Personal Location Agent for Communicating Entities (PLACE)

Justin Lin[1], Robert Laddaga[1], and Hirohisa Naito[2]

[1] MIT Artificial Intelligence Laboratory
{JL79, rladdaga}@ai.mit.edu
[2] Fujitsu Laboratories Ltd.
{naitou}@jp.fujitsu.com

Abstract. Traditionally, location systems have been built bottom-up beginning with low-level sensors and adding layers up to high-level context. Consequently, they have focused on a single location-detection technology. With sharing of user location in mind, we created Personal Location Agent for Communicating Entities (PLACE), an infrastructure that incorporates multiple location technologies for the purpose of establishing user location with better coverage, at varying granularities, and with better accuracy. PLACE supports sensor fusion and access control using a common versatile language to describe user locations in a common universe. Its design provides an alternative approach towards location systems and insight into the general problem of sharing user location information.

1 Introduction

The pervasive computing movement calls for a wealth of context information in order to perform the appropriate tasks for each person in each of their possible situations [1, 2]. Of this context, location has become one of the most popular topics among research groups. This increase in location technologies has produced a large number of mobile (location) applications that fall under two categories: (1) users performing context-based tasks with his/her own location and (2) users performing context-based tasks with the locations of other users. The first category includes reminder services [3, 4], navigational services [5, 6], location-aware information delivery [4], environment customizations [7], location-based tour guides [8, 9], and general context engines [10]. However, the second category remains largely unexplored. Some research groups have ventured into this domain (e.g. ActiveMap [11]), but possibly due to the daunting privacy concerns, there has been relatively little research in this area. In this paper, we demonstrate the potential behind sharing one's location with others, given a location system infrastructure designed specifically for this purpose.

In addition to location applications, the exploding interest in location context has resulted in a multitude of location detection technologies developed by many different groups, each of which seems to contribute something unique and interesting to the field [12]. Individually, these systems are useful, yet each faces the limitations of their methods. Working together, they can obtain location information with better coverage, at varying granularities, and with better accuracy - three qualities particularly useful in sharing user locations. Consider how people generally solve such problems. When we need to locate some object or place, we refer to several sources of location information, by asking multiple people, referring to multiple maps, using GPS receivers, and other methods. In doing so we transparently interpret multiple coordinate systems in terms of each other, resolve disparate symbolic references, and gauge the probabilities associated with each bit of evidence. These things that we do with little thought or effort are incredibly difficult for our computer systems to do. In light of this complexity, we can see why there has been little effort to bring the unique and interesting features of various location systems together into one collective system.

This paper aims to provide insight into an alternative approach toward user location systems, namely building an infrastructure with location-sharing in mind from its conception. We discuss the Personal Location Agent for Communicating Entities (PLACE), an infrastructure that utilizes (1) a semantic representation of location as a means to attain fusion of sensors at varying granularities for better coverage and accuracy, and (2) an intermediary software agent between location devices and location services that performs sensor fusion and access control on behalf of the user. We begin with a brief generalization of current location systems and the details and motivations behind our general design. Then, we move to a discussion on how to achieve a common universe of locations to allow for clear communication of location information between a world of entities. Finally, we provide brief overviews of applications that capitalize on our design, after which we give some concluding thoughts.

2 Related Work

There exists a multitude of location systems today. Each system typically chooses a location representation that suits the technology used in determining one's location. For example, the Global Positioning System (GPS) uses multiple satellites that provide specially coded signals to a GPS receiver, which then can use these signals to compute its location in latitude-longitude coordinates. Further, the Active Badges location system places a base per room that detects via infrared communication when badges enter the room, thereby providing a symbolic location representation, namely which room a user is in [14].

Each location system generally has both unique features and limitations in terms of technology, accuracy, scalability, and cost. Because these systems are usually developed in separate research facilities with different views on the ideal location system, there appears to be little effort to have two or more location systems coexist and collaborate with each other. Consequently, applications built to utilize location technologies have generally limited themselves to one location system, as shown in Fig. 1. Thus, most applications that exist today face the limitations of their chosen location systems.

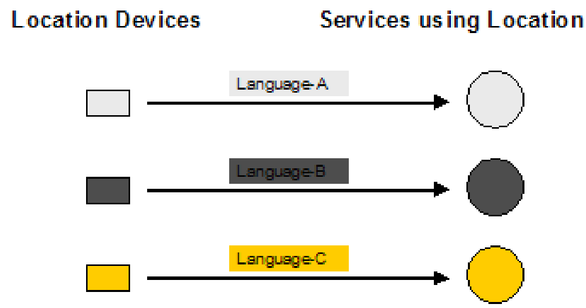

Fig. 1. Conventional Method

Korkea-aho and Tang [15] have made a similar observation that each location application seems to create its own representation. They aimed to “create a simple lowest common denominator data set that as many location information sources and applications in the Internet as possible could use.” The result was called the “Common Spatial Location Data Set,” consisting of geodetic latitude, longitude and altitude, accuracy and time of measurement, speed, direction, course, and orientation. This is a sensible solution for location systems in general; however, in the context of pervasive computing and describing locations of users, this geographic location representation would not be ideal. People generally tend to reason about locations as semantic places rather than as coordinates.

Furthermore, various specifications have been proposed. First, as specifications for describing geographical objects, maps, etc. G-XML [16] and GML [17] should be noted. They create a standard for encoding geospatial data into XML documents. Likewise, POIX [18] and NVML [5] are specifications for describing positions of moving and static objects and encoding this information into XML. Both sets of specifications are based on common location expressions and add additional vocabulary depending upon their objective. There are some ongoing activities attempting to standardize coordinate representations for location expressions. However, they have not attempted to standardize symbolic representation for expressions of location information.

SignalSoft Corporation [19] has also attempted to utilize multiple location sources in their product suite. The Location Manager receives data from multiple location sources and delivers the data to location applications. They have successfully created a commercial product but restrict themselves to cell-phone-based location detection technologies. Furthermore, these cell-phone-based technologies use a spatial location representation.

Multisensor data fusion is a field that has a huge amount of attention and numerous published results [20]. However, most work in the field is based on a signal processing, bottom up approach. We favor an approach that proceeds top down from a relevant decision problem, as an approach to filter away much of the irrelevant ambiguous or conflicting data.

Access control is another highly developed and explored area, with a great deal of literature. For specification of access control for purposes of maintaining military secrets, see [21]. For general discussion of security, confidentiality and information integrity issues, see [22]. Our own efforts are more influenced by Role Based Access Control, [23].

3 Sharing Location - Conceptual Requirements

In the context of this paper, location sharing is the sharing of information about locations of persons for the purpose of performing some task effectively, such as joining them for a meeting, delivering a package, predicting their near future behavior, and determining location dependent privileges. Sharing location information can be categorized into: (1) systems sharing location information with other systems, (2) systems sharing location information with humans, and (3) humans sharing location information with other humans. Wherever machines (systems) are involved, we need to be concerned about interoperable data representations. Wherever humans are involved, we need to be concerned about issues of privacy, and the ability to present information in an understandable manner. Note that humans may well have privacy concerns about data residing physically (persistently) on certain machines. Further, even in machine to machine interactions, humans may need to diagnose and debug failures, and hence need to understand data involved in the interactions.

3.1 Common Language

Ideally, all location devices and services would use a common, universal location representation and language. Unfortunately, this is not the reality. The vast contrast between location representations that exist today is not merely a consequence of philosophical differences, but rather of the inherent unique features that are specific to each location technology. GPS provides latitude-longitude coordinates while Active Badges provide symbolic spaces, but both technologies prove to be useful in different scenarios.

Our goal is to allow for bidirectional translation to and from a common language that expresses location in a representation chosen and designed specifically for user-location services. This language serves as a standard, if you will, between communicating entities, but bidirectional translation provides versatility when required. Because a universal language does not exist today, communication of location information is restricted to be from the location device to the location service (again, designed specific to the location device). With a common user-location language, location information can be freely exchanged between all communicating entities, device to service and service to service.

3.2 Common Map of Locations

With a common language, entities can communicate location information, but cannot necessarily fully understand each other without having the same map of locations. Upon finding out that a user is in “John Hancock Building,” in order to be able to realize and use this location (e.g. locate on a map, get directions, build higher-level context), one must have at least a similar understanding of “John Hancock Building” as the provider of the information. For the understandings to be similar enough, we require that the processes by which each user binds an object or location to the symbol, result in roughly the “same” object or location.

3.3 Sensor Fusion

Deducing a user’s location is an attempt to capture some information about the world. The multitude of location technologies that exist today provides many paths one can take to obtain this information. However, each location technology has its limitations; this motivates us to combine technologies and to “take all that we can get.” With this approach comes the problem of sensor fusion.

Perhaps an even more important purpose of sensor fusion and multiple location technologies is making location systems more robust to environmental conditions. Any given location system can fail for any number of reasons, and it will often be the case that backup systems require multiple sources to recover the coverage or accuracy of the primary system. Multiple systems for providing location information, and a flexible architecture that allows for arbitrary combinations of information from disparate sensor sources and other information, is central to a more robust and self-adaptive architecture [24, 25].

3.4 Access Control

The presence of an intermediary between location devices and location services provides us with a convenient place to perform access control. Traditionally, many location applications have not implemented access control because they do not involve sharing location information with other users. Those that do involve sharing of location information are often based on the user carrying a location device that the user can detach or disable, and thus address the access control issue (e.g. [11]). This sidestepping of access control issues does not adequately deal with common situations in which the user wants some people to have access to location information under certain circumstances. It also fails to deal adequately with non-personal location devices such as cameras, and with location information needed for emergency reasons. A less binary solution, and one that is less device dependent, is called for. Just as P3P [13] seeks to offer users complete control over their personal information on Web sites they visit, a location system must offer users complete control over their location information.

4 PLACE

We propose a design and implementation exploring the possibilities of a unifying location system infrastructure that allows user-location to be commonly understood among communicating entities. Our design goal is to create a user-location infrastructure supporting *sensor fusion* and *access control* using a *common versatile language* to describe locations in a *common universe*.

The ultimate goal behind our unifying location system is to allow for all communicating entities to fully understand each other, with respect to location information. In order to accomplish this, these entities must all share a common universe of locations. Sensor fusion cannot operate unless sensors are able to collaborate. Users cannot perform access control based on location context without precisely matching their perceptions of the locations. Communicated location information carries meaning exactly to the extent that there is a common universe of locations shared by the communicating parties.

In the next few sections we describe our software environment, and our design for language, representation, access control, and a common universe of locations.

4.1 Software Platform

The ubiquitous computing environment within which our location service is designed to function is the MIT AI Lab Intelligent Room [26]. The Intelligent Room is based on a distributed agent infrastructure called Metaglue [27], which provides directory and brokering services to agents. Metaglue supports lightweight agents, and provides an architecture that organizes agents into societies, in such a fashion as to allow intersociety communication. Because the service of sensor fusion occurs on behalf of each user, a distributed, agent-based software platform is quite appropriate, even if the Metaglue architecture did not dictate such a solution.

In the Metaglue environment, agents work within a society, and each user has a society of agents. Each user's society represents the computing environment devoted to performing services for the user. As seen in Fig. 2, we place the responsibility of sensor fusion and access control (described below) on an agent called the Personal Location Agent (PLA) that acts as an intermediary between location detection devices and location services.

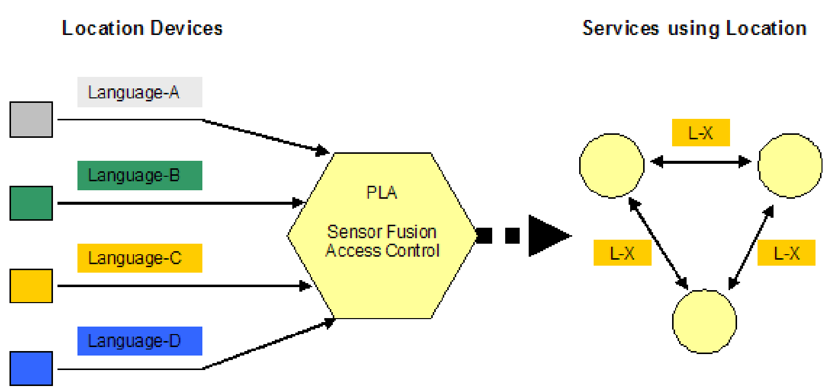

Fig. 2. General Design of PLACE

A much less obvious, yet perhaps more important motivation to using an agent-based platform for location is the vision of user-centric computation. Traditionally, location applications have simply obtained location information directly from the users' location devices, giving essentially no regard to the user as an entity, or the ultimate consumer of location information. This approach runs counter to the new trend away from application-centric computation towards user-centric computation. We propose that by performing location detection and treating location as a property of a user entity, regardless of what services actually use the information, we achieve the notion of computation acting on behalf of the user, as opposed to the application. In doing so, we separate location detection from location utilization.

Indeed, the Metaglue environment in which we have embedded the location services is itself a component of a larger ubiquitous computing environment: The MIT Oxygen prototype system, within which the Intelligent Room functions as a prototype of Oxygen's Enviro21 [28].

4.2 Language

The first step to bringing all entities into the same universe of locations is to have them understand each others' languages. There is a spectrum of possible solutions between two extreme solutions to the language problem. First, all entities can speak their own language as long as they offer translation to and from all other languages. This represents the more cumbersome (but also more flexible) of the two solutions because it places the responsibility of deciphering all incoming information on each entity. Alternatively, all entities could use the same location representation and avoid translation altogether. However, forcing one representation onto all entities would entail losing some valuable information that perhaps can only be encapsulated with a different representation.

Our current solution is to create a common default language with which all communicating entities communicate, but offers translation when necessary. The translation is carried out within the Personal Location Agent, in order to localize the translation burden. Even when using translation, it helps to reduce the problem to translating to and from a common representation in the PLA, rather than use a cross-product translation approach. Though a universal location representation for all entities and for all situations would inevitably be inadequate, a common language for all entities but for only communicating user-location seems much more reasonable. We choose to represent locations with semantic names for places and the various relations between places, simply because people seem to usually reason about locations in this way. Reasoning about location information is central to our purpose.

4.3 Relation-Map

Semantic names alone offer little with respect to a unifying location system. It is the semantic (e.g. geographical, physical) relations existing between the places that offers great value to our system. We name this knowledge representation for locations comprised of semantic names connected by semantic relations a *relation-map*.

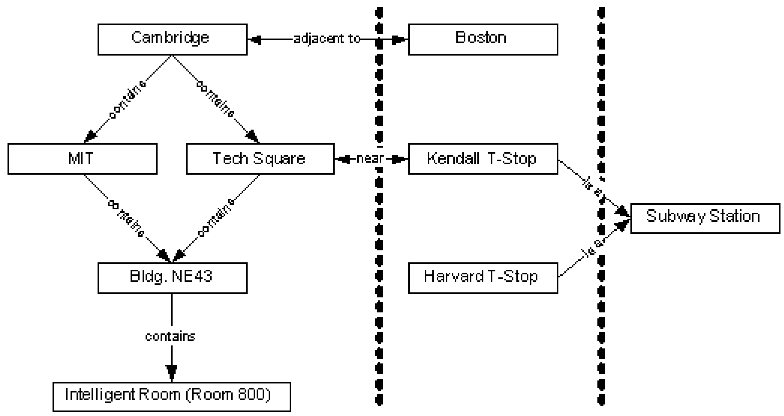

Fig. 3. Example of a Relation-Map

Collaborating Sources at Varying Granularities Without relations, a unifying location system cannot interpret location information at varying granularities as supporting or conflicting information. Consider Fig. 3 showing a portion of a relation-map. If two location devices return "Intelligent Room" and "Bldg. NE43," the system can determine that the clues support each other only if it knows that "Bldg. NE43" contains "Intelligent Room." Otherwise, "Bldg. NE43" and "Intelligent Room" compete with each other, and the user will be concluded to be either in one or the other.

Building Context Furthermore, PLACE uses a relation map not only to handle multiple collaborating sources at varying granularities, but also to serve as a database of information with which further inferencing can take place. Given that a user is in "Intelligent Room," we can use the relation-map shown in Fig. 3 to make several inferences, including, for example: "the user is at Cambridge," or "the user is near Kendall T-Stop." In this way, we combine sensor information with general knowledge about the world to build rich location context.

A Note on Context in a Relation-Map In Fig. 3, the dominant relation verb is "contains." Containment is the most common relation used among the various location representations that exist today. The primary advantage of using a containment hierarchy is that it provides the most desired organization for a world of places. However, many location representations either require a strict containment hierarchy of places or make it very difficult to represent multiple parent places. Brumitt and Shafer [29] recognize that the real world cannot be neatly "partitioned into unambiguous, non-overlapping pieces that exhaustively cover the area in question." The structure of a relation-map allows for easy representation of these multiple hierarchies.

While most location representations tend to focus on containment alone, a relation-map supports other location verbs such as "near" and "adjacent to." As shown above, these verbs are very useful in establishing further context by allowing for proximity and adjacency representations along with containment. Furthermore, a relation-map admits the possibility of annotation with additional entities and relations (such as "is a"). Such annotation would typically be done by and for applications that need the additional relational information. This could very easily extend to other context analysis such as behavior or activity deduction (i.e. context building not relevant to location) by simply adding relational verbs. If, in order to provide one of these features, we need to add a particular verb, the appropriate application will add it, as an annotation.

4.4 Access Control

Within the Intelligent Room and E21 projects, access control is treated as a separate and separable service, orthogonal to services like location awareness. Consequently, our design does not solve any access control problems per se, but simply uses access control mechanisms supplied as a service by the intelligent environment. However, there still remains significant design work on the location side, because we must specify the expressive power that we require of the access control system in order to allow the flexible access control that we need.

Location tracking is a delicate issue. The ability to find other people at any time is a very powerful utility; however, the ability for other people to find you at any time is a rather uncomfortable notion. People do not want to be stalked by strangers, suffocated by friends and family, or always locatable by business associates. Therefore, when using a PLACE service, it is important for users to have the functionality to precisely fine-tune the access control to their location information. Though widely used as a standard access control mechanism, Access Control Lists (ACLs) alone do not suffice because they do not take advantage of contextual information that users may find pertinent in their distribution of location information [26].

PLACE approaches this problem with access control based on who, why, when, how, what, and where. The following highlights potential uses for each type of control.

- Who is asking?
  - no strangers - keep undesired people from tracking you.
  - family only - parents feel more secure knowing the locations of their children.
  - travel groups - prevents individual members from getting lost.
- Why does someone want to know?

  - emergencies - be locatable by anyone during emergency situations that are relevant to you.
  - technical support - IT employees only want to be found if there is a legitimate technical support issue.
- When does someone want to know?
  - during scheduled meetings times - let other members of meetings know how far you are from the meeting.
  - during business hours (9am-5pm) only - the ability to quickly find resources is conducive to productivity and efficiency.
- How are they asking, or how is the location information provided?
  - give abstract location information, but not output of a camera.
  - give location information if they are asking over a secure network path.
- What am I doing?
  - sleeping/taking a vacation - regardless of where, do not disclose location.
  - giving lecture - allow all who are interested in your lecture to find out where you are speaking.
- Where am I?
  - not in library - prevent people from distracting you while studying.
  - at home - know when friends get home from school.
- Combination
  - no business contacts outside work hours - provide increased productivity at work while not allowing work to travel with you after work hours.
  - friends during designated social hours - allow friends to socialize with you only when desired.

The information of who, why, etc. simply contributes to user context, and PLACE utilizes user context to deduce permissions on user information, such as location. If additional context is available, users can add further precision to the access control of their information. Thus, as shown in the medical database domain [30], users can enhance their access control by incorporating context as additional variables and constraints.

With our software platform having an agent society per user, we can allow users to constantly obtain location information but control distribution of this information to services and other users. More specifically, the Personal Location Agent (PLA) working in a user's society can keep track of the user's location at all times, but control distribution of this information to those outside of the user's society. The primary benefit to this design is that the user can allow services working on his/her behalf to access his/her location at all times, while restricting access to foreign services.

Semantic Expressions The ability to establish location at varying granularities with a relation-map enables a user to perform such access control with greater ease. Using the convenient organization of places portrayed by a relation map, users can create more intuitive distribution rules. For example, instead of creating a rule stating, "If a user is in room-1, room-2, . . . , or room-n, then tell those in room-1, room-2, . . . , or room-n that the user is at work," we can instead create a rule stating, "If a user is in his work-building, then tell those in his work-building that the user is at work."

Furthermore, by customizing a relation map and inserting one's own perception of places and relations between places, a user can create rules that more adequately portrays his/her personal life. For example, with the customized relation map shown in Fig. 4, a user can create access control rules using semantic terms like "social place," "work place," "lunch place," etc.

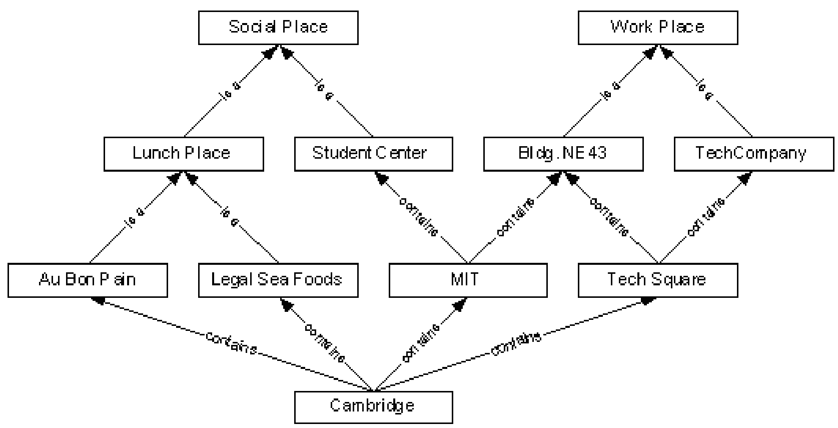

Fig. 4. Relation-Map with User Customizations

Distribution at Varying Granularities Users may also use the additional context available in a relation map to distribute location information at varying granularities. For example, a user may wish to have the following distribution:

- Let family members know the user's exact location.
- Let friends know the user's location at building granularity.
- Let acquaintances know the user's location at city granularity.
- Let strangers know the user's location at planet granularity.

4.5 Uniquely Identifiable Places

In the above examples, the semantic names for the places are casually assumed to be unique. Uniqueness is essential when communicating location because entities must not confuse two places with identical names. For example, more than one building can have a room called "Room 800." More than one institution can have a building called "Bldg. NE43." In the relation-map shown in Fig. 3, we do not explicitly address the potential ambiguities.

One solution would be to include a specification in each name, such as "Cambridge.MIT.Bldg-NE43.Room150." However, this approach has two problems. First, the specification must address all granularities in the universe of locations in order to ensure that there is not a duplicate place specification. That is, we must ensure that the universe of locations does not include more than one room with the name "Cambridge.MIT.Bldg-NE43.Room800." Second, the specifications imply a strict hierarchy of containment, an unrealistic perception of the world.

The only solution likely to work in the long run, solving most problems, and scaling to very large systems, is one that has the ability to actively translate symbolic information from an outside representation into an internal one, and vice versa. In order for this ability to work, it is necessary for systems to conduct investigatory dialogues to resolve ambiguities, remove conflicts and derive new concepts and associations. This is of course far too difficult, except in very crude approximation, for today's technology.

Instead, our solution is to have a distributed common map of locations among communicating entities. This allows for unique identification numbers to be attached to each location. All entities may communicate using these unique IDs and understand each other because they all possess the identical mapping from the IDs to the symbolic locations. At the same time, this common map of locations also solves the problem of having communicating entities possess the same vocabulary of locations.

We want to be clear that this solution is only a stopgap. It won't scale well, and carries with it enormous problems about generating, distributing, and coordinating information. We are only interested in solving those problems as much as is required to support our stopgap. Furthermore, it simply sweeps under the rug the fact that existing and new applications that depend on other representations for object and locations will still need to be semantically coordinated with the unique identifiers.

5 Applications

Our general design was inspired by the issues that we found inherent in the problem of sharing user locations. The following describes a series of applications that demonstrate the motivations behind the features of our design.

5.1 Patient Tracker

Patient Tracker uses location information of patients to track a patient in a medical facility and to provide doctors with appropriate and relevant information when approaching a patient. Whereas doctors in the ER might want to immediately know the patient's vital signs (including temperature, blood pressure, heart rate, respiratory rate, etc.), those in the recovery department might want to immediately know the patient's laboratory values (cell blood count, electrolytes, etc.) or study results (x-rays, CAT scans, etc.). In this way, Patient Tracker not only utilizes symbolic representation of physical location but also a relation-map semantic relation that categorizes the rooms in the hospital (see Fig. 5). In any case, whether given in certain departments or all departments,

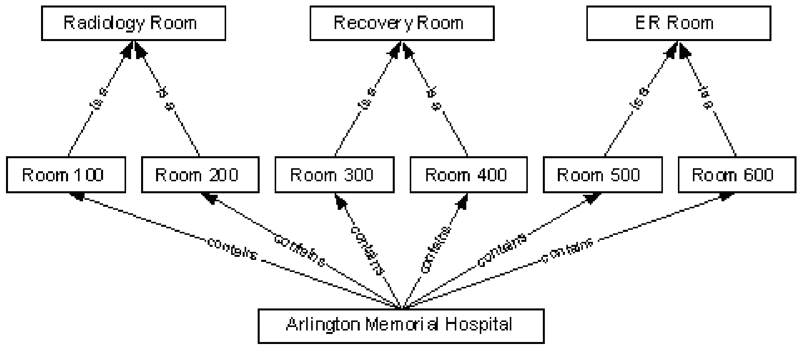

Fig. 5. Relation-Map used in Patient Tracker

patient information given to doctors before or upon arrival in preparation for treatment can save doctors time and energy. Furthermore, in terms of privacy issues, patients do not want to show their location to persons other than doctors and nurses responsible for their care, and appropriate visitors.

5.2 Expert Finder

Expert Finder uses location information to find people declared as experts of a specified topic. In some cases, it is much easier to ask someone for help rather than search through manuals and documents. For example, one might want to find the nearest IT employee for immediate on-site technical support. Unlike the Patient Tracker, Expert Finder potentially deals with people outside the scope of a building. Upon medical emergency, one might want to find the nearest doctor in the area. Thus, in addition to semantic representation and relation maps, the ability to interpret multiple location technologies and to perform sensor fusion is necessary to ensure that one is able to understand the doctor's location.

Even experts do not want to show their location while they are in a private location. However, in a significant enough emergency, that privacy concern may need to be over-ridden. Flexible, Role Based Access Control can solve this type of problem. In this case, when an expert comes close to a user in public space, the system will tell that to the user.

6 Conclusion

Driven by varying motivations, research groups have developed numerous location technologies and services that exist today. Traditionally, applications are rapidly built in response to the new and exciting location technologies available. As a consequence, applications are usually specific to one location technology.

We present a general design of a location system aimed to capitalize on multiple location technologies for the purpose of selectively sharing location information. Our design does not hope to serve as the general location infrastructure for all location services, but for services specific to user-location. Using a location representation consisting of semantic names and relations, we permit collaboration between location sources with different, yet related perspectives of the world. Similarly, we allow for further context inferencing that enables PLACE to distribute location at varying granularities and with intuitive access control. Finally, we demonstrate the utility of such features with a series of applications.

We hope the issues addressed in this paper will stimulate discussion about sharing location information, and will eventually enable us to communicate about locations seamlessly. We also think this idea can apply to other categories of context [31, 32].

7 Acknowledgements

The authors are grateful for support and advice from Howie Shrobe, for advice and suggestions from Krzysztof Gajos, and for support from Rod Brooks and the MIT Oxygen Alliance Partners: Acer, Delta, Hewlett-Packard, Nokia, Philips and NTT.

References

1. Dertouzos, M.: The Future of Computing. Scientific American, July 1999.
2. Ark, S. , Selker, T.: A Look at Human Interaction with Pervasive Computers. IBM Systems Journal, Vol.38, No.4, 1999.
3. Dey, A., Abowd, G.: CyberMinder: A Context-Aware System for Supporting Reminders. Proceedings of International Symposium on Handheld and Ubiquitous Computing (HUC 99), 1999.
4. Marmasse, N., Schmandt, C.: Location-Aware Information Delivery with ComMotion. In Proceedings of International Symposium on Handheld and Ubiquitous Computing (HUC 99), 1999.
5. Sekiguchi, M., Takayama K., Naito, H., Maeda Y., Horai H., Toriumi M.: NaVigation Markup Language (NVML). World Wide Web Consortium (W3C) Note, 6 Aug. 1999. http://www.w3.org/TR/1999/NOTE-NVML-19990806.
6. Long, S., Aust, D., Abowd, G., Atkeson, C.: Cyberguide: Prototyping context-aware mobile applications Proceedings of ACM CHI'96 Project Note, 1996.
7. Snoeren, A., Balakrishnan, H.: An End-to-End Approach to Host Mobility. Proceedings of the 6th Annual International Conference on Mobile Computing and Networking (MOBICOM 2000), Boston, MA, Aug. 2000.
8. Jose, R., Davies, R.: Scalable and Flexible Location-Based Services for Ubiquitous Information Access. Proceedings of International Symposium on Handheld and Ubiquitous Computing (HUC 99), 1999.
9. Oppermann, R., Specht, M.: A Context-Sensitive Nomadic Exhibition Guide. Proceedings of International Symposium on Handheld and Ubiquitous Computing (HUC 99), 1999.
10. Schmidt, A., et al. Advanced Interaction in Context. In Proceedings of International Symposium on Handheld and Ubiquitous Computing (HUC 99), 1999.
11. McCarthy, J., Meidel, E.: ActiveMap: A Visualization Tool for Location Awareness to Support Informal Interactions. Proceedings of International Symposium on Handheld and Ubiquitous Computing (HUC 99), 1999.

12. Hightower, J., Borriello, G.: Location Systems for Ubiquitous Computing. Computer, Aug. 2001, pp.57-66.
13. World Wide Web Consortium: Platform for Privacy Preferences Project. http://www.w3.org/P3P/.
14. Want, R., et al.: The Active Badge Location System. ACM Trans. Information Systems, Jan. 1992, pp. 91-102.
15. Korkea-aho, M., Tang, H.: Experiences of Expressing Location Information for Applications in the Internet. Workshop Proceedings of Ubicomp 2001: Location Modeling for Ubiquitous Computing, Sept. 2001.
16. Database Promotion Center: G-XML (Geospatial-eXtensible Markup Language) Protocol 2.0, March 2001. http://gisclh.dpc.or.jp/gxml/contents/index.htm
17. Open GIS Consortium: Geography Markup Language (GML) 1.0, Dec. 1999. http://www.opengis.org/techno/rfc11info.htm
18. MOSTEC: MObile Information Standard TEchnical Committee: POIX: Point Of Interest eXchange language, version 2.0 Document Revision 1, July 1999. http://www.w3.org/TR/poix/.
19. SignalSoft Corp. Wireless Location Services. http://www.signalsoftcorp.com.
20. Hall, D., Llinas, J.: Handbook of Multisensor Data Fusion, CRC Press, Boca Raton, June 2001.
21. Trusted Computer Security Evaluation Criteria, DOD 5200.28-STD. Department of Defense, 1985.
22. Krause, M., Tipton, H. F.: Handbook of Information Security Management. CRC Press LLC, 1998.
23. Baldwin, R.W. Naming and Grouping Privileges to Simplify Security Management in Large Databases. In IEEE Symposium on Computer Security and Privacy, 1990.
24. Robertson, P., Laddaga, R., and Shrobe, H.: Proceedings of the First International Workshop on Self-Adaptive Software Lecture Notes in Computer Science 1936, Springer-Verlag, 2001.
25. Laddaga, R.: Creating Robust Software Through Self-Adaptation. IEEE Intelligent Systems, Vol 14, May/June 1999.
26. Hanssens, N., Kulkarni, A., Tuchinda, R., Horton, T.: Building Agent-Based Intelligent Workspaces. In submission.
27. Coen, M., Phillips, B. Warshawsky, N., Weisman, L., Peters, S., Finin, P. Meeting the Computational Needs of Intelligent Environments: The Metaglue System. Proceedings of MANSE'99, Dublin, Ireland, 1999.
28. Gajos, K. and Kulkrani, A.: FIRE: An Information Retrieval Interface for Intelligent Environments. Proceedings of the International Workshop on Information Presnetations and Natural Multimodal Dialogue (IPNMD 2001) Verona, Italy, 2001.
29. Brumitt, B., Shafer, S.: Topological World Modeling Using Semantic Spaces. Workshop Proceedings of Ubicomp 2001: Location Modeling for Ubiquitous Computing, Sept. 2001.
30. Tzelepi, S., Koukopoulos, D., Pangalos, G.: A Flexible Content and Context-Based Access Control Model for Multimedia Medical Image Database Systems. Proceedings of ACM MMSig'01, 2001.
31. Naito, H., Takayama, K., Maeda, Y.: Situated Information. In the Proceedings of the International Workshop on AI in mobile Systems (AIMS2001), IJCAI 2001, pp.43-48, 2001.
32. Shafer, S.A.N., Brumitt, B., Cadiz, J.: Interaction Issues in Context-Aware Interactive Environments. Journal of Interactions, to appear.