Fully Distributed Location-Aware Computing

John Light[1], Eric Pattison[2], Trevor Pering[1], Murali Sundar[1], Roy Want[1]

[1]Intel Research
2111 NE 25th Ave. JF3-375., Hillsboro OR 97124
503 264 9080
first.last@intel.com

[2]University of Calgary
Dept. of Computer Science
2500 University Drive NW, Calgary, AB T2N 1N4
ericp@cpsc.ucalgary.ca

ABSTRACT

There have been many proposals for location-aware computing that involve centralized infrastructure. Using cell-phone systems or GPS databases requires considerable fixed resources to maintain location information and provide services. We propose the Ubiquitous Walkabout, which, by using independent Information Beacons and an "always on" mobile device, provides a richer and more flexible user experience. The approach provides more capability, privacy and scalability than a cellular approach (and may also be less expensive), and more capability and flexibility than a GPS approach.

Author Keywords

Location-aware computing, ubiquitous computing, Personal Server, Information Beacons.

THE UBIQUITOUS WALKABOUT

The Ubiquitous Walkabout, part of the Ubiquity project in Intel Research, is an investigation into distributed location-aware computing. It involves Information Beacons, which wirelessly broadcast specific information about a location to a small (30 ft) vicinity, and a Personal Server, which can receive the beacon messages and process them as they are received. Walking down a street equipped with this technology, a user would receive the information from many Information Beacons in turn, each of which would transmit information about a specific service or other offering in its vicinity. Based on previously expressed policies and preference provided by the user, the Personal Server might report the service or offering to the user, respond to the offering automatically, log the information for later use, or ignore it.

THE PERSONAL SERVER

The Personal Server [1] is a capability that can be part of any small mobile device such as a cell phone or PDA. It can run continuously for a long time (one or more days) and provide considerable computational and storage resources to its user during that time. It uses advanced power management capabilities to provide the appearance of being "always on", though it may in fact enter sleep modes from which it can quickly return. In its current form, the Personal Server communicates with the world through one or more wireless (radio) interfaces. It can talk to Personal Computers, public displays, PDAs, cell phones, or personal I/O devices like a wireless watch.

INFORMATION BEACONS

An Information Beacon can be as simple as a radio chip, microcontroller, and power supply. It need only be able to store a small amount of information that it broadcasts repeatedly, and can be manufactured for under $20 in large quantities. To establish the content of a beacon, it can simply be plugged into a standard Personal Computer.

We are currently using Berkeley Motes as Information Beacons. They use a simple radio technology to broadcast a short (300 byte) message repeatedly. We have recently used mote technology based on a Bluetooth radio (iMote) to provide larger messages.

FULLY DISTRIBUTED

The Ubiquitous Walkabout provides a location-aware computing experience that differs from some previous descriptions.

- The Personal Server can store and utilize considerable user context, both explicit and inferred, as well as volumes of acquired content.
- The Personal Server can run agents of considerable complexity on behalf of the user, and the user can freely experiment with choice and configuration of agents.
- The beacon owner has complete control over content, and the Personal Server owner has complete control over response to it.
- All preferences, policies, actions, and recorded data are limited to the Personal Server (or other user specified devices), so privacy is maximized.
- All of the above capabilities can be delivered in a fully scalable manner that doesn't involve a centralized resource.

COMPARISON TO OTHER APPROACHES

We will compare the proposed approach to Location-Aware Computing (LAC) to existing directions. The primary directions being pursued now are:

1. Cell phone-based infrastructure approach. This is based on knowing what cell you are in, which the cellular provider can use to base services on.

2. Cell phone-based database approach. Since the cell you are in is available at the handset, an application running in the handset can provide services based on a local or online database mapping cells to services.
3. 911-based infrastructure approach. Soon most cellular systems will be able to track users to within a few meters (many already do). This information is typically only available to the cellular provider (and its designees).
4. GPS database approach. If the mobile device contains a GPS receiver, it can access a local or online database to find out about services in the vicinity.
5. Hotspot database approach. Since hotspots are common in some areas, a directory mapping hotspot IDs to their locations can be used to access the same databases that the GPS approach uses.

We will discuss these approaches and the proposal in the following dimensions:

1. Scalability.
2. Services.
3. Performance.
4. Reliability.
5. Privacy.
6. Security.
7. Economics.

All of the current approaches have one thing in common: they depend on some form of infrastructure. In the case of the three cellular approaches, the infrastructure includes the cellular system itself. The three database approaches require creation and maintenance of a publicly available database. While the distribution of these databases can be decentralized, their maintenance and control may need to be centralized. The cell-based and 911-based infrastructure approaches involve databases privately held by the cellular providers (and its designees).

Consider the potential magnitude of these central databases. If half the stores in the United States and a third of its citizens wanted to have a LAC presence, the size of these databases would be enormous. You can think about it as the union of all the yellow and white phone books in the U.S. If the LAC is to meet its potential, much of this information would need to be dynamic (unlike a phone book), and the update process would be daunting (and likely slow). Think of it as if someone would try to implement the web as a centralized database. (The databases could be decentralized, a la Yellow Pages, but even online YPs can take over a week to update.)

The proposed approach eliminates central infrastructure. Just as with the web, it allows the problem to be naturally distributed. Just as with the web, each user communicates directly with the information provider, without an intermediary. As individuals acquire mobile devices that work with the proposal and as shops and individuals acquire information beacons, the system scales naturally.

What services can be offered with the various approaches? We submit that any service can ultimately be offered by any of these approaches, existing or proposed. If the hypothesis is that there is no difference among them in this respect, then we are faced with proving a negative. We invite you to provide counterexamples.

The quality of the services offered may vary, however. The database approaches may suffer update latencies, making it difficult to provide highly targeted changes. The update latency consists of two factors: updating the database itself (including distribution delays) and updating the mobile device with changes. Minimizing both latencies seems difficult. Updating the private databases in the cell-based approaches is not much easier, though the delay updating the mobile device is not present since the LAC functionality is implemented in the cellular provider's back end, which is near the database.

Another aspect of service quality is the level of innovation that can be brought to bear on the problem. Can vendors try different approaches and change approaches to meet perceived needs. The database approaches allow plenty of innovation using the distributed database data, but they may restrict innovation that requires new data formats or types of information. Cellular approaches require innovation either by the cellular provider or by third parties with access to the providers back end. History has shown cellular providers to be slow innovators, and supportive of third party innovation only at a steep price. The proposed approach, since it is fully distributed, allows innovation on multiple levels since new ideas can be tried out locally and migrated elsewhere at little cost.

Apparent performance will likely be seen as the timeliness of getting up-to-date information from the source (e.g., store) to the user. This time includes both the time to determine location and the timeliness of the database updates. The cellular approaches and the proposal can determine location quickly. The GPS approach can suffer acquisition delays or failures. The database update problem was discussed above.

All of the approaches depend to some extent on real-time radio traffic, which is inherently unpredictable. Cellular providers have spent vast sums to ensure reliable operation under most circumstances. GPS systems are known to fail in urban areas and seldom work indoors. The proposed approach depends on untested (in this use) radio technology, so its reliability is yet to be determined. We believe it can be made locally reliable, but many problems will have to be worked out.

Privacy is critical to user acceptance of Location-Aware Computing. There are many aspects to privacy, and the various approaches affect different aspects differently.

Here are some relevant privacy issues:

1. Stealth. This is the ability to be somewhere without anyone knowing you are there. GPS preserves stealth since the device doesn't radiate. Cellular approaches are not stealthy, but we assume that no one is going to use the information. Information beacons can preserve stealth if they don't require a response from the mobile device.
2. Anonymity. This is freedom from being identified. As long as you have stealth privacy, you have anonymity, but if the stealth is not preserved, you might still be able to acquire beacon information without disclosing your identity or some information that could be used to establish your identity (such as a MAC address).
3. Tracking. This is the ability for someone to track your comings and goings through your use of the technology. The cellular network can track you, and with the advent of 911 technology, it can track you in great detail. The GPS user is largely safe from being tracked, except by observing downloads of the services database. Information beacons allowing stealth operation can't be used to track you, but if they require non-anonymous ID, or even a consistent anonymous ID, they can.
4. Interests. Independent of knowing where you are, there is the question of learning about your interests. Can an outsider learn about your preferences and interests by observing some aspect of your behavior? Observation of selective incremental downloads of a database could be used for that purpose. The proposed approach provides this class of privacy if using stealth operation.
5. Data. This is specific information about you, typically collected by you for your own use. Access by others is considered a violation of privacy. Protection of this data is usually implemented as a form of computer security, which is discussed in a later paragraph.

All of these privacy issues depend to a large extent on how well you trust the support infrastructure that is required to provide it. Most people trust their cellular provider with all these forms of privacy. We trust that the provider isn't tracking us or listening in on our conversations with the bank. Public databases are another thing. Who is going to manage and monitor those databases? Only the current proposal in stealth mode preserves all of these, with some question about data privacy.

Security is a difficult question because all security is relative to expectations. We trust the cellular provider to keep information about us secure, and mostly it is. The public database proposals, along with the current proposal, share the issue that the security of the mobile device determines our overall security since they contain considerable information about us. Can mobile devices be made secure? We suspect they can be secured against access through their normal I/O mechanisms, but what happens when the device falls into a miscreant's hands? As more and more important information is kept on mobile devices, the importance of this question grows.

What are the economics of the various approaches? The cost of GPS capability in a mobile device will continue to drop, so we assume that it will eventually be nearly free. The cost of cellular service is well understood, but the cost of backend services is not. Some reports on European experience imply that cellular providers are loath to part with location information, making it intolerably expensive. This begs the question of what their actual costs of providing these services might be. The cost of maintaining and delivering public service databases is unknown, but it will probably involve a monthly fee, since the information must be kept fresh. The proposed approach involves a small incremental cost to a standard mobile device (for the Personal Server capability) and a cost of $20-50 to purchase each information beacon, and some cost to run it. We suspect that the one-time cost of an information beacon is much smaller than the fees to support the ongoing costs of maintaining public or private service databases.

An additional question for all these approaches is how you start them up, since they suffer to varying extents from the "chicken and egg" problem.

CONCLUSION

We have presented an alternative approach for Location-Aware Computing that offers several advantages over current approaches. We discussed what those approaches are and their advantages and disadvantages relative to the proposed alternative.

We hope we have successfully conveyed that the proposal is worthy of more extensive consideration by the larger research community.

REFERENCES

1. Want R., Trevor Pering, Gunner Danneels, Muthu Kumar, Murali Sundar and John Light. *The Personal Server: changing the way we think about ubiquitous computing*,. Proceedings of UbiComp 2002, Springer LNCS #2498, Goteborg, Sweden, pp194-209.

BIOGRAPHY

John is a member of the Ubiquity Strategic Research Project in Intel Research, working with Roy Want and others to investigate the applicability of the Personal Server concept to various aspects of Ubiquitous and Proactive Computing. Prior to that, John worked on Information Visualization and Knowledge Management research at Intel, including the Miramar 3D workspace, and at other companies on graphics, operating system, and CAD products.