A Survey and Taxonomy of Location Systems for Ubiquitous Computing

Jeffrey Hightower and Gaetano Borriello
University of Washington, Computer Science and Engineering
Box 352350, Seattle, WA 98195

Technical Report UW-CSE 01-08-03

August 24, 2001

Abstract

Emerging mobile computing applications often need to know where things are physically located. To meet this need, many different location systems and technologies have been developed. In this paper we present a the basic techniques used for location-sensing, describe a taxonomy of location system properties, present a survey of research and commercial location systems that define the field, show how the taxonomy can be used to evaluate location-sensing systems, and offer suggestions for future research. It is our hope that this paper is a useful reference for researchers and location-aware application builders alike for understanding and evaluating the many options in this domain.

1 Introduction

To serve us well, emerging mobile computing applications will need to know the physical location of things so that they can record them and report them to us: Are we almost to the campsite? What lab bench was I standing by when I prepared these tissue samples? How should our search-and-rescue team move to quickly locate all the avalanche victims? Can I automatically display this stock devaluation chart on the large screen I am standing next to? Where is the nearest cardiac defibrillation unit?

Researchers are working to meet these and similar needs by developing systems and technologies that automatically locate people, equipment, and other

This technical report is an extended version of the article *Location Systems for Ubiquitous Computing* [22] tangibles. Indeed, many systems over the years have addressed the problem of automatic location-sensing. Because each approach solves a slightly different problem or supports different applications, they vary in many parameters, such as the physical phenomena used for location determination, the form factor of the sensing apparatus, power requirements, infrastructure versus portable elements, and resolution in time and space.

To make sense of this domain, we have developed a taxonomy to help developers of location-aware applications better evaluate their options when choosing a location-sensing system. The taxonomy may also aid researchers in identifying opportunities for new location-sensing techniques.

In Section 2, we present the basic techniques, such as triangulation, used for location-sensing. Section 3 defines a taxonomy by examining issues in location system implementations. Section 4 then surveys several important commercial and research location systems and places them in the taxonomy. In Section 5, we give an example of applying the taxonomy to choose a location-sensing system for an application. Finally, Section 6 describes future research directions and Section 7 concludes.

2 Location Sensing Techniques

Triangulation, scene analysis, and proximity are the three principal techniques for automatic location-sensing. Location systems may employ them individually or in combination. For each technique we describe its basic concepts, list some implementation technologies, and give examples of location systems which use the technique.

2.1 Triangulation

The triangulation location-sensing technique uses the geometric properties of triangles to compute object locations. Triangulation is divisible into the subcategories of *lateration*, using distance measurements, and *angulation*, using primarily angle or bearing measurements.

2.1.1 Lateration

We define the term lateration to mean for distance measurements what angulation means for angles. Lateration computes the position of an object by measuring its distance from multiple reference positions. Calculating an object's position in two dimensions requires distance measurements from 3 non-collinear points as shown in Figure 1. In 3 dimensions, distance measurements from 4 non-coplanar points are required. Domain-specific knowledge may reduce the number of required distance measurements. For example, the Active Bat Location System measures distance from indoor mobile tags, called Bats, to a grid of ceiling mounted ultrasound sensors [21]. A Bat's 3-dimensional position can be determined using only 3 distance measurements because the sensors in the

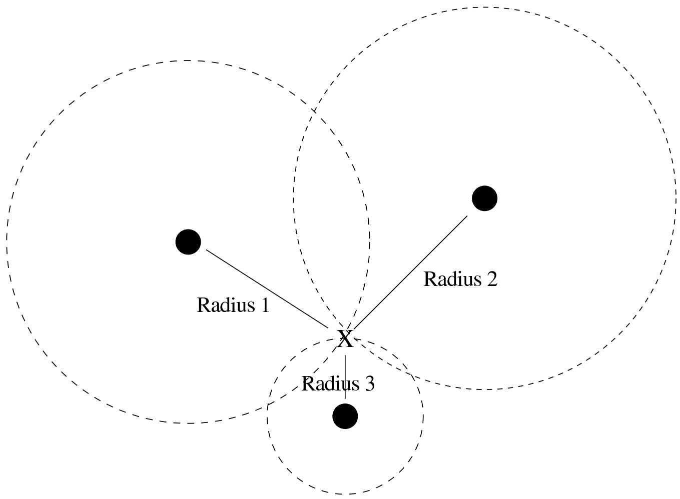

Figure 1: Determining 2D position using lateration requires distance measurements between the object 'X' and 3 non-collinear points.

ceiling are always above the receiver. The geometric ambiguity of only 3 distance measurements can be resolved because the Bat is known to be below the sensors and not in the alternate possible position on the next floor or roof above the sensor grid.

There are three general approaches to measuring the distances required by the lateration technique.

1. **Direct.** Direct measurement of distance uses a physical action or movement. For example, a robot can extend a probe until it touches something solid or take measurements with a tape measure. Direct distance measurements are simple to understand but difficult to obtain automatically due to the complexities involved in coordinating autonomous physical movement.

2. **Time-of-Flight.** Measuring distance from an object to some point P using time-of-flight means measuring the time it takes to travel between the object and point P at a known velocity. The object itself may be moving, such as an airplane traveling at a known velocity for a given time interval, or, as is far more typical, the object is approximately stationary and we are instead observing the difference in transmission and arrival time of an emitted signal. For example, sound waves have a velocity of approximately 344 meters per second in 21°C air. Therefore, an ultrasound pulse sent by an object and arriving at point P 14.5 milliseconds later allows us to conclude that the object is 5 meters away from point P. Measuring the time-of-flight of light or radio is also possible but requires clocks with much higher resolution (by six orders of magnitude) then those used for timing ultrasound since a light pulse emitted by the object has a velocity

of 299,792,458 meters per second and will travel the 5 meters to point P in 16.7 nanoseconds. Also, depending on the capabilities of the object and the receiver at point P, it may be necessary to measure a round-trip delay corresponding to twice the distance.

Ignoring pulses arriving at point P via an indirect (and hence longer) path caused by reflections in the environment is a challenge in measuring time-of-flight since direct and reflected pulses look identical. Active Bats and others statistically prune away reflected measurements by aggregating multiple receivers' measurements and observing the environment's reflective properties.

Another issue in taking time-of-flight measurements is agreement about the time. When only one measurement is needed, as with round-trip sound or radar reflections, "agreement" is simple because the transmitting object is also the receiver and must simply maintain its own time with sufficient precision to compute the distance. However, in a system like GPS, the receiver is not synchronized with the satellite transmitters and thus cannot precisely measure the time it took the signal to reach the ground from space. Therefore, GPS satellites are precisely synchronized with each other and transmit their local time in the signal allowing receivers to compute the *difference* in time-of-flight. GPS receivers can compute their 3-dimensional position (latitude, longitude, and elevation) using 4 satellites. The satellites are always above the receivers so only 3 satellites would normally be required to provide distance measurements in order to estimate a 3D position. However in GPS a fourth satellite measurement is required to allow us to solve for the forth unknown, the error between the receiver clock and the synchronized satellite clocks – a system of four equations (4 satellite signals) and four unknowns (X, Y, Z, and transmission time). Refer to [16] for an excellent summary of GPS theory. To maintain synchronization, each of the 27 GPS satellites contains four cesium/rubidium atomic clocks which are locally averaged to maintain a time accuracy of 1 part in 10^{13} seconds. Furthermore, each satellite gets synchronized daily to the more accurate atomic clocks at US Naval Observatory by US Air Force GPS ground control.

Time-of-flight location-sensing systems include GPS, the Active Bat Location System [21], the Cricket Location Support System [32], Bluesoft [27], and PulsON Time Modulated Ultra Wideband technology [36].

3. **Attenuation.** The intensity of an emitted signal decreases as the distance from the emission source increases. The decrease relative to the original intensity is the attenuation. Given a function correlating attenuation and distance for a type of emission and the original strength of the emission, it is possible to estimate the distance from an object to some point P by measuring the strength of the emission when it reaches P. For example, a free space radio signal emitted by an object will be attenuated by a factor proportional to $1/r^2$ when it reaches point P at distance r from the

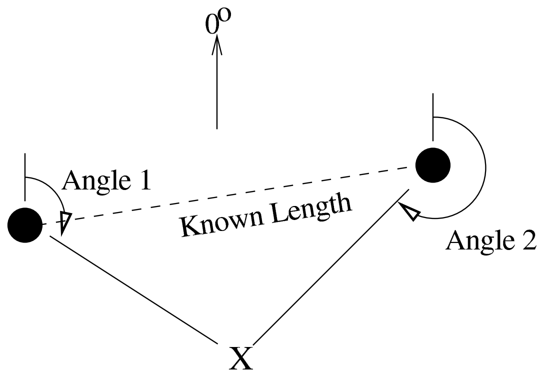

Figure 2: This example of 2D angulation illustrates locating object 'X' using angles relative to a 0° reference vector and the distance between two reference points. 2D angulation always requires at least two angle and one distance measurement to unambiguously locate an object.

object.

In environments with many obstructions such as an indoor office space, measuring distance using attenuation is usually less accurate than time-of-flight. Signal propagation issues such as reflection, refraction, and multipath cause the attenuation to correlate poorly with distance resulting in inaccurate and imprecise distance estimates.

The SpotON ad hoc location system implements attenuation measurement using low-cost tags. SpotON tags use radio signal attenuation to estimate inter-tag distance [23] and exploits the density of tag clusters and correlation of multiple measurements to mitigate some of the signal propagation difficulties.

2.1.2 Angulation

Angulation is similar to lateration except, instead of distances, angles are used for determining the position of an object. In general, two dimensional angulation requires two angle measurements and one length measurement such as the distance between the reference points as shown in Figure 2. In three dimensions, one length measurement, one azimuth measurement, and two angle measurements are needed to specify a precise position. Angulation implementations sometimes choose to designate a constant reference vector (e.g. magnetic north) as 0°.

Phased antenna arrays are an excellent enabling technology for the angulation technique. Multiple antennas with known separation measure the time of arrival of a signal. Given the differences in arrival times and the geometry of the receiving array, it is then possible to compute the angle from which the emission originated. If there are enough elements in the array and large enough separations, the angulation calculation can be performed.

The VHF Omnidirectional Ranging (VOR) aircraft navigation system is a different example of the angulation technique. As any pilot knows, VOR stations are ground-based transmitters in known locations which repeatedly broadcast 2 simultaneous signal pulses. The first signal is an omnidirectional reference containing the station's identity. The second signal is swept rapidly through 360° like the light from a lighthouse at a rate such that the signals are in phase at magnetic north and 180° out of phase to the south. By measuring the phase shift, aircraft listening to a VOR station can compute their "radial," the compass angle formed by the direct vector to the VOR station and magnetic north, to 1°. Aircraft location can be computed via angulation using 2 VOR stations. VHF radio signals are limited to line-of-sight reception and the range of the transmitted signals is 40-130 nautical miles.

2.2 Scene Analysis

The scene analysis location-sensing technique uses features of a scene observed from a particular vantage point to draw conclusions about the location of the observer or of objects in the scene. Usually the observed scenes are simplified to obtain features that are easy to represent and compare (e.g., the shape of horizon silhouettes such as Figure 3 as seen by a vehicle mounted camera [5]). In *static* scene analysis, observed features are looked up in a predefined dataset that maps them to object locations. In contrast, *differential* scene analysis tracks the difference between successive scenes to estimate location. Differences in the scenes will correspond to movements of the observer and if features in the scenes are known to be at specific positions, the observer can compute its own position relative to them.

The advantage of scene analysis is that the location of objects can be inferred using passive observation and features that do not correspond to geometric angles or distances As we have seen, measuring geometric quantities often requires motion or the emission of signals, both of which can compromise privacy and can require more power. The disadvantage of scene analysis is that the observer needs to have access to the features of the environment against which it will compare its observed scenes. Furthermore, changes to the environment in a way that alters the perceived features of the scenes may necessitate reconstruction of the predefined dataset or retrieval of an entirely new dataset.

The scene itself can consist of visual images, such as frames captured by a wearable camera [35], or any other measurable physical phenomena, such as the electromagnetic characteristics that occur when an object is at a particular position and orientation. The Microsoft Research RADAR location system is an example of the latter. RADAR uses a dataset of signal strength measurements created by observing the radio transmissions of an 802.11 wireless networking device at many positions and orientations throughout a building [3]. The location of other 802.11 network devices can then be computed by performing table lookup on the prebuilt dataset. The observed features, signal strength values in

Figure 3: Horizon shapes extracted from a visual scene can be used statically to look up the observer's location from a prebuilt dataset or dynamically to compute movement of the vehicle mounted camera.

this case, correlate with particular locations in the building but do not directly map to geometric lengths and angles describing those locations.

2.3 Proximity

A proximity location-sensing technique entails determining when an object is "near" a known location. The object's presence is sensed using a physical phenomenon with limited range. There are three general approaches to sensing proximity:

1. **Detecting physical contact.** Detecting physical contact with an object is the most basic sort of proximity sensing. Technologies for sensing physical contact include pressure sensors, touch sensors, and capacitive field detectors. Capacitive field detection has been used to implement a Touch Mouse [25] and Contact, a system for intra-body data communication among objects in direct contact with a person's skin [31].

2. **Monitoring wireless cellular access points.** Monitoring when a mobile device is in range of one or more access points in a wireless cellular network is another implementation of the proximity location technique and is illustrated by Figure 4. Examples of such systems include the Active Badge Location System [37] and the Xerox ParcTAB System [39], both using diffuse infrared cells in an office environment, and the Carnegie Mellon Wireless Andrew [24] using a campus-wide 802.11 wireless radio network.

3. **Observing automatic ID systems.** A third implementation of the proximity location-sensing technique uses automatic identification systems

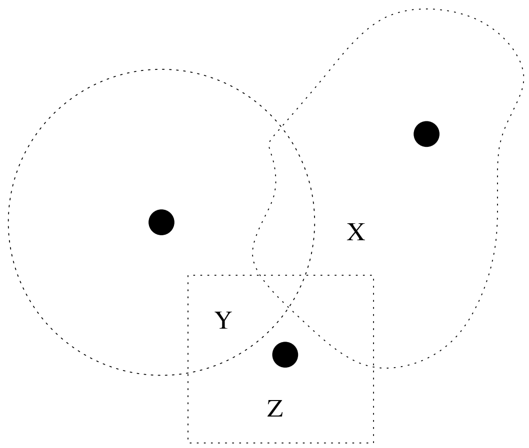

Figure 4: Objects 'X', 'Y', and 'Z' are located by monitoring their connectivity to one or more access point in a wireless cellular network. The cell geometry is an artifact of the wireless technology technology used in the implementation. For example, a radio cellular network cell may have the shape of the region containing object 'X' while diffuse infrared in a room is constrained by the walls resulting in a square shape.

such as credit card point-of-sale terminals, computer login histories, land-line telephone records, electronic card lock logs, and identification tags such as electronic highway E-Toll systems, UPC product codes, and injectable livestock identification capsules [38]. If the device scanning the label, interrogating the tag, or monitoring the transaction has a known location, the location of the mobile object can be inferred.

Proximity approaches may need to be combined with identification systems if they do not include a method for identification in the proximity detection. For example, the Contact system [31] enables communication between objects a user is touching and all these objects can exchange identification information over the same communication channel. Livestock tags have unique signatures identifying individual animals. Similarly for cell phones. In contrast, the Touch Mouse and pressure sensors, require an auxiliary identification system since the method used to detect proximity does not provide identification directly.

3 Location System Properties

A broad set of issues arises when we discuss and classify location system implementations. These issues are generally independent of the technologies or techniques a system uses. Although certainly not all orthogonal, nor equally applicable to every system, the classification axes we present do form a reasonable taxonomy for characterizing or evaluating location systems.

The Global Positioning System is perhaps the most widely publicized location-sensing system. GPS provides an excellent lateration framework for determining geographic positions. The worldwide satellite constellation has reliable and ubiquitous coverage and, assuming a differential reference or use of the Wide Area Augmentation System, allows receivers to compute their location to within 1 to 5 meters [12]. Aircraft, hikers, search-and-rescue teams, and rental cars all currently use GPS. Given its celebrity, we use GPS as a running example to introduce our classifiers.

3.1 Physical Position and Symbolic Location

A location system can provide two kinds of information: physical and symbolic. GPS provides physical *positions*. For example, our building is situated at at $47°39'17"N$ by $122°18'23"W$, at a 20.5-meter elevation. In contrast, *symbolic* location encompasses abstract ideas of where something is: in the kitchen, in Kalamazoo, next to a mailbox, on a train approaching Denver.

A system providing a physical position can usually be augmented to provide corresponding symbolic location information with additional information, infrastructure, or both. For example, a laptop equipped with a GPS receiver can access a separate database that contains the positions and geometric service regions of other objects to provide applications with symbolic information [7]. Linking real-time train positions to the reservation and ticketing database can help locate a passenger on a train. Applications can also use the physical position to determine a range of symbolic information. For example, one application can use a single GPS physical position to find the closest printer, while another may link it with calendar information to provide information about that person's current activity.

The distinction between physical position and symbolic location is more pronounced with some technologies than others. GPS is clearly a physical-positioning technology. Point-of-sale logs, bar code scanners, and systems that monitor computer login activity are symbolic location technologies mostly based on proximity to known objects. However, some systems such as Cricket can be used in either mode, depending on their specific configuration.

The resolution of physical-positioning systems can have implications for the definitiveness of the symbolic information they can be used to derive. For example, knowing where a person is inside a building, to within 10 meters, may be ineffective in placing that person in a specific room because of the position of walls within that 10-meter range. Purely symbolic location systems typically provide only very coarse-grained physical positions. Using them often requires multiple readings or sensors to increase accuracy – such as using multiple overlapping proximity sensors to detect someone's position within a room.

3.2 Absolute versus Relative

An *absolute* location system uses a shared reference grid for all located objects. For example, all GPS receivers use latitude, longitude, and altitude – or their equivalents, such as Universal Transverse Mercator (UTM) coordinates – for reporting location. Two GPS receivers placed at the same position will report equivalent position readings, and $47°39'17"N$ by $122°18'23"W$ refers to the same place regardless of GPS receiver.

In a *relative* system, each object can have its own frame of reference. For example, a mountain rescue team searching for avalanche victims can use hand-held computers to locate victims' avalanche transceivers. Each rescuer's device reports the victims' position relative to itself.

An absolute location can be transformed into a relative location – relative to a second reference point, that is. However, a second absolute location is not always available. In reverse, we can use triangulation to determine an absolute position from multiple relative readings if we know the absolute position of the reference points. But we often can't know these positions if the reference points are themselves mobile. Thus, the absolute versus relative distinction denotes primarily what information is available and how the system uses it rather than any innate capabilities.

3.3 Localized Location Computation

Some systems provide a location capability and insist that the object being located actually computes its own position. This model ensures privacy by mandating that no other entity may know where the located object is unless the object specifically takes action to publish that information. For example, orbiting GPS satellites have no knowledge about who uses the signals they transmit. Online map servers such as Expedia [28] and old-fashioned road atlases and print maps also fall into this category.

In contrast, some systems require the located object to periodically broadcast, respond with, or otherwise emit telemetry to allow the external infrastructure to locate it. The infrastructure can find objects in its purview without directly involving the objects in the computation. Personal-badge-location systems fit into this category, as do bar codes and the radio frequency identification tags that prevent merchandise theft, track shipments, and help identify livestock in the field [14] [26]. Placing the burden on the infrastructure decreases the computational and power demands on the objects being located, which makes many more applications possible due to lower costs and smaller form factors.

The policy for manipulating location data need not be dictated by where the computation is performed. For example, system-level access control can provide privacy for a movement history in a personal-location system while still allowing the infrastructure to perform the location computation. Doing so, however, imposes a requirement of trust in the access control.

3.4 Accuracy and Precision

A location system should report locations accurately and consistently from measurement to measurement. Some inexpensive GPS receivers can locate positions to within 10 meters for approximately 95 percent of measurements. More expensive differential units usually do much better, reaching 1- to 3- meter accuracies 99 percent of the time. These distances denote the *accuracy*, or grain size, of the position information GPS can provide. The percentages denote *precision*, or how often we can expect to get that accuracy.

Obviously, if we can live with less accuracy, we may be able to trade it for increased precision. Thus, we really must place the two attributes in a common framework for comparison. To arrive at a concise quantitative summary of accuracy and precision, we can assess the error distribution incurred when locating objects, along with any relevant dependencies such as the necessary density of infrastructural elements. For example, "Using five base stations per 300 square meters of indoor floor space, location-sensing system X can accurately locate objects within error margins defined by a Gaussian distribution centered at the objects' true locations and having a standard deviation of 2 meters."

Sensor fusion seeks to improve accuracy and precision by integrating many location or positioning systems to form hierarchical and overlapping levels of resolution. Statistically merging error distributions is an effective way to assess the combined effect of multiple sensors.

The ad hoc sensor networking and smart dust community [29] often addresses the related issue of *adaptive fidelity*. A location system with this ability can adjust its precision in response to dynamic situations such as partial failures or directives to conserve battery power.

Often, we evaluate a location-sensing system's accuracy to determine whether it is suitable for a particular application. Motion-capture installations that support computer animation [11] feature centimeter-level spatial positioning and precise temporal resolution, but most applications do not require this level of accuracy. GPS tags might suffice for species biologists concerned about the position of a migrating whale pod to a precision of 1 square kilometer. A personal location system for home or office applications might need enough accuracy to answer the query, "Which room was I in around noon?" but not "Where, to the nearest cubic centimeter, was my left thumb at 12:01:34 p.m.?"

3.5 Scale

A location-sensing system may be able to locate objects worldwide, within a metropolitan area, throughout a campus, in a particular building, or within a single room. Further, the number of objects the system can locate with a certain amount of infrastructure or over a given time may be limited. For example, GPS can serve an unlimited number of receivers worldwide using 24 satellites plus three redundant backups. On the other hand, some electronic tag readers cannot read *any* tag if more than one is within range.

To assess the scale of a location-sensing system, we consider its coverage area per unit of infrastructure and the number of objects the system can locate per unit of infrastructure per time interval. Time is an important consideration because of the limited bandwidth available in sensing objects. For example, a radio-frequency-based technology can only tolerate a maximum number of communications before the channel becomes congested. Beyond this threshold, either latency in determining the objects' positions will increase or a loss in accuracy will occur because the system calculates the objects' positions less frequently.

Systems can often expand to a larger scale by increasing the infrastructure. For example, a tag system that locates objects in a single building can operate on a campus by outfitting all campus buildings and outdoor areas with the necessary sensor infrastructure. Hindrances to scalability in a location system include not only the infrastructure cost but also middleware complexity – it may prove difficult to manage the larger and more distributed databases required for a campus-sized deployment.

3.6 Recognition

For applications that need to recognize or classify located objects to take a specific action based on their location, an automatic identification mechanism is needed. For example, a modern airport baggage handling system needs to automatically route outbound and inbound luggage to the correct flight or claim carousel. A proximity-location system consisting of tag scanners installed at key locations along the automatic baggage conveyers makes recognition a simple matter of printing the appropriate destination codes on the adhesive luggage check stickers. In contrast, GPS satellites have no inherent mechanism for recognizing individual receivers.

Systems with recognition capability may recognize only some feature types. For example, cameras and vision systems can easily distinguish the color or shape of an object but cannot automatically recognize individual people or a particular apple drawn from a bushel basket.

A general technique for providing recognition capability assigns names or globally unique IDs (GUID) to objects the system locates. Once a tag, badge, or label on the object reveals its GUID, the infrastructure can access an external database to look up the name, type, or other semantic information about the object. It can also combine the GUID with other contextual information so it can interpret the same object differently under varying circumstances. For example, a person can retrieve the descriptions of objects in a museum in a specified language. The infrastructure can also reverse the GUID model to emit IDs such as URLs that mobile objects can recognize and use [4].

3.7 Cost

We can assess the cost of a location-sensing system in several ways. Time costs include factors such as the installation process's length and the system's administration needs. Space costs involve the amount of installed infrastructure and the hardware's size and form factor.

Capital costs include factors such as the price per mobile unit or infrastructure element and the salaries of support personnel. For example, GPS receivers need an antenna of sufficient size for adequate satellite reception and may need a second antenna to receive the land-based differential signal. Support personnel at the US Air Force GPS command station must regularly monitor the status of the GPS satellites. Further, building and launching the satellites required a major capital investment by the US government.

A simple civilian GPS receiver costs around \$100 and represents the incremental cost of making a new object positionable independently of its global location. A system that uses infrared beacons for broad-casting room IDs requires a beacon for every room in which users want the system to find them. In this case, both the infrastructure and the object the system locates contribute to the incremental cost.

3.8 Limitations

Some systems will not function in certain environments. One difficulty with GPS is that receivers usually cannot detect the satellites' transmissions indoors. This limitation has implications for the kind of applications we can build using GPS. For example, because most wired phones are located indoors, even if its accuracy and precision were high enough to make it conceivable, GPS does not provide adequate support for an application that routes phone calls to the land-line phone nearest the intended recipient. A possible solution that maintains GPS interaction yet works indoors uses a system of GPS repeaters mounted at the edges of buildings to rebroadcast the signals inside.

Some tagging systems can read tags properly only when a single tag is present. In some cases, colocated systems that use the same operating frequency experience interference. In general, we assess functional limitations by considering the characteristics of the underlying technologies that implement the location system.

4 A Survey of Location Systems

We can use our taxonomy to survey some of the research and commercial location technologies that are representative of the location-sensing field. Tables 1 and 2 summarize the properties of these technologies. In Table 1, the open circles indicate that the systems can be classified as either absolute or relative, and the checkmarks indicate that localized location computation (LLC) or recognition applies to the system. Physical-symbolic and absolute-relative are paired alternatives, and a system is usually one or the other in each category.

Technology Name	**Properties**						
	Technique	**Phys**	**Symb**	**Abs**	**Rel**	**LLC**	**Recognition**
GPS	Radio time-of-flight lateration	•		•		√	
Active Badges	Diffuse infrared cellular proximity		•	•			√
Active Bats	Ultrasound time-of-flight lateration	•		•			√
MotionStar	Scene analysis, lateration	•		•			√
VHF Omnidirectional Ranging (VOR)	Angulation	•		•		√	
Cricket	Proximity, lateration		•	○	○	√	
MSR RADAR	802.11 RF scene analysis & triangulation	•		•			√
PinPoint 3D-iD	RF lateration	•		•			√
Avalanche Transceivers	Radio signal strength proximity	•			•		
Easy Living	Vision, triangulation		•	•			√
Smart Floor	Physical contact proximity	•		•			√
Automatic ID Systems	Proximity		•	○	○		√
Wireless Andrew	802.11 cellular proximity		•	•			√
E911	Triangulation	•		•			√
SpotON	Ad hoc lateration	•			•		√

Table 1: Location system properties.

Technology Name	Classification Criteria			
	Accu & Prec	**Scale**	**Cost**	**Limitations**
GPS	1-5 meters (95-99%)	24 satellites worldwide	Expensive infrastructure, $100 receivers	Not indoors
Active Badges	Room size	1 base per room, badge per base per 10 sec	Administration costs, cheap tags & bases	Sunlight & fluorescent interference with infrared
Active Bats	9cm (95%)	1 base per $10m^2$, 25 computations per room per sec	Administration costs, cheap tags & sensors	Required ceiling sensor grids
MotionStar	1mm, 1ms, 0.1° (nearly 100%)	Controller per scene, 108 sensors per scene	Controlled scenes, expensive hardware	Control unit tether, precise installation
VHF Omnidirectional Ranging (VOR)	1° radial (≈100%)	Several transmitters per metropolitan area	Expensive infrastructure, inexpensive aircraft receivers	30-140 nautical miles line of sight
Cricket	4x4 ft. regions (≈100%)	≈ 1 beacon per 16 sq. ft.	$10 beacons & receivers	No central management, receiver computation
MSR RADAR	3-4.3m (50%)	3 bases per floor	802.11 network installation, ≈$100 wireless NICs	Wireless NICs required
PinPoint 3D-iD	1-3m	Several bases per building	Infrastructure installation, expensive hardware	Proprietary, 802.11 interference
Avalanche Transceivers	Variable, 60-80m range	1 transceiver per person	≈$200 per transceiver	Short radio range, unwanted signal attenuation
Easy Living	Variable	3 cameras per small room	Processing power, installed cameras	Ubiquitous public cameras
Smart Floor	Spacing of pressure sensors (100%)	Complete sensor grid per floor	Installation of sensor grid, creation of footfall training dataset	Recognition may not scale to large populations
Automatic ID Systems	Range of sensing phenomenon (RFID typically ¡ 1m)	Sensor per location	Installation, variable hardware costs	Must known sensor locations
Wireless Andrew	802.11 cell size (≈100m indoor, 1km free space)	Many bases per campus	802.11 deployment, ≈$100 wireless NICs	Wireless NICs required, RF cell geometries
E911	150-300m (95%)	Density of cellular infrastructure	Upgrading phone hardware or cell infrastructure	Only where cell coverage exists
SpotON	Depends on cluster size	Cluster at least 2 tags	$30 per tag, no infrastructure	Attenuation less accurate than time-of-flight

Table 2: Location system classification criteria.

Figure 5: Olivetti Active Badge (right) and a base station (left) used in the system's infrastructure.

4.1 Active Badge

The first and arguably archetypal indoor badge sensing system, the Active Badge location system, which was developed at Olivetti Research Laboratory, now AT&T Cambridge [37], consists of a cellular proximity system that uses diffuse infrared technology. Each person the system can locate wears a small infrared badge like that shown in Figure 5. The badge emits a globally unique identifier every 10 seconds or on demand. A central server collects this data from fixed infrared sensors around the building, aggregates it, and provides an application programming interface for using the data.

The Active Badge system provides absolute location information. A badge's location is symbolic, representing, for example, the room – or other infrared constraining volume – in which the badge is located. The Cambridge group also designed one of the first large software architectures for handling this type of symbolic location data [20].

As with any diffuse infrared system, Active Badges have difficulty in locations with fluorescent lighting or direct sunlight because of the spurious infrared emissions these light sources generate. Diffuse infrared has an effective range of several meters, which limits cell sizes to small- or medium-sized rooms. In larger rooms, the system can use multiple infrared beacons.

4.2 Active Bat

In more recent work, AT&T researchers have developed the Active Bat location system, which uses an ultrasound time-of-flight lateration technique to provide more accurate physical positioning than Active Badges [21]. Users and objects carry Active Bat tags shown in Figure 6. In response to a request the controller sends via short-range radio, a Bat emits an ultrasonic pulse to a grid

Figure 6: Mobile “Bat” of the At&T Cambridge Active Bat location system. Image courtesy of AT&T Laboratories Cambridge.

of ceiling-mounted receivers. At the same time the controller sends the radio frequency request packet, it also sends a synchronized reset signal to the ceiling sensors using a wired serial network. Each ceiling sensor measures the time interval from reset to ultrasonic pulse arrival and computes its distance from the Bat. The local controller then forwards the distance measurements to a central controller, which performs the lateration computation. Statistical pruning eliminates erroneous sensor measurements caused by a ceiling sensor hearing a reflected ultrasound pulse instead of one that traveled along the direct path from the Bat to the sensor.

The system, as reported in 1999, can locate Bats to within 9cm of their true position for 95 percent of the measurements, and work to improve the accuracy even further is in progress. It can also compute orientation information given predefined knowledge about the placement of Bats on the rigid form of an object and allowing for the ease with which ultrasound is obstructed. Each Bat has a GUID for addressing and recognition.

Using ultrasound time of flight this way requires a large fixed-sensor infrastructure throughout the ceiling and is rather sensitive to the precise placement of these sensors. Thus, scalability, ease of deployment, and cost are disadvantages of this approach.

4.3 Cricket

Complementing the Active Bat system, the Cricket Location Support System uses ultrasound emitters to create the infrastructure and embeds receivers in the object being located [32]. This approach forces the mobile objects to perform all their own triangulation computations. Cricket uses the radio frequency signal not only for synchronization of the time measurement, but also to delineate the time region during which the receiver should consider the sounds it receives. The system can identify any ultrasound it hears after the end of the radio frequency packet as a reflection and ignore it. A randomized algorithm allows multiple uncoordinated beacons to coexist in the same space. Each beacon also transmits a string of data that describes the semantics of the areas it delineates using the short-range radio.

Like the Active Bat system, Cricket uses ultrasonic time-of-flight data and a radio frequency control signal, but this system does not require a grid of ceiling sensors with fixed locations because its mobile receivers perform the timing and computation functions. Cricket, in its currently implemented form, is much less precise than Active Bat in that it can accurately delineate 4x4 square-foot regions within a room, while Active Bat is accurate to 9cm. However, the fundamental limit of range-estimation accuracy used in Cricket should be no different than Active Bat, and future implementations may compete with each other on accuracy.

Cricket implements both the lateration and proximity techniques. Receiving multiple beacons lets receivers triangulate their position. Receiving only one beacon still provides useful proximity information when combined with the semantic string the beacon transmits on the radio.

Cricket's advantages include privacy and decentralized scalability, while its disadvantages include a lack of centralized management or monitoring and the computational burden – and consequently power burden – that timing and processing both the ultrasound pulses and RF data place on the mobile receivers.

4.4 RADAR

A Microsoft Research group has developed RADAR, a building-wide tracking system based on the IEEE 802.11 WaveLAN wireless networking technology [3]. RADAR measures, at the base station, the signal strength and signal-to-noise ratio of signals that wireless devices send, then it uses this data to compute the 2D position within a building. Microsoft has developed two RADAR implementations: one using scene analysis and the other using lateration.

The RADAR approach offers two advantages: It requires only a few base stations, and it uses the same infrastructure that provides the building's general-purpose wireless networking. Likewise, RADAR suffers two disadvantages. First, the object it is tracking must support a wireless LAN, which may be impractical on small or power-constrained devices. Second, generalizing RADAR to multifloored buildings or three dimensions presents a nontrivial problem.

RADAR's scene-analysis implementation can place objects to within about 3 meters of their actual position with 50 percent probability, while the signal-strength lateration implementation has 4.3-meter accuracy at the same probability level. Although the scene-analysis version provides greater accuracy, significant changes in the environment, such as moving metal file cabinets or large groups of people congregating in rooms or hallways, may necessitate reconstructing the predefined signal-strength database or creating an entirely new database.

Several commercial companies such as WhereNet [15] and Pinpoint [13] sell wireless asset-tracking packages, which are similar in form to RADAR. Pinpoint's 3D-iD performs indoor position tracking using proprietary base station and tag hardware to measure radio time of flight. Pinpoint's system achieves 1- to 3-meter accuracy and, by virtue of being a commercial product, offers easier deployment and administration than many research systems.

The 3D-iD system suffers the disadvantage that each antenna has a narrow cone of influence, which can make ubiquitous deployment prohibitively expensive. Thus, 3D-iD best suits large indoor space settings such as hospitals or warehouses. It has difficulty interoperating with the 802.11 wireless networking infrastructure because of radio spectrum collision in the unregulated Industrial, Scientific, and Medical (ISM) band.

4.5 MotionStar Magnetic Tracker

Electromagnetic sensing offers a classic position-tracking method [33]. The large body of research and products that support virtual reality and motion capture for computer animation often offer modern incarnations of this technology. For example, Ascension offers a variety of motion-capture solutions, including Flock of Birds and, shown in Figure 7, the MotionStar DC magnetic tracker [2]. These tracking systems generate axial DC magnetic-field pulses from a transmitting antenna in a fixed location. The system computes the position and orientation of the receiving antennas by measuring the response in three orthogonal axes to the transmitted field pulse, combined with the constant effect of the earth's magnetic field.

Tracking systems such as MotionStar sense precise physical positions relative to the magnetic transmitting antenna. These systems offer the advantage of very high precision and accuracy, on the order of less than 1mm spatial resolution, 1ms time resolution, and $0.1°$ orientation capability. Disadvantages include steep implementation costs and the need to tether the tracked object to a control unit. Further, the sensors must remain within 1 to 3 meters of the transmitter, and accuracy degrades with the presence of metallic objects in the environment.

Many other technologies have been used in virtual environments or in support of computer animation. A CDMA radio ranging approach has been suggested [6], and many companies sell optical, infrared, and mechanical motion-capture systems. Like MotionStar, these systems are not designed to be scalable for use in large, location-aware applications. Rather, they capture position in one precisely controlled environment.

4.6 Easy Living

Several groups have explored using computer vision technology to figure out where things are. Microsoft Research's Easy Living provides one example of this approach. Easy Living uses the Digiclops real-time 3D cameras shown in Figure 8 to provide stereo-vision positioning capability in a home environment

Figure 7: MotionStar DC magnetic tracker, a precision system used in motion capture for computer animation, tracks the position and orientation of up to 108 sensor points on an object or scene. Key components include (left and right) the magnetic pulse transmitting antennas and (center) the receiving antennas and controller. Image courtesy of Ascension Technology Corporation.

[34]. Although Easy Living uses high-performance cameras, vision systems typically use substantial amounts of processing power to analyze frames captured with comparatively low-complexity hardware.

State-of-the-art integrated systems [17] demonstrate that multimodal processing – silhouette, skin color, and face pattern – can significantly enhance accuracy. Vision location systems must, however, constantly struggle to maintain analysis accuracy as scene complexity increases and more occlusive motion occurs. The dependence on infrastructural processing power, along with public wariness of ubiquitous cameras, can limit the scalability or suitability of vision location systems in many applications.

4.7 Smart Floor

In Georgia Tech's Smart Floor proximity location system [30], embedded pressure sensors capture footfalls, and the system uses the data for position tracking and pedestrian recognition. This unobtrusive direct physical contact system does not require people to carry a device or wear a tag. However, the system has the disadvantages of poor scalability and high incremental cost because the floor of each building in which Smart Floor is deployed must be physically altered to install the pressure sensor grids.

Figure 8: Digiclops color 3D camera, made by Point Grey Research and used by the Microsoft Research Easy Living group to provide stereo-vision positioning in a home environment. Image courtesy of Point Grey Research Inc.

4.8 E911

The US Federal Communications Commission's E911 telecommunication initiatives require that wireless phone providers develop a way to locate any phone that makes a 911 emergency call [10]. E911 is not a specific location-sensing system, but we include it because the initiatives have spawned many companies that are developing a variety of location systems to determine a cellular phone's location.

Location systems developed to comply with the E911 initiatives will also support new consumer services. For example, a wireless telephone can use this technology to find the nearest gas station, post office, movie theater, bus, or automated teller machine. Data from many cellular users can be aggregated to identify areas of traffic congestion. Many business speculators tout this model of mobile consumerism, or *mCommerce*, as being the "next big thing."

To comply with E911, vendors are exploring several RF techniques, including antenna proximity, angulation using phased antenna arrays, lateration via signal attenuation and time of flight, as well as GPS-enabled handsets that transmit their computed location to the cellular system [1]. To meet the FCC requirement, positioning must be accurate to within 150 meters for 95 percent of calls with receiver-based handset solutions such as GPS, or to within 300 meters with network-transmitter-based approaches.

5 Applying the Taxonomy

In addition to simply reasoning about a location-sensing system, our taxonomy can be applied to evaluate the characteristics of a location system needed by a particular application or the suitability of an existing location system for the application. To illustrate, consider choosing a location-sensing system for a personal ubiquitous jukebox. The jukebox allows each user to customize their own audio stream to accompany them as they move throughout a home or office environment. The audio stream is generated by the infrastructure from both a fixed repository of the user's personal digital audio files and streaming content such as internet radio stations. Audio stream playback takes advantage of fixed speakers the user encounters in the environment. Stream content is mediated when multiple users are in physical proximity of each other and must share the same speakers.

Tables 1 and 2 summarize the "fingerprint" of how location-sensing systems fits into the taxonomy. Given this application specification, we can use the taxonomy to create the fingerprint of a location-sensing system that would meet the needs of this ubiquitous jukebox. For example:

1. **Physical versus Symbolic.** The jukebox requires symbolic locations. The user needs to able to be located in regions in the environment corresponding to areas served by audio speakers. Knowing the physical (x,y,z) position of the user is not directly useful.

2. **Absolute or Relative.** Because the application uses fixed speakers driven by the infrastructure, absolute locations are needed.

3. **Localized Local Computation.** The infrastructure is already managing and keeping private each user's repository of digital audio files so allowing it to compute the user's locations and then protecting this information with access-controls is reasonable.

4. **Recognition.** The jukebox requires the capability to recognize and distinguish individual people in order to pipe users' audio streams to the correct speakers.

5. **Accuracy and Precision.** Accuracy must be sufficient to distinguish the regions in which various speakers may be heard – probably on the order of $4-6m^2$ regions. Jukebox precision should be very high, on the order of 99%.

6. **Cost.** A low cost location-sensing system is always desirable, but a location-sensing system with incremental costs in the infrastructure may be acceptable in this case since expanding to serve new areas already requires installing additional speaker infrastructure.

7. **Limitations.** The location-sensing system must function in the indoor environment.

From this fingerprint, we can immediately rule out certain location systems and analyze the suitability of others. GPS will not work indoors. MotionStar provides excessive accuracy and lacks the necessary scalability. The Cricket system, while providing enough accuracy and being low cost, may not be the ideal choice in the case because the infrastructure is already managing the audio streams and file storage so forcing the mobile tags to compute which speakers they are near is an unnecessary divergence from an already infrastructure-centric application model.

The best candidate for this jukebox application drawn from the systems we have seen in this paper is probably the Active Badge system. Active Badges may require more installed infrastructure then is desirable but because diffuse infrared and audio are both generally constrained by the same physical barriers such as walls of a room, the coverage of an Active Badge bastation corresponds nicely with the region served by a set of speakers. Symbolic speaker location information is potentially easy to manage using the Active Badge system.

Much like a modern software development cycle, this entire evaluation process is circular – specify the application, construct the fingerprint of a location-sensing system meeting the needs of the application, determine the existing systems which comes closest to matching that fingerprint, evaluate how well those systems supports the application, respecify the application, and repeat until the application is fully specified and either a location system is chose or the decision is made to construct a new location-sensing system.

6 Research Directions

Location sensing is a mature enough field to define a space within a taxonomy that is generally populated by existing systems, as Tables 1 and 2 have shown. As such, future work should generally focus on lowering cost, reducing the amount of infrastructure, improving scalability, and creating systems that are more flexible within the taxonomy. This does not imply, however, that location-sensing is a solved problem or that further advancements are simply a matter of rote technology improvement. Rather, location-sensing is now entering an exciting phase in which cross-pollination with ideas from other computer science and engineering disciplines motivates future research.

6.1 Sensor fusion

Defined as the use of multiple technologies or location systems simultaneously to form hierarchical and overlapping levels of sensing, sensor fusion can provide aggregate properties unavailable when using location systems individually.

For example, integrating several systems with different error distributions may increase accuracy and precision beyond what is possible using an individual system. The more independent the techniques, the more effectively they can be combined.

Figure 9: Robots have many on-board sensors for use in localization, multi-robot collaboration, and zero-knowledge map building.

An example of current sensor fusion research, multisensor collaborative robot localization and map building presents a problem usually divided into two subproblems:

- tracking location as the environment changes or the robot moves, and
- determining robot location from a zero-knowledge start state.

Autonomous robots, such as those shown in Figure 9, employ a myriad of on-board sensors including ultrasound and laser range finders, inertial trackers, and cameras. The robots use Markov and Bayesian statistical techniques and multi-robot collaboration to accomplish sensor fusion [19]. These techniques provide important starting points for combining location systems for ubiquitous computing.

6.2 Ad Hoc Location Sensing

This approach to locating objects without drawing on the infrastructure or central control borrows ideas from the ad hoc networking research community. In a purely ad hoc location-sensing system, all of the entities become mobile objects with the same sensors and capabilities. To estimate their locations, objects cooperate with other nearby objects by sharing sensor data to factor out overall measurement error. In this way, a cluster of ad hoc objects converges to an accurate estimate of all nearby objects' positions. Objects in the cluster are located relative to one another or absolutely if some objects in the cluster occupy known locations.

Figure 10: Prototype SpotON radio tag. These tags use radio signal attenuation to perform ad hoc lateration. Ad hoc clusters of tags cooperate to factor out measurement errors for all tag positions.

The techniques for building ad hoc systems include triangulation, scene analysis, or proximity. The work of Doherty et al. [18] and Bulusu et al. [8] explores ad hoc proximity systems that consider variants of the following question: Given a set $\mathcal{S}$ of tiny sensor devices and a proximity model of radio connectivity, such as a sphere or circle with a fixed radius, if we know that $s_0 \ldots s_n$, $s_i \subset \mathcal{S}$ are subsets of sensors in proximity to one another, how accurately can we infer the relative location of all sensors in set $\mathcal{S}$?

Doherty et al. present an algorithmic approach to this problem as well as a framework for describing error bounds on the computed locations. Bulusu et al. extend this basic connectivity notion by adding an ideal theoretical model of outdoor radio behavior and a regular grid of reference nodes at known locations.

The SpotON system implements ad hoc lateration with low-cost tags. SpotON tags use radio signal attenuation to estimate intertag distance [23]. They exploit the density of tags and correlation of multiple measurements to improve both accuracy and precision. Figure 10 shows a prototype SpotON tag.

Sensing object locations with no fixed infrastructure represents a highly scalable and low-cost approach. In the future, infrastructural systems could incorporate ad hoc concepts to increase accuracy or reduce cost. For example, it might be possible for a system like Active Bat to use a sparser ceiling-mounted ultrasound receiver grid if Bats could also accurately measure their distance from other Bats and share this information with the infrastructure.

6.3 Location-Sensing-System Accuracy: A Challenge

Comparing the accuracy and precision of different location-sensing systems can be an arduous task because many system descriptions lack a concise summary of these parameters. We therefore suggest that future quantitative evaluations of location-sensing systems include the *error distribution*, summarizing the system's accuracy and precision and any relevant dependencies such as the density of infrastructural elements. For example, "Using five base stations per 300 square meters of indoor floor space, location-sensing system X can accurately locate objects within error margins defined by a Gaussian distribution centered at the objects' true location and a standard deviation of 2 meters." We strongly encourage the location-sensing research and development community to investigate how to best obtain and represent such error distributions.

In addition to its comparison value, researchers could use a location-sensing system's accurately described error distribution as partial input for simulating a system – even a hypothetical one. Prototyping an application with a simulator avoids the cost of purchasing, deploying, and configuring a hardware infrastructure when the goal is simply to evaluate the suitability of a certain location-sensing system. Preliminary work on this idea has begun. For example, Bylund and Espinoza have built a simulator for a campus-sized position-sensing system that uses a Quake III gaming arena [9].

7 Conclusion

In this paper, we have presented the basic techniques used for location sensing, taxonomized location system properties, and surveyed research and commercial location systems that define the field. We applied our taxonomy to a ubiquitous jukebox application to illustrate its value in evaluating the requirements of a location system needed by a particular application or the suitability of an existing location system for an application. Finally, we have observed that since the space defined by our taxonomy is generally populated, location-sensing field is entering an exciting time where the cross-pollination of ideas amongst existing systems and from other disciplines of computer science and engineering is motivating future research such as sensor fusion and ad-hoc location sensing.

With decreasing costs of silicon and wireless connectivity, location systems will become increasingly common. Increased attention and effort will foster improvements in various aspects of the design space. We offer our approach to comparing these systems to help researchers make better choices for the location systems they use in ubiquitous applications.

8 Acknowledgments

The authors thank Trevor Pering from Intel Research and Ken Yasuhara, Neil Spring, and Vibha Sazawal from the University of Washington for their editorial feedback on this paper. We also thank Dieter Fox and Larry Arnstein at UW for providing valuable insights that helped clarify our presentation.

References

[1] Airbiquity. Website, 2001. http://www.airbiquity.com/.

[2] Ascension Technology Corporation, PO Box 527, Burlington, VT 05402. *Technical Description of DC Magnetic Trackers*, 2001.

[3] Paramvir Bahl and Venkata Padmanabhan. RADAR: An in-building RF-based user location and tracking system. In *Proceedings of IEEE INFOCOM*, volume 2, pages 775–784, March 2000.

[4] John Barton and Tim Kindberg. The cooltown user experience. Technical Report 2001-22, HP Laboratories, Palo Alto, CA, February 2001.

[5] J. Ross Beveridge, Christopher R. Graves, and Christopher E. Lesher. Local search as a tool for horizon line matching. In *Image Understanding Workshop*, pages 683–686, Los Altos, CA, February 1996. ARPA, Morgan Kaugmann.

[6] Steven R. Bible, Michael Zyda, and Don Brutzman. Using spread-spectrum ranging techniques for position tracking in a virtual environment. In *Second IEEE Workshop on Networked Realities*, Boston, MA, October 1995.

[7] Barry Brumitt, John Krumm, B. Meyers, and S. Shafer. Ubiquitous computing and the role of geometry. In *Special Issue on Smart Spaces and Environments*, volume 7-5, pages 41–43. IEEE Personal Communications, October 2000.

[8] Nirupama Bulusu, John Heidemann, and Deborah Estrin. GPS-less low cost outdoor localization for very small devices. *IEEE Personal Communications*, 7(5):28–34, October 2000. Special Issue on Smart Spaces and Environments.

[9] Markus Bylund and Fredrik Espinoza. Using quake III arena to simulate sensors and actuators when evaluating and testing mobile services. In *CHI 2001 Extended Abstracts*, pages 241–242. ACM, March-April 2001. Short Talk.

[10] Federal Communications Commission. Fcc wireless 911 requirements fact sheet, January 2001. http://www.fcc.gov/e911/.

[11] Ascension Technology Corporation. Website, 2001. http://www.ascension-tech.com/.

[12] Garmin Corporation. About GPS. Website, 2001. http://www.garmin.com/aboutGPS/.

[13] PinPoint Corporation. Website, 2001. http://www.pinpointco.com/.

[14] Sensormatic Electronics Corporation. Website, 2001. http://www.sensormatic.com/.

[15] WhereNet Corporation. Website, 2001. http://www.widata.com/.

[16] Peter H. Dana. Global positioning system overview. Website, 2000. http://www.colorado.edu/geography/gcraft/notes/gps/gps.html.

[17] T. Darrell, G. Gordon, M. Harville, and J. Woodfill. Integrated person tracking using stereo, color, and pattern detection. In *Conference on Computer Vision and Pattern Recognition*, pages 601–608. IEEE Computer Society, June 1998.

[18] Lance Doherty, Kristofer S. J. Pister, and Laurent El Ghaoui. Convex position estimation in wireless sensor networks. In *Proceedings of IEEE Infocom 2001*, volume 3, pages 1655–1663. IEEE, IEEE Computer Society Press, April 2001.

[19] Dieter Fox, Wolfram Burgard, Hannes Kruppa, and Sebastian Thrun. A probabilistic approach to collaborative multi-robot localization. *Autonomous Robots*, 8(3):325–244, June 2000.

[20] Andy Harter and Andy Hopper. A distributed location system for the active office. In *IEEE Network*, pages 62–70. IEEE Computer Society Press, January/February 1994.

[21] Andy Harter, Andy Hopper, Pete Steggles, Any Ward, and Paul Webster. The anatomy of a context-aware application. In *Proceedings of the 5th Annual ACM/IEEE International Conference on Mobile Computing and Networking (Mobicom 1999)*, pages 59–68, Seattle, WA, August 1999. ACM Press.

[22] Jeffrey Hightower and Gaetano Borriello. Location systems for ubiquitous computing. *Computer*, 34(8):57–66, August 2001.

[23] Jeffrey Hightower, Roy Want, and Gaetano Borriello. SpotON: An indoor 3d location sensing technology based on RF signal strength. UW-CSE 00-02-02, University of Washington, Department of Computer Science and Engineering, Seattle, WA, February 2000.

[24] Alex Hills. Wireless andrew. *IEEE Spectrum*, 36(6):49–53, June 1999.

[25] Ken Hinckley and Mike Sinclair. Touch-sensing input devices. In *Proceedings of the 1999 Conference on Human Factors in Computing Systems (CHI 1999)*. ACM, 1999.

[26] Axcess Incorporated. Website, 2001. http://www.axsi.com/.

[27] Bluesoft Incorporated. Website, 2001. http://www.bluesoft-inc.com/.

[28] Expedia Maps. Website, 2001. http://maps.expedia.com/.

[29] DARPA Information Technology Office. Sensor information technology. Website, 2001. www.darpa.mil/ito/research/sensit/.

[30] Robert J. Orr and Gregory D. Abowd. The smart floor: A mechanism for natural user identification and tracking. In *Proceedings of the 2000 Conference on Human Factors in Computing Systems (CHI 2000)*, The Hague, Netherlands, April 2000. ACM.

[31] Kurt Partridge, Larry Arnstein, Gaetano Borriello, and Turner Whitted. Fast intrabody signaling. Demonstration at Wireless and Mobile Computer Systems and Applications (WMCSA), December 2000.

[32] Nissanka B. Priyantha, Anit Chakraborty, and Hari Balakrishnan. The cricket location-support system. In *Proceedings of MOBICOM 2000*, pages 32–43, Boston, MA, August 2000. ACM, ACM Press.

[33] F. Raab, E. Blood, T. Steiner, and H. Jones. Magnetic position and orientation tracking system. *IEEE Transactions on Aerospace and Electronic Systems*, 15(5):709–717, September 1979.

[34] Microsoft Research. Easy living. Website, 2001. http://www.research.microsoft.com/easyliving/.

[35] Thad Starner, Bernt Schiele, and Alex Pentland. Visual context awareness via wearable computing. In *International Symposium on Wearable Computers*, pages 50–57. IEEE Computer Society Press, Pittsburgh, PA, October 1998.

[36] Time Domain Corporation, 7057 Old Madison Pike, Huntsville, AL 35806. *PulsON Technology: Time Modulated Ultra Wideband Overview*, 2001.

[37] Roy Want, Andy Hopper, Veronica Falcao, and Jon Gibbons. The active badge location system. *ACM Transactions on Information Systems*, 10(1):91–102, January 1992.

[38] Roy Want and Dan M. Russell. Ubiquitous electronic tagging. *IEEE Distributed Systems Online*, 1(2), September 2000.

[39] Roy Want, Bill Schilit, Norman Adams, Rich Gold, Karin Petersen, David Goldberg, John Ellis, and Mark Weiser. The parctab ubiquitous computing experiment. In Tomasz Imielinski, editor, *Mobile Computing*, chapter 2, pages 45–101. Kluwer Publishing, February 1997. ISBN 0-7923-9697-9.