

The Real-World Wide Web Browser: An Interface for a Continuously Available, General Purpose, Spatialized Information Space

Rob Kooper and Blair MacIntyre

*College of Computing, GVI Center
Georgia Institute of Technology
Atlanta, GA, USA*

{blair, kooper}@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 present the assumptions we make about the characteristics of such a system, and discuss the implications of those assumptions for an AR interface. In particular, we focus on the implications of continuous use, context-awareness, and distributed publishing.

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

1 Introduction

The idea of a wearable augmented reality (AR) system that envelops the mobile user in a spatially-registered 3D information space is an exciting one, and has been experimented with, in different forms, by many people over the years. While some systems have been designed with general-purpose information spaces in mind [9, 25, 26], most researchers have focused on task-specific systems [e.g., 2, 6, 7, 10, 18]. This is understandable, both because a significant set of AR applications are task-specific (i.e. maintenance, repair and training), and because many of the hard problems that must be solved when creating an AR system (e.g., tracking, registration, and interaction) become tractable in the context of a specific application or task domain [e.g., 7, 18].

Despite the difficulties in creating a general-purpose AR system, we are interested in how to use AR as an interface to a general purpose, spatially-located, global information space. We envision a future where head-worn displays are as light and unobtrusive as a pair of glasses and can therefore be worn continuously, allowing users to be immersed in an augmented information space as they move through their daily lives. In addition to presenting information tied to a user's specific location (e.g., related to machinery, restaurants, or people near the

user), such a system can be used to provide any information that might be pertinent to the user's context (e.g., calendar reminders, waiting times at restaurants if it is near dinner time, real-time messages from friends, or reference documents related to current activities).

The goal of our research, using a spatialized AR information space to display context sensitive information, can be traced back to the writings of Bush and Engelbart on augmenting human capabilities by enabling new forms of information access and processing [4, 8]. The larger goal of our research is to merge personal AR devices¹, such as head-worn auditory and visual displays, with other "off-the-desktop" interaction paradigms, including ubiquitous computing [29], context-aware computing [1] and intelligent environments [3]. We refer to these computational environments collectively as augmented environments (AEs) and believe that personal AR devices will play a key role in their success. Among other benefits, personal AR devices will allow hands-free, personalized interfaces to these complex, 3D information spaces.

Much of the proposed benefit of AEs derives from the possibility of using ubiquitous computation, communication and sensing to enable the presentation of context-sensitive information and services to the user. One significant feature of such systems is their ability to generate a wide variety of potentially useful information and present it to the user in a way that leverages on their ability to synthesize and extract useful data [17, 29, 30]. Since personal AR devices can be worn (and therefore perceived) continuously, we can place information at spatially relevant locations in the environment and can tailor the information to the wearer, such that the AR devices present the widest variety of information to an individual.

1. We refer to head-worn displays as person displays because they present information that can only be perceived by the wearer. This is in contrast to other possible AR displays, such as projectors [19], that can be perceived simultaneously by multiple people.

The central elements of this paper are twofold. First, in Section 3, we will describe the characteristics we assume a general purpose spatialized information system will exhibit and summarize the implications of those characteristics on the human computer interface. Second, in Section 4, we will present a set of interface guidelines and design goals for a continuously worn, spatialized AR interface to such a space, based on the characteristics presented in Section 3, and also present and discuss our current design of a Real World Wide Web browser.

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 [26], 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 [1]. 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 [26]. 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 believe in extending 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. In WorldBoard, only geo-spatial location is used as context; information is placed anywhere in the world, in 1 meter cubes and specified in UTM coordinates [27]. Furthermore, 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. [2, 6, 7, 9, 10, 11, 18, 25, 26]). 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. As we will discuss in Section 4.1, there are significant differences in how task-specific (i.e. worn to augment the user’s primary activity) and general-purpose (i.e. often not related to the user’s primary activity) interfaces are designed. The most relevant spatialized AR systems to this work are Audio Aura [17] and the Touring Machine [10]. Audio Aura is an audio-only AR system designed to provide knowledge workers with a continuous sense of generally relevant information (e.g., the state of their E-mail, the activity of their coworkers), combined with context-sensitive information (e.g. is the person whose office we are looking at in today? How long is the printer queue for the printer we are near?). Through careful design, Audio Aura created an audio space that could be worn continuously. 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), does not support non-location-based concepts of context, and requires the use of a handheld tablet for interaction and to display much of its data.

Both Audio Aura and the Touring Machine can also be thought of as ambient, or peripheral, displays [30]. Their intent was to present information in such a way that it resides in the periphery of the user’s attention, and becomes integrated with their awareness of their surroundings. Similarly, our goal is to present a sense of the information available to the user in their current context in a peripheral or ambient fashion. However, unlike other ambient displays, we also want to allow users to easily “dig down” and obtain more detailed information in the same display space.

Like AR, wearable computing researchers have explored both task specific systems (e.g., [28]) and general purpose, continuous use systems (e.g., [21]). However, unlike AR, general purpose, continuous-use systems have received significant attention in the wearable community (e.g., [24]). 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 place 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 [14, 15], 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 works 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 Implications

3.1 The RWWW: Systems Assumptions

As mentioned above, we believe that the current WWW will evolve over time into the RWWW, as more and more of the information available on web servers has some kind of context associated with it. The RWWW will likely retain its structure as a loosely organized, heterogeneous collection of information servers, but will serve up semantically richer data, and take advantage of new kinds of servers to help structure the information space. We also believe that, because of the often serendipitous content on the Web, the RWWW will be most useful if it is presented continuously to the user.

If the RWWW evolves out of the WWW, it will not suddenly “appear”, but will rather be embedded in the existing notion of the 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 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 (such as GML [12] or WBXML [31]) gain acceptance, sites that already allow location-based queries (such as MapQuest [16] and Restaurant.com [20]) will become far more common.

The availability of more complex contextual information will also emerge soon, on the tails of initiatives (such as Microsoft’s .NET) that are pushing for semantically meaningful structuring of information on the WWW. As more and more of the new information on the web is tagged in semantically meaningful ways, tools will emerge that help us semi-automatically associate meta-data (stored on our personal servers) with information on the web, such as time (e.g., when did I access this data, when should I be reminded of it), people (e.g., who was I near when this data was created, who is

it associated with), projects (e.g., what data is related to what, what was I doing when it was created), and places (e.g., where was I when I created this data); such meta-data yet another type of contextual information.

As the RWWW grows, centralized indices (such as google.com) and authored portals (such as yahoo.com, restaurant.com and citysearch.com) 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). (WorldBoard’s idea of channels is closest of to today’s authored portals.) Both 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 will be necessary to ensure authoritative access to common kinds of information. For example, as I walk through downtown New York, an index may return an eclectic collection of information about the current neighborhood, published by individuals around the world (based on their own visits, cultural criticisms, etc.). In contrast an authored portal, such as restaurant.com, will return current, consistent, and (relatively) trustworthy information about local restaurants. Alternatively, a tour portal might provide a high-level guide through the neighborhood, complemented by a mixture of the other 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. we could place a brick wall on the road).
- *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 foreseeable future. This implies a need to handle a mixture of location-based, contextually-tagged and untagged data.

We also assume that, based on the state of current tracking technology, that at least coarse location information will be available of the user, and that accurate head and body orientation will be available.

3.2 The RWWW Browser: Interface Assumptions

The most significant design choice we have made is that the RWWW Browser is intended to be worn continuously. 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 doing.
- *minimal interaction*: The user should rarely (if ever) be required to interact with the system when they do not want to. However, continuously worn systems will still require occasional interaction with the user; these interactions should require minimal effort.

The first two implications in the previous section, continuously changing data and safety, are also significantly more important for 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.

In the context of a continuously worn system, each of these assumptions have significant design implications. In particular, the browser interface must behave fundamentally different than the interfaces for task-specific AR systems. When designing a task-specific interface, such as a maintenance and repair system [11], we implicitly assume that the user is interested in the data being displayed because it is relevant to their current task. When the display is no longer relevant, the user can simply turn off or remove the AR system.

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 [30]); 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 them when they are focused on other tasks.

4 A RWWW Browser

As mentioned above, we envision a RWWW Browser that immerses the user in a spatialized information space, merged with their perception of the real world. Since it is continuously worn, the design will need to integrate all the implications laid out in the previous section.

We will eventually include more than just head-worn visual displays in the RWWW browser interface, but for we initially focus on the use of see-through HMDs. as the primary display medium. We expect that when we do incorporate multiple displays and other input/output modalities, we will have a greater range of display and interaction options, such as using audio for some forms of peripheral awareness and choosing which display to

use for a given task based on the affordances of the device.

In the two following subsections, we will present our design principles for a RWWW Browser, and then the prototype implementation we have built to explore this problem.

4.1 Design Principles for a RWWW Browser

Our initial prototype, is based on the following design principles, which derive from the implications listed in the sections above:

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, and can adapt accordingly (*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*).

4.2 A RWWW Browser Prototype

In this section, we will present our prototype of a RWWW browser based on the design principles in the previous section. Our browser currently supports only HTML content, allowing us to use an existing HTML rendering engine for display. (We are working on adding support for 3D and audio content.)

The browser allows the user to look around the environment and see an indication of the available information nodes, represented as either small twinkling stars (*link awareness*, the default mode) or small, fixed-

size thumbnails (*thumbnail awareness*, for easier browsing). The star view gives the user a quick feel for how much information is available, and where it is located, without cluttering the space (principle 1).

To help the user organize information, we have adapted the channel notion from WorldBoard. As in WorldBoard, a channel provides a logical organization for the RWWW content; individual channels can be turned on and off to control what information is displayed (principle 2). Like WorldBoard, our channels will often contain authored information, such as a restaurant guide or building information. However, our channels are also used as the basis for controlling context-based access to the RWWW by allowing the channels to encapsulate a dynamic search parameterized by the current context. For example, a channel could be defined to create a node for each person in your current environment (a common form of context) that contains information obtained from an appropriate search engine. In this prototype, information nodes can contain a web page or a set of context keywords which in turn map to a web page. We currently use Google [13] to map the context keywords to a relevant web page, but will eventually use richer mechanisms.

As mentioned above, the default view through the HMD places small twinkling stars at the spatial location of each information node. Each currently available channel (see Section 4.3 for a discussion of how channels are made available) is listed along the bottom of the display (principle 3). To provide a lightweight method of obtaining additional information about a node, we have implemented a simple approximation to two-level gaze selection: a glance-select and a gaze-select (principle 4). At any time, if any nodes are close to the center of the display (indicated by a small square target in the center of the display), the closest node will change from a star to a small rendering of the node. Additionally, a description of the node will be attached to the square target in the center of the screen, and a leader line will connect the two. For example, the node with the red arc around it in Figure 1 has just been glance-selected. As a user pans their view around a scene, the glance-selected node is changed immediately when a new node is closer to the center of the display.

Once a node is glance-selected, a 2 second countdown begins, as indicated by the proportion of the red circle displayed around the node. When the red circle is complete, the information node is gaze-selected; a yellow circle is drawn around the information node to indicate the gaze-selection status.

The selected information node will be gaze-selected, and have the yellow circle drawn around it, until another node is gaze-selected. When a node is gaze-selected, a larger version of the information node will be displayed in the upper left corner of the screen, with the small descriptive text rendered next to it, as shown in Figure 2



Figure 1. The browser in default *link awareness* mode (for clarity, a screen-capture from the HMD is displayed on the left, in addition to a small image taken through the HMD on the right). If any links are near the center of the display, the one closest to the center has its description displayed. The link changes to a small thumbnail (10cm in world coordinates) to provide the user with more information about its content, and a countdown begins till it is gaze-selected.



Figure 2. When an item is gaze-selected, its link is highlighted and the link description plus a larger thumbnail are attached to the HUD.

(principle 3). At the same time, a significantly larger rendering of the node is positioned in the user's *body space* (a space centered in the middle of the users body at their waist, which is oriented in the direction their body is facing), directly in front of their waist, as shown in Figure 3. By looking down, the user can read the information node that has been gaze-selected, and it will remain suspended in front of them until a new node is gaze-selected (principle 4).

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 the "push" button, 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




Figure 3. When an item is gaze-selected, the detailed view of that item is automatically attached to a temporary holding location in the user's body space (directly in front of the user's waist). In this image, the user is looking down toward the floor in front of them to see the larger image of the node.




Figure 4. The small image to the left 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.

those that are saved in the body space, we draw a green circle around the nodes saved in the body space (see Figure 4).

To aid the user when they are looking for a particular piece of information, they can toggle the browser in and out of *thumbnail awareness mode* (see Figure 5). In this mode the browser will render all the information in small thumbnails, rather than as stars. These thumbnails will not be large enough for users to read the pages content, but are large enough that a familiar page may be visually recognizable.

Each channel icon displays the channel name and a number indicating the number of nodes currently in that channel. Our goal is to use context to decide which channels are presented to the user (principle 5), allowing the user the ability to change which channel is displayed



Figure 5. The browser in *thumbnail awareness mode*. The nodes appear as thumbnails, to give the user a greater ability to scan a space for something they are interested in.

and explicitly toggle channels on and off. When the system (or the channel itself) determines that a channel's content has changed significantly or contains especially important information (i.e. the weather channel informing the user of a thunderstorm in the region), the system adds the channel to, or removes it from, the channel bar. Furthermore, the system must alert the user that a channel may be of interest; currently, we cause the channel button to flash.

The final action a user can perform in our prototype is to control the content of the descriptive text that is displayed when a node is glanced at: the user can toggle between the description of the node, and the URL of the information. We believe this may help users to decide how trustworthy the information is.

To satisfy principle 6 (non-spatialized data should be displayed in a reasonable location), we are experimenting with the idea of spatial information servers that can provide, for any space, locations to display non-spatial data. For example, in our lab, we could display non spatial information on a virtual bulletin board on a blank area of the wall.

We have designed the interface to require six buttons to operate: one button to cycle right through the available channels, one button to toggle a channel on or off, one pair of buttons to push and pop items into and out of the body space, one button to toggle between awareness modes, and one button to switch between URL display and descriptive text display.

4.3 System Overview

Our RWWW browser implementation is built using Java and Java3D. The browser itself runs on a Sun workstation and uses a stereo Sony Glasstron for display. The user's head and body are tracked using an InterSense IS-600 tracker.

The application is split into two separate programs: the first is the prototype browser, and the second is a Wizard of Oz interface that allows an observer to generate

spatialized data, context data and portals. The browser maintains a list of channels and the information nodes associated with each channel.

We use an existing HTML rendering engine for the information nodes, so our browser can display any content that this engine can render, taking the 2D output and texture mapping it onto polygons in the 3D world. The information nodes encapsulate the different representations, including stars and thumbnails. The thumbnail is the content of the URL, rendered as a 32 x 32 pixel texturemap on a 10cm x 10cm polygon. The information nodes are also responsible for translating the set of context keywords to a URL. Finally, the title of the web page is used as the descriptive text for each information node.

We have experimented with different-sized thumbnails, including having thumbnails that change size as a function of their distance from the user. More distant thumbnails were rendered at larger sizes, so that farther thumbnails appeared to shrink more slowly than they otherwise would. Our hope was that we would then be able to possibly see the thumbnails of more distant objects: at 10cm x 10cm, the thumbnails quickly become too small to see. However, when we changed the size in this way, users who tried the system reported two problems. First, it became very hard to judge the relative location of thumbnails. Second, the thumbnails appeared to move away from the users as they tried to approach them (the size cue overrode the stereo cue).

The Wizard of Oz interface provides the browser with the names of the channels, which channels are selected, which channels should attract the user's attention, the current representation of the information nodes are, what information nodes exist in each channel and where information nodes are to be placed. This Wizard of Oz interface also mimics the controller as a set of six buttons.

5 Discussion and Future Work

In this paper we list five implications of a continuously worn interface to an information space like the RWWW: non-interference, minimal interaction, safety, continuously changing data, and heterogeneous data. These five implications are then used to develop a list of principles to which a RWWW browser should adhere.

During the design and implementation of our prototype browser, initial tests with other researchers and visitors to our lab suggested that our design principles are reasonable, and aided in the iterative design of the interface itself.

We are currently investigating possible interaction devices. Although the Twiddler is a good device for textual input, we believe a simpler device would be more appropriate, providing quick access to each browser function. It is likely that we will use a range of input

devices and modalities (speech, buttons, gestures) that support the interactive tasks we define in different situations. For example, a belt worn button controller may not be appropriate while driving a car, since the users hands will be busy, but such a controller may be ideal when walking in a noisy environment.

Contextual information is currently entered by hand through the Wizard of Oz interface. We are exploring the use of systems that will be able to provide the browser with dynamically generated context (i.e. [5, 22]). We are also interested in exploring environments with a dense sensing infrastructure (such as a lab environment) and those with a sparse infrastructure (such as an outdoor environment that we do not control).

Some of these environments will have additional displays available for browser display. To minimize the amount of information displayed on the HMD, we are interested in taking advantage of these additional displays and moving some information to them from the HMD. Attempting to do this will raise interesting privacy and security issues (which displays are we allowed to use, and who can see them, for example).

We have only begun to address the issue of how to place non-spatialized information in the AR space in a reasonable way. One of the ideas we are pursuing is to introduce the notion of *physical space servers*, which would provide information about the physical environment. Such servers could contain pre-defined databases (containing models of the space, specifications of who owns what areas, etc.) but could also be generated dynamically (e.g., as described in [23]). Building on the character of the WWW, a set of these servers would allow anyone to define the characteristics of the physical environment, such as marking a specific area (e.g., a blank part of the wall) as a bulletin board on which the RWWW Browser should place non-spatialized data. Other non-spatialized data, like to-do notes and names, may be attached to the user's *body space*, or to related people and objects. This implies a rich, dynamic database of contextual and spatial data.

As described in Section 4.1, we currently use the WorldBoard notion of channels to bring order to the potentially huge amounts of available information. Eventually, we will need to explore how mobile users can create, destroy and edit channels in a reasonable way. This is especially true with channels that are defined as context-based searches: users may want to be able to split, merge or otherwise manipulate the information sets generated by the current context. In addition to being able to control the content of a channel, users should be able to control how the information nodes are displayed. For example, a user may decide that their personal to-do channel should always display the raw node content in space, rather than as stars or thumbnails.

For such a browser to be truly lightweight and usable, context must be used not only to generate content for

channels, but control the interface. Property 5 (in Section 4.1) states that context should not be used to actively change the content of display. However, as with all guidelines, there are cases when users may want to do otherwise. For example, when the user is in their car, moving quickly, the system should notice that many of the links (especially those close to the user) would be moving so fast on the display that showing them would be significantly distracting and provide little information. However, information attached to them, or to their car, would still be useful.

We have described a system of gaze selection based on the center of the display. We use this simplification based on the lack of eye tracking in the prototype system. However, if we had accurate eye tracking with depth information, we would probably want the gaze selection to work differently: currently, because the “gaze point” is attached to a fixed location on the screen, it is something that users can use predictably and adapt to. For example, if they are talking to someone with links around them, they can turn their head slightly and thus not have the links around their companion constantly activating. A true eye tracking system would not allow this. However, it would allow the system to tell what user is looking at, including differentiating between looking at the display and looking at the environment. This would allow a new set of interaction techniques, such as having information closer than the user’s gaze location fade out because the user is looking “through” it.

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

The authors would like to acknowledge the members of the Augmented Environments Lab and Future Computing Environments Group at Georgia Tech for their influence on this work. This work was supported by Siemens via a GVU Industrial Affiliate Grant, and ONR under Grant N000140010361.

7 References

- [1] Abowd, G.D., Dey, A.K., Brotherton, J., and Orr, R.J. “Context-awareness in Wearable and Ubiquitous Computing”, In *Virtual Reality Society International Journal*, Vol. 3, 1999, pp. 200-211.
- [2] Billinghurst, M., Bowskill, J., Jessop, M., and Morphet, J. “A Wearable Spatial Conferencing Space”, In *ISWC '98*, pp. 76-93.
- [3] Bobick, A., Intille, S., Davis, J., Baird, F., Pinhanez, C., Campbell, L., Ivanov, Y., Schütte, and A., Wilson, A. “The KidsRoom: A Perceptually-Based Interactive and Immersive Story Environment” In *PRESENCE*, 8(4), pp. 367-391.
- [4] Bush, V. “As We May Think” In *The Atlantic Monthly* 176 No. 1, 101-108 (July 1945).
- [5] Cooltown, www.cooltown.hpl.hp.com
- [6] Cater, J.P., and Huffman, S.D. “Use of Remote Access Virtual Environment Network (RAVEN) for Coordinated IVA-EVA Astronaut Training and Evaluation”, In *PRESENCE*, 4(2), pp. 103-109.
- [7] Curtis, D., Mizell, D., Gruenbaum, P., and Janin, A. “Several Devils in the Details: Making AR App Work in the Airplane Factory”, In *IWAR '98*, 1998.
- [8] Engelbart, D.C. “Augmenting Human Intellect: A Conceptual Framework”, Stanford Research Institute, Report AFOSR-3233 (October 1962)
- [9] Feiner, S., MacIntyre, B., Haupt, M. and Solomon, E.. “Windows on The World: 2D Windows for 3D Augmented Reality”, In *UIST '93*, pp. 145-155.
- [10] Feiner, S., MacIntyre, B., Hollerer, T., and Webster, T. “A Touring Machine: Prototyping 3D mobile augmented reality 28 systems for exploring the urban environment”, In *Personal Technology*, 1(4): 208-217, 1997.
- [11] Feiner, S., MacIntyre, B. and Seligmann, D. “Knowledge-Based Augmented Reality”, In *Communications of the ACM*, 36(7), pp. 53-62, July, 1993.
- [12] Geography Markup Language (GML), www.w3.org/Mobile/posdep/Presentations/GMLIntro/index.htm
- [13] Google, www.google.com.
- [14] GUIDE project, www.comp.lancs.ac.uk/computing/research/mpg/most/guide.html.
- [15] Long, S., Kooper, R., Abowd, G.D., and Atkeson, C. “Rapid Prototyping of Mobile Context-Aware Applications: The CyberGuide Case Study”, *MobiCom*, 1996.
- [16] MapQuest, www.mapquest.com.
- [17] Mynatt, E.D. Back, M., Want, W., Baer M., and Ellis, J.B. “Designing audio aura”, In *CHI '98*, pp. 566-573.
- [18] Navab, N., Bascle, B., Appel, M., and Cubillo, E. “Scene augmentation via the fusion of industrial drawings and uncalibrated images with a view to marker-less calibration”, In *IWAR '99*, 1999.
- [19] Raskar, R., Welch, G., and Fuchs, H. “SAR: Spatially Augmented Reality”, In *IWAR '98*, 1998.
- [20] Restaurant.com, www.restaurant.com.
- [21] Rhodes, B.J. “The Wearable Remembrance Agent: A system for Augmented Memory”, In *ISWC '97*, pp. 146-155.
- [22] Salber, D., Dey, A.K., and Abowd, G.D. “The Context Toolkit: Aiding the Development of Context-Enabled Applications”, In *CHI '99*, pp. 434-441.
- [23] Shafer, S., Krumm, J., Brumitt, B., Meyers, B., Czerwinski, M., and Robbins, D. “The New EasyLiving Project at Microsoft Research”, In *1998 DARPA / NIST Smart Spaces Workshop*, pp.7-127 - 7-130.
- [24] Starner, T., Mann, S., Rhodes, B., Levine, J., Healey, J., Kirsch, D., Picard, R., and Pentland A. “Augmented Reality Through Wearable Computing” In *PRESENCE*, 6(4), 1997.
- [25] Sutherland, I., “A Head-Mounted Three-Dimensional Display”, In *Fall Joint Computer Conference*, Thompson Books, Washington, D.C. (1968), pp. 757-764.
- [26] Spohrer, J.C. “Information in places”, In *Systems Journal*, vol. 38, pp. 602-628, 1999.
- [27] Terry, N.G. “How to Read the Universal Transverse Mercator (UTM) Grid”, In *GPS World*, 32, 1996.
- [28] Thompson, C., Ockerman, J.J., Najjar, L., and Rogers, E. “Factory Automation Support Technology (FAST): A New Paradigm of Continuous Learning and Support Using a Wearable”, In *ISWC '97*, pp. 31-38
- [29] Weiser, M. “Some Computer Science Issues in Ubiquitous Computing”, In *Communications of the ACM*, 36(7), pp. 75-84 (July 1993).
- [30] Wisneski, C., Ishii, H., Dahley, A., Gorbet, M., Brave, S., Ullmer, B. and Yarin, P. “Ambient Displays: Turning Architectural Space into an Interface between People and Digital Information”, In *CoBuild '98*, pp. 22-32.
- [31] WorldBoard XML (WBXML), www.worldboard.org.