

A Touring Machine: Prototyping 3D Mobile Augmented Reality Systems for Exploring the Urban Environment

Steven Feiner, Blair MacIntyre, Tobias Höllerer

Department of Computer Science
Columbia University
New York, NY 10027
{feiner,bm,htobias}@cs.columbia.edu
<http://www.cs.columbia.edu/graphics/>

Anthony Webster

Graduate School of Architecture, Planning and
Preservation
Columbia University
New York, NY 10027
acw18@columbia.edu
<http://www.cc.columbia.edu/~archpub/BT/>

Abstract

We describe a prototype system that combines together the overlaid 3D graphics of augmented reality with the untethered freedom of mobile computing. The goal is to explore how these two technologies might together make possible wearable computer systems that can support users in their everyday interactions with the world. We introduce an application that presents information about our university's campus, using a head-tracked, see-through, head-worn, 3D display, and an untracked, opaque, handheld, 2D display with stylus and trackpad. We provide an illustrated explanation of how our prototype is used, and describe our rationale behind designing its software infrastructure and selecting the hardware on which it runs.

Keywords: Augmented Reality, Virtual Environments, Mobile Computing, Wearable Computing, GPS.

1. Introduction

Recent years have seen significant advances in two promising fields of user interface research: *virtual environments*, in which 3D displays and interaction devices immerse the user in a synthesized world, and *mobile computing*, in which increasingly small and inexpensive computers and wireless networking allow users to roam the real world without being tethered to stationary machines. We are interested in how virtual environments can be combined with mobile computing, with the ultimate goal of supporting ordinary users in their interactions with the world.

To experiment with these ideas, we have been building the system described in this paper. The kind of virtual environment technology with which we have been working is *augmented reality*. Unlike most virtual environments, in which a virtual world *replaces* the real world, in augmented reality a virtual world *supplements* the real world with additional information. This concept was pioneered by Ivan Sutherland [27], and is accomplished through the

use of tracked “see-through” displays that enrich the user's view of the world by overlaying visual, auditory, and even haptic, material on what she experiences.

The application that we are addressing is that of providing users with information about their surroundings, creating a personal “touring machine.” There are several themes that we have stressed in this work:

- Presenting information about a real environment that is integrated into the 3D space of that environment.
- Supporting outdoor users as they move about a relatively large space on foot.
- Combining multiple display and interaction technologies to take advantage of their complementary capabilities.

Our prototype assists users who are interested in our university's campus, overlaying information about items of interest in their vicinity. As a user moves about, she is tracked through a combination of satellite-based, differential GPS (Global Positioning System) position tracking and magnetometer/inclinometer orientation tracking. Information is presented and manipulated on a combination of a head-tracked, see-through, headworn, 3D display, and an untracked, opaque, handheld, 2D display with stylus and trackpad.

Our emphasis in this project has been on developing experimental user interface software, not on designing hardware. Therefore, we have used commercially available hardware throughout. As we describe later, this has necessitated a number of compromises, especially in the accuracy with which the user's 3D position and orientation is tracked. These have in turn affected the design of our user interface, which relies on approaches that require only approximate, rather than precise, registration of virtual and real objects.

In Section 2 we present related work. Section 3 describes a scenario in our application domain, including pictures generated by a running testbed implementation. In Section 4, we describe both our high-level approach in designing our system and the specific hardware and software used. Finally, Section 5 presents our conclusions and

the directions that we will be taking as we continue to develop the system.

2. Related Work

Previous research in augmented reality has addressed a variety of application areas including aircraft cockpit control [12], assistance in surgery [26], viewing hidden building infrastructure [10], maintenance and repair [9], and parts assembly [5, 29]. In contrast to these systems, which use see-through headworn displays, Rekimoto [23] has used handheld displays to overlay information on color-coded objects. Much effort has also been directed towards developing techniques for precise tracking using tethered trackers (e.g., [16, 2, 28, 25]).

Work in mobile user interfaces has included several projects that allow users to explore large spaces. Loomis and his colleagues have developed an application that makes it possible for blind users to navigate a university campus by tracking their position with differential GPS and orientation with a magnetometer to present spatialized sonic location cues [18]. Petrie et al. have field-tested a GPS-based navigation aid for blind users that uses a speech synthesizer to describe city routes [22]. The CMU Wearable Computer Project has developed several generations of mobile user interfaces using a single handheld or untracked headworn display with GPS, including a campus tour [24]. Long et al. have explored the use of infrared tracking in conjunction with handheld displays [17]. Mann [20] has developed a family of wearable systems with headworn displays, the most recent of which uses optical flow to overlay textual information on automatically recognized objects.

Our work emphasizes the combination of these two streams of research: augmented reality and mobile user interfaces. We describe a prototype application that uses tracked see-through displays and 3D graphics without assuming precise registration, and explore how a combination of displays and interaction devices can be used together to take advantage of their individual strengths.

Prior to the development of VRML, several researchers experimented with integrating hypertext and virtual environments [7, 8, 1]. All investigated the advantages of presenting hypertext on the same 3D display as all other material, be it headworn or desktop. In contrast, our current work exploits the different capabilities of our displays by presenting hypertext documents on the relatively high-resolution 2D handheld display, which is itself embedded within the 3D space viewed through the lower-resolution headworn display.



Figure 1. Prototype campus information system. The user wears a backpack and headworn display, and holds a handheld display and its stylus.

3. Application Scenario

Consider the following scenario, whose figures were created using our system. The user is standing in the middle of our campus, wearing our prototype system, as shown in Figure 1. His tracked see-through headworn display is driven by a computer contained in his backpack. He is holding a handheld computer and stylus.

As the user looks around the campus, his see-through headworn display overlays textual labels on campus buildings, as shown in Figures 2 and 3. (These images were shot through the headworn display, as described in Section 4.3, and are somewhat difficult to read because of the low brightness of the display and limitations of the recording technology.) Because we label buildings, and not specific building features, the relative inaccuracy of the trackers we are using is not a significant problem for this application.

At the top of the display is a menu of choices: “Columbia:”, “Where am I?”, “Depts?”, and “Buildings?”. When selected, each of these choices sends a URL to a web browser running on the handheld computer. The browser then presents information about the campus, the user’s current location, a list of departments, and a list of buildings, respectively. The URL points to a custom HTTP server on the handheld computer that generates a page on the fly containing the relevant information. The generated pages contain links back to the server itself and to regular web pages



Figure 2. View shot through the see-through headworn display, showing campus buildings with overlaid names. Labels increase in brightness as they near the center of the display.

elsewhere. (An additional menu item, “Blank”, allows the headworn display to be blanked when the user wants to view the unaugmented campus.) Menu entries are selected using a touchpad mounted on the back of the handheld computer. The touchpad’s x coordinates are inverted to preserve intuitive control of the menus.

Labels seen through the headworn display are grey, increasing in intensity as they approach the center of the display. The one label closest to the center is highlighted yellow. If it remains highlighted for more than a second, it changes to green, indicating that it has been selected, and a second menu bar is added below the first, containing entries for that building. A selected building remains selected until the user’s head orientation dwells on another building for more than a second as indicated by the color change. This approximation of gaze-directed selection can be disabled or enabled through a menu item.

When a building is selected, a conical green compass pointer appears at the bottom of the headworn display, oriented in the building’s direction. The pointer turns red if the building is more than 90 degrees away from the user’s head orientation. This allows the user to find the building more easily if they turn away from it. The pointer is especially useful for finding buildings selected from the handheld computer. This is made possible by our custom HTTP server, which can tell the backpack computer to select a building on the headworn display.

The building’s menu bar contains the name of the building, plus additional items: “Architecture”, “Departments”, and “Miscellaneous”. Selecting the name of the building from the menu using the trackpad sends a relevant URL to the handheld computer’s browser. Selecting any of the remaining menu entries also sends a URL to the browser and creates a collection of items that are positioned near the building on the headworn display.



Figure 3. A view of the Philosophy Building with the “Departments” menu item highlighted.

To call the user’s attention to the new material on the handheld computer, when menu items that send URLs are selected, a copy of the menu item is translated down to and off the bottom of the headworn display. For example, Figure 3 shows the Philosophy Building with the “Departments” menu item highlighted prior to selection. When the item is selected, the building is surrounded with the names of the departments that it contains, as shown in Figure 4. The automatically-generated web page displayed on the handheld is shown in Figure 5(a).

There are two ways to access information about the selected building. On the headworn display, the user can cycle through the surrounding items with the trackpad and select any to present relevant information about it on the handheld display. Alternatively, the user can select a corresponding item from the automatically-generated web page. For example, Figure 5(b) shows the regular web page for one of the departments in the Philosophy Building, accessed through the system. The lists of buildings and departments produced by the top-level menu items on the headworn display can also be used to access this information; e.g., to find out about a building or department whose name is known.

4. System Design

While we wanted our system to be as lightweight and comfortable as possible, we also decided to use only off-the-shelf hardware to avoid the expense, effort, and time involved in building our own. Consequently we often settled for items that were far bulkier than we would like them to be, in return for the increased flexibility that they offered. The combined weight of the system is just under 40 pounds.



(a)




(b)

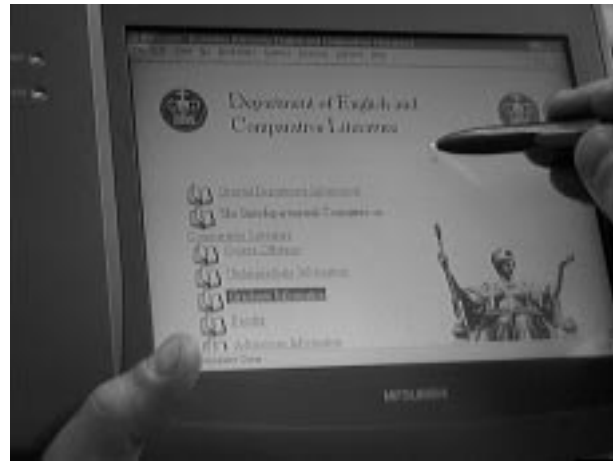


(c)

Figure 4. After the “Departments” menu item is selected, the department list for the Philosophy Building is added to the world, arrayed about the building. The three figures show the label animation sequence: (a) a fraction of a second after selection, (b) approximately half a second later, and (c) after the animation has finished.



(a)



(b)

Figure 5. (a) Selecting the “Departments” menu item causes an automatically-generated URL to be sent to the web browser on the handheld computer, containing the department list for the Philosophy Building. (b) Actual home page for the English and Comparative Literature department, as selected from either the generated browser page or the department list of Figure 4.

The following subsections describe some of the hardware and software choices that we made in designing our system, whose hardware design is diagrammed in Figure 6.

4.1. Hardware

Backpack computer. It was important to us that our main computer not only be portable, but also capable of working with readily available peripherals, including high-performance 3D graphics cards. We chose a Fieldworks 7600, which includes a 133MHz Pentium, 64Mbyte memory, 512K cache, 2GB disk, and a card cage that can hold 3 ISA and 3 PCI cards. While this system is our biggest compromise in terms of weight and size, it has significantly simplified our development effort.




Figure 6. Hardware design of our prototype campus information system.

Graphics card. We use an Omnicomp 3Demon card, which is based on the Glint 500DTX chipset, including hardware support for 3D transformations and rendering using OpenGL.

Handheld computer. Our handheld computer is a Mitsubishi Amity, which has a 75MHz DX4, 640x480 color display, 340MB disk, 16MB main memory, PCMCIA slot, and integral stylus. Control of the headworn display menu is accomplished through a Cirque GlidePoint trackpad that we mounted on the back of the handheld computer. (We originally considered having the handheld computer stylus control the headworn display's menu when it was within a designated physical area of the handheld computer's display. We decided against this, however, because it would be difficult to remain in that area when the user was not looking at the handheld display.)

Headworn display. Video see-through displays currently provide a number of advantages over optical see-through displays, particularly with regard to registration and proper occlusion effects [26]. However, video-based systems restrict the resolution of the real world to that of the virtual world. While we believe that this is a good trade-off in

many applications, we feel that augmented reality systems will become commonplace only when they truly add to reality, rather than subtract from it. In our work we have selected the relatively lightweight Virtual I/O *i-glasses* head-worn display. This is a 60,000 triad color display. We are also experimenting with a Virtual I/O 640x480 resolution greyscale display.

Orientation tracker. We use the built-in tracking provided with our headworn display. This includes a magnetometer, which senses the earth's magnetic field to determine head yaw, and a two-axis inclinometer that uses gravity to detect head pitch and roll.

Position tracking. We use a Trimble DSM GPS receiver to obtain position information for its antenna, which is located on the backpack above the user's head. While normal GPS generates readings that are accurate only within about 100 meters, it can be routinely coupled with correction information broadcast from another receiver at a known location that contains information about how far it is off. We subscribe to a differential correction service provided by Differential Corrections Inc., which allows us to achieve about one-meter accuracy.

Network. To provide communication with the rest of our infrastructure we use NCR WaveLan spread-spectrum 2Mbit/sec radio modems in both the backpack and handheld PCs, which operate with a network of base stations on campus.

Power. With the exception of the computers, each of the other hardware components has relatively modest power requirements of under 10 watts each. We run them all using an NRG Power-MAX NiCad rechargeable battery belt. It has the added advantage of allowing a fully charged replacement powerpack to be plugged in prior to unplugging the depleted powerpack, without interrupting power.

4.2. Software

Infrastructure. We use COTERIE [19], a system that provides language-level support for distributed virtual environments. COTERIE is based on the distributed data-object paradigm for distributed shared memory. Any data object in COTERIE can be declared to be a shared object that either exists in one process, and is accessed via remote-method invocation, or is replicated fully in any process that is interested in it. The replicated shared objects support asynchronous data propagation with atomic serializable updates, and asynchronous notification of updates. COTERIE runs on Windows NT/95, Solaris, and IRIX, and includes the standard services needed for building virtual environment applications, including support for assorted trackers, etc. This software is built on top of Modula-3 [14] and Obliq [4].

Graphics package. We use a version of Obliq-3D [21], a display-list based 3D graphics package, which we have modified both to provide additional features needed for vir-

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.