The Location Stack: Multi-sensor Fusion in Action

Jeffrey Hightower and Gaetano Borriello

Dep't of Computer Science & Engineering
University of Washington
Box 352350
Seattle, WA 98195
+1 206 543 1695
{jeffro,gaetano}@cs.washington.edu

Intel Research Seattle
1100 NE 45th Street, Suite 600
Seattle, WA 98105
+1 206 633 6555
{jeffrey.r.hightower,gaetano.borriello}@intel.com

ABSTRACT

The Location Stack is a set of design abstractions and sensor fusion techniques for location systems. It employs novel probabilistic techniques such as particle filters to fuse readings from multiple sensor technologies while providing a uniform programming interface to applications. Our implementation is publicly available and supports many location sensor technologies. Specifically, our live demonstration tracks multiple people using statistical sensor fusion of RFID proximity tags and ultrasonic distance measurement badges. Participants are invited to don tracking badges and watch a projected visualization of the real-time probabilistic estimates of all participants' locations.

Keywords

Location sensing, sensor fusion, particle filters

INTRODUCTION

Location is essential information for many ubiquitous computing systems: We want our home to learn and respond to its inhabitants' movements. We want to capture and optimize workflow in a factory. We need directions from one place to another. We want to interact naturally with input-output devices casually encountered in the environment. Yet, to meet these goals, existing location-aware ubicomp systems can be improved in two areas:

1. Solid design abstractions can provide a common vocabulary for comparative evaluation of location systems.
2. Fusing readings from multiple different sensor technologies can exploit the advantages of each technology while presenting a single application programming interface that probabilistically represents location information.

Our contribution is in both of these areas. Based on lessons from a previous survey of location systems [1] we created the Location Stack, a common vocabulary and general framework for multi-sensor location-aware ubiquitous computing. In this demonstration, we highlight our Fusion layer's use of Bayesian filter techniques, more specifically, particle filters and multi-hypothesis tracking to estimate people's locations in real multi-sensor environments. Our implementation supports sensor fusion of many location sensor technologies including infrared proximity badges, passive RFID tags, ultrasonic ranging badges, active radio proximity tags, global positioning system receivers, infrared laser range-finders, 802.11b wireless clients, and, more importantly, any combination of these. Our architecture consists of scalable distributed services communicating with asynchronous XML messages and remote procedure calls, similar to many modern ubiquitous computing systems.

LOCATION STACK ABSTRACTIONS

The Location Stack codifies a set of layered abstractions based on properties identified in a previous survey of location systems [1] and the design experiences of several projects [2]. Figure 1 shows the Location Stack.

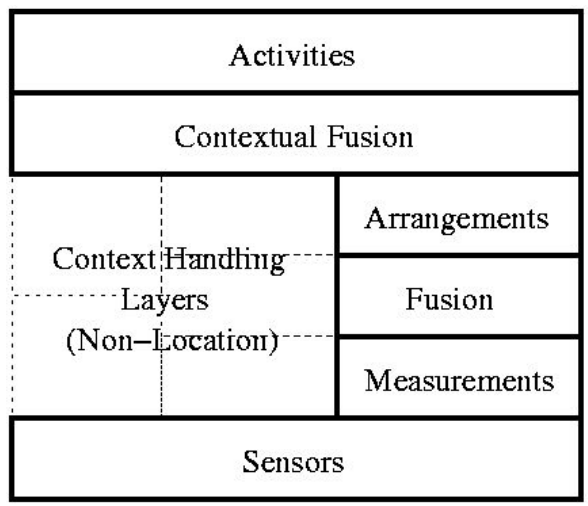

Figure 1: The Location Stack abstractions are a general framework and common vocabulary for location-aware ubiquitous computing systems.

We briefly discuss the layers and the interfaces they provide with particular emphasis on the fusion layer – the thrust of this demonstration.

Sensors

The Sensors layer consists of the sensing hardware for detecting a variety of physical phenomena. Our implementation has drivers for many common location technologies including infrared proximity badges, passive

RFID tags, ultrasonic ranging badges, active radio proximity tags, global positioning system receivers, infrared laser range-finders, and 802.11b wireless clients. Information is pushed up the stack as sensors generate new information about the changing state of the physical world. This demonstration uses passive RFID tags and ultrasonic ranging badges.

Measurements

Each sensor driver discretizes and classifies the data produced into measurements of type Distance, Angle, Proximity, or Position as well as several aggregate types such as Scan (a distance-angle combination). For example, infrared badges and RFID sensors both produce proximity measurements with likelihood models based on the power of the infrared emitters and the range and antenna characteristics of the radio. These measurement likelihood models describe the probability of observing a measurement given a location of the person or object. Such a model consists of two types of information: First, the sensor noise and, second, a map of the environment. The problem of constructing maps of indoor environments receives substantial attention in the robotics research community and is not our focus in this work.

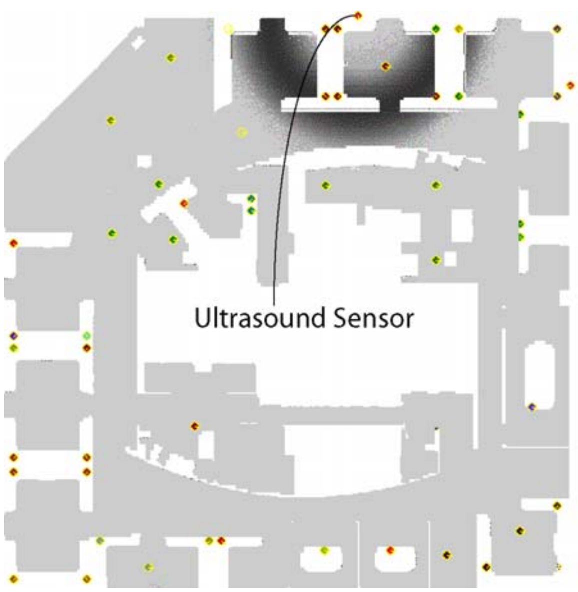

Figure 2: Measurement likelihood model ultrasound tags. Darker areas represent higher likelihood.

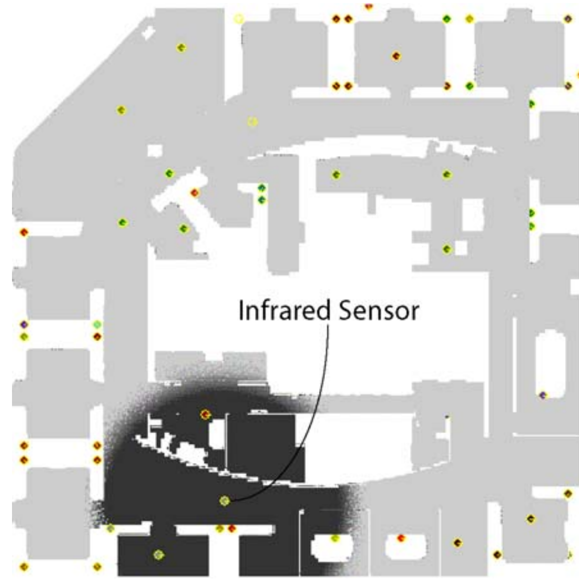

Figure 3: Measurement likelihood model for infrared proximity badges. Darker areas represent higher likelihood.

Figure 2 shows the likelihood model at all locations in our lab for a specific 4.5 meter ultrasound distance measurement. The likelihood function is a ring around the location of the sensor where the width of the ring is the uncertainty in the measured distance. Such noise may be represented by a Gaussian distribution centered at the measured distance. Furthermore, since ultrasound sensors frequently produce measurements that are far from the true distance due to reflections, all locations in the environment have some likelihood, as indicated by the gray areas in the map. White areas are blocked by obstacles. Figure 3 illustrates the sensor model for the infrared badge systems. Infrared sensors provide only proximity information, so likelihood is a circular region around the receiver. RFID tags are also a proximity technology and behave similarly.

Fusion

The Fusion layer continually merges measurements into a probabilistic representation of objects' locations and presents a uniform programming interface to this representation. In this demonstration we illustrate estimating the location of multiple people where each person wears an RFID tags and ultrasonic ranging badge. Due to these sensors' low accuracy (relative to robotics and motion capture sensors like precision scanning laser range finders), the belief over each person's location is typically very uncertain and often multi-modal, hence we apply a Bayesian filter techniques called particle filters which is commonly used in robot localization and is optimized for this type of scenario. Particle filters can naturally integrate information from different sensors. Refer to [3] for a general survey of Bayesian filtering techniques for location estimation or [4] for an in depth treatment of particle filters and Monte Carlo statistical techniques.

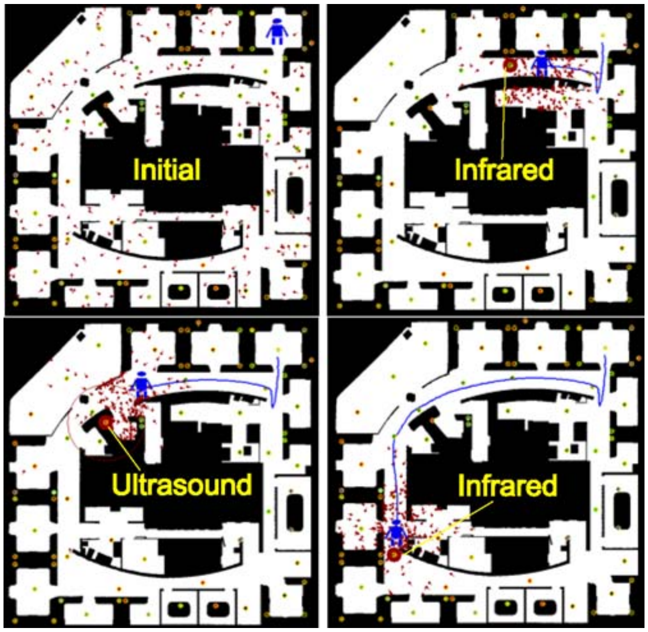

Figure 4: Sensor fusion of infrared and ultrasound sensors. Density of the particles reflects the probability posterior of the person's location.

Figure 4 shows snapshots from a typical sequence projected onto a map of the environment. In this example, the person is wearing an infrared badge and ultrasound tag and starts in the upper right corner as indicated by the icon. Since the start location is unknown to the system, the particles are spread uniformly throughout the free-space of the environment. The second picture (top right) shows the location probability after the person has moved out of the cubicles and into the upper hallway. At this point, the samples are spread over different locations. After an ultrasound sensor detects the person, their location can be estimated more accurately, as shown in the third (bottom left) picture in Figure 4. Later, after moving down the hallway on the left, the samples are spread over a larger area, since this area is only covered by infrared sensors that only provide very coarse location information (bottom right).

A single sensor fusion service running on a modern PC (1.8 GHz Pentium 4 with 512 MB memory) can perform real-time multi-sensor probabilistic tracking of more than 40 objects at a sustainable rate of 2 measurements per second per object. Objects are tracked in 7 dimensions (x, y, z, pitch, roll, yaw, and linear velocity). Higher performance (more objects or a faster measurement rate) can be realized by reducing the state space to two dimensions or through more advanced techniques such as our technique of constraining the particle filters to Voronoi graphs of the environment discussed below. Another way to increase performance is to distribute computation across multiple fusion services, although applying certain Arrangements layer operators then poses additional challenges.

There are two pieces of additional research we have contributed but are not highlighting in this demonstration. First, we have shown how particle filters can be used more efficiently by constraining possible locations of a person to locations on a Voronoi graph of free space that naturally represents typical human motion along the main axes of the environment. In experiments we found that such Voronoi graph tracking results in better estimates with less computation. Furthermore, the Voronoi graph structure can be used to learn high-level motion patterns of a person. For example, the graph can capture information such as "Rebecca goes into room 22 with probability 0.67 when she walks down hallway 9." More details on using Voronoi graphs with particle filters and on applying high-level learning can be found in [5,6]. Second, although also not shown in this demonstration, other work of ours at the Fusion layer has addressed the problem of estimating objects' identities in situations where explicit identity information is not provided by the all the sensors. In particular, we have introduced a technique to combine highly accurate anonymous sensors like scanning infrared laser range finders with less accurate identity-certain location technologies like infrared and ultrasonic badges [7].

Arrangements

We provide two operators to relate the locations of multiple objects. We provide a test for multi-object proximity given a distance and a test for containment with a map region. Because we operate directly on the location probability posteriors of each object, the results of these tests can also be probabilistic. For example the proximity test produces a pairwise confidence matrix that a given group of objects are within 4 meters of one another. Taken together, these operators provide a probabilistic implementation of the "programming with space" metaphor as used with great success in AT&T Sentient Computing project [8]. Future work in our implementation of the Arrangements layer is to provide an additional operator to test for more general geometric formations of multiple objects.

Context and Activities

The Contextual Fusion layer combines location information with other contextual information such as personal data (schedules, email threads, contact lists, task lists), temperature, and light level while the Activities layer categorizes contextual information into semantic states defining an application's interpretation of the world. Our implementation of the Context and Activities layers is in its infancy because few ubiquitous computing systems have been deployed which take sensor information all the way up to the level of human activity inference. To make inroads, we are collaborating with the Assisted Cognition research group, a group seeking to create novel computer systems that will enhance the quality of life of people suffering from Alzheimer disease and similar cognitive disorders [9]. Our goal for this collaboration is to design general interfaces for the Context and Activities layer based on usage patterns of the existing Fusion and Arrangements layers in support of these higher level learning tasks.

SUMMARY

Our demonstration highlights the primary capabilities of our Location Stack implementation: We show a highly flexible system which can track multiple people using statistical sensor fusion of information from multiple sensor technologies, in this case, RFID proximity tags and ultrasonic distance measurement badges.

The Location Stack abstractions structure location systems into a layered architecture with robust separation of concerns allowing us to partition the work and research problems appropriately. Our implementation is a publicly available Java package containing a complete framework for operating a multi-sensor location system in a ubiquitous computing environment. The implementation is typical of a modern ubiquitous computing system: a set of reliable distributed services communicating using asynchronous XML messages and linked using dynamic service discovery capability in the middleware. The Location Stack is deployed in our laboratory and workspace at Intel Research Seattle, operates nearly 24x7, and is used by other research projects as a reliable source of location information.

REFERENCES

1. Hightower, J. and Borriello, G. Location systems for ubiquitous computing. *Computer*, 34(8):57–66, August 2001.
2. Hightower, J., Brumitt, B., and Borriello, G. The location stack: A layered model for location in ubiquitous computing. *Proceedings of the 4th IEEEWorkshop on Mobile Computing Systems & Applications (WMCSA 2002)*, pages 22–28, Callicoon, NY, June 2002. IEEE Computer Society Press.
3. Fox, D., Hightower, J., Liao, L., Schulz, D., and Borriello, G. Bayesian Filtering for Location Estimation. *IEEE Pervasive Computing*, vol. 2, no. 3, pp. 24-33, IEEE Computer Society Press, July-September 2003
4. Doucet, A., and de Freitas, N., Gordon, N. editors. *Sequential Monte Carlo in Practice*. Springer-Verlag, New York, 2001.
5. Liao, L., Fox, D., Hightower, J., Kautz, H., and Schulz, D. Voronoi tracking: Location estimation using sparse and noisy sensor data. *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2003.
6. Rabiner, L. R. A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE*. IEEE, 1989. IEEE Log Number 8825949.
7. Schulz, D., Fox, D., and Hightower, J. People Tracking with Anonymous and ID-Sensors using Rao-Blackwellised Particle Filters. *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI)*, 2003
8. Addlesee, M., Curwen, R., Hodges, S., Newman, J., Steggles, P., Ward, A., and Hopper, A.. Implementing a sentient computing system. *Computer*, 34(8):50–56, August 2001.
9. Kautz, H., Etzioni, O., Fox, D., Weld, D., and Shastri, L. Foundations of assisted cognition systems. *UW-CSE 03-AC-01*, University of Washington, Department of Computer Science and Engineering, Seattle, WA, March 2003.

Demonstrations Supplement
Ubiquitous Computing Conference

CONTACT INFO

First Name	Jeffrey
Last Name	Hightower
Organization	University of Washington
Street Address	Box 352350
City	Seattle
State/Province	WA
Country	USA
Postal Code	98195
Daytime Telephone	206-545-2517
Email	jeffro@cs.washington.edu
URL	www.cs.washington.edu/homes/jeffro/

DEMONSTRATIONS DESCRIPTION

Title: The Location Stack: Multisensor Fusion in Action

Project Description (100 words max):

The Location Stack uses novel statistical techniques such as particle filtering to perform real-time location sensing by fusing measurements from multiple sensor technologies.

Envisioned Interaction:

UbiComp attendees will have the opportunity to experience live multisensor probabilistic location sensing. A space will be equipped with several RFID detectors, ultrasonic sensors, and perhaps additional sensor technologies as prep time allows (we support many types of sensor technologies). Participants will be invited to don tracking badges. A visualization showing the real-time probabilistic locations of all participants will be projected on a large screen showing a 2D top view of the environment.

TECHNICAL REQUIREMENTS

SPACE

Ample free space is very important for an effective location sensing demo. As such, this demo requires a minimum of 15'x15'x7' (L/W/H). An isolated room with a larger size (e.g. 30'x30') is desirable to allow more simultaneous participants, more freedom of motion for each participant, and simpler logistics. If not in a dedicated room, a corner is desirable to simplify sensor placement on the walls. Finally, it is very desirable to be informed as to the exact dimensions and floor plan of the allocated space before arriving at the conference.

ACOUSTICAL

Some of our sensors use ultrasound, but audible sound (created by us or heard from other demos) is not a constraint for this demo.

LIGHTING

No lighting constraints exist except it should be dark enough to be able to clearly display a visualization using an LCD projector.

TIME

The demo should run at the opening demo reception because it requires an overseer to distribute sensors to participants and collect the sensors as they leave. The demo can be experienced by multiple participants limited by the size of the space allocated.

COMPUTATIONAL EQUIPMENT

The demo involves 3-4 PCs, 2 LCD projectors, and several pieces of custom sensor hardware, all of which we will bring. UbiComp could provide a projection screen if the walls are not suitable for projection.

NETWORKING

We use a small isolated network of machines. External network access is not necessary nor will 802.11b be required.

RADIO FREQUENCIES

Sensors emit ultrasound, 916.5MHz RF (Berkeley motes), and multiple RFID emissions of moderate intensity around 915MHz.

POWER

3-4 PCs, 2 projectors, and ~12 wall warts for sensors. It is desirable to have power strips distributed on each side or corner of the space.