphics library [29]. EMMIE is distributed over several machines. Its primary structure is a simple replicated object directory implemented in Repo, similar to the one described in [28]. Each EMMIE process has a copy of a shared directory, and when any process adds or removes an object, the others are notified. In this way, all the processes are peers, with no centralized master process required to coordinate the application. The directories are replicated in each process to ensure fast access to the objects when needed for real time graphics generation. 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 the object components is a Repo-3D scene graph that defines the appearance of the object. This object is constructed of a hierarchy of Repo-3D choice groups, each of which allows the various processes in EMMIE to choose between the various possible 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 various interaction techniques and object representations can be experimented with simply and cleanly.

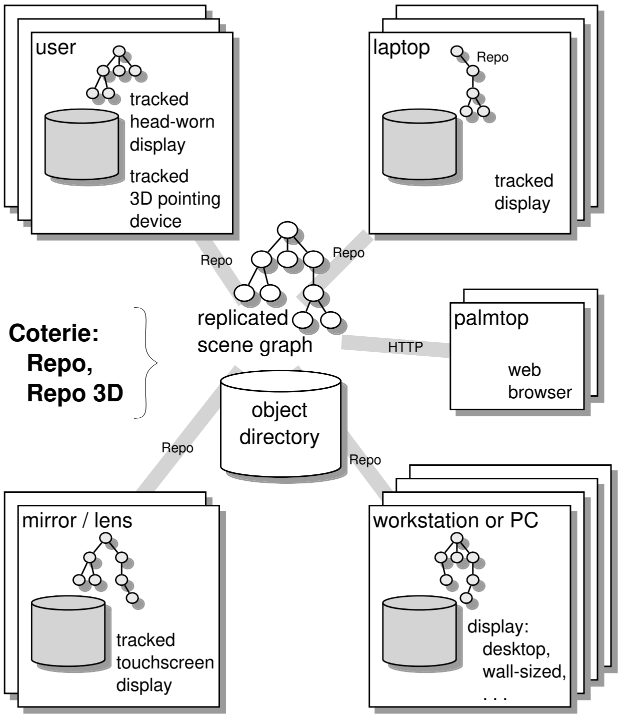

Figure 8: Architecture of the EMMIE system

Figure 8 shows a diagram of the 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, which 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 display's orientation is tracked with a built-in magnetometer and inclinometer. The magic mirror and lens are implemented on a Wacom PL-300 LCD panel with pen-input facilities, driven by a PC, and tracked by a Logitech 6DOF ultrasonic tracker. Other workstations and laptops join the environment by running a background thread implementing EMMIE's drag & drop functionality, allowing them to be fully integrated in the environment. While we assume the workstation displays stay in fixed positions, the laptop displays are tracked with Logitech 6DOF trackers. Hand-held devices such as the 3Com Palm Pilot 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.

Future Work: Automated 3D Layout Assistance

The next avenue we would like to explore is the use of dynamic, context-sensitive techniques that can be added to an environment manager to help users manage information more effectively. There are additional analogs of 2D window management techniques that can be adapted to this environment. A range of new techniques could also help the user deal with the dynamic, 3D nature of augmented environments.

First, let us consider the 3D analogs of some techniques used in 2D window managers to help users both position objects when they are created, and keep them organized. For example, when a new item is created, the environment manager should 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, as not only virtual but real items must be considered. We do not want to place a new virtual item on top of some important physical item such as a telephone, for example, unless it is meaningfully associated with it.

Other techniques could help the user deal with the dynamic nature of augmented environments. For example, returning to the analogy between desktop window managers and environment managers such as EMMIE, we note that window managers do exactly what the user tells them; no more, no less. Such an approach is fine for the desktop interface, as the desktop itself is generally static and lifeless, with activity happening only inside statically positioned windows. By contrast, an environment manager must be suited to the dynamic nature of an augmented environment, interacting with the user's actions and assisting the user in managing their ever-changing environment.

On a traditional 2D desktop UI, the organization of information typically changes in response to some user action, such as starting or stopping programs, moving windows, and so forth. In an augmented environment, information is positioned in 3D, often in relation to objects that may move. As the user, other people, and objects in the environment move about, the relationships between the virtual and real world change. For example, if a group of objects is suspended in space above a table, and two users of the system wish to talk to one another, they most likely do not want the objects blocking their view of each other. Rather than forcing the users to move to accommodate the virtual objects, 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 infer that it is important for a user to be able to see all other users and relevant displays. Therefore, the system could constrain the virtual objects so they do not interfere with the users' views of each other and of their relevant displays.

Conclusions

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. Many of the interaction and user interface techniques presented are variations of ones that have been proposed before. Our major contribution is their coordination within a single cohesive framework for collaborative use. By merging virtual, physical, and hybrid interaction 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.

Acknowledgements

This research is supported by a stipend to Andreas Butz from the German Academic Exchange Service (DAAD), ONR Contract N00014-97-1-0838, the Advanced Network & Services National Tele-Immersion Initiative, and gifts from Intel, Microsoft, and Mitsubishi Electric Research Laboratory.

References

1. C. Ahlberg, C. Williamson, and B. Shneiderman. Dynamic queries for information exploration: An implementation and evaluation. In *Proceedings of CHI '92*, pages 619–626. ACM press, 1992.

2. D. Amselem. A window on shared virtual environments. *Presence: Teleoperators and Virtual Environments*, 4(2):130–145, 1995.

3. I. G. Angus and H. A. Sowizral. VRMosaic: WEB access from within a virtual environment. In N. Gershon and S. Eick, editors, *Proc. IEEE Visualization '96*, pages 59–64. IEEE Computer Society Press, October 1996.

4. R. T. Azuma. A survey of augmented reality. *Presence: Teleoperators and Virtual Environments*, 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 *Proceedings of CHI 94*, pages 358–364. ACM press, 1994.

6. M. Billinghurst, J. Bowskill, N. Dyer, and J. Morphett. An evaluation of wearable information spaces. In *Proc. IEEE VRAIS '98*, pages 20–27, March 14–18 1998.

7. M. Billinghurst, J. Bowskill, M. Jessop, and J. Morphett. A wearable spatial conferencing space. In *Proc. of the 2nd International Symposium on Wearable Computers*, pages 76–83, 1998.

8. M. Billinghurst, S. Weghorst, and T. Furness. Shared space: An augmented reality approach for computer supported collaborative work. *Virtual Reality*, 3(1):25–36, 1998.

9. D. A. Bowman and L. F. Hodges. An evaluation of techniques for grabbing and manipulating remote objects in immersive virtual environments. In *1997 Symposium on Interactive 3D Graphics*, pages 35–38, Providence, RI, 1997.

10. A. Butz, C. Beshers, and S. Feiner. Of vampire mirrors and privacy lamps: Privacy management in multi-user augmented environments [technote]. In *Proceedings of UIST'98*. ACM SIGGRAPH, 1998.

11. C. Carlsson and O. Hagsand. DIVE — A platform for multi-user virtual environments. *Computers and Graphics*, 17(6):663–669, Nov.–Dec. 1993.

12. E. Churchill and D. Snowdon. Collaborative virtual environments: An introductory review of issues and systems. *Virtual Reality*, 3(1):3–15, 1998.

13. C. Cruz-Neira, D. J. Sandin, and T. A. DeFanti. Surround-screen projection-based virtual reality: The design and implementation of the CAVE. In *Computer Graphics (Proc. ACM SIGGRAPH '93)*, Annual Conference Series, pages 135–142, Aug. 1993.

14. P. Dykstra. X11 in virtual environments. In *Proc. IEEE 1993 Symposium on Research Frontiers in Virtual Reality*, pages 118–119, San Jose, CA, October 25–26 1993.

15. Ericsson, IBM, Intel, Nokia, and Toshiba. Bluetooth mobile wireless initiative. http://www.bluetooth.com, 1998.

16. S. Feiner, B. MacIntyre, M. Haupt, and E. Solomon. Windows on the world: 2D windows for 3D augmented reality. In *Proceedings of UIST '93*, pages 145–155, 1993.

17. 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.

18. S. Feiner and A. Shamash. Hybrid user interfaces: Breeding virtually bigger interfaces for physically smaller computers. In *Proceedings of UIST '91*, pages 9–17. ACM press, 1991.

19. K. P. Fishkin, T. P. Moran, and B. L. Harrison. Embodied user interfaces: Towards invisible user interfaces. In *Proc. of EHCI '98*, Heraklion, Greece, 1998.

20. G. W. Fitzmaurice. Situated information spaces and spatially aware palmtop computers. *CACM*, 36(7):38–49, July 1993.

21. E. Foxlin, M. Harrington, and G. Pfeifer. Constellation: A wide-range wireless motion-tracking system for augmented reality and virtual set applications. In *Proc. SIGGRAPH '98*, pages 372–378, July 19–24 1998.

22. S. Gottschalk and J. Hughes. Autocalibration for virtual environments tracking hardware. In *Proc. SIGGRAPH '93*, pages 65–72, Anaheim, August 1–6 1993.

23. T. Holmer, L. Lacour, and N. Streitz. i-LAND: An interactive landscape for creativity and innovation. In *Proceedings of ACM CSCW'98 Conference on Computer-Supported Cooperative Work*, Videos, page 423, 1998.

24. D. Hopkins. Directional selection is easy as pie menus! *login: The Usenix Association Newsletter*, 12(5), Sept. 1987.

25. H. Ishii, M. Kobayashi, and K. Arita. Iterative design of seamless collaboration media. *Communications of the ACM*, 37(8):83–97, Aug. 1994.

26. H. Ishii and N. Miyake. Toward an open shared workspace: Computer and video fusion approach of Teamworkstation. *CACM*, 34(12):37–50, December 1991.

27. B. MacIntyre and S. Feiner. Future multimedia interfaces. *Multimedia systems*, 1996(4):250–268, 1996.

28. B. MacIntyre and S. Feiner. Language-level support for exploratory programming of distributed virtual environments. In *Proc. UIST '96*, pages 83–94, Seattle, WA, November 6–8 1996.

29. 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.

30. The PIT: Protein interactive theater. URL: http:// www.cs.unc.edu/ Research/ graphics/ GRIP/ PIT.html, 1998.

31. 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. of SIGGRAPH '98*, pages 179–188, 1998.

32. J. Rekimoto. Pick-and-drop: A direct manipulation technique for multiple computer environments. In *Proceedings of UIST '97*, pages 31–39. ACM Press, 1997.

33. J. Rekimoto. A multiple device approach for supporting whiteboard-based interactions. In *Proceedings of CHI 98*, pages 344–351. ACM Press, 1998.

34. J. Rekimoto and K. Nagao. The world through the computer: Computer augmented interaction with real world environments. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, Virtual and Augmented Realities, pages 29–36, 1995.

35. S. Robertson, C. Wharton, C. Ashworth, and M. Franzke. Dual device user interface design: PDAs and interactive television. In *Proceedings of CHI 96*, pages 79–86. ACM Press, 1996.

36. D. Seligmann and J. Edmark. Automatically generated 3D virtual environments for multimedia communication. In *Proceedings of the Fifth International Conference in Central Europe on Computer Graphics and Visualization (WSCG97)*, February 10–14 1997.

37. 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.

38. Z. Szalavari, D. Schmalstieg, A. Fuhrmann, and M. Gervautz. "Studierstube": An environment for collaboration in augmented reality. *Virtual Reality*, 3(1):37–48, 1998.

39. B. Ullmer, H. Ishii, and D. Glas. mediaBlocks: Physical containers, transports, and controls for online media. In M. Cohen, editor, *SIGGRAPH 98 Conference Proceedings*, Annual Conference Series, pages 379–386. ACM SIGGRAPH, Addison Wesley, July 1998. ISBN 0-89791-999-8.

40. R. Want, K. Fishkin, A. Gujar, and B. Harrison. Bridging physical and virtual worlds with electronic tags. Technical report, Xerox Palo Alto Research Center, Palo Alto, CA, USA, Sept. 1998. (in press).

41. M. Weiser. The computer for the 21st century. *Scientific American*, 3(265):94–104, 1991.