Building Applications to Establish Location Awareness: New Approaches to the Design, Implementation, and Evaluation of Mobile and Ubiquitous Interfaces

D. Scott McCrickard[1], Miten Sampat, Jason Chong Lee
Center for HCI and Department of Computer Science
Virginia Polytechnic Institute and State University (Virginia Tech)
Blacksburg VA 24061-0106 USA

ABSTRACT
An emerging challenge in the design of interfaces for mobile devices is the appropriate use of information about the location of the user. This chapter considers tradeoffs in privacy, computing power, memory capacity, and wireless signal availability that accompany the obtaining and use of location information and other contextual information in the design of interfaces. The increasing ability to integrate location knowledge in our mobile, ubiquitous applications and their accompanying tradeoffs requires that we consider their impact on the development of user interfaces, leading to an Agile Usability approach to design borne from agile software development and usability engineering. The chapter concludes with three development efforts that make use of location knowledge in mobile interfaces.

INTRODUCTION
A key challenge in the emerging field of ubiquitous computing is in understanding the unique user problems that new mobile, wearable, and embedded technology can address. This chapter focuses on problems related to location determination—different ways to determine location at low cost with off-the shelf devices and emerging computing environments, and novel methods for integrating location knowledge in the design of applications. For example, many web sites use location knowledge from IP addresses to automatically provide the user with relevant weather and traffic information for the current location. There is significant opportunity in the use of location awareness for HCI researchers to explore information-interaction paradigms for the uncertainty and unpredictability that is inherent to many location detection systems—particularly indoor systems that use Wifi signals which can be blocked by roofs, walls, shelves, and even people!

The prior knowledge of location to make such decisions in the presentation of information affords it to be categorized as *context awareness*, the use information that can be used to identify the situation of an entity to appropriately tailor the presentation of and interaction with information to the current situation (Dey, 2001). While context awareness can include a wide variety of information—including knowledge of who is in your surrounding area, events that are happening, and other people in your vicinity—this chapter focuses on the identification and use of location information, perhaps the most cheaply and readily available type of context information. This chapter considers the tradeoffs in privacy, computing power, memory capacity, and wireless (Wifi) signal availability in building interfaces that help users in their everyday tasks. We discuss our

[1] Contact author: mccricks@cs.vt.edu, http://people.cs.vt.edu/~mccricks/

own SeeVT system, which uses Wifi signals in location determination (Sampat et al., 2005). The SeeVT system provides the backbone for supplying location information to mobile devices on a university campus. Numerous interfaces built on SeeVT provide timely and appropriate location information to visitors in key areas of the campus.

The increasing ability to integrate location knowledge in our mobile, ubiquitous applications requires that we consider its impact on the development of user interfaces. This chapter describes the merging of agile software development methods from software engineering with the scenario-based design (SBD) methodology from usability engineering to create a rapid iteration design approach that is heavy in client feedback and significant in its level of reusability. Also presented are three interfaces developed using our Agile Usability methodology, focusing on the benefits found in using the Agile Usability approach and the tradeoffs made in establishing location awareness.

BACKGROUND

From the early days, navigation has been central to progress. Explorers who set sail to explore the oceans relied on measurements with respect to the positions of celestial bodies. Mathematical and astronomical techniques were used to locate oneself with respect to relatively stationery objects. The use of radio signals proved to be fairly robust and more accurate, leading to the development of one of the first modern methods of navigation during World War II, called LORAN (LOng RAnge Navigation). LORAN laid the foundation of what we know as the Global Positioning System, or GPS (Pace et al., 1995). Primarily commissioned by the United States Department of Defense for military purposes, GPS relies on 24 satellites that revolve around the Earth to provide precision location information in three dimensions. By relying on signals simultaneously received by four satellites, GPS provides much higher precision than previous techniques. GPS navigation is used in a wide range of applications from in-car navigation, to Geographic Information System (GIS)-mapping, to GPS-guided bombs.

GPS has become the standard for outdoor location-awareness as it provides feedback in a familiar measurement metric. Information systems like in-car navigators have adopted GPS as the standard for obtaining location, since it requires little or no additional infrastructure deployments and operates worldwide. However, GPS has great difficulty in predicting location in dense urban areas, and indoors, as the signals can be lost when they travel through buildings and other such structures. With an accuracy of about 100 meters (Pace et al., 1995), using GPS for indoor location determination does not carry much value. Along with poor lateral accuracy, GPS cannot make altitude distinctions of three to four meters—the average height of a story in a building—thus making it hard to determine, e.g., whether a device is on the first floor or on the second floor. Despite continued progress through technological enhancements, GPS has not yet evolved sufficiently to accommodate the consumer information-technology space. This chapter primarily focuses on technologies making inroads for indoor location determination.

While GPS has clear advantages in outdoor location determination, there have been other efforts focused around the use of sensors and sensing equipment to determine location within buildings and in urban areas. Active Badges was one of the earliest efforts at indoor location determination (Want et al., 1992). Active Badges rely on users carrying badges which actively emitted infrared signals that are then picked up by a network of embedded sensors in and around the building. Despite concerns about badge size and sensor deployment costs, this and other early efforts inspired designers to think about the possibilities of information systems that could utilize location-information to infer the context of the user, or simply the context of use. One notable related project is MIT's Cricket location system, which involved easy-to-install motes that acted as beepers instead of as a sensor network (Priyantha, Chakraborty, and Balakrishnan, 2000). The user device would identify location based on the signals received from the motes rather than requiring a broadcast from a personal device. Cricket was meant to be easy to deploy, pervasive and privacy observant. However, solutions like Cricket require deployment of a dense sensor network—reasonable for some situations, but lacking the ubiquity necessary to be an inexpensive, widely available, easy-to-implement solution.

To provide a more ubiquitous solution, it is necessary to consider the use of existing signals—many of which were created for other purposes but can be used to determine location and context. For example, mobile phone towers, IEEE 802.11 wireless access points (Wifi), and fixed Bluetooth devices (as well as the previously mentioned GPS) all broadcast signals that have identification information associated with them. By using that information, combined with the same sort of triangulation algorithms used with GPS, the location of a device can be estimated. The accuracy of the estimation is relative to the number and strength of the signals that are detected, and since one would expect that more "interesting" places would have more signals, accuracy would be greatest at these places—hence providing best accuracy at the most important places. Place Lab is perhaps the most widespread solution that embraces the use of pre-existing signals to obtain location information (LaMarca et al., 2005). Using the broadcasted signals discussed previously, Place Lab allows the designer to establish location information indoors or outdoors, with the initiative of allowing the user community to contribute to the overall effort by collecting radio environment signatures from around the world to build a central repository of signal vectors. Any client device using Place Lab can download and share the signal vectors for its relevant geography—requiring little or no infrastructure deployment. Place Lab provides a location awareness accuracy of approximately 20 meters.

Our work focuses specifically on the use of Wifi access networks, seeking to categorize the benefits according to the level of access and the amount of information available in the physical space. We propose three categories of indoor location determination techniques: *sniffing* of signals in the environment, *web-services access* to obtain information specific to the area, and *smart algorithms* that take advantage of other information available on mobile devices. In the remainder of this chapter, we describe these techniques in more detail, and we discuss how these techniques have been implemented and used in our framework, called SeeVT using our Agile Usability development process.

CATEGORIZING INDOOR LOCATION DETERMINATION TECHNIQUES

When analyzing location awareness, it is clear that the goal is not just to obtain the location itself, but information associated with the location—eventually leading to full context awareness to include people and events in the space, as described in (Dey, 2001). For example, indoor location awareness attributes such as the name of the building, the floor, surrounding environments, and other specific information attributed with the space are of particular interest to designers. Designers of systems intended to support location awareness benefit not only from location accuracy, but also from the metadata (tailored to the current level of location accuracy) that affords several types of cross-interpretations and interpolations of location and other context as well.

Access to this information can be stored with the program, given sufficient computing power and memory. This approach is reasonable for small areas that change infrequently—a library or a nature walk could be examples. Information about the area can be made accessible within the application with low memory requirements and rapid information lookup. However, changes to the information require updates to the data, a potentially intolerable cost for areas where location-related changes occur frequently. For example, a reconfigurable office building where the purpose and even the structure of cubicles change frequently would not be well served by a standalone application. Instead, some sort of web-based repository of information would best meet its needs. Taking this model another step, a mobile system could request and gather information from a wide range of sources, integrating it for the user into a complete picture of the location. As an example, a university campus or networked city would benefit from a smart algorithm that integrated indoor and outdoor signals of various types to communicate a maximally complete picture of the user's location.

Of course, each added layer of access comes with additional costs as well. Simple algorithms may sense known signals from the environment (for example, GPS and wireless signals) to determine location without broadcasting presence. However, other solutions described previously might require requesting or broadcasting of information, revealing the location to a server, information source, or rogue presence—potentially resulting in serious violations of privacy and security. The remainder of this section describes the costs and benefits for three types of indoor location determination approaches: sniffing, web services, and smart algorithms.

Sniffing

As the name suggests, sniffing algorithms sense multiple points of a broadcast environment, using the points to interpret the location of a device. The radio environment is generally comprised of one or more standard protocols that could be used to interpret location: modern environments include radio signals including Wifi, Bluetooth, microwaves, and a host of other mediums; creating interesting possibilities for location interpolation. Sniffing is also desirable because all location interpolation and calculations are performed on the client device—eliminating the need for a third-party service to perform the analysis and produce results. As mentioned previously, there are some benefits and disadvantages to this approach.

Performing the location determination on the client device eliminates the need for potentially slow information exchange over a network. This approach gives designers the flexibility they need in order to perform quick and responsive changes to the interfaces as well as decision matrices within their applications. For example, a mobile device with a slow processor and limited memory will need a highly efficient implementation to achieve a speedy analysis. A limiting factor for this approach is the caching of previously known radio vectors. Since most analysis algorithms require a large pool of previously recorded radio-signal vectors to interpolate location, it translates into large volumes of data being pre-cached on the client device. A partial solution for this exists already, pre-caching only for regions that the user is most likely to encounter or visit. Though this is not a complete solution to the resource crunch, it is a reasonable approach for certain situations, with periodic updates or fetches when radio-vectors are upgraded or the system encounters an unknown location.

Herecast is an example of a system using the sniffing model (Paciga and Lutfiyya, 2005). It maintains a central database of known radio vectors, which are then published to client devices on a periodic basis. The clients are programmed to cache only a few known locations that the user has encountered, and relies largely on user participation to enter accurate location information when they enter new areas that the system has not encountered before. The accuracy for these systems is generally acceptable, but there is always the worry of not having a cache of an area that the application is about to encounter. The lack of linking to a service also means that other contextual information associated with the location is hard to integrate with this approach due to device caching constraints and metadata volatility.

Web-services model

Keeping with the fundamental idea of mobile devices facing a resource crunch, this approach has client devices and applications use a central service for location determination. This means that the client device simply measures or "sees" the radio environment and reports it to the central service. The service then performs the necessary computation to interpolate the user location (potentially including other timely information) and communicates it back to the client. This also allows the client to store a minimal amount of data locally and to perform only the simplest of operations—important for mobile devices that often trade off their small size for minimal resources.

The approach is elegant in many ways, but faces several challenges in its simplistic approach such as the problem of network latency leading to lengthy times to perform the transactions. However, as the speed and pervasiveness of mobile networks is on the rise, as is the capabilities of silicon integration technologies for mobile platforms, designing large-scale centralized systems based on the web-services model will be a reasonable approach for many situations. Mobile online applications such as Friend Finders and child tracking services for parents are classic examples of tools that require central services to allow beneficial functionality to the end user.

Our own SeeVT system uses the web services model by allowing its clients to perform Simple Object Access Protocol (SOAP) web service calls to a standard web interface, and submit a radio vector for analysis. It then performs the necessary location determination using a probablistic algorithm and returns the location to the client. SeeVT provides the interface designer access to functionality on the service end as well. It allows the designers to control sessions and monitor the progress of clients by using a logging feature, and provides handles to integrate other widgets as well. For example, if an application wants to perform a search based on the users current location, SeeVT allows the designer to add functionality to its modules to perform further server-side computations.

Smart Algorithms

Looking ahead, algorithms that span large and diverse geographic areas will require the integration of many signals, information requests, and additional inputs. Place Lab attempts to address this issue for all radio signals (LaMarca et al., 2005). Currently it can compute location using mobile phone tower signals, Wifi, fixed Bluetooth devices, and GPS. However, we expect that other information will be used for location determination in the near future. For example, the ARDEX project at Virginia Tech seeks to use cameras—quickly becoming commonplace on mobile devices—to create a real-time fiducial-based system for location determination based on augmented reality algorithms (Jacobs, Velez, and Gabbard, 2006). The goal of the system is to integrate it with SeeVT such that anyone at defined hot spots can take a picture of their surrounding area and obtain information about their location. In an interesting twist on this approach, the GumSpots positioning system allows users to take a picture of the gum spots on the ground in urban areas and performs image recognition on them to return user location (Kaufman and Sears, 2006). Other information recording devices could be used in similar ways to help determine or enhance the understanding of our current location.

BUILDING INTERFACES FOR LOCATION-KNOWLEDGEABLE DEVICES

This section begins with a discussion on possible application scenarios that can leverage location knowledge in mobile devices. This section first describes *Agile Usability*, an extension of agile software development methodologies to include key aspects of usability engineering—resulting in an interface building technique that is well suited to ubiquitous and location-knowledgeable computing devices, both from the standpoint of interaction as well as development processes. Next, three case studies illustrate real world applications that have been built using these processes. Each case study describes key aspects of the application, illustrating one of the indoor location determination techniques and highlighting key lessons learned from the use of Agile Usability.

Agile Software Development, Usability Engineering, and Agile Usability

Ubiquitous and pervasive systems are often introduced to augment and support everyday tasks in novel ways using newly developed technology or by using existing technology in different ways. Since end-user needs are often ill-defined for ubiquitous systems, development needs to quickly incorporate stakeholder feedback so the systems can be iteratively improved to address new and changing requirements. This section discusses the use of an agile development methodology to build ubiquitous systems. Based on our own work (Lee et al., 2004; Lee, Chewar, and McCrickard, 2005) and on prior investigation of agile development methods (Beck, 1999; Koch, 2004; Constantine, 2001), we present a usability engineering approach for the construction of interfaces for mobile and ubiquitous devices.

Agile software development methodologies have been developed to address continuous requirements and system changes that can occur during the development process. They focus on quick delivery of working software, incremental releases, team communication, collaboration and the ability to respond to change (Beck et al., 2001). One stated benefit of agile methods is a flattening of the cost of change curve throughout the development process. This makes agile methods ideally suited to handle the iterative and incremental development process needed to effectively engineer ubiquitous systems. One shortcoming of many agile methods is a lack of consideration for the needs of end users (Constantine 2001). Current agile development methodologies have on-site clients to help guide the development process and ensure that all required functionality is included. However, many ubiquitous and pervasive systems require continuous usability evaluations involving end-users to ensure that such systems adequately address their needs and explore how they are incorporated in people's daily tasks and affect their behavior. Researchers, including Miller (2005), Constantine (2001), and Beyer et al. (2004) have developed ways to integrate system and software engineering with usability engineering. We present our approach to agile usability engineering, henceforth referred to as Agile Usability, with the added benefit of usability knowledge capture and reuse.

Our approach combines the software development practices of Extreme Programming (XP) with the interaction design practices of Scenario-Based Design (SBD) (Beck, 1999; Rosson and Carroll, 2002). The key features of this process are an incremental development process supported by continual light-weight usability evaluations, close contact with project stakeholders, an agile interface architecture model, known as a central design record (CDR), that bridges interface design and system implementation issues, and proactive knowledge capture and reuse of interface design knowledge (See Figure 1).

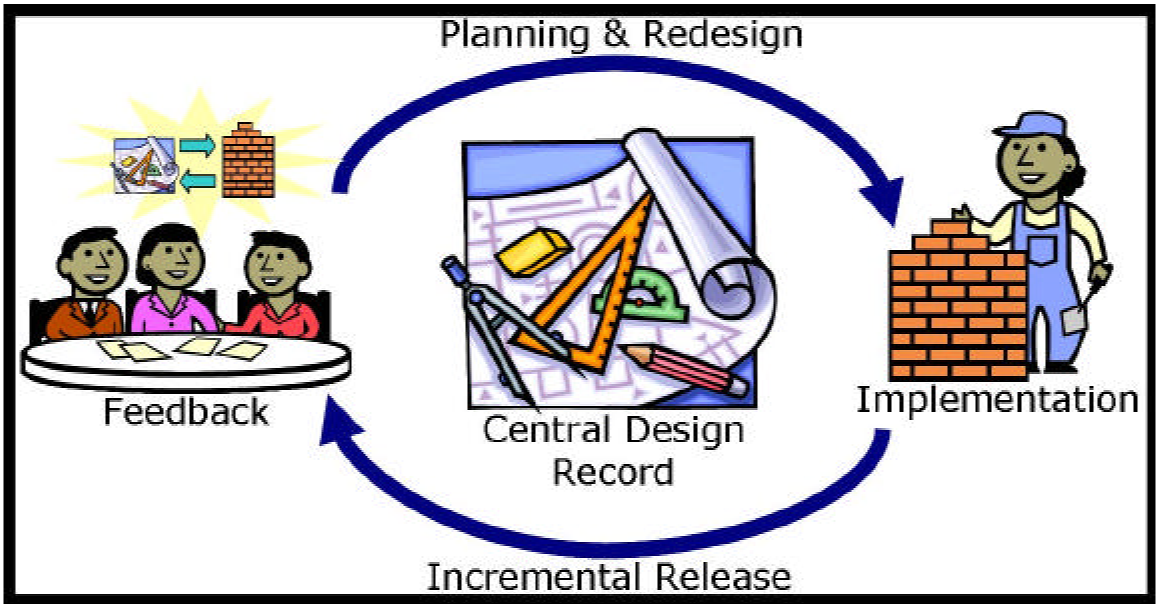

Figure 1: The Agile Usability process. The Central Design Record bridges interface design with implementation issues. This enables incremental improvement incorporating feedback from project stakeholders and usability evaluations.

Running a large-scale requirements analysis process for developing ubiquitous systems is not as beneficial as when designing other types of systems as it can be very difficult to envision how a ubiquitous system will be used in a specific situation or how the introduction of that system will affect how people behave or use it. In this type of development process, portions of the system are developed and evaluated by end users on a continual basis. This helps developers in uncovering new requirements and dealing with changing user needs as development proceeds. This type of development process requires some amount of discipline and rigor in terms of the types of development practices to follow. Specific details of these XP programming practices are detailed in Beck's book on the subject (Beck, 1999). Our use of these practices are elaborated in a technical report (Lee et al., 2005).

An incremental development process necessitates close collaboration with customers and end users to provide guidance on what features are needed and whether the system is usable. Ideally, representatives from these groups will be onsite with the developers working in the same team. Our customers were not strictly onsite, although they were in the same general location as the developers. Regularly scheduled meetings and continual contact through email and IM were essential to maintaining project velocity.

The key design representation is the central design record (CDR), which draws on and makes connections between design artifacts from XP and SBD. Stories, which describe individual system features, are developed and maintained by the customer with the help of the developers. They are prioritized by the customer and developed incrementally in that order. These include all features needed to develop the system including underlying infrastructure such as databases, networking software or hardware drivers. Scenarios, which are narratives describing the system in use, are used to communicate interface design features and behaviors between project stakeholders. Claims, which describe the positive and negative psychological effects of interface features in a design, are developed from the scenarios to highlight critical interaction design features. Story identification and development may lead to changes to the scenarios and claims. The reverse may also be true. This coupling between interface design and system implementation is critical for ubiquitous systems as developers must deal with both interactional and technological issues when deploying a system to the population.

In addition to acting as a communication point between stakeholders and highlighting connections between interface design and implementation, the CDR is important as a record of design decisions. As developers iterate on their designs, they often need to revisit previous design decisions. The explicit tradeoffs highlighted in the claims can be used by developers and clients to determine how best to resolve design issues that come up. Perhaps most important, Agile Usability drives developers to explore key development techniques in the development of location-based interfaces—the techniques used advance the field and can be reused in other situations.

Agile Usability has been applied in numerous situations, three of which are highlighted in this chapter as case studies. Each case study describes how the user tasks were identified, how stakeholder feedback was included, how our agile methodology was employed, and how appropriate location detection technologies were integrated. The discussion portion of this section will compare and contrast the lessons learned in the different case studies—highlighting specific usability engineering lessons and advancements that can be used by others.

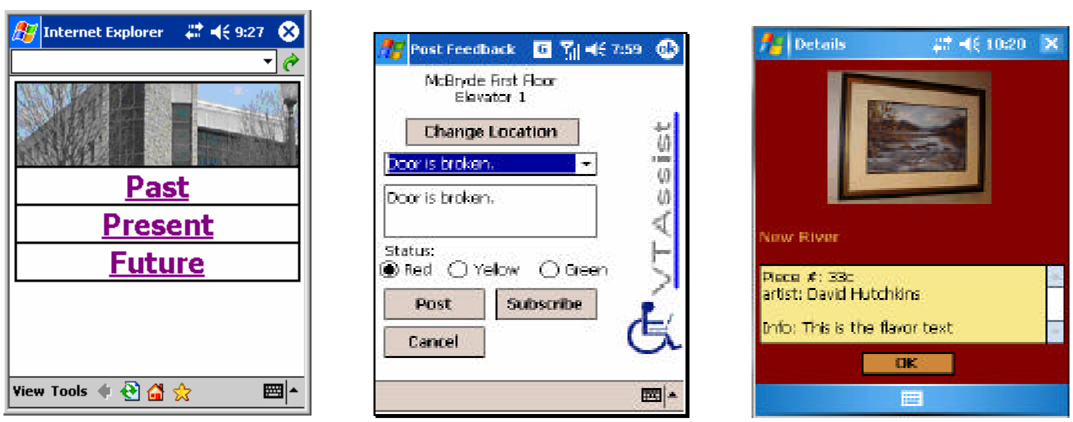

Figure 2: Screenshots from three applications built on SeeVT. From the left, the Alumni Tour Guide, VTAssist, and SeeVT-ART.

Case Study 1: Alumni Tour Guide

The alumni tour guide application was built for visitors to the Virginia Tech campus. The system notifies users about points-of-interest in the vicinity as and when they move about the VT campus (Nair et al., 2006). This image-intensive system provides easy-to-understand views of the prior and current layout of buildings in the current area. By focusing on an almost exclusively image-based presentation, users spend little time reading text and more time reflecting on their surroundings and reminiscing about past times in the area. See Figure 2 for a screenshot of the guide.

The earliest prototypes of the tour guide proposed a complex set of operations, but task analyses and client discussions performed in the Agile Usability stages indicated that many alumni—particularly those less familiar with handheld and mobile technology—would be unlikely to want to seek out solutions using the technology. Instead, later prototypes and the final product focused on the presentation and contrast of historical and modern images of the current user location. For example, alumni can use the tour guide to note how an area that once housed some administrative offices in old homes has been rebuilt as a multi-story technology center for the campus. This pictorial comparison, available at any time with only a few clicks, was well received by our client as an important step in connecting the campus of the past with the exciting innovations of the present and future.

As the target users are alumni returning to campus, most are without access to the wireless network, and the logistics are significant in providing access to the thousands of people who return for reunions, sporting events, and graduations. As such, the Alumni Tour Guide uses the sniffing location detection method to identify current location. This method fits well with the nature of the tasks of interest to alumni: they care most about the general space usage and the historical perspectives of a location that change little over time.

Case Study 2: VTAssist

Building interfaces is often difficult when the target audience has needs and skills different than those of the developer: for example, users with mobility impairments. It often takes many iterations to focus on the most appropriate solutions—a perfect candidate for Agile Usability. A pair of developers used our methodology to build VTAssist, a location-aware application to enable users with mobility impairments, specifically users in wheelchairs, to navigate a campus environment (Bhatia et al., 2006). VTAssist helps people in wheelchairs navigate in an environment more conducive to those who are not restricted in movement. Unlike typical handhelds and Tablet PC applications (the two platforms for which VTAssist was created), the VTAssist system must attract the user's attention at times of need or danger, guide them to alternate paths, and provide them with a means to obtain personal assistance when necessary. Perhaps most importantly, VTAssist allows users to quickly and easily supply feedback on issues and difficulties at their current location—both helping future visitors and building a sense of community among those who traverse the campus. See Figure 2 for a screenshot of the VTAssist.

In developing VT-Assist using Agile Usability, we found that needs and requirements changed over time, requiring that the methodology account for those changes. For example, the original design was intended to help wheelchair users find location accessible resources and locations, but later the need was identified to keep that information constantly updated, resulting in the addition of the collaborative feedback feature. It was this feature that was deemed most important to the system—the feature that would keep the information in VT-Assist current, and would enable users to take an active role in maintaining the information, helping others, and helping themselves.

Due to the importance of the feedback feature in maintaining up-to-date information for those in wheelchairs, VTAssist uses the web-services model. Certainly it would be possible to obtain some benefit from the sniffing model, but the client reaction indicated the importance of user feedback in maintaining an accurate database of problems and in providing feedback channels to frustrated users looking for an outlet for their comments. In addition, the server-side computations of location and location information (including comments from users and from facility administrators) results in faster, more up-to-date reports about the facilities.

Case Study 3: Conference Center Guide

The conference center guide, known as SeeVT-ART, addresses the desires of visitors and alumni to our area in coming to, and generally in returning to, our university campus—specifically the campus alumni and conference center (Kelly et al., 2006). SeeVT-ART provides multimodal information through images, text, and audio descriptions of the artwork featured in the center. Users can obtain alerts about interesting regional and university-specific features within the center, and they can be guided to related art by the same artist or on the same topic. The alerts were designed to be minimally intrusive, allowing users to obtain more information if they desired it or to maintain their traversal through the center if preferred. See Figure 2 for a screenshot of SeeVT-ART.

Agile Usability was particularly effective in this situation because of the large amount of input from the client, who generated a lot of ideas that, given unlimited time and resources, would have contributed to the interface. Agile Usability forced the developers to prioritize—addressing the most important changes first while creating placeholders illustrating where additional functionality would be added. Prioritization of changes through Agile Usability also highlighted the technological limitations of the underlying SeeVT system, specifically those related to the low accuracy of location detection, and how that influenced the system design. For example, when a user enters certain areas densely populated with artistically interesting objects, SeeVT-ART requires the user to select from a list of the art pieces as it is impossible to determine with accurate precision where the user is standing or (with any precision) what direction the user is facing. These limitations suggested the need for smart algorithms that use information about the area and that integrate additional location determination methods.

Smart algorithms that store location data over time and use it to improve location detection can be useful in determining data such as the speed at which a user is walking and the direction a user is facing. SeeVT-ART can use this data to identify the piece of art at which a user most likely is looking. Our ongoing work is looking at integrating not only the widely accessible broadcast signals from GPS, cellular technology, and fixed Bluetooth, but also RFID, vision algorithms, and augmented reality (AR) solutions. Our early investigation into a camera-based AR solution combines information about the current location with image processing by a camera mounted on the handheld to identify the artwork and augment the user's understanding of it with information about the artist, provenance, and so forth. These types of solutions promise a richer and more complete understanding of the importance of a location than any one method could accomplish alone.

CONCLUSIONS AND FUTURE DIRECTIONS

The three location-knowledgeable SeeVT applications described in this document offer a glimpse into the possibilities for location-knowledgeable mobile devices. The increasing presence of wireless networks, improvements in the power and utility of GPS, and development of other technologies that can be used to determine location portends the ubiquity of location-knowledgeable applications in the not-too-distant future. Delivery of location-appropriate information in a timely and useful manner with minimal unwanted interruption will be the goal of such systems. Our ongoing development efforts seek to meet this goal.

In support of our development efforts, we explore new usability engineering approaches particularly appropriate for location-knowledgeable applications. The use of stories and the knowledge capturing structures of Agile Usability combined with its rapid multiple iterations enable convergence on solutions to the most important issues faced by emerging application areas. We repeatedly found that designers are able to identify issues of importance to the target users, while keeping in perspective the design as a whole. Our ongoing work seeks ways to capture and share the knowledge produced from designing these applications not only within a given design but across designs, leading to the systematic scientific advancement of the field.

In the future, these developing Agile Usability techniques will be supported by specific tools and toolkits for leveraging the location-awareness needs of on-the-go users. An early contribution that can be drawn from this work is the novel methods for supporting location awareness in users—browseable historical images of the current location, rapid feedback methods for reporting problems, new map presentation techniques—all methods that should be captured in a toolkit and reused in other location awareness situations.

REFERENCES

Bahl, P. and Padmanabhan, V. N. (2000). RADAR: An In-Building RF-Based User Location and Tracking System. In *Proceedings of IEEE INFOCOM* (volume 2, 775-784). Tel Aviv, Israel.

Beck, K. (1999). Embracing Change with Extreme Programming. *IEEE Computer*. Oct. '99. Volume 32, Issue 10, 70-77.

Beck, K., Beedle, M., van Bennekum, A., Cockburn, A., Cunningham, W., Fowler, M., Grenning, J., Highsmith, J., Hunt, A., Jeffries, R., Kern, J., Marick, B., Martin, R.C., Mellor, S., Schwaber, K., Sutherland, J., and Thomas, D. (2001). The Agile Manifesto. Available at http://agilemanifesto.org

Beyer, H. R., Holtzblatt, K., Baker, L. (2004). An Agile Customer-Centered Method: Rapid Contextual Design. In *Proceedings of Extreme Programming and Agile Methods 2004 (XP/Agile Universe)* (50-59). Calgary Canada.

Bhatia, S., Dahn, C., Lee, J. C., Sampat, M., and McCrickard, D. S. (2006). VTAssist-A Location-based Feedback Notification System for the Disabled." In *Proceedings of the ACM Southeast Conference (ACMSE '06)* (512-517). Melbourne FL.

Constantine, L. L. (2001). Process Agility and Software Usability: Toward Lightweight Usage-Centered Design." *Information Age*, *8* (2). Reprinted in L.Constantine (Ed.), Beyond Chaos: The Expert Edge in Managing Software Development. Addison-Wesley, Boston.

Dey, A. K. (2001). Understanding and Using Context. *Personal and Ubiquitous Computing* Volume 5, Number 1, 4-7.

Jacobs, J., Velez, M, and Gabbard, J. (2006). ARDEX: An Integrated Framework for Handheld Augmented Reality. In *Proceedings of the First Annual Virginia Tech Center for Human-Computer Interaction Research Experience for Undergrads Symposium*, (6). Blacksburg VA.

Kaufman, J. and Sears, J. (2006). GSPS: GumSpots Positioning System. In *IPT 2006 Spring Show*, available at http://itp.nyu.edu/show/detail.php?project_id=539.

Kelly, S., Hood, B., Lee, J. C., Sampat, M., and McCrickard, D. S. (2006) Enabling Opportunistic Navigation Through Location-Aware Notification Systems. Pending paper submission.

Koch, A. S. (2004). Agile Software Development: Evaluating the Methods For Your Organization. Artech House Publishers.

LaMarca, A., Chawathe, Y., Consolvo, S., Hightower, J., Smith, I., Scott, J., Sohn, T., Howard, J., Hughes, J., Potter, F., Tabert, J., Powledge, P., Borriello, G., and Schilit, B. (2005). Place Lab: Device Positioning Using Radio Beacons in the Wild. In *Proceedings of the 3rd International Conference on Pervasive Computing (Pervasive 2005)* (134-151). Munich, Germany.

Lee, J. C., Lin, S., Chewar, C. M., McCrickard, D. S., Fabian, A., and Jackson, A. (2004). From Chaos to Cooperation: Teaching Analytic Evaluation with LINK-UP. In *Proceedings of the World Conference on E-Learning in Corporate, Government, Healthcare, and Higher Education (E-Learn '04)* (2755-2762) Washington DC.

Lee, J. C., Chewar, C. M., and McCrickard, D. S. (2005). Image is Everything: Advancing HCI Knowledge and Interface Design Using the System Image. In *Proceedings of the ACM Southeast Conference (ACMSE '05)* (2-376-2-381). Kennesaw GA.

Lee, J. C., Wahid, S., Chewar, C. M., Congleton, B., McCrickard, D. S. (2005). Spiraling Toward Usability: An Integrated Design Environment and Management System. Technical Report TR-05-15, Computer Science, Virginia Tech.

Miller, L. (2005). Case Study of Customer Input for a Successful Product. In *Proceedings of the Agile 2005 Conference* (225-234). Denver CO.

Nair, S., Kumar, A., Sampat, M., Lee, J. C., and McCrickard, D. S. (2006). Alumni Campus Tour: Capturing the Fourth Dimension in Location Based Notification Systems. In *Proceedings of the ACM Southeast Conference (ACMSE '06)* (500-505). Melbourne FL.

Pace, S., Frost, G. P., Lachow, I., Frelinger, D., Fossum, D., Wassem, D., and Pinto, M. M. (1995). The Global Positioning System: Assessing National Policies, Rand Corporation, Reference Number: MR-614-OSTP.

Paciga, M. and Lutfiyya, H. (2005). Herecast: An Open Infrastructure for Location-Based Services Using WiFi , In *Proceedings of Wireless and Mobile Computing, Networking, and Communications (WiMoB 2005)* (21-28). Montreal, Canada.

Priyantha, N. B., Chakraborty, A., and Balakrishnan, H. (2000). The Cricket Location-Support system. In *Proceedings of the Sixth Annual International Conference on Mobile Computing and Networking (MOBICOM 2000)* (32-43). Boston, MA.

Rosson, M .B. and Carroll, J. M. (2002). Usability Engineering: Scenario-Based Development of Human-Computer Interaction, Morgan Kaufman, New York, NY.

Sampat, M., Kumar, A., Prakash, A., and McCrickard, D. S. (2005). Increasing Understanding of a New Environment using Location-Based Notification Systems. Poster paper in *Proceedings of 11th InternationalConference on Human-Computer Interaction (HCII '05)* (5 pages). Las Vegas NV.

Sciacchitano, B., Cerwinski, C., Brown, I., Sampat, M., Lee, J. C., and McCrickard. D. S. (2006). Intelligent Library Navigation using Location-aware Systems. In *Proceedings of the ACM Southeast Conference (ACMSE '06)* (371-376). Melbourne FL.

Tom, H. 1994. The geographic information systems (GIS) standards infrastructure. StandardView Volume 2, Number 3, 33-142.

Want, R., Hopper, A., Falcao, V., Gibbons, J. (1992). The Active Badge Location System. ACM Transactions on Information Systems, Volume 40, Number 1, 91-102.

Youssef, M. and Agrawala , A. K. (2004). Handling Samples Correlation in the Horus System. In *Proceedings of IEEE INFOCOM* (9 pages). Hong Kong, China.