Augmenting Intelligent Environments:
Augmented Reality as an Interface to Intelligent Environments

Blair MacIntyre
Department of Computer Science, Columbia University
1214 Amsterdam Ave., Mailcode 0401, New York, NY 10027
phone: (212) 222-5315, fax: (212) 666-0140
e-mail: bm@cs.columbia.edu

Elizabeth D. Mynatt
Xerox Palo Alto Research Lab
3333 Coyote Hill Road, Palo Alto, CA 94304
phone: (650) 812-4945, fax: (650) 812-4890
email: mynatt@parc.xerox.com

Abstract

We believe that *augmented reality* techniques, that use computer generated information to enhance the user's perception of the world, provide a solid foundation for imbuing intelligent environments with ubiquitous display capabilities. Visual and auditory *personal augmented reality displays* that are worn or carried by individual users, combined with accurate position and orientation of the user, allow information to be presented *in context* on a *per user* basis. Augmented reality is useful for both explicit and direct interaction for providing foreground information as well as for implicit and passive interaction for providing background interaction. We believe that both auditory and visual augmented reality play an important role, and that the combination will yield far more compelling interfaces than either could alone.

Introduction

To quote from the call for the workshop, when we consider interfaces to Intelligent Environments, we agree that

> ... interaction with these environments should be in terms of forms that people are naturally comfortable with. Their user-interface primitives are not menus, mice and windows but gesture, speech, context, and affect. At least in theory, it should be possible for a person to forget she is using a computer while interacting with one.

However, much of the research on user-interfaces for intelligent environments has focused on supporting natural input (determining what the user is trying to communicate to the environment.) An equally important challenge is how the environment communicates with the user in natural and meaningful ways. We believe that *augmented reality* techniques, that use computer generated information to enhance the user's perception of the world, provide a solid foundation for imbuing the environment with ubiquitous display capabilities. In this paper, we focus on visual and auditory *personal augmented reality displays* that are worn or carried by individual users. As we are able to accurately sense the location and orientation of the occupants of the environment (perhaps to also support gesture recognition or contextual assistance based on location and activity), the auditory and visual augmentations can be presented *in context* on a *per user* basis. When combined with more traditional devices, a rich information space is created where combinations of public and private information can be presented to users simultaneously using a combination of displays. This augmenting and leveraging of the real world makes possible a wider variety of interaction techniques and ways of organizing information.

Our research has two main thrusts: the development of new UI paradigms and supporting software for computer environments emphasizing augmented reality, and the creation of prototype applications that serve as testbeds for these ideas.

Background

Our groups are investigating the development of augmented reality systems as a way of presenting additional information to users about their environment. At Columbia, the goal has been to explore augmenting a user's view of the world using see-through head-mounted displays, whereas at PARC the focus has been on augmenting the auditory space around the user.

At Columbia, we have built a number of self-contained, single-user augmented reality prototype systems over the years that demonstrate a variety of uses for visual augment reality. One of our earliest systems, KARMA (Knowledge-based Augmented Reality for Maintenance Assistance) [Feiner et al 1993b] used a see-through head-worn display to explain simple end-user maintenance for a laser printer. A knowledge-based graphics design system interactively designed the graphics and simple textual callouts that the

user saw. In another early project, we explored how 2D X Windows could be used within augmented reality environment [Feiner et al 1993a] by mapping the "desktop" onto part of a virtual sphere surrounding the user. A simple hypermedia system allowed links to be made between 2D windows and real world objects, allowing information about the environment to be displayed on the head-mounted display. We have continued to explore other uses for single-user augmented reality over the past four years, and are currently developing a rapid prototyping testbed, called COTERIE, to explore more complex multi-user, multi-display scenarios [MacIntyre and Feiner, 1996b].

At PARC, our current work explores using audio to connect a person's activities in the physical world with information culled from the virtual world[1]. In Audio Aura [Mynatt et al 1997] our goal is to provide serendipitous information, via background auditory cues, that is tied to people's physical actions in the workplace. For example, dropping by someone's empty office results in hearing an *auditory footprint*, a qualitative cue for how long they have been gone. Our current system combines three known technologies: active badges, distributed systems and digital audio delivered via portable, wireless headphones. An active badge [Want et al 1992] is a small electronic tag designed to be worn by a person. It repeatedly emits a unique infrared signal that is detected by a low-cost network of IR sensors placed around a building. The information from the IR sensors is collected and combined with other data sources such as online calendars and email queues. Audio Aura services (such as the auditory footprint) specify what information they require to provide that service. Services are triggered by changes in this Audio Aura database resulting in auditory cues being sent to the user's wireless headphones.

Research Directions

Our common research direction is to use augmented reality as an interface to intelligent environments. We are investigating the intersection of technology, work practice and design aesthetics issues in designing augmented realities.

Current sensing and display technologies offer powerful affordances but also impose challenging constraints. The accuracy and responsiveness of sensors and their accompanying infrastructure constrain the experiences that one can create. For example, in Audio Aura, the delays inherent in the active badge infrastructure require that a user linger at the doorway of an office to hear an auditory footprint. We cannot support the equivalent of glancing in the door while walking without trying to predict the user's movements in advance. If we also wish to explore techniques that take advantage of accurately sensing the user's location and where they are looking, such as overlaid graphics or spatialized audio, we need to use

---

1. Virtual world refers to cyberspace, information in networked computer systems.

technologies that can provide information about the position and orientation of a user's head with significantly higher accuracy and lower latency. Unfortunately, current accurate sensing technologies do not support many people or cover large areas, unlike active badges, so we can only explore these techniques within small areas containing a few people. However, even in this restricted context, there are new interaction and information organization challenges that need to be addressed. For example, when a see-through, head-worn display is used to overlay graphics on the environment, relatively small changes in the user's head position or orientation change the perceived relationship between the head-worn display, non head-worn displays and the other objects in the physical world. The information presented on the head-worn display (and perhaps on other displays as well) must be continuously updated and reorganized to reflect this ever-changing context.

In envisioning uses for augmented reality systems, we are trying to take advantage of more than just the constant availability of these displays, and find ways to use augmented reality to compliment and enrich existing work practices. For example, one goal of our work is to orchestrate hand-held, head-worn, desk-top, and wall-mounted units to work together to create a hybrid information space. In such an environment, multiple users wearing see-through head-worn displays might view common public information on physical displays in the environment, while each user's head-worn display could privately overlay their personal annotations on the public data, either visually or aurally. We are particularly interested in developing techniques to manage the relationship between the physical world and the possibly large amount of virtual information (on both the head-mounted and non head-mounted displays), a problem we refer to as *environment management* in analogy to window management [MacIntyre and Feiner 1996a].

Of course, high-precision sensing is not required to provide useful and enriching augmentations. One Audio Aura service, the *group pulse*, attempts to support distributed work groups by turning activity information about group members into a background, dynamic auditory cue akin to hearing the bustle of a co-located team. In general, our search for ways of enriching the periphery is part of PARC's efforts in creating *calm technology* [Weiser and Brown, 1997]. The design of the sonic ecology created in Audio Aura is key to its eventual success or failure. We are experimenting with a variety of sound types (music, sound effects, voice) as well as techniques for layering sounds in a dynamic soundscape.

Conclusion

In summary, we believe that augmented reality will provide a powerful and flexible user-interface for intelligent environments. In this paper, we illustrate different uses of augmented reality as well as the underlying media used to create them. Augmented reality is useful for both explicit and direct interaction for

providing foreground information as well as for implicit and passive interaction for providing background interaction. In both cases, augmented reality can enrich the connections between people and their environment (including other people). We believe that both auditory and visual augmented reality play an important role, and that the combination will yield far more compelling interfaces than either could alone. An underlying question in all of this work is what type of information should be provided, in what media should information be presented, how do we understand a user's context, and can we predict the actions and needs of the users.

Acknowledgments

References

Feiner, S., MacIntyre, B., Haupt, M. and Solomon, E. 1993a. Windows on The World: 2D Windows for 3D Augmented Reality. In *Proc. UIST '93 (ACM Symp on User Interface Software and Technology)*, pages 145–155. Atlanta, Georgia..

Feiner, S., MacIntyre, B. and Seligmann, D. 1993b. Knowledge-Based Augmented Reality. *Communications of the ACM* 36(7): 53–62.

MacIntyre, B. and Feiner, S. 1996a. Future Multimedia User Interfaces. *Multimedia Systems*, 4(5): 250–268.

MacIntyre, B. and Feiner, S. 1996b. Language-level Support for Exploratory Programming of Distributed Virtual Environments. In *Proc. UIST '96 (ACM Symp on User Interface Software and Technology)*, pages 83–94, Seattle, WA.

Mynatt, E. D., Back, M., Want, R. and Frederick, R. 1997. Audio Aura: Light-Weight Audio Augmented Reality. In *Proc. UIST '97 (ACM Symp on User Interface Software and Technology)*, pages 210-212, Banff, CA.

Want, R., Hopper, A., Falcao, V. and Gibbons, J. 1992. The Active Badge Location System. *ACM Transactions on Information Systems*. 10(1): 91-102.

Weiser, M. and Brown, J. S. 1997. Designing Calm Technology. At *CHI '97 Workshop on Research Issues in Ubiquitous Computing*, Atlanta, GA.