

iCam: Precise at-a-distance Interaction in the Physical Environment

Shwetak N. Patel¹, Jun Rekimoto², and Gregory D. Abowd¹

¹ College of Computing & GVU Center
Georgia Institute of Technology
801 Atlantic Drive, Atlanta GA 30332-0280 USA
{shwetak, abowd}@cc.gatech.edu

² Interaction Laboratory
Sony Computer Science Laboratories, Inc.
3-14-13 Higashigotanda
Shinagawa-ku, Tokyo 141-0022 Japan
rekimoto@csl.sony.co.jp

Abstract. Precise indoor localization is quickly becoming a reality, but application demonstrations to date have been limited to use of only a single piece of location information attached to an individual sensing device. The localized device is often held by an individual, allowing applications, often unreliably, to make high-level predictions of user intent based solely on that single piece of location information. In this paper, we demonstrate how effective integration of sensing and laser-assisted interaction results in a handheld device, the iCam, which simultaneously calculates its own location as well as the location of another object in the environment. We describe how iCam is built and demonstrate how location-aware at-a-distance interaction simplifies certain location-aware activities.

1 Introduction and Motivation

We have seen great progress in our research community towards the goal of practical, precise indoor localization. A variety of techniques, including those that introduce new infrastructure (e.g. ultrasound [20, 2], camera tracking [28], ultra-wideband [24]) and those that leverage existing infrastructure (e.g. 802.11 [3, 12], GSM [18], Bluetooth [14]) show that we are not far off from having everyday devices that know where they are in the physical world.

Location-aware applications are limited, however, when using only knowledge of a single object's location. This location information is usually that of the device itself, and applications assume that the device is in the possession of an individual. In that case, device location relates to the owner, and services provided are dependent on the owner's assumed location. This technique is unsatisfactory for applications, such as the canonical mobile tour guide, in which focus of attention, not just location, may be the desired trigger for delivering information (e.g., for what object is the tourist wanting further information) [1, 8].

One augmented reality solution is to gain precise location and orientation information of the individual and project information about the world in the field of view of that individual. However, it can be difficult to place content in the environment, that is, to create the link between physical location and virtual information. Another solution is to tag the environment with glyphs and recognize those tags, usually through some form of computer vision or active tagging approach. Tagging is not always feasible or aesthetically desirable, though this solution does allow for objects to be moved in the environment and does not require precise localization.




Figure 1: The iCam handheld. The back of the device (shown on the bottom) houses the laser pointer and the camera in addition to the collection of sensors inside the housing.

In this paper, we provide another alternative: augmenting a precisely located device with a laser-based range finder so that it can also accurately determine the location of an object at-a-distance. We describe the iCam handheld device as a demonstration of this concept (shown in Figure 1). iCam integrates a commercially available indoor ultra-wideband positioning system [24] with a magnetic compass, an accelerometer, a camera, and a laser pointer. This portable handheld device tracks its own location and orientation to within 7 cm and 1.5 degrees of rotation. By using the laser pointer, a user can point to any object or surface in the physical environment and iCam calculates that object's absolute 3D location to 20 cm accuracy from a distance of up to 5 meters for our prototype. With this capability, we demonstrate how one can use iCam to place information at any arbitrary location in the environment and how it can greatly simplify both the authoring of location-triggered content as well as calibration. The same iCam handheld features a camera which allows users to view any digital content through an augmented viewfinder.

The important contribution of this paper is a demonstration of how simultaneous knowledge of two pieces of location information simplifies aspects of the location-aware experience, both from the user and developer perspectives. Among those tasks that are made easier are overall system calibration, creation of a map of the physical environment, and attaching virtual information to physical locations.

In this paper, we first describe the related work and then illustrate the user experience this system enables. We present a location-based tour guide application that leverages iCam's unique capabilities, describe the detailed implementation of iCam, and then finish up with future work and conclusions.

2 Related Work

Although comprehensive surveys of location technology and location-aware applications are not practical for this paper, we will highlight some notable examples. Researchers have extensively studied indoor location technologies. Hightower provides an overview of the various location technologies and techniques [9]. The two basic approaches are to build the entire sensing infrastructure from the ground up (*e.g.*, ActiveBadge [26], Cricket[20], and Active Bat[2]) or to leverage existing infrastructure that can yield localization, either through triangulation or fingerprinting (*e.g.*, 802.11 work such as RADAR[3] and Place Lab[12], GSM [18], and Bluetooth [14]). Applications that leverage some of these location technologies started with the original location research at Olivetti and Xerox PARC, where researchers used office tasks as motivation to create experiences such as Audio Aura, auto call forwarding, and desktop migration [16, 26]. Location information has also been used to trigger events and reminders in applications like Forget-me-not and CybreMinder [6, 11]. Later, researchers utilized location-awareness in tourist applications to help people navigate and explore unfamiliar spaces, such as with CyberGuide [1] and GUIDE [4].

Location has been used both implicitly and explicitly to attempt to describe and interact with the physical world. Researchers have explored selection of physical objects in an environment for various augmented reality tasks (*e.g.*, NaviCam [23]), which includes ways to tag the physical world using printed barcodes, 2-dimensional

glyphs [22, 23], RFID, and active beacons [25] to connect the physical and electronic worlds. The applications created with this technique demonstrate the potential of leveraging knowledge of the world beyond just the location of the device or individual. However, the current solutions that use static labels (such as barcodes or 2-dimensional glyphs) are limited in distance due to the camera resolution and perception techniques used to decipher the glyphs. Static labels are also not practical to deploy in highly interactive spaces because of the difficulty inherent to placing labels on every object.

Previous research projects have also explored laser pointer interaction, both for interaction at-a-distance with large displays [5, 10, 15, 17] and for the selection tasks of physically tagged objects [19, 21]. A popular laser pointer interaction scheme is to use a camera focused on a region of a wall or object in which a laser spot may appear. Simple computer vision techniques locate the red laser dot and follow it around the interaction region. Such a scheme is appealing for meetings or presentations, during which one can interact with a display from a distance by simply pointing at it with an ordinary laser pointer. Two-way laser pointing techniques have been proposed using active tags placed in the environment, but these are subject to the same scalability limitations as static tags mentioned previously.

The XWand system demonstrates how individuals can use laser pointing techniques with a collection of other sensors to support selection of and interaction with devices through a special-purpose device [28]. Although this camera-based tracking solution provides very flexible ways to point and select objects within the environment, it requires significant overhead in terms of camera infrastructure. An extension of XWand, called the World Cursor, removes the vision requirement by using the XWand to steer a remotely controlled laser pointer around a room [27]. The remotely controlled laser pointer has a model of where it is pointing in 3D-space and has sufficient geometric information to know where its red laser dot is pointing. The drawback of this is that it is limited to objects that are within the line-of-sight of the steerable laser. For example, it would be difficult to interact with the sides of an object placed against the wall. iCam's mobility addresses some of these problems.

The drawback of these approaches is the time-consuming setup, calibration, and the difficulty of adding or editing digital content. Most systems assume a model of the space and attributes of the objects that are preprogrammed. Although iCam requires a 3D model of the space, we provide a very quick and easy means to define that model.

3 Demonstrating the iCam Experience

3.1 The iCam Experience

iCam supports two basic modes of interaction. The first is a simple seeking or gathering of information from the physical and virtual space. The second is defining the geometry of the physical space and the digital content that may appear within it. This

includes the calibration of the infrastructure, the mapping and defining of the physical space, and the authoring and attaching of virtual content.

A laser pointer mounted on the handheld produces a bright red dot that iCam uses to target objects or places of interest. For example, a user can point the visible laser dot at a light switch or trace the outline of a door by moving the dot around the door-frame (see Figure 2). Additionally, users can interact with the space between the handheld device and the position where the dot lands by moving a virtual cursor on the viewfinder up and down the beam produced by the laser pointer. This movement of the cursor allows interaction within the free space where there may not be physical artifacts for the dot to actually land on. If the user is interested in scanning the space for information, iCam can be used similarly to a video camera (such as when recording a video of a scene) to find the relevant information in the area.

3.1.1 Seeking Information

For seeking information, iCam is used much like a video camera in that the user can scan the space and see augmentations of digital content, such as textual information or pictures, over the live view of the physical space (see Figure 6). A user can also interact with the physical space by, for example, actuating a light switch by pointing at it and pressing a virtual button on the handheld. More complicated widgets may also be available in the space. In the light switch example, the physical switch could be augmented with additional control options (*e.g.*, timing, mood lighting, or dimming capabilities) that appear over the physical switch when the user points at it. The user can then use the virtual cursor in the augmented view to interact with the virtual switch.

iCam also supports zooming to obtain different levels of detail in the scene. If the user is interested in viewing the available virtual content for a space, he can use the zoom out feature to view a large amount of available content quickly and to reduce the amount of movement for scanning the space. Likewise, zooming into a scene can reveal more detailed information. Zooming also enables long distance interaction. If the laser dot is not easily viewable, the user can zoom in and use the virtual cursor to do fine-grained movements from a distance.

3.1.2 Defining the Physical Space and Placing Content

iCam provides an easy method for defining the physical layout of a space (see Figure 2). Users can trace parts of the physical space, similar to outlining with a pen, by using the laser dot as the visual feedback. The viewfinder provides feedback of the trail or “ink” left behind with the trace, which allows users to view exactly what was marked. The accuracy of the system allows users to provide detailed selections of the areas of interest. This selection method allows a person to produce a geometric model quickly by tracing the outlines of the walls, doors, and windows of a room.

We provide a very simple content addition mechanism called *beaming*. Beaming involves the attachment of authored content, such as a note, anywhere in the physical space. Users can create content using the iCam interface or import it from another

computer system. After producing the content, the user can either beam it to where the laser dot lies or place the authored content in free space by moving the virtual cursor displayed in the viewfinder along the laser beam.




Figure 2: Example of a user tracing the door frame to create a map of the space. Notice the virtual ink left behind through the viewfinder for visual feedback.

3.2 Tour Guide Application

To demonstrate the use and capabilities of the iCam system, we revisit the canonical tour guide example. Although this is a popular application in many location-aware systems, an often unaddressed but important task for these applications is creating the digital representation of the physical space. Especially for precise indoor solutions, some level of a geometric model of the space is a necessary component. Thus, we focus on a tour guide application that provides an easy way to create the tour itself.

Our system provides a very easy way to define the map of the physical space by using a special “learning” or mapping mode. A user can walk around the space and add annotations to physical objects. The tour creator would create the content on the iCam and beam it to the appropriate physical artifact by pointing at it. Arrows can also be placed in the environment to suggest where to go next from each exhibit. In this case, the tour creator would find an appropriate location and draw a line where an arrow should appear. The creator can also make free hand annotations by using iCam as a pen to mark and produce callouts. Additionally, a user can take a snapshot of the current view, use a stylus to draw directly on the display, and then place the annotation back into the environment. This is desirable since it may be difficult to control the laser for certain strokes, especially from very far distances.

Figures 3-6 show the use of a simple content creation interface with iCam. The interface consists of a blank note on which the user can place any combination of pictures, text, free-hand writing, or short audio clips. The user selects the type of content he wants to add from the toolbar and places it on the note. He can insert pictures by looking through a list of image files on the device, enter text by using the handheld's virtual keyboard, add free-hand writing by directly writing on the screen with a finger or stylus, or add an audio note by speaking into the onboard microphone. He may also import pre-authored content from a remote computer and attach it to the notes created with iCam.

During a tour, the iCam displays the annotations and virtual content through the viewfinder while it is pointed toward the relevant parts of the physical space. Every user will see the same content from every vantage point, similar to how it would look if it were a physical artifact in the space. A user can view all the available annotations in an area by scanning the space with the viewfinder and using the zooming capabilities. Knowing the user intent and focus is an important feature in a mobile tour guide and here it is made explicit without needing to infer based solely location.



Figure 3: The user scans the physical environment to find a place to add new content.




Figure 4: The user authors the content using the iCam's notes interface.



Figure 5: The user beams the created note to the desired location on the bookshelf. Note the red laser pointer on the green book.



Figure 6: The user views the same note from a different vantage point.

4 iCam Implementation Details

iCam is a handheld device (shown in Figure 1) that can accurately locate its position and determine its absolute orientation when it is indoors. It can also determine the 3D position of any object to which points with its onboard laser pointer. iCam is built using a Sony Vaio Type-U handheld instrumented with a variety of onboard sensors. The device can localize its position in 3D space to within 7 cm of accuracy, determine its orientation (azimuth, tilt, and roll) within 1.5 degrees, and determine the 3D location of other objects and surfaces within 20 cm in all directions.¹ It accomplishes this by integrating a modified version of a commercially available, ultra-wideband location system (from Ubisense [24]), a 3-axis magnetic compass, a 2-axis accelerometer, and a laser pointer tracked by the handheld's camera. The net location update rate is 15 Hz, which is enough for most interactive applications.

¹ The 20 cm is an average estimate for the error in our instrumented space of 10 m x 15 m and interacting at distances of up to 5 meters. This value would increase for much larger spaces because of potential angular errors. It is also affected by the resolution of the camera, as we later describe.

4.1 3D Position

The handheld determines its location using the Ubisense ultra-wideband location system installed in the environment. We chose to use the Ubisense system because of its commercial availability and its ability to handle large spaces with mild to moderate occlusions. The handheld device is instrumented with an active tag that constantly emits its identity at very high frequency bands (5.8-7.2 GHz). Sensors placed in the environment detect these signals and triangulate the tag's location based on the signal's time and angle of arrival. The advertised average accuracy of the system is 10-15 cm. However, after modifying the tag's antenna (replacing it with a copper cube antenna with larger surface area) and strategically placing 6 sensors in the space, we obtained average accuracies of 7 cm in all directions. The accuracy was measured by averaging differences between iCam's calculated positions from actual known locations which were measured using a standard tape measure. We instrumented a 10 m x 15 m space using six ultra-wideband sensors. We placed more sensors in the areas where there were more occlusions and thus a greater potential for multi-path reflections (e.g., around desks, structural supports, etc).

4.2 3D Orientation

The iCam system accomplishes absolute orientation with an Aichi Mi AMI201 3-axis magnetic compass and an Analog Devices ADXL series 2-axis accelerometer. The compass determines the handheld's azimuth or bearing angle and the accelerometer determines the tilt and roll angles. Because magnetic compasses are only accurate when held parallel to the ground (perpendicular to the gravitation axis), the accelerometer and the third magnetic axis provide a means to compensate for the tilt and roll angles and keep the compass electronically gimballed regardless of level. This allows the user to move freely with the handheld and still obtain very accurate bearing information.

A common problem indoors is the magnetic interference produced by some consumer electronic devices such as televisions and computers. The iCam mitigates this problem by constantly monitoring the magnetic values in all directions and dynamically subtracts out any abnormal readings. Since it is difficult to calibrate out all time-varying interferences, especially small changes, we only focus on significant magnetic deviations. iCam accomplishes its dynamic calibration by comparing the current magnetic values with that of the initial calibration data and the history of readings up to that point. The initial calibration is done in an area that is known to have minimal or no magnetic interference. The device also stores the recent component values from the compass (sampled at 15 Hz). iCam detects magnetic disturbances by comparing the magnitude value from the calibration sequence to the current magnitude. Any significant magnetic interference in the environment will cause a magnetic spike in all three directional components of the compass, thus greatly increasing its net magnitude. When this phenomenon occurs, iCam replaces the current magnetic values with the most recent unaffected value. This dead reckoning approach works well for short-lived magnetic disturbances.

4.3 3D Range-finding

By coupling the handheld's position and orientation with its distance from an object, it is possible to determine the 3D position of that object. We do this by using the camera to track the red dot produced by the onboard laser pointer. We chose the laser tracking approach over ultrasound for two reasons. First, we already had plans to use a laser pointer for visual feedback, so it was natural to leverage its capability. Second, ultrasound is not collimated enough to aim precisely at a small surface.

The laser diode is mounted at a known fixed position parallel to the camera (see Figure 7). The onboard camera has a resolution of 640x480 pixels and horizontal field of view of approximately 50° and vertical field of view of approximately 40°. The iCam projects a laser beam onto an object or surface in the field of view of the camera. The onboard camera captures the dot from the laser along with the rest of the scene and a simple algorithm runs over the image looking for the brightest pixels. We then calculate the range to the object based on where along the vertical axis of the image this laser dot falls. The closer the dot is to the center of the image, the further away the object is. The distance is calculated using the angle, ϕ , between the camera's central focal point and the position of the dot in the camera's view (see Figure 7).

The equation used for this distance is:

$$D = \frac{h}{\tan \phi},$$

$$\phi = C * P$$

where h is the distance between the camera and laser diode centers, and ϕ is the number of pixels from the laser dot to the center of focus (P) multiplied by the radians per pixel constant (C). The radians per pixel constant (C) is an approximation that we derived in the lab through a series of calibration sequences with known distances D . The values are then averaged to determine the final approximation for C. Another reason for the approximation is the curvature of the lens with respect to the focal plane. Depending on which part of the image is used, the radian per pixel value (C) changes slightly. However, this change is very slight, and the approximation works well for our system.

Since the laser dot only appears on the vertical axis near the center of the image, we limit the detection algorithm to only that region. This helps prevent false positives from other bright light sources that may appear in the image. This approach works reasonably well indoors with about 20 cm accuracy for distances between 1 and 5 meters. The accuracy and range are limited by the resolution of the image. In our simple setup, we spread 240 pixels vertically across the 20-degree half field of view, because the laser is parallel to the optical axis of the camera. By using 3 sub-pixel analysis we are able to obtain 20 cm accuracy. Mounting the laser horizontally instead of vertically to the camera would improve the accuracy due to the slight in-

crease in horizontal resolution (320 pixels vs 240 pixels). This produces an increase to about 16 cm for the 1-5 m range. The 20 cm accuracy currently requires 3 sub-pixel analysis and increasing that to 10 sub-pixels would provide about 7 cm accuracy for the 240 pixel vertical resolution case. In addition, an optical zoom capability would also greatly improve accuracy and give the user more precise control for interacting with distant objects. However, this requires manual operation of the lens and optical system by the user.




Figure 7: Left: Camera and laser diode mounted parallel to each other behind the iCam handheld. Right: Diagram showing the components of the distance equation stated previously.

4.4 Calibration

The system requires two calibration steps: the orientation sensors on the handheld and for the location system in the environment.² To calibrate the orientation sensors, iCam prompts the user to hold the handheld level to the ground and then fully rotate it around all three axes.

Surveying and calibrating the Ubisense location sensors is typically a fairly time-consuming and tedious task. Since the sensors are mounted at an elevated position, it is difficult to measure absolute positions accurately between multiple sensors. Often, to find the proper setup it requires testing several different configurations. Thus, the overall process can be a very long task, involving the following steps:

- 1) placing the sensors in the environment;
- 2) surveying the location of the sensors (both relative to the space and each other);
and
- 3) calibrating each of the sensors to known points in the environment

² The ranging system does not need to be calibrated by the end user because the needed values are determined during construction and do not change in different environments.

By leveraging the capabilities of iCam, we can accelerate these steps and greatly ease the overall setup burden.

The Ubisense calibration software requires two pieces of information: the location of the wall-mounted location sensors and mobile sensor readings from known points in space. To survey the locations of the wall-mounted sensors, we initially pick one pair that can be seen by iCam from the same location. Keeping the iCam in the same location in the room, we measure the (X, Y, Z) locations of these two Ubisense sensors using the iCam's range-finding and orientation capabilities. The laser pointer provides visual feedback to assist in aiming. In addition, all the Ubisense sensors have a square marking on the foreshore which serves as a consistent reference point for aiming. Then we pick a new wall-mounted sensor and move iCam to a location where this sensor and one of the previous pair are both visible. We measure the (X, Y, Z) locations of these two from the new vantage point and proceed in the same pairwise fashion around the room until each wall-mounted sensor has been paired and measured with a previously measured one. In the end, we resolve the relative translations between all the pairs into 3D coordinates in a common coordinate system. For this procedure, we can assume that all the coordinate frames are parallel in orientation, since the wall-mounted Ubisense sensors are effectively omni-directional when laid out.

While carrying out these surveying measurements, we also record readings on the mobile Ubisense sensor mounted on iCam. After resolving all the 3D measurements into one coordinate system, these measurements serve as the mobile sensor readings for the other half of the Ubisense calibration input (step 3 from above). Additional readings can also be recorded afterwards to improve position accuracy.

4.5 Overall system architecture

Figure 8 shows the overall system architecture. The position sensors in the environment are connected through ethernet to a PC. The location software runs on that PC and transmits all resulting location information back to the iCam handheld device via 802.11b. The handheld computes its orientation and its range to other objects locally and then wirelessly transmits these values back to the PC. The handheld's application is written in Java and parts of the user interface (viewfinder and 3D overlays) are created using the OpenGL GL4Java extension. The geometric model and attributes of the physical and virtual space are stored on the PC, but are cached locally on the handheld during startup to speed up interaction. As the user modifies aspects of the space, the system updates the locally stored model and notifies the PC of the appropriate changes.




Figure 8: Overall system architecture of the iCam system.

5 Conclusions and Future Work

We have presented a handheld device called iCam that precisely determines its own location and orientation as well as the position of objects and surfaces in the physical environment. In addition to localizing itself and other objects, iCam serves as an interaction device for augmented reality applications. iCam offers several advantages over other similar devices. It provides accurate position information of both the handheld itself and a point in space indicated by a laser pointer. It uses a practical infrastructure of commercially available equipment to achieve this dual localization. iCam's simple mechanism for interaction between the physical and virtual space facilitates an easy way to map the physical space. Finally, iCam simplifies the calibration of the location infrastructure by making it easier and faster. We have not previously seen a single device that can simultaneously and affordably accommodate calibration and content placement while also serving as a primary end user interaction device.

iCam's ability to localize objects accurately from a distance greatly simplifies the mapping process by allowing the users to define exactly what they intend *in situ*. We have demonstrated this by explaining how iCam would be used to place content for a location-aware tour guide, but this is certainly just the first step in exploring its capabilities. A system like iCam enables many other augmented reality applications that leverage the user's explicit intent and actions with the physical space. One can imagine other practical applications that take advantage of these capabilities. For example, suppose someone is moving into a new space and is trying to decide where to place furniture. He has hired a moving company to handle the actual moving of the furniture, but the person is unable to be there when the actual moving takes place. To ensure that the movers place the furniture where the person needs it to be, the user can use iCam to mark spaces around the room where he would like the furniture placed (*e.g.*, trace a box around the south side of the room and attach a note with the

text "black bookcases goes here"). When the movers come to the space, they can then use iCam to view where he has designated the furniture to be placed.

The tour guide and the last example fall into the larger class of groupware of applications that support asynchronous, collocated collaboration. Inventory location for customers (*e.g.*, helping someone find a book in a bookstore or a bottle of wine at a store) is another example. For individuals with mobility challenges, the precise, at-a-distance interaction can effectively extend their reach. These individuals, for example, could use iCam to indicate to caregivers where to find or place certain items around their home.

Authoring of location-aware content can also be generalized to the description of behaviors that are associated to locations. Imagine developing a remote control that can point to a device and then control it. The "programmer" of the remote would need to indicate what physical space is occupied by the device, so that when a user points to that space subsequent control commands are sent to the correct logical device. This suggests a model of "programming the environment" in a very literal sense, and we see promise for this kind of use of location-enhanced, at-a-distance interaction that can extend other recent work in programming simplifications for Ubicomp [7, 13].

These applications require simultaneous knowledge of the user interface's own location and orientation as well as the position of objects and surfaces in the physical environment. Our iCam device enables exploration of this kind of sophisticated location-aware experience, both from the user and developer perspectives. Among those tasks that are made easier are overall system calibration, creation of a map of the physical environment, and attaching virtual information to physical locations.

Acknowledgements

The authors thank all the members of the Ubicomp Research Group, and in particular Julie Kientz, Khai Truong, Gillian Hayes, Giovanni Iachello, and Jay Summet, as well as members of the Interaction Lab at the Sony Computer Science Laboratory in Japan. We also thank John Krumm from Microsoft Research for his help shepherding the final draft of this paper. This work is sponsored in part by the National Science Foundation (Grant No. 0513111).

References

1. Abowd, G. D., Atkeson, C.G., Hong, J., Long S., Kooper, R., and Pinkerton, M. Cyberguide: A Mobile Context-Aware Tour Guide. *ACM Wireless Networks*. Volume 3, pp 421-433. 1997.
2. Active Bat. The BAT Ultrasonic Location System. <http://www.uk.research.att.com/bat/>. 2006.
3. Bahl, P. and Padmanabhan, V. RADAR: An In-Building RF-Based User Location and Tracking System. In the proceedings of *IEEE Infocom 2000*. Los Alamitos. pp. 775-784. 2000.

4. Cheverst, K., Davies, N., Mitchell, K., Friday, A., and Efstratiou, C. Developing a Context-aware Electronic Tourist Guide: Some Issues and Experiences. In the proceedings of *Conference on Human Factors in Computing Systems (CHI) 2000*. Netherlands. pp 17-24. April 2000.
5. Cooperstock, J. R., Fels, S.S., Buxton, W. and Smith, K.C. Reactive Environments: Throwing Away Your Keyboard and Mouse. *Communications of the ACM*. Volume 40, pp 65-73. 1997.
6. Dey, A.K. and Abowd, G.D. CybreMinder: A Context-Aware System for Supporting Reminders. In the proceedings of *The 2nd International Symposium on Handheld and Ubiquitous Computing (HUC2K)*. Bristol, UK. pp. 172-186. September 2000.
7. Dey, A.K., Hamid, R., Beckmann, C., Li, I., and Hsu, D. aCAPpella: Programming by Demonstration of Context-Aware Applications. In the proceedings of *Conference on Human Factors in Computing Systems (CHI) 2004*. pp. 33-40. April 2004
8. Dow, S., Lee, J., Oezbek, C., MacIntyre, B., Bolter, J.D., and Gandy, M. Exploring Spatial Narratives and Mixed Reality Experiences in Oakland Cemetery. In the proceedings of *ACM SIGCHI Conference on Advances in Computer Entertainment (ACE 2005)*. Valencia, Spain. June 2005.
9. Hightower, J. and Borriello, G. A Survey and Taxonomy of Location Systems for Ubiquitous Computing, University of Washington Tech Report CSC-01-08-03. 2001.
10. Kirstein, C. and Mueller, H. Interaction with a Projection Screen using a Camera-tracked Laser Pointer. In the proceedings of *The International Conference on Multimedia Modeling*. Lausanne, Switzerland. 1998
11. Lamming, M. and Flynn, M. Forget-me-not: Intimate Computing in Support of Human Memory. In the proceedings of *The Symposium on Next Generation Human Interfaces*. Tokyo, Japan. 1994.
12. LaMarca, A., Chawathe, Y., Consolvo, S., Hightower, J., Smith, I., Scott, I., Sohn, T., Howard, J., Hughes, J., Potter, F., Tabert, J., Powledge, R., Borriello, G., and Schilit, B. Place Lab: Device Positioning Using Radio Beacons in the Wild. In the proceedings of *Pervasive 2005*. Munich, Germany. pp. 116 – 133. 2005.
13. Li, Y., Hong, J. I., and Landay, J. A. 2004. Topiary: a tool for prototyping location-enhanced applications. In the proceedings of *The ACM Symposium on User interface Software and Technology (UIST 2004)*. Santa Fe, NM. pp 217-226. October 2004.
14. Madhavapeddy, A. and Tse, T. Study of Bluetooth Propagation Using Accurate Indoor Location Mapping. *The Seventh International Conference on Ubiquitous Computing (UbiComp 2005)*. Tokyo, Japan. pp 105-122. September 2005.
15. Myers, B. A., Bhatnagar, R., Nichols, J., Peck, C.H., Kong, D., Miller, R., and Long, A.C., Interacting at a Distance: Measuring the Performance of Laser Pointers and Other Devices. In the proceedings of *Conference on Human Factors in Computing Systems (CHI 2002)*. Minneapolis, Minnesota. 2002.
16. Mynatt, E.D., Back, M., Want, R., Baer, M., and Ellis, J.B. Designing Audio Aura. In the proceedings of *Conference on Human Factors in Computing Systems (CHI 1998)*. pp. 566-573. April 1998.
17. Olsen, D. R. and Nielsen, T. Laser Pointer Interaction. In the proceedings of *Conference on Human Factors in Computing Systems (CHI 2001)*. Seattle, Washington. 2001.
18. Otsason, V., Varshavsky, A., LaMarca A., and de Lara, E. Accurate GSM Indoor Localization. In the proceedings of *The Seventh International Conference on Ubiquitous Computing (UbiComp 2005)*. Tokyo, Japan. September 2005.
19. Patel, S. N. and Abowd, G.D. A 2-way Laser-assisted Selection Scheme for Handhelds in a Physical Environment. *The Fifth International Conference on Ubiquitous Computing (UbiComp 2003)*. Seattle, WA. pp 200 – 207. 2003.

20. Priyantha, N. B., Chakraborty, A., and Balakrishnan, H. The Cricket Location-Support System. In the proceedings of *The International Conference on Mobile Computing and Networking (Mobicom 2000)*. Boston, MA. August 2000.
21. Ringwald, M. Spontaneous Interaction with Everyday Devices Using a PDA. Workshop on Supporting Spontaneous Interaction in Ubiquitous Computing Settings. UbiComp 2002. Goeteborg, Sweden. 2002
22. Rekimoto J. and Ayatsuka Y. CyberCode: Designing Augmented Reality Environments with Visual Tags. In the proceedings of *Designing Augmented Reality Environments (DARE 2000)*. Elsinore, Denmark. pp 1 – 10. 2000.
23. Rekimoto, J and Katashi, N. The World through the Computer: Computer Augmented Interaction with Real World Environments. In the proceedings of the *ACM Symposium on User Interface Software and Technology (UIST 1995)*. Pittsburgh, PA. pp 29-36. 1995.
24. Ubisense. <http://www.ubisense.net>. 2005.
25. Want, R., Fishkin, K., Gujar, A., and Harrison, B. Bridging physical and virtual worlds with electronic tags. In the proceedings of *Conference on Human Factors in Computing Systems (CHI 2001)*, Pittsburgh, PA. 1999.
26. Want, R., Hopper, A., Falcao, V., and Gibbons, J. The active badge location system. *ACM Transactions on Information Systems*. Volume 10. pp. 91-102. January 1992.
27. Wilson, A. and Pham, H. Pointing in Intelligent Environments with the WorldCursor. In the proceedings of *Interact 2003*, Zurich, Switzerland. September 2003.
28. Wilson, A. and Shafer, S. XWand: UI for intelligent spaces. In the proceedings of *Conference on Human Factors in Computing Systems (CHI 2003)*. Ft. Lauderdale, Florida. pp 545 – 552. April 2003.