Rapid Prototyping of Mobile Context-Aware Applications: The Cyberguide Case Study

Sue Long, Rob Kooper,
Gregory D. Abowd, and Christopher G. Atkeson
College of Computing
Georgia Institute of Technology
Atlanta, GA 30332-0280
{suelong,kooper,abowd,cga}@cc.gatech.edu
http://www.cc.gatech.edu/fce

ABSTRACT

We present the Cyberguide project, in which we are building prototypes of a mobile context-aware tour guide that provide information to a tourist based on knowledge of position and orientation. We describe features of existing Cyberguide prototypes and discuss research issues that have emerged in our context-aware applications development in a mobile environment.

Keywords: Mobile computing applications, context-awareness, location-dependent applications, hand-held devices

1 Introduction

The project we report on in this paper, Cyberguide, has as its main focus the rapid prototyping of hand-held mobile applications in order to assess the utility of context-awareness in mobile devices. The challenge we are addressing in Cyberguide is how to build mobile applications that make use of the context of the user. Initially, we are concerned with only a small part of the user's context, specifically location and orientation.

The application which drives the development of Cyberguide is a that of a general tour guide. More specifically, the initial prototypes of Cyberguide were designed to assist visitors in a tour of the GVU (Graphics, Visualization and Usability) Center during our monthly open houses. Visitors to a GVU open house are typically given a map of the various labs and an information packet describing all of the projects that are being demonstrated at various sites. In building Cyberguide, we wanted to support the tasks of the visitor to the GVU open house. Moving all of the paper-based information into a hand-held intelligent tour guide that knows where you are and what you are looking at and can answer typical visitor questions provides a testbed for research questions on mobile, context-aware application development.

Our short-term goal was to prototype versions of Cyberguide on commercially available PDAs and pen-based PCs. Context awareness initially meant the current physical position and orientation of the Cyberguide unit (and since it is hand-held, this locates the user as well). Position information improves the utility of a tour guide application. As the prototypes of Cyberguide evolve, we will be able to handle more of the user's context, such as where she has been, what was seen and heard there, as well as where others are and have been.

1.1 Overview

In this paper, we discuss the evolution of the Cyberguide design and prototype as well as what future research areas our experience has uncovered. We begin by describing scenarios for the use of context-aware mobile applications. We then compare our research to other related work. Next we outline a conceptual design of Cyberguide followed by a description of some of our prototypes. A discussion of issues that arose during our prototype development follows. We conclude with an outline of our plans for future work.

To appear in the Proceedings of the 2^{nd} *ACM International Conference on Mobile Computing And Networking, 1996 (MobiCom'96).*

2 Scenarios For a Mobile Context-Aware Application

This section outlines some possible uses for future mobile context aware devices. Some of these uses are currently being implemented and some are futuristic. The ideal hand-held device will have a screen and pen/finger interface, access to substantial storage resources, possibly through an internal device such as a CD drive, or through substantial communication and networking resources (cell phone, pager, data radio interface) providing access to other storage servers (such as the Web), an audio input and output interface with speech generation and potentially sophisticated voice recognition, and a video input and output interface. The video input (a video camera) could be pointed at the user to interpret user gestures, or pointed at the environment to interpret objects or symbols in the environment. The video output could be integrated into the main screen or be a separate video display device, such as an attached screen or heads up display on glasses worn by the user.

One major application of mobile context aware devices are personal guides. Museums could provide these devices and allow users to take personalized tours seeing any exhibits desired in any order, in contrast to today's taped tours. Walking tours of cities or historical sites could be assisted by electronic guidebooks. The hand-held devices could use position measurement systems such as indoor beacons or the Global Positioning System (GPS) to locate the user, and an electronic compass or inertial navigation system to find user orientation. Objects of interest could be marked with visual markers or active beacons or recognized using computer vision. Some objects, such as animals at a zoo or aquarium, might be difficult to mark but could be recognized with simple computer vision and some assistance from the environment (indications that this is the elephant cage, for example). The personal guide could also assist in route planning and providing directions. Some of these functions are currently being provided by automobile based navigation systems. We describe our implementation of a mobile tour guide for our "open houses" in a later section.

There are other ways to assist users. Consider a traveler in Japan that does not speak or read Japanese. The hand-held device could act as a pocket multilingual dictionary, actually speaking the appropriate phrase with the appropriate pronunciation to a taxi driver, for example (or even showing the appropriate Kanji and an associated map on the screen). A device that included video input or a scanner could assist in reading signs or menus. A device that could show stored images might be able to show a shopkeeper the desired object or favorite meal. Another more futuristic use is to assist the user by recognizing faces at a cocktail party and reminding the user who people are.

Real time communication allows a personal device to act as an agent for the user. A personal guide to Disneyworld could make reservations at particular rides, and alert the user when the reservation was available. The device could also tell the user which rides had the shortest lines. Similar approaches are currently being used for automobile traffic management in Atlanta.

An important application of context aware devices is enhanced reality. A heads up display could provide "X-ray" vision for the user. While surveying a building for renovation, the location of hidden plumbing or electrical conduits could be indicated to the user, based on information from sensors and/or building plans. At an archeological site a visitor could be provided with various overlays indicating what used to be above the current ground level as well as what is below the current ground level.

Context-aware devices can also be used as tools. Simple sonar devices are used to make room measurements today. It would not take much to have a hand-held device that both videotaped and mapped a room along with user commentary. An ecological field study or an archeological dig could be assisted by a device that automatically recorded the context of a particular find, including noting the surrounding objects. Consider an electronic field guide that assisted the user in recognizing plants or insects.

One of the most interesting applications of context aware devices is to support group interaction on a tour or in a classroom, for example. Participants in a live demonstration of some new technology could use their personal device to help steer the demo using majority voting or consensus among the viewers. Each participant could run a personalized version of the same demo by expressing their own choices. In this case context is which demo a participant is participating in or attending to, and the personal machine may switch to another context if it detects the user is attending to that context instead.

3 Related Work

In thinking about and developing a location-aware application, we were greatly influenced by the work on the PARCTab at Xerox PARC [10], the InfoPad project at Berkeley [5], the Olivetti Active Badge system [9] and the Personal Shopping Assistant proposed at AT&T [2]. We wanted to build useful applications that might take advantage of the hardware developed in the PARCTab and InfoPad projects. We did not want to build our own hardware, so we have a different focus from all of these projects. There are a number of commercially available and relatively inexpensive hand-held units that would suffice for our purposes, such as the Apple MessagePad with the Newton Operating System[1], a MagicCap[2] machine or a pen-based palmtop/tablet PC. We chose to work initially with the Apple MessagePad 100 with Newton 1.3 and pen-based PCs running Windows for Pen Computing 1.0.

For positioning, we considered the Active Badge system, but rejected it for reasons of cost and long-term objectives. The Active Badge system combines position detection with communication. For room-level granularity of position, this is reasonable since the communications range is on par with the position resolution. With Cyberguide, it is not clear that positioning and communication systems should always share physical resources. Certain versions of our prototype did; other prototypes did not. We provided for the separation of the wireless communications capabilities from the positioning system, so we could seek out more cost-effective solutions for both.

We tried to pay attention to the higher level conceptual design of Cyberguide, but we have not been as general in our handling of context-aware mobile objects as has Schilit [8].

4 Architecture of Cyberguide

From the beginning, we have been focusing on developing a family of systems and not just a single application, so it made sense to think about a conceptual design, or architecture, that captured the essence of Cyberguide. We have divided the system into several independent components, or modules, and have found it useful to present those components in terms of the people a tourist would like to have available while exploring unfamiliar territory. The overall system serves as a tour guide, but we can think of a tour guide as playing the role of cartographer, librarian, navigator and messenger. The services provided by these components are:

- **Cartographer (Mapping)** This person has intimate knowledge of the physical surroundings, such as the location of buildings, interesting sights within a building, or pathways that the tourist can access. This component is realized in our systems by a map (or maps) of the physical environments that the tourist is visiting.
- **Librarian (Information)** This person provides access to all of the information about sights that a tourist might encounter during their visit. This would include descriptions of buildings or other interesting sights and the identities of people associated with the areas. The librarian can answer specific question about certain sights (“Who works in that building?” or “What artist painted that picture?” or “What other demonstrations are related to what I am looking at?”). This component is realized as a structured repository of information relating to objects and people of interest in the physical world.
- **Navigator (Positioning)** The interests of the tourist lie relatively close to their physical location. Therefore, it is important to know exactly where the tourist is, in order to show the immediate surroundings on the map or answer questions about those surroundings (“What am I looking at?”). The navigator is responsible for charting the location of the tourist within the physical surroundings. This component is realized by a positioning module that delivers accurate information on tourist location and orientation.
- **Messenger (Communications)** A tourist will want to send and receive information, and so the messenger provides a delivery service. For example, when visiting an exhibit or demonstration, the tourist might want to speak with the owner of the exhibit. If the owner is not present, the tourist can leave a message. In order to find out where other tourists are located, each tourist can communicate her current location to some central service that others can access. It might also be desirable to broadcast information to a set of tourists (“The bus will be leaving from the departure point in 15 minutes.”). This component is realized as a set of (wireless) communications services.

The utility of this architectural decomposition for Cyberguide is that it provides an extensible and modular approach to system development. It is extensible because we can always add further services. For example, we have considered adding an historian whose purpose is to document where the tourist has been and what her reactions were to the things she saw. It is modular because it has allowed us to change the implementation of one component of the system with minimal impact on the rest of the system. For example, we have implemented different versions of the navigator and the librarian without having to alter the other components. Of course, these components are related in some ways; for instance, position information ultimately has to be translated into a location on the physical map. Defining standard interfaces between the components is the means by which we achieve separation between and coordination among the various components.

---

[1] MessagePad and the Newton Operating System are registered trademarks of Apple Computer, Inc.

[2] MagicCap is a registered trademark of General Magic, Inc.

5 The Indoor Cyberguide

In this section, we describe how each of the separate modules in the conceptual architecture have evolved in the indoor version of Cyberguide on the Apple MessagePad.

5.1 Map Module

Developed for use by vistors to GVU open houses, the map module, shown on the left side of Figure 1, contains a map of the entire GVU Center. Passageways and demonstration stations (stars in Figure 1) are shown. Only a limited view of the lab can be seen at any given time. The user can scroll the map around and zoom in and out to see alternative views. There is an icon to show the user's location on the map. Using information from the positioning module, we implemented automatic scrolling of the map. If desired, the user's position is updated automatically and the map is scrolled to ensure that the user's current position remains on the visible portion of the map.

5.2 Information Module

The information module (shown on the right side of Figure 1) contains information about each of the demos on display at GVU open house. This includes abstracts of the project being demoed, background information on those involved with the project, as well as where to get further information. The location of each demo is marked on the map by a star. The user selects the star icon for a demo to reveal its name. Selecting the name brings up the information page for that demo. The user can also go directly into the information module and search for information for specific demo pages either by category or by project name.

One version of the information module was hard-coded, providing very fast response but requiring a recompilation every time demo information needed to be updated. Another implementation used Newton files, called soups, to store information. The use of soups avoided hard-coding data into the application and simplified demo information updates, but did not have adequate response time. Our third implementation of the information module used Newton Books, the Newton platform documentation viewer, to store the demo information. The use of Newton Books improved our access time considerably, allowing for an automated information update process without requiring data be hard-coded directly into the application. Throughout all three versions of our information module, we were able to modify the information module independent of the development efforts of the other modules, validating the modularity of part of our design.

5.3 Communication Module

Our initial implementation of a communication module consisted of a wired Internet Appletalk connection from a Apple MessagePad through a Unix Appletalk Gateway. Although there are future plans to release TCP/IP for the Newton platform, currently it is limited to Appletalk network connections. We built an application level protocol on top of a public domain implementation of the Appletalk protocol for Solaris[4]. This allows us to open a connection-based Appletalk stream from the Apple MessagePad to a UNIX platform. We then invoked our gateway application to repacketize Appletalk packets into TCP/IP packets for transmission over the Internet. This allowed for TCP/IP connectivity from a Apple MessagePad via an Appletalk connection. We could then fetch HTML documents as well as send and receive e-mail. We utilized this functionality within Cyberguide by providing a questionnaire for users to complete, which was sent to the developers as an e-mail message. (see Figure 2)

5.4 Position Module

Position is the obvious starting point for a context aware mobile device. We considered several methods for sensing the user location. The outdoor version can use GPS. Indoors GPS signals are weak or not available. We considered RF for indoor position measurement, but found off the shelf solutions too expensive.

One solution for an indoor positioning system was to use infrared (IR). Our first positioning system was based on using TV remote control units as active beacons, and using a special IR receiver tuned to the carrier frequency ($\approx$ 40kHz) of those beacons (Figure 3). A microcontroller (Motorola 68332) interfaced the IR receiver to the serial port of the Apple MessagePad. We deployed an array of remote controls hanging from the ceiling (Figure 3 right), each remote control acting as a position beacon by repeatedly beaming out a unique pattern. The 68332 translates the IR pattern into a unique cell identifier that is sent to the Apple MessagePad's serial port. As the tourist moves around the room and passes into the range of a new cell, the position (indicated by an arrowhead) is updated on the map. Keeping track of the last recorded cell location provides a good guess as to the location the tourist is heading, so we indicate an assumed orientation by pointing the position icon accordingly.

The remote control system is too expensive for large scale use as the cost of the 68332 microcontroller is added to the cost of the Newton. Our second positioning system combines position measurement with communication in using the built in Newton IR port. Two ways to utilize the IR port are to use the proprietary Newton IR protocol or to use RS232. The proprietary

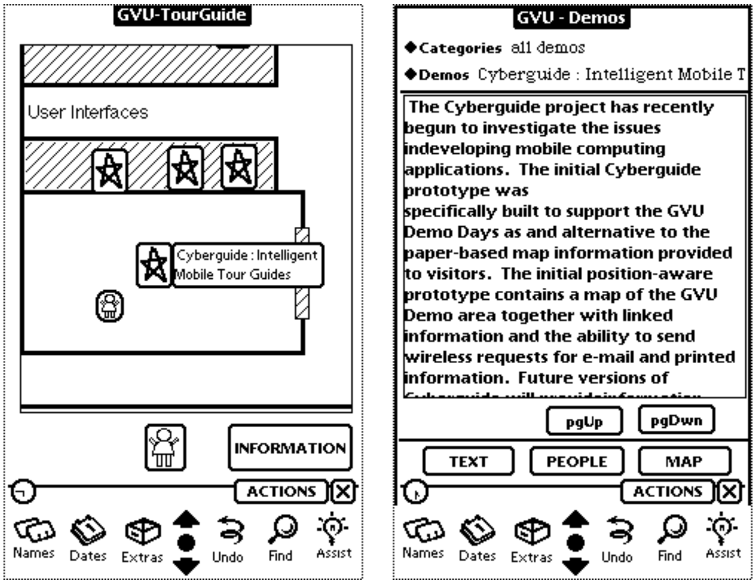

The Cyberguide project has recently begun to investigate the issues indeveloping mobile computing applications. The initial Cyberguide prototype was
specifically built to support the GVU Demo Days as and alternative to the paper-based map information provided to visitors. The initial position-aware prototype contains a map of the GVU Demo area together with linked information and the ability to send wireless requests for e-mail and printed information. Future versions of

Figure 1: Newton screen dumps: map view on left, information view on right.

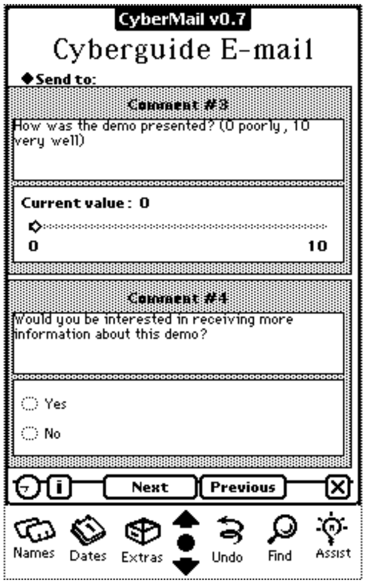

How was the demo presented? (0 poorly, 10 very well)

Would you be interested in receiving more information about this demo?

○ Yes

○ No

Figure 2: Questionnaire using communications module for delivery.

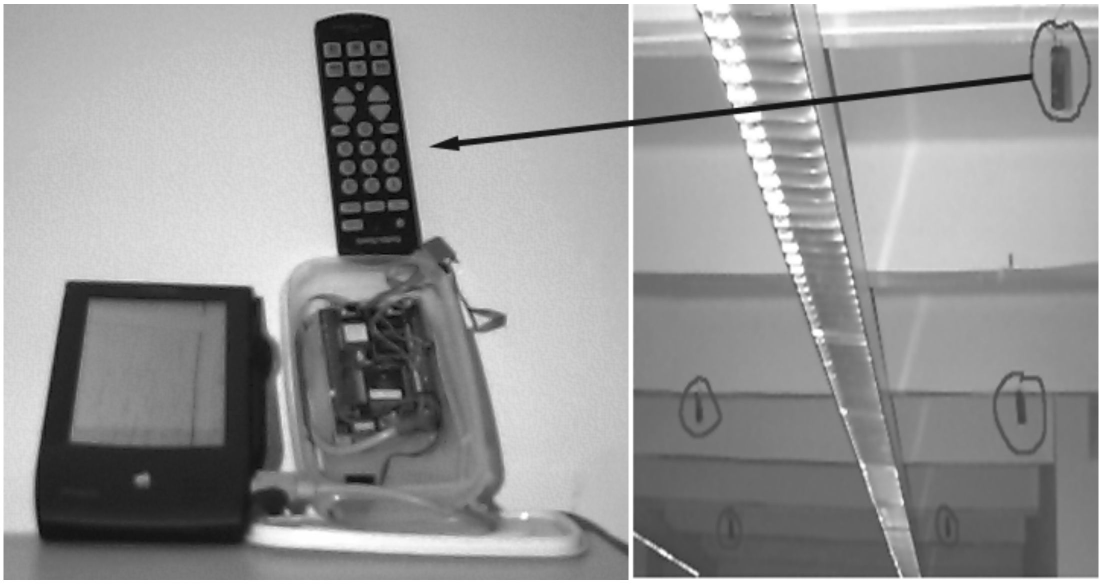

Figure 3: IR positioning prototype (left) and the array of positioning beacons in the GVU Lab (right).

Newton IR protocol not robust in the presence of multiple senders and receivers during mobile use. We used the IR port in serial IR mode where access to the IR port is through a half duplex connection encoding data using raw RS232. The serial IR mode is appropriate for our positioning system, but has a range of only three feet.

We are currently implementing a new IR system based on the same IR units used inside the Apple MessagePad (from Sharp) for beacons and communication stations. These units relay data from a serial interface on a networked workstation out through the IR interface using RS232 encoding. The networked workstations control the data that is passed between IR units. Instead of a secondary processor, pen-based PCs will use an IR transceiver attached to their serial port and Apple MessagePads will use their built in IR port.

With these IR units in place, we will also be able to implement crude low range wireless communications. We are no longer limited to sending single beacon numbers, but can send relatively large packets of information. The drawbacks are that our range is only three feet and we are communicating over an RS232 connection without error detection or correction. Given these limitations, we plan to implement relatively simple transmissions and broadcasts.

6 The Outdoor Cyberguide

There were several motivations for building a Cyberguide prototype for outdoor use (Figure 4). First, we wanted to use Cyberguide over a wider area than the relatively small GVU Center. We also wanted to test the modularity of our design by having to change critical features. The two features that were changed on this prototype were the underlying map and the physical positioning system. We obtained a different map and inserted that into the map module without any problems. For positioning, we replaced the IR positioning module with a Trimble GPS unit attached to the Apple MessagePad serial port. (see right side of Figure 4). The GPS unit sends a position in latitude and longitude which was then translated into a pixel coordinate representing the user's current position on the map.

The outdoor positioning system has been tested by two prototypes. We first built a proof of concept tour of the Georgia Tech campus (shown in Figure 4). We are currently working on a more full-scale outdoor Cyberguide, CyBARguide, to guide a tourist in pursuit of refreshment at neighborhood establishments in Atlanta. This involves handling multiple maps, querying of a large amount of information and some minimal routing facilities. We also plan to make the data within the information module modifiable so the user can add personalized information including personal impressions that may be useful for future reference, a type of virtual grafitti.

7 PC Implementations

In order to verify the platform independence of our conceptual design, we initiated two separate efforts build-

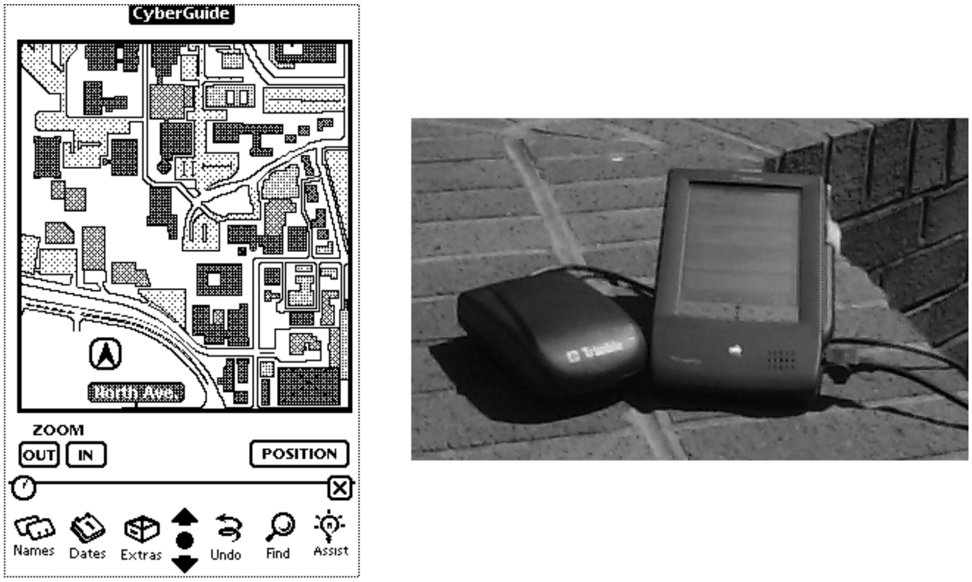

Figure 4: The outdoor Cyberguide (left) with GPS unit (right).

ing pen-based PC versions of Cyberguide. These limited functionality PC versions were written using Borland's Delphi environment and Microsoft's Visual Basic. Both were installed on Dauphin DTR-1 palmtops running Pen for Windows Computing 1.0. The functionality of these Cyberguides are summarized below.

The Visual Basic prototype implemented the map, positioning and information modules, but not the communication module. A commercial web browser is used as the information module. The positioning module is identical to the indoor IR module on the Apple MessagePad. The map module displays maps, like the Apple MessagePad, using a bitmap image on which additional icons are displayed.

The Delphi version has implemented the map and information module. Web pages containing demo information are stored locally as database objects using a stand-alone Borland database engine. Information is viewed using a public domain Delphi HTML viewer. Though faster than the Visual Basic counterpart, it is a long-term disadvantage to have the information base stored locally. In contrast to both the Apple MessagePad and Visual Basic versions, the Delphi version uses vector based maps. This allows for arbitrary scaling and rotation of the map and well as path generation (not implemented). Currently there is no communication module and only a manual positioning capability, similar to the first Apple MessagePad prototype.

8 Issues

Our experience over the last year developing versions of Cyberguide with different features on different platforms has given us a certain amount of insight into the important issues in developing mobile, context-aware applications. In this section, we summarize some of those issues.

Our prototyping efforts were iterative, modifying both hardware and software to improve functionality. Our primary focus, to assess the impact of mobile technology for a specific task, necessitated rapid and inexpensive prototyping. This motivated the use of inexpensive commercially available hardware. When choosing our hardware platform, we considered several mobile handheld devices before deciding on the Apple MessagePad. One of the biggest obstacles was the requirement to program the Apple MessagePad in its own proprietary language. However, given the price and product availability we developed all three iterations of our initial prototype on the Apple MessagePad.

Over the last nine months we have used Cyberguide for many GVU open houses. We gathered informal feedback through informal surveys, formal questionnaires and informal user comments. We also observed visitors as they tried to use Cyberguide to maneuver around the lab. During each iteration we incorporated the user feedback and our own reactions to what was good and what was bad into the next iteration on the design. While our major focus was to prototype a context-aware mobile application rapidly, we realize that little can be determined concerning the impact of such technology unless the technology is put in the hands of real users.

8.1 Coupling of Positioning and Communication

There is an interesting relationship between the positioning and communication systems. Systems such as PARCTab and the Active Badge rely on a close coupling of positioning and communication. This is because the location of any mobile entity is determined by the beacon which receives a communication from the entity. In Cyberguide, the indoor and outdoor positioning system worked by having the beacon inform the mobile entity where it was located. The disadvantage of this latter approach for Cyberguide is that only the mobile entity knows its location. For applications in which you want objects to know about the position of other objects, there must be some sort of communication. However, it can be impossible or undesirable to couple positioning and communication together. For example, if position is coming from GPS, then a separate means of communication must be used. In our current version of indoor IR positioning using the Sharp IR units, we can couple positioning and communication, but the range of the IR link is so limited (3 feet) that communication will be cumbersome. It makes sense to use a short range IR positioning system because position information can be localized to objects of interest. Communication, on the other hand, needs to be uniform throughout some space.

8.2 Communication Medium

We have been trying to implement communications services on commercial hand-held units. While communication is important to Cyberguide, there is no obvious appropriate choice for a wireless communication medium to suit our needs. It clearly is not a priority among the manufacturers of these units to provide high bandwidth, cost-effective wireless communications. There are many potential solutions for communications (IR, spread-spectrum RF, cellular packet, cellular modem), but the variety and quality of these services changes so much that manufacturers tend not to build communications devices into the units, for fear of premature obsolescence. Instead, they rely on third party communications solutions with standard interfaces (e.g., PCMCIA). In our experience, there is a need to have high bandwidth communication to the mobile unit and low bandwidth connection back to the network. We do not assume that the hand-held unit will always carry around with it the entire information base or map associated with the area the tourist is visiting. Rather, that information should be provided on demand and relative to the position/orientation of the tourist.

8.3 Map Representation

We have experimented with bitmap and vector-based maps. The bitmap representation is easy to obtain (scanning) for any area and is relatively inexpensive to store and display. Scaling and rotation, however, are cumbersome with that representation. Since we were decorating the map to highlight places of interest that were not on the original bitmap, it was difficult to control the display of the decorations after scaling for a zoom in or out. Another problem with the bitmap representation is accuracy with respect to the real-world. In the outdoor version of Cyberguide, we noticed a drift on the positioning system for a certain region of the map attributed to the map itself being out of scale. Also, a bitmap representation is not suited to doing higher level map services, such as generating a path to direct the tourist to a location of interest.

A vector-based representation, on the other hand, is easier to handle in terms of manipulation and additional services such as way-finding, but was not a feasible solution on the interpreted platforms (Visual Basic and Apple MessagePad 100 with Newton 1.3 using the Newton Toolkit 1.5) because it was computationally overwhelming to manage the display (this may be solved with compiling capabilities of later versions of the Newton Toolkit). It is also more difficult to obtain a vector-based map for a large and detailed area. For example, when we built the Delphi prototype, it was not difficult to build a map tool to construct the vector-based map, but it would take a very long time to create a map of the GVU Lab with the detail we already had in the bitmap version. For outdoor use, there are already commercially available structured map databases for large areas, and these are being used in navigational systems in rental cars. However, the size of the map database prohibits local storage on hand-held units, so there is an even stronger argument for high bandwidth downstream wireless connectivity.

8.4 Cross-Platform Issues

We developed prototypes on multiple platforms to validate our platform independent conceptual design. It was encouraging to see the same indoor positioning system work for both the Apple MessagePad and Visual Basic prototypes. We will soon have this working for the Delphi platform as well. The latest version of the indoor positioning system will work for any device with a serial port. Outdoor positioning is similarly cross-platform, relying on either a serial or PCMCIA interface.

For communication, we also see the need to standardize the interface and protocol. For information services, it is natural to want support for wireless TCP/IP to enable full Internet capability. The information browsing can then be treated as a Web browsing task, for example. There are efforts already to support TCP/IP for platforms such as the Newton platform, but the bandwidth does not yet support delivery of complex graphics, as we would need for map delivery.

9 Future Work

We have worked for a year on Cyberguide. We have always tried to keep in mind the very long term goals of this kind of application, and many of those ideas have been expressed by others and by us in Section 2. The following is a list of features that we feel are most important for near-term future versions of mobile, context-aware applications such as Cyberguide.

9.1 Modifiable Information Base

As a tourist visits a place, she may read about prepared information, but she may also have her own thoughts and reactions to what she sees or may overhear someone else interpret an exhibit in an interesting way that she would like to record. Capturing relevant information along the way and adding that to the information base would be useful. There has been some recent work on capturing information at meetings [6] or in the classroom [1, 7] and it seems a natural extension to add this kind of capture facility for the tourist.

9.2 Increased Communications

Cyberguide is communications poor, and that has prevented us from doing more interesting activities, such as broadcasting messages, informing a single user about other tourists, sending messages between tourists and moving the information sources off of the hand-held unit. This is a high priority in our future work and we will most likely investigate third party wireless modems with TCP/IP support on the hand-held units. Since we are currently limited to what is in place from commercial vendors, we might now take the time to investigate research platforms such as the InfoPad.

9.3 Improved Context Awareness

One way to view the capturing activity above is as a way to augment the system's understanding of the tourist's context by remembering what was interesting. We currently have a very limited notion of context in Cyberguide — physical location and crude orientation. We have experimented with capturing historical context (what sights have been visited already), but there are a number of other aspects of the tourist's context that can be useful. For instance, knowing where everyone else is located might suggest places of potential interest. Knowing the tourist's reaction to some exhibits would help in suggesting other related places of interest. Being aware of time of day or day of week may aid in more intelligent user queries of interesting attractions or activities. We feel that the secret to context-awareness is in doing it behind the scenes. The more that can be automatically captured and turned into context, the better. If the user has to explicitly inform the system about context information ("I am currently located here." or "This exhibit is boring to me." or "The museum is currently closed.") then the context is unlikely to be fully utilized.

9.4 Leveraging the Web

In the touring example, we found that much of the information we wanted to display was already available on the Web. Web browsing is now a natural information browsing metaphor. It makes sense to leverage this available information resource and mode of interaction for all of the information needs. The map module then provides hooks or links to the information on the Web, and it is delivered on demand, similar to when a user selects a URL on a browser. There have already been some research [3] and commercial attempts at providing mobile Web browsers.

9.5 Use of Vision

In the extreme case, we can think of a communion between the physical and electronic worlds, as suggested by work in augmented reality. Replace the hand-held unit with a pair of goggles and as the user wanders around, information about certain exhibits can be summoned up and overlaid on top of the actual image. Vision techniques can be used to augment the positioning system to inform the system and tourist what the tourist is looking at. We have experimented with vision systems as an extension to Cyberguide. Ultimately, we want to move toward personalized vision systems.

9.6 Ubiquitous Positioning System

Our current prototypes are exclusively indoor or outdoor, not both. This is mainly because we had no one positioning system that worked in both conditions. GPS is unreliable indoors and the IR-based beacon system is impractical for us to implement outdoors. We intend to integrate both positioning systems into one application, to allow a tourist to wander in and out of buildings and have Cyberguide automatically switch the positioning system.

9.7 Increased Multimedia Support

In the context of the GVU open houses, visitors often gather around a demo and stretch to watch the activity on a desktop machine as the researcher describes the related research. We would like to have Cyberguide units in the vicinity of the demonstration be able to pick up a live feed from the demonstration for display on their own unit. Initially, this could be achieved by a simple mirroring of the demonstration machine's display onto the hand-held unit. Ultimately, we envision the visitor "plugging in" to the demonstration and being able to control it from their hand-held unit.

10 Conclusions

We have described research we have been conducting over the past year prototyping mobile, context-aware applications. The main focus of our development effort has been commercially available hardware platforms and with the specific application focus of a tour guide. We described several iterations on the Cyberguide application, built to support tours through various venues. Our experience prototyping several variations of Cyberguide on various platforms has raised some significant research issues in the continued development of mobile, context-aware applications.

Acknowledgments

The Cyberguide project is very much a group development effort. The authors would like to acknowledge the support of the several members of the Future Computing Environments Group at Georgia Tech, particularly Dietmar Aust, without whom this project would surely have failed. Gregory Abowd would like to acknowledge the numerous undergraduate students in his software engineering courses who put forth effort in defining and building many of the prototype systems. Cyberguide was initially supported by a seed grant from Peter Freeman, Dean of the College of Computing at Georgia Tech. Subsequent effort on the project was support by BellSouth Intelliventures through the Industrial Affiliates Program in the GVU Center at Georgia Tech. Some of our equipment has been loaned to us from Apple Computer, Inc. and we would like to thank Jim Spohrer in particular for his continued support of our work.

Sue Long now works for Wink Communications in Alameda, California. She can be contacted via e-mail as sue.long@wink.com.

References

[1] G. D. Abowd, C. Atkeson, A. Feinstein, C. Hmelo, R. Kooper, S. Long, N. Sawhney, and M. Tan. Teaching and learning as multimedia authoring: The classroom 2000 project. In *Proceedings of the ACM Conference on Multimedia — Multimedia'96*, 1996. To appear.

[2] A. Asthana, M. Cravatts, and P. Krzyzanouski. An indoor wireless system for personalized shopping assistance. In L.-F. Cabrera and M. Sattyanarayanan, editors, *Workshop on Mobile Computing Systems and Applications*, pages 69–74. IEEE Computer Society Press, December 1994.

[3] J. F. Bartlett. W4 — the wireless World Wide Web. In L.-F. Cabrera and M. Sattyanarayanan, editors, *Workshop on Mobile Computing Systems and Applications*, pages 176–178. IEEE Computer Society Press, December 1994.

[4] B. Krupczak, M. Ammar, and K. Calvert. Multi-subsystem protocol architectures: Motivation and experience with an adaptor-based approach. In *Proceedings of IEEE INFOCOM'96*, 1996.

[5] S. Long, D. Aust, G. D. Abowd, and C. G. Atkeson. Rapid prototyping of mobile context-aware applications: The cyberguide case study. In *Proceedings of the 1995 conference on Human Factors in Computing Systems — CHI'96*, 1996. Short paper.

[6] A. C. Long, Jr., S. Narayanaswamy, A. Burstein, R. Han, K. Lutz, B. Richards, S. Sheng, R. W. Brodersen, and J. Rabaey. A prototype user interface for a mobile multimedia terminal. In *Proceedings of the 1995 conference on Human Factors in Computing Systems — CHI'95*, 1995. Interactive experience demonstration.

[7] S. Minneman, S. Harrison, B. Janseen, G. Kurtenbach, T. Moran, I. Smith, and B. van Melle. A confederation of tools for capturing and accessing collaborative activity. In *Proceedings of the ACM Conference on Multimedia — Multimedia'95*, November 1995.

[8] N. N. Sawhney, G. D. Abowd, and C. Atkeson. Can electronic notebooks enhance the classroom? Technical Report GIT–96–06, Graphics, Visualization and Usability Center, Georgia Institute of Technology, February 1996.

[9] W. N. Schilit. *System architecture for context-aware mobile computing*. PhD thesis, Columbia University, 1995.

[10] R. Want, A. Hopper, V. Falcao, and J. Gibbons. The active badge location system. *ACM Transactions on Information Systems*, 10(1):91–102, January 1992.

[11] R. Want, B. Schilit, N. Adams, R. Gold, K. Petersen, J. Ellis, D. Goldberg, and M. Weiser. The PARCTAB ubiquitous computing experiment. Technical Report CSL-95-1, Xerox Palo Alto Research Center, March 1995.