

Cyberguide: A mobile context-aware tour guide

Gregory D. Abowd^a, Christopher G. Atkeson^a, Jason Hong^a, Sue Long^{a,b}, Rob Kooper^a and Mike Pinkerton^a

^a Graphics, Visualization and Usability Centre, College of Computing, Georgia Institute of Technology, Atlanta, GA 30332-0280, USA

^b Wink Communications, Alameda, CA 94501, USA

Future computing environments will free the user from the constraints of the desktop. Applications for a mobile environment should take advantage of contextual information, such as position, to offer greater services to the user. In this paper, we present the Cyberguide project, in which we are building prototypes of a mobile context-aware tour guide. Knowledge of the user's current location, as well as a history of past locations, are used to provide more of the kind of services that we come to expect from a real tour guide. We describe the architecture and features of a variety of Cyberguide prototypes developed for indoor and outdoor use on a number of different hand-held platforms. We also discuss the general research issues that have emerged in our context-aware applications development in a mobile environment.

1. Introduction

Future computing environments promise to free the user from the constraints of stationary desktop computing, yet relatively few researchers are investigating what applications maximally benefit from mobility. Current use of mobile technology shows a slow evolution from our current desktop paradigm of computing, but the history of interaction shows that the adoption of new technology usually brings about a radical revolution in the way humans use and view technology [11]. Whereas the effective use of mobile technology will give rise to an interaction paradigm shift, it is difficult to predict what that shift will be. We follow the advice of Alan Kay, therefore, and choose to predict the future by inventing it. Our approach is to think first about what activities could be best supported by mobile technology and then determine how the technology would have to work. This applications focus is important to distinguishing our work in mobile computing.

In April 1995, we formed the Future Computing Environments (FCE) Group within the College of Computing and the Graphics, Visualization and Usability (GVU Center) at Georgia Tech to promote such an applications focus. Our group is committed to the rapid prototyping of applications that benefit from the use of emerging mobile and ubiquitous computing technologies. Quick development of these futuristic applications allows us to predict and shape what our everyday lives will be like when today's novel technology becomes commonplace.

Applications for a mobile environment should take advantage of contextual information, such as position, to offer greater services to the user. In this paper, we present the Cyberguide project, a series of prototypes of a mobile, hand-held context-aware tour guide. Initially, we are concerned with only a small part of the user's context, specifically location and orientation. Knowledge of the user's current location, as well as a history of past locations, are used to provide more of the kind of services that we come

to expect from a real tour guide. We describe the architecture and features of a variety of Cyberguide prototypes developed for indoor and outdoor use on a number of different hand-held platforms. We also discuss the general research issues that have emerged in our experience of developing context-aware applications in a mobile environment. Some of these research issues overlap with those that we have considered in applying other applications of ubiquitous computing technology.

The general application domain which has driven the development of Cyberguide is tourism, but we have found it necessary to be even more focused in our research. The initial prototypes of Cyberguide, therefore, were designed to assist a very specific kind of tourist – a visitor in a tour of the GVU Center Lab 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. Moving all of the paper-based information into a hand-held, position-aware unit provided a testbed for research questions on mobile, context-aware application development.

The long-term goal is an application that knows where the tourist is, what she is looking at, can predict and answer questions she might pose, and provide the ability to interact with other people and the environment. Our short-term goal was to prototype versions of Cyberguide on commercially available PDAs and pen-based PCs in which context-awareness simply 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 have been able to handle more of the user's context, such as where she and others have been, and we have increased the amount in which the tourist can interact and communicate with the place and people she is visiting.

1.1. Overview

This paper is an extended version of an earlier report on Cyberguide [7], we discuss the evolution of the Cyberguide design and prototype as well as what future research areas our experience has uncovered. We begin in section 2 by describing scenarios for the use of context-aware mobile applications. In section 3, we provide context for our research within the area of applications-centered mobile computing. The generic architecture of Cyberguide is explained in section 4. We will describe in section 5 the initial realization of the generic components of the Cyberguide architecture, a series of prototypes developed for the Apple MessagePad. We will then describe in section 6 how the initial indoor prototypes were extended for use outdoors and for greater interaction with the environment. We conclude in sections 7 and 8 with a discussion of significant issues for context-aware applications development and how our past experience will influence our future development plans.

2. Scenarios for a mobile context-aware application

This section outlines some possible uses for future mobile context-aware applications. Some of these uses are currently being implemented and some are futuristic. We begin with our initial assumptions about what technology we expect Cyberguide to use. Tourists are usually quite happy to carry around a book that describes the location they are visiting, so a reasonable packaging would be in the form of a hand-held device. 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. In fact, many museums now provide portable devices for just such a purpose, but what we are envisioning is a device that would allow the tourist to go anywhere she pleases and be able to receive information about anywhere she is. Walking tours of cities or historical sites could be assisted by these 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 on-board navigation systems.

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 a theme park 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 major cities.

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.

Many tourists take records of some sort of their traveling experiences, either by taking pictures or videos or by composing a travel diary. Imagine the possibilities if the recording of these experiences could be more efficiently and accurately recorded. A drive across the country could result in a trail superimposed upon a map, and clicking on the trail would reveal an image of what you could see at that moment – an automatically-generated spatial index into your memories.

These are but a few of the possibilities we can imagine that a context-aware application can provide for the tourist. We have investigated many of these possibilities already and report on them later.

3. Related work

In thinking about and developing a location-aware application, we were greatly influenced by work such as the PARCTab at Xerox PARC [15], the InfoPad project at Berkeley [8], the Olivetti Active Badge system [14] and the Personal Shopping Assistant proposed at AT&T [3]. 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 System¹, a MagicCap² 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. Each platform was available for \$150–500 with relatively powerful development environments. This low cost of hardware was critical to the success of Cyberguide because it made it possible to put a number of units in the hands of many students, all with different ideas that they were allowed to investigate.

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.

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

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

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 [13].

4. Architecture of Cyberguide

From the beginning, we have viewed Cyberguide as a family of prototypes and not just a single prototype, so it made sense to think about a conceptual design, or architecture, that captured the essence of the mobile tour guide. We have divided the system into several independent components, or building, and have found it useful to present those components both in terms of the generic function and personified 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 (map component)*. 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 component)*. 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 component)*. 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 component)*. 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




Figure 1. The map (left) and information (right) interfaces of the initial Cyberguide MessagePad prototype.

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.

5. The indoor Cyberguide

In this section, we describe how each of the separate modules in the conceptual architecture have been realized in the initial series of prototypes developed on the Apple MessagePad for use indoors during GVV open houses.

5.1. Map component

The initial map module, shown on the left side of figure 1, contains a map of the entire GVV 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 component

The information module (shown on the right side of figure 1) contains information about each of the demos on display at the GVV 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

Figure 2. Questionnaire using communications module for delivery.

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 component

Our initial implementation of a communication module consisted of a wired Internet Appletalk connection from an Apple MessagePad through a Unix Appletalk Gateway. We designed 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 repackage Appletalk packets into TCP/IP packets for transmission over the Internet. This allowed for TCP/IP connectivity from an 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 component

Position is the obvious starting point for a context-aware mobile device. We considered several methods for sensing the user location. Outdoors, continuous services, such as GPS, can be used. Indoors, however, 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 ($\approx 40\text{kHz}$) 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 roughly equivalent to that of the MessagePad.

6. Extending the initial prototype

The first Cyberguide prototypes were completed within 6 months. To test out the genericity of our architectural approach, we decided to develop further prototypes that altered one or more of the major components described in section 4 and increased overall functionality. We describe these extended prototypes here.

6.1. Outdoor positioning

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 also developed a more functional outdoor prototype that covered three surrounding neighborhoods of the campus, described later.

6.2. Alternate platforms

In order to verify the platform independence of our conceptual design, we initiated two separate efforts building 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 initially installed on Dauphin DTR-1 palmtops running Pen for Windows Computing 1.0.

Explore Litigation Insights

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

Real-Time Litigation Alerts



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

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

Advanced Docket Research



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

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

Analytics At Your Fingertips



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

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

API

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

LAW FIRMS

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

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

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.