An Interface for a Continuously Available, General Purpose, Spatialized Information Space

Rob Kooper and Blair MacIntyre
GVU Center and College of Computing
Georgia Institute of Technology
Atlanta, GA 30332 USA
`{kooper,blair}@cc.gatech.edu`

Abstract

In this paper we describe an augmented reality (AR) system that acts as continuously available interface to a spatialized information space based on the World Wide Web. We call such an information space the Real-World Wide Web (RWWW). We present the assumptions we make about the characteristics of such a system, discuss the implications of those assumptions for an AR interface, and describe a RWWW browser we are building.

Keywords: augmented reality, human-computer interaction, adaptive interfaces, mobile computing, wearable computing.

1 Introduction

In our previous work, we have focused on task-specific AR systems (e.g., (Feiner, MacIntyre, Hollerer, & Webster, 1997)). Recently, we have begun exploring how to create AR interfaces for general-purpose 3D information spaces, as suggested by (Spohrer, 1999). In our current research, we are investigating the use of 3D augmented reality (AR) to envelop mobile users in contextually-relevant, spatially-registered information spaces. We envision an enhanced version of the World Wide Web, where each information object (i.e., “web page”, containing 2D and 3D visual and auditory information) may have contextual meta-data associated with it, such as a location, person, or activity. This meta-data would be used to decide when and where to present the information objects to users. We refer to such a space as the *Real*-World Wide Web (RWWW), and interfaces to it as RWWW Browsers. In a RWWW Browser, web pages would be registered with real-world locations and presented to the user at the appropriate place and time.

We are particularly interested in RWWW Browsers that can be worn *continuously*, providing users with constant awareness of the information spaces they move through in their daily lives. We are interested in presenting the user not only with information that is relevant to their current location (i.e. restaurants or people near by) but any information that might be pertinent to their context (i.e. to-do items’s, calendar entries, messages from friends) (Salber, Dey, & Abowd, 1999; Starner, 1999). Unlike most research in continuously-worn wearable computing, where researchers use 2D heads-up-displays (e.g., (Rhodes, 1997; Starner, 1999)), the ability to place information in 3D around the user raises new opportunities and challenges. In particular, the interface design must balance the conflicting requirements of minimizing the volume of information displayed (to avoid distracting the user and cluttering their visual field) with the need to provide rich context (to capitalize on the users ability to rapidly scan and synthesize data).

2 Background and Related Work

Our work is related to four main areas of research: augmented reality, ambient or peripheral displays, wearable computing and context-aware computing. Our expectation, and that of others (Spohrer, 1999), is that the architecture supporting the general purpose spatialized information space will evolve from the current World Wide Web (WWW), retaining its structure as a loosely organized, heterogeneous collection of information servers. The key difference is that this future WWW will serve up virtual information that contains contextual data describing its association with the real world. Such data might describe a specific geo-spatial location for the information, but might also be more abstract, such as associating the information with a person, time, situation or activity (Abowd, Dey, Brotherton, &

Orr, 1999). Because such a system effectively embeds the WWW in the real world, we refer to such a context-based information space as the "Real-World" Wide Web (RWWW).

Our vision of the RWWW is closest to Spohrer's notion of the WorldBoard (Spohrer, 1999). Like Spohrer, our ultimate vision is of a global information space that merges the real and the virtual, with a variety of content, and accessible to anyone. Both WorldBoard and the RWWW extend the current WWW infrastructure by enhancing the data provided by existing servers, and by creating new servers that deliver specialized content. Our initial vision differs from the initial vision of WorldBoard primarily in the granularity of context that can be associated with information. To simplify the interface, WorldBoard focuses on the use of carefully authored information *channels*. We believe that these simplifications cause WorldBoard to lose much of the character of the WWW that has made it so successful, namely the ability for anyone to author, structure and publish information as they wish; the focus on authored content restricts what can be encountered by the user, and the use of a coarse coordinate system as the only context restricts where and (perhaps more importantly) why information can be encountered.

There have been numerous other AR systems built over the years (e.g, (Feiner, MacIntyre, & Seligmann, 1993)). We are not aware of any other attempts to create a continuously worn, general purpose interface to a spatialized information space such as the RWWW. The most relevant spatialized AR systems to our work are Audio Aura (Mynatt, Back, Want, M., & Ellis, 1998) and the Touring Machine (Feiner et al., 1997). Audio Aura is an audio-only AR system designed to provide knowledge workers with a continuous sense of generally relevant, context-sensitive, information (e.g., the state of their E-mail, the activity of their coworkers). The Touring Machine (and other tour systems) exhibits some of the characteristics we desire, but it is designed for a specific kind of data (tours of outdoor sites). In addition, the interface design assumes the user's primary task is taking the tour, and are focussed on the data being presented. Therefore, the Touring Machine interface is not subject to the same constraints we illustrate in Section 3.

Like AR, wearable computing researchers have explored both task specific systems (e.g., (Thompson, Ockerman, Najjar, & Rogers, 1997)) and general purpose, continuous use systems (e.g., (Rhodes, 1997)). However, unlike AR, general purpose, continuous-use systems have received significant attention in the wearable community (e.g., (Starner et al., 1997)). The key difference between AR and typical wearable computers is information spatialization: current wearables are not capable of complex 3D graphics, so they generally use 2D heads-up displays (HUDs) which place the information in a fixed location in 2D. This has the advantage that the user can predict where information will and will not appear, and can adapt accordingly. However, it has the disadvantage that information is limited to a small space and cannot be associated directly with relevant parts of the world. This limits the users ability to take advantage of their spatial memory and perception, and significantly decreases the amount of information that can be displayed.

Finally, our work can be viewed as a context-aware system. It has close similarities to context-aware tour guides (Long, Kooper, Abowd, & Atkeson, 1996), which attempt to present information to the user based on some notion of their current context. Most context aware systems that we know about use hand-held displays, which are suitable because they are aimed at more task-specific, often query-reply, interactions. Unlike these systems, we are interested in presenting the information to the user continuously, and therefore use a see-through, head-worn display. However, many of the issues related to defining, obtaining, and acting on contextual information are common across all of these systems.

3 Assumptions and Design Goals

As the RWWW evolves out of the WWW, it will not suddenly "appear", but will rather be integrated with the existing WWW for the foreseeable future; some WWW pages will have detailed context associated with them, some will have limited contextual information, and some will have none at all. The availability of more complex contextual information will begin to emerge soon, on the tails of initiatives (such as Microsoft's .NET) that are pushing for semantically meaningful structuring of information on the WWW. The corollary to this gradual evolution is that the RWWW will begin to appear "any time now". Given the current interest in creating context-aware applications for mobile phones and palm-sized devices, a rudimentary form of the RWWW is already being developed, with the initial contextual tagging being simple location information. As geo-spatial tagging standards gain acceptance, sites that already allow location-based queries (such as MapQuest and Restaurant.com) will become far more common.

As the RWWW grows, centralized indices (such as google.com), authored portals (such as yahoo.com) and information services (such as Fodors Restaurant Guide) will adapt and allow information to be retrieved automatically, based not only on content, but also on standard contextual cues (such as location or person identity).

WorldBoard's idea of channels is closest of to today's authored portals, but we believe all three kinds of sites will be important: automatically created indices that crawl the web like today's search engines will be necessary to ensure that as much information as possible is made available to the user, based on their current context, and authored portals and information services will be necessary to ensure authoritative access to common kinds of information.

The implications of these assumptions about the RWWW on any interface to the RWWW are:

- *continuously changing data:* The combination of context awareness and widely distributed servers implies that data may change at any time, and may also be changing continuously.
- *safety:* The distributed authorship of the RWWW implies that both trustworthy and untrustworthy data can be mixed together. Directly displaying unfiltered information from the RWWW may be dangerous (e.g. a malicious person could place a virtual brick wall on a busy highway).
- *heterogeneous data*: The gradual evolution of the RWWW implies that old and new (i.e., tagged and untagged) data will be mixed together for the forseeable future. This implies a need to handle a mixture of location-based, contextually-tagged and untagged data.

As described above, the most significant design choice is that the RWWW Browser is intended to be worn continuously, as in (Starner, 1999). This decision implies two interface characteristics, common to most continuously-available user interfaces:

- *non-interference:* The continuous availability of the system implies that the interface should support continuous awareness of the virtual space without interfering with other tasks the wearer may be engaged in.
- *minimal interaction*: The user should rarely (if ever) be required to interact with the system when they do not want to. While continuously worn systems will require occasional interaction with the user, these interactions should require minimal effort.

Satisfying the first two implications, continuously changing data and safety, is vital for a continuous-use systems; if the system were designed to be worn briefly from time to time, these points would be interesting observations, but might have little practical impact on the interface design.

The fundamental differences in interface design between task-specific and continuous use AR systems are analogous to the differences between applications designed for desktop computers (such as a word processor) and ambient displays (such as the Ambient Room (Wisneski et al., 1998)); it is easy to turn away from a desktop display or iconify the application, but it is a far more significant act for a user to leave an Ambient Room if it is interfering with a task they are trying to accomplish. For a RWWW Browser (and the Ambient Room) to be successful, it must present the user with a continuous awareness of their information space, and do it in such a way that does not adversely impact their focal tasks.

Based on the assumption described above, we have come up with a list of design goals for our initial prototype RWWW Browser.

1. The default state of the browser is to display information nodes as simple, consistent icons, instead of presenting the "raw," unfiltered content from the RWWW spatially in the 3D world. This avoids visual confusion, as well as badly timed or maliciously-placed content (*non-interference, minimal interaction, safety, continuously changing data*).
2. The user should be able to organize the information at a coarse level, so they can quickly control what information is displayed (*non-interference, minimal interaction, safety*).
3. Unless the user explicitly requests otherwise, detailed information (especially if it is to remain displayed for any length of time) should be placed in a fixed location relative to the user (such as relative to their head or body). The user can then predict where information will be (*non-interference, minimal interaction*).
4. The user should be able to easily access and dismiss different levels of detail of the information space, ranging from the basic information about a node to the full details of all nodes in the space (*minimal interaction*).
5. The browser should not significantly change the content of the display based on context or location changes. Instead, the browser should peripherally notify the user about potentially significant changes in the information space and provide them with a simple way to explore the new content (*non-interference, safety, continuously changing data*).
6. Non-spatialized data should be displayed at reasonable locations in the environment (*heterogeneous data*).

4 The Initial prototype RWWW Browser

Based on the assumptions and goals stated in the previous section, we have developed an initial interface for our browser that focuses on providing an awareness of the amount of information available, with a coarse indication of the information content (Figure 1). First, web pages are organized by grouping them into related information channels (following (Spohrer, 1999)). The user can control which channels are displayed, giving them fast (but coarse) control of the amount of information in their visual space. Second, unlike previous work where detailed information such as labels are placed in the world (e.g., (Feiner et al., 1997)), web pages are displayed using spatially-located anchors (currently, small "twinkling stars" that are colored to indicate which channel they belong to). This approach gives the user an indication of the amount, location and kind of information available, while minimizing visual clutter. Information that does not have detailed spatial location information, such as a page associated with a time or activity, is automatically positioned in space at a reasonable location by the browser.

Additional information about an anchor is obtained using a simple form of two-level gaze selection, extending the approach we developed in (Feiner et al., 1997) (i.e., the object closest to the center of the screen is selected). The first level (*glance*) immediately displays the title of the selected anchor near the center of the screen, replaces the anchor with a thumbnail rendering of the web page, and begins to draw a red circle clockwise around the thumbnail being glanced at. The second level (*gaze*) is activated by gazing at the same anchor for a few seconds (the time it takes to draw the complete red circle). Gazing causes a yellow circle drawn around the selected thumbnail, and a larger rendering to be placed in a fixed location in the upper left corner of the screen. This approach allows an area to be scanned quickly to obtain more detailed information. Detailed information is always placed at a fixed location (the upper left corner of the display in Figure 1 and Figure 2) so that users can predict where information will appear, and therefore locate it or work around it when necessary. A node remains gaze-selected until a new node is gaze-selected.

To support the situation where a user needs to simultaneously view multiple nodes, or simply wishes to retain a node for future reference, we have added a temporary storage area in the user's body space by allowing the gaze-selected page to be "pushed" to the right around the user's waist. When the user pushes a node, all pages in their body space are rotated to the right (both the saved pages, and the gaze-selected page). To distinguish between nodes that are gaze-selected and those that are saved in the body space, we draw a green circle around the nodes saved in the body space (see Figure 2).

The system alerts the user of significant changes to the set of "relevant information" by flashing the channel indicator along the bottom of the display, but new channels are not displayed without an explicit user action. Currently, our Browser can handle any traditional 2D web page that can be rendered by the Java HTML renderer.

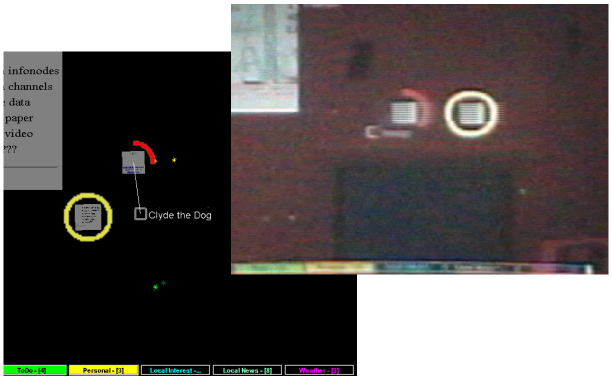

Figure 1 The prototype RWWW browser. For clarity, a screen-capture from the HMD is displayed on the left, in addition to an image taken through the HMD on the right. The link closest to the center has a text description and small thumbnail (10cm in world coordinates) displayed, connected by a leader line. A red circle is slowly drawn around the thumbnail (over the course of a few seconds). After a short gaze (the time taken to draw the red circle), a larger image of the web page appears in the upper left corner of the screen, and the circle turns yellow to indicate which node is gaze-selected and displayed in the upper left.

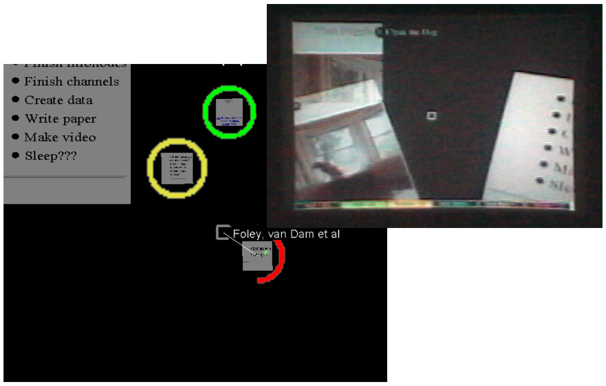

Figure 2 The small image to the right shows what the user sees when they look down and slightly to the right. The node that is currently gaze-selected (the picture) is to the left, and another node (a checklist) has been saved and is to the right. The large image is a screen shot showing what would be displayed on the HMD if the user looked straight ahead at the information nodes. The green circle indicates the saved node, the yellow circle indicates the gaze-selected node.

5 Discussion and Future Work

In this paper we list five important issues that must be addressed when creating a continuously worn interface to an information space like the RWWW: non-interference, minimal interaction, safety, continuously changing data, and heterogeneous data. These five issues are then used to develop a list of design goals to which we believe a RWWW browser should adhere. During the design and implementation of our prototype browser, initial tests with other researchers and visitors to our lab aided in the iterative design of the interface, and suggested that our design principles are reasonable.

We believe that coarse organization in channels, as is done currently in the browser, is a useful start point for organizing the large information space. However, channels alone are not sufficient without the use of contextual information (such as location, time, activity or the identities of nearby people) to prompt the user as to which channels might be most relevant to their current context. We are currently integrating context-sensing infrastructure, developed by others in our research group (Salber et al., 1999), into our prototype to allow us to begin experimenting with the use of context, both to select appropriate channels, and to prioritize the information in the channels. We are also beginning to investigate how knowledge of which physical and virtual objects the user is looking at should affect the display of the information space.

In addition to experimenting with the Browser itself, we are creating a simplified RWWW infrastructure that supports a richer set of content types (such as 3D objects and audio), adds contextual meta-data to each document, and provides the Browser with additional information such as models of the physical environment (to support more controlled spatialized display). These models will not only contain a coarse layout of the room (where walls and windows are) but can also contain hints as to where non-spatialized information should be placed.

In this paper, we have discussed our first steps toward implementing a Real-World Wide Web browser. We have developed a set of guiding principles for our experiments, and implemented an initial prototype of a Browser. We hope to take the lessons we have learned, and take the next step toward the eventual deployment of a RWWW Browser.

6 Acknowledgments

This work was supported by Siemens via a GVU Industrial Affiliate Grant, DARPA/ITO under the Information Technology Expeditions, Ubiquitous Computing, Quorum, and PCES programs, ONR under Grant N000140010361, and equipment and software donations from Sun Microsystems and Microsoft.

7 References

Abowd, G. D., Dey, A. K., Brotherton, J., & Orr, R. J. (1999). Context-awareness in Wearable and Ubiquitous Computing. *Virtual Reality Society International Journal, 3*, 200-211.

Feiner, S., MacIntyre, B., Hollerer, T., & Webster, T. (1997). A Touring Machine: Prototyping 3D mobile augmented reality 28 systems for exploring the urban environment. *Personal Technology, 1*(4), 208-217.

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

Long, S., Kooper, R., Abowd, G. D., & Atkeson, C. (1996). *Rapid Prototyping of Mobile Context-Aware Applications: The CyberGuide Case Study.* Paper presented at the MobiCom.

Mynatt, E. D., Back, M. J., Want, W., M., B., & Ellis, J. B. (1998). *Designing Audio Aura.* Paper presented at the Proc. of ACM Computer Human Interaction Conference (CHI '98).

Rhodes, B. J. (1997). *The Wearable Remembrance Agent: A system for Augmented Memory.* Paper presented at the ISWC '97.

Salber, D., Dey, A. K., & Abowd, G. D. (1999). *The Context Toolkit: Aiding the Development of Context-Enabled Applications.* Paper presented at the Proc. of ACM Computer Human Interaction Conference (CHI '99).

Spohrer, J. C. (1999). Information in places. *Systems Journal, 38*(4), 602-628.

Starner, T. (1999). *Wearable Computing and Context Awareness.* Unpublished PhD thesis, MIT, Cambridge, MA.

Starner, T., Mann, S., Rhodes, B., Levine, J., Healey, J., Kirsch, D., Picard, R., & A., P. (1997). Augmented Reality Through Wearable Computing. *Presence: Teleoperators and Virtual Environments, 6*(4).

Thompson, C., Ockerman, J. J., Najjar, L., & Rogers, E. (1997). *Factory Automation Support Technology (FAST): A New Paradigm of Continuous Learning and Support Using a Wearable.* Paper presented at the ISWC '97.

Wisneski, C., Ishii, H., Dahley, A., Gorbet, M., Brave, S., Ullmer, B., & Yarin, P. (1998). *Ambient Displays: Turning Architectual Space into an Interface between People and Digital Information.* Paper presented at the CoBuild '98.