

Research Article

Accuracy of iPhone Locations: A Comparison of Assisted GPS, WiFi and Cellular Positioning

Paul A Zandbergen

*Department of Geography
University of New Mexico*

Abstract

The 3G iPhone was the first consumer device to provide a seamless integration of three positioning technologies: Assisted GPS (A-GPS), WiFi positioning and cellular network positioning. This study presents an evaluation of the accuracy of locations obtained using these three positioning modes on the 3G iPhone. A-GPS locations were validated using surveyed benchmarks and compared to a traditional low-cost GPS receiver running simultaneously. WiFi and cellular positions for indoor locations were validated using high resolution orthophotography. Results indicate that A-GPS locations obtained using the 3G iPhone are much less accurate than those from regular autonomous GPS units (average median error of 8 m for ten 20-minute field tests) but appear sufficient for most Location Based Services (LBS). WiFi locations using the 3G iPhone are much less accurate (median error of 74 m for 58 observations) and fail to meet the published accuracy specifications. Positional errors in WiFi also reveal erratic spatial patterns resulting from the design of the calibration effort underlying the WiFi positioning system. Cellular positioning using the 3G iPhone is the least accurate positioning method (median error of 600 m for 64 observations), consistent with previous studies. Pros and cons of the three positioning technologies are presented in terms of coverage, accuracy and reliability, followed by a discussion of the implications for LBS using the 3G iPhone and similar mobile devices.

1 Introduction

Reliable location information has become a cornerstone of many applications, including emergency services, navigation, commercial services, recreation, tracking and network-

Address for correspondence: Paul A. Zandbergen, Department of Geography, Bandelier West Room 111, MSC01 1110, 1 University of New Mexico, Albuquerque, NM 87131, USA. E-mail: zandberg@unm.edu

ing. Global Positioning Systems (GPS) have emerged as the leading technology to provide location information to these Location Based Services (LBS). A GPS receiver provides accurate location, speed and time to a user, anywhere in the world and under any weather condition. Improvements in GPS receiver technology have resulted in very reliable and affordable GPS receivers for a wide range of applications. Most newer model cell phones are GPS-enabled, resulting in the widespread adoption of consumer applications that rely on GPS.

The GPS technology adopted in most cell phones employs a server-side component for the processing of the GPS signal and is referred to as Assisted GPS (A-GPS). While high-sensitivity GPS chipsets have been adopted in recent years, A-GPS does not work well indoors and as a result complementary positioning systems are employed under these conditions. Metropolitan-scale WiFi positioning has become a reality in most urban areas in the U.S. for WiFi enabled devices for both indoor and outdoor locations. Cell phones can also fall back on the more established cellular network positioning techniques. These three different techniques will be reviewed, followed by a discussion of their implementation in Apple's popular 3G iPhone. The empirical part of this study consists of an evaluation of the accuracy of locations obtained using these three positioning modes of the 3G iPhone.

1.1 Assisted GPS

Most GPS-enabled cell phones, including the 3G iPhone, employ a technology known as Assisted GPS (A-GPS). With A-GPS many of the functions of a full GPS receiver are performed by a remote GPS location server. This remote server provides the A-GPS mobile device with satellite orbit and clock information; the initial position and time estimate; satellite selection, range and range rate; and position computation. The mobile device contains a very basic GPS receiver that needs to synchronize to given satellites that are visible and transfer pseudo range information to the location server over the cellular network. With A-GPS the mobile device does not need to decode the GPS messages for each satellite or perform an extensive search for visible satellites when the system is turned on. This results in reduced power consumption and rapid time-to-first-fix. Most cellular service providers have adopted A-GPS as the technology of choice to meet the U.S. Federal Communications Commission (FCC) requirements for location information to support E911 services. The FCC requires that handset-based systems locate the caller to within 50 m for 67% of calls and to 150 m for 95% of calls.

GPS and A-GPS, however, do not work very well in high density urban areas due to limited satellite visibility, and typically do not work at all indoors due to signal obstructions. As a result, the availability of reliable positioning is limited in areas where people spend most of their time. This is changing somewhat with the adoption of high-sensitivity GPS (HSGPS) chip sets. Using a large number of correlators these chips are able to obtain a position fix using very weak GPS signals. While HSGPS does work under challenging conditions (e.g. urban canyons, indoors), positional accuracy and time-to-first-fix are often much worse than under ideal conditions. Many newer A-GPS systems employ HSGPS chip sets but even these receivers are not always able to obtain a position fix indoors, particularly inside buildings with a lot of steel in their construction or through several layers of concrete or brick.

1.2 Alternative Positioning Systems

A number of indoor positioning systems have been developed to overcome the limitations of GPS. These are based on terrestrial beacons and use cellular network signals, WiFi signals, Bluetooth, infrared, ultrasound or other radio frequencies. Good recent overviews of these systems are provided by and Kolodziej and Hjelm (2006) and Bensky (2008). Examples of systems include Active Badges based on infrared (Want et al. 1992) and Active Bats based on ultrasound (Harter et al. 1999), both developed by AT&T, RADAR developed by Microsoft Research (Bahl and Padmanabhan 2000), Ekahau (Ekahau 2008) and AeroScout (AeroScout 2008) all based on WiFi signals, MIT Cricket (Priyantha et al. 2000) based on radio signals and ultrasound, and MoteTrack (Lorincz and Welsh 2007) based on radio signals. Several of these systems have shown very promising results in terms of achieving high positional accuracy in highly controlled indoor environments. However, their widespread adoption is limited by the fact that implementation is typically very resource intensive, including a high density of base stations and extensive calibration efforts. Several systems also require setting up specialized beacons and/or the use of special tags to track mobile devices. As a result these systems have been mostly targeted at relatively small indoor sites, such as a single building. Several systems allow for the continuous tracking of assets (people, devices, goods) within this controlled environment. While these positioning systems are obviously of great interest for certain applications, they do not lend themselves very well to complement GPS in order to achieve seamless indoor-outdoor positioning for large metropolitan-scale areas. For widespread implementation on commercial mobile devices, WiFi and cellular positioning have emerged as the most viable alternatives at the present time. The Rosum TV system could become a third alternative but implementation is still in a pilot stage (Rosum 2008).

1.3 WiFi Positioning

WiFi positioning uses terrestrial based WiFi access points (APs) to determine location. Over the past several years tens of millions of APs using the 802.11 standard have been deployed by individuals, homeowners, businesses, academic institutions, retail stores and public buildings. All of these APs repeatedly broadcast a signal announcing their existence to the surrounding area. These signals typically travel several hundred meters in all directions. The density of APs in urban areas is so high that the signals often overlap, creating a natural reference system for determining location. WiFi positioning software identifies the existing WiFi signals within range of a WiFi enabled mobile device and calculates the current location of the device.

Coverage of WiFi positioning is best in heavily populated areas. WiFi APs are deployed for private and public use to provide high speed wireless coverage inside buildings and for selected outdoor areas. As a result, WiFi positioning in theory has excellent coverage and performance indoors. These attributes distinguish it from GPS which struggles to deliver positioning information in indoor environments. WiFi positioning does not require that a connection be established to the WiFi network: the WiFi signals are only recorded in the form of their unique MAC address and signal strength at a particular location. This allows WiFi positioning to use potentially very weak signals, as well as encrypted signals, without having to establish a connection.

Several positioning algorithms have been developed for WiFi positioning. These fall into the broad categories of geometric techniques, statistical techniques, fingerprinting and particle filters. Fingerprinting is also referred to in the literature as radio mapping, database correlation or pattern recognition. For a review of these techniques, see Hightower and Borriello (2004) and Gezici (2008). While originally developed for indoor positioning, these have been extended for outdoor use with some modifications. Many of the geometric and statistical techniques rely on knowledge of the exact location of APs and/or the ability to model signal strength as a function of distance from the AP's location. This is not feasible for metropolitan-scale implementation. First, there are simply too many APs to consider, typically in the thousands for a single city. Second, modeling signal strength in a complex and highly variable environment (buildings, structures, vegetation, vehicles, etc.) is very challenging. As a result, fingerprinting techniques have emerged as the preferred method for metropolitan-scale WiFi positioning, since they do not require the exact location of APs and do not attempt to model signal strength. Instead, fingerprinting employs a calibration phase (also referred to as the training or offline phase) in which WiFi signals are observed at known locations. The set of APs and their respective signal strengths presents a "fingerprint" that is unique to that location. For large areas these data are collected using a technique known as "wardriving" – a mobile device with a WiFi receiver (typically a laptop) is hooked up to a GPS device, and the WiFi signals and GPS coordinates are recorded as the device moves through an area (typically in a vehicle). In the positioning phase (or online phase) the observed WiFi signals at an unknown location are compared to the database of previously recorded fingerprints to determine the closest match. Several matching techniques have been developed for this, including k-nearest neighbor estimation, support vector regression, smallest M-vertex polygon, Bayesian modeling, neural networks and kernelized distance estimation (see Roos et al. 2002, Youssef et al. 2003, Kaemarungsi and Krishnamurthy 2004, Kolodziej and Hjelm 2006, Yim 2008). K-nearest neighbor estimation has been most widely used, in part due to its computational simplicity and in part because it performs well relative to other techniques (Lin and Lin 2005)

Most of the knowledge on the performance of WiFi positioning has been gained from studies in well-controlled indoor environments with a very high AP density (e.g. Mok and Retscher 2007, Wallbaum 2007, Liao and Kao 2008, Manodham et al. 2008, Yin et al. 2008, Wayne et al. 2009). Performance varies with AP density and distribution, reliability of the positional reference database, and the positioning algorithm employed, among other factors. For a single building with a substantial number of APs, median horizontal accuracies of between 1 and 5 m have been achieved (Mok and Retscher 2007, Wallbaum 2007, Swanguang and Krishnamurthy 2008, Wayne et al. 2009).

The same approach used for indoor WiFi positioning can be employed for outdoor positioning. While the average AP density (in units per km²) for metropolitan-scale areas is much lower than a typical indoor environment, AP density in many urban areas is large enough so that signals from different APs overlap, creating the possibility of a seamless indoor-outdoor positioning system based on WiFi signals. A pioneering effort in this regard was made by the Place Lab project of the Intel Corporation. Place Lab developed a working prototype for Seattle, WA and published several studies on the performance of WiFi positioning. For a review of Place Lab, see LaMarca et al. (2005) and Hightower et al. (2006). For well-calibrated areas, Place Lab was able to achieve median positioning errors of between 15 to 40 m (Cheng et al. 2005). In a comparison of three different algorithms (centroid, fingerprinting and particle filter) the differences in accuracy

between algorithms was smaller than between different neighborhoods within the study area (Cheng et al. 2005). Positional accuracy was relatively robust to changes in the APs within the study area, to noise in the GPS data in the calibration stage and to a reduction in the density of the calibration data (Cheng et al. 2005). In short, Place Lab demonstrated the feasibility of metropolitan-scale WiFi positioning with moderate positional accuracy using the existing infrastructure of 802.11 APs. The Place Lab project was terminated in 2005, but the software and documentation remain available.

Several WiFi positioning systems are currently in operation. The most well established of these is developed by Skyhook Wireless, which is the system employed on the 3G iPhone. Alternatives include Navizon, WeFi and PlaceEngine. All these systems work in a similar manner. First, an application needs to be installed on a WiFi enabled device (or can be integrated into the device's firmware as is the case with the 3G iPhone). Upon activation the application records the available WiFi signals and sends this information to a remote location server. The location server compares the recorded signals to those in a database and the estimated location is then reported back to the mobile device.

Both Navizon and WeFi rely on the user community to populate the database of WiFi signals. Users are encouraged to input their (known) location when available, using a GPS signal or other means, and this is uploaded to a community database. Coverage is in theory global, but in reality very sporadic based on the contributions of users. PlaceEngine is a prototype with coverage limited to selected cities in Japan. Skyhook Wireless by contrast is the only fully commercialized system and employs its own fleet of data collectors. Coverage includes most urban areas in the U.S., Canada and Western Europe as well as selected cities in Asia. Skyhook was founded in 2003 and started creating a commercial WiFi positioning system, building on the efforts of Intel's Place Lab project. Skyhook refined the WiFi positioning technology (it holds several patents) and in 2008 released its hybrid positioning system (XPS) which combines GPS, WiFi and cellular positioning. Skyhook's documentation indicates that its WiFi positioning algorithms are based on fingerprinting (Figure 1), although the specific algorithms are proprietary. The specific algorithms behind XPS, however, are not described. For example, Skyhook claims that XPS improves the positional accuracy of WiFi positioning by 50% by leveraging signals from just two GPS satellites. These hybrid positioning algorithms, however, are not disclosed and Skyhook's patents also do not provide further insights into their workings. The technology developed by Skyhook has been endorsed by several equipment manufacturers, including Apple (which adopted Skyhook's system for the iPhone in 2008) and SiRF (a leading GPS chip manufacturer). Several online services (e.g. AOL, MapQuest) have also partnered with Skyhook to provide location-aware web services. Despite the widespread endorsement by industry partners, Skyhook's WiFi positioning systems has recently been criticized because it is vulnerable to location spoofing and location database manipulation attacks (Tippenhauer et al. 2009).

Limited information has been published on the performance of the existing metropolitan-scale WiFi positioning systems. In fact, since the Place Lab project was terminated in 2005, most peer reviewed publications on WiFi positioning have been limited to controlled indoor environments. This is in sharp contrast to the dramatic growth in the sale of WiFi devices and the rapidly expanding infrastructure to deliver wireless network service. Navizon (2007) states in its documentation that it achieves an accuracy of 20 to 40 m, but no actual test data have been published. PlaceEngine includes the following statement in its documentation: "it is difficult to determine the precise accuracy of the PlaceEngine service, but to give a rough estimation we believe it to be on

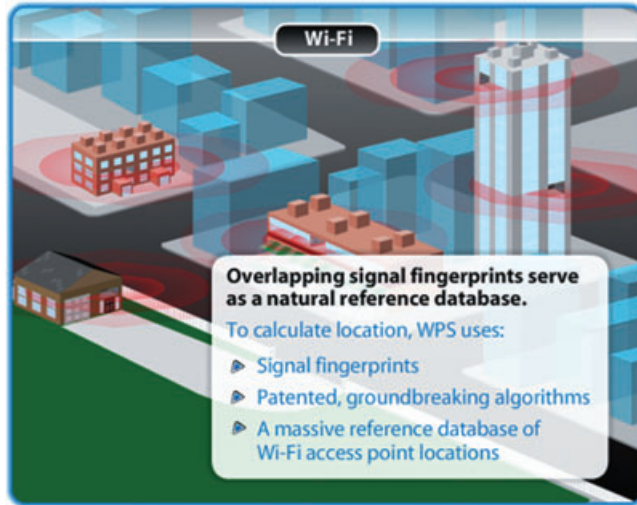


Figure 1 Description of the positioning method behind the WiFi positioning system (Source: Skyhook Wireless 2008)




Figure 2 Statements by Skyhook Wireless regarding the performance of its WiFi positioning system (Source: Skyhook Wireless 2008)

the order of 5 to 100 m”. WeFi does not make any performance information available. Skyhook states in its documentation that their WiFi positioning is accurate to within “20 m, indoors and outdoors” (Figure 2). The only published performance test results are provided in the form of a white paper (Skyhook Wireless 2008). Testing of positional accuracy consisted of static point tests (at undisclosed locations) and driving tests in several cities. Ground truth was determined “using digital maps and aerial photography” (p. 7). While no systematic accuracy metrics for different test conditions are reported, the

general conclusion was a median accuracy in urban canyons of 20 to 30 m with availability >97% and time-to-first-fix <1 second. In several tests the WiFi positioning system was found to be more accurate than a handheld GPS receiver (Garmin eTrex with SiRF III chip).

1.4 Cellular Positioning

Cellular networks have quickly developed into an extensive wireless communication infrastructure with almost worldwide coverage. Cellular service areas are divided into cells and each of these cells has a base station (cellular tower) associated with it. Techniques have been developed to track a mobile client when it is moving through the network. These location management techniques rely on the two-way communication between the mobile device and the network. When a user connects to the network, the mobile device is allocated to the base station transmitting with the strongest field strength. The most basic form of cellular positioning is to use the (known) location of this base station. This method is known as cell identification (cell ID). Location accuracy depends solely on the cell size, but this can be enhanced with support of other techniques. For example, some cells are divided into different sectors by directional base station antennas which can substantially reduce positional error. Further improvements can be achieved by using the received signal strength, although signal strength can vary considerably due to fading, topography, obstacles and other factors. Finally, cell ID techniques can be improved by using the timing advance that is calculated by the base station. These techniques are referred to as enhanced cell ID or E-CID.

When a mobile device is within range of multiple base stations, more complex positioning algorithms can be employed, mostly relying on Time Difference of Arrival (TDOA). For a review see Chapter 8 in Bensky (2008). Positional accuracy of these techniques depends on the density of base stations and the reliability of time of arrival measurements. The latter factor is in turn dependent on the bandwidth of the cellular signal, making the Global System for Mobile Communications (GSM) signal potentially more accurate than the Code Division Multiple Access (CDMA) signal. Angle of Arrival (AOA) is another algorithm adopted by cellular service providers, but is not as widely implemented. Fingerprinting techniques have also been developed for cellular positioning similar to those discussed for WiFi positioning (e.g. Juurakko and Backman 2004, Chen et al. 2006). Due to the effort involved in calibration, however, fingerprinting has not yet been adopted by any of the major cellular service providers.

Whatever specific positioning algorithm is used, the positional accuracy of cellular positioning depends greatly on the density of base stations. Horizontal error has therefore been found to vary greatly across urban-rural gradients, with a median error in the order of 50 to several hundred meters in urban areas and in the order of several hundred meters to several kilometers in rural areas (e.g. Weiss 2003, Lin and Juang 2005, Mohr et al. 2008). In one of the more comprehensive recent studies by Mohr et al. (2008) using three different cellular operators in the U.K., the median error was 246 m in a dense urban setting and 626 m in a rural setting.

Notwithstanding recent refinements in positioning algorithms, there are fundamental limits to the accuracy that can be achieved with typical densities of cell towers (Gustafsson and Gunnarsson 2005). This has been widely recognized by cellular providers and hence the adoption of A-GPS as the technology of choice to meet the FCC requirements for positioning. Nevertheless, cellular positioning remains of interest since

it may be available when other signals are not, as well as for older devices that are not A-GPS and/or WiFi enabled. As a result, the refinement of algorithms for cellular positioning remains an area of active research (e.g. Schwaighofer et al. 2003, Otsason et al. 2005, Chen et al. 2006).

The cellular positioning implemented on Apple's 3G iPhone is based on Google Mobile Maps. Limited information is published on the performance of this particular version of cellular positioning, but Google's description suggests it is based on cell ID.

1.5 iPhone Locations

When the iPhone was first released on 29 June 2007, it was met with rave reviews and quickly became a commercial success. While it was not the first smart phone on the market, its unique design, multi-touch screen, virtual keyboard and built-in functionality certainly raised the expectations as to what a smart phone should be capable of. Sales of the iPhone reached 17.4 million by December 2008, making Apple the third largest mobile manufacturer after Nokia and Samsung.

One of the substantive criticisms of the original iPhone was the fact it had no built-in A-GPS capability. This changed with the release of the 3G iPhone on 11 July 2008 which included a hybrid positioning system consisting of A-GPS, WiFi and cellular positioning. While each of these techniques was previously available separately for mobile devices, the 3G iPhone was the first consumer device to employ the hybrid positioning system. To implement the 3G iPhone's hybrid positioning system, Apple teamed up with Skyhook Wireless and Google. Previously Apple and Skyhook had teamed up to deliver WiFi positioning for the regular iPhone (without A-GPS).

From a user's perspective the positioning system of the 3G iPhone switches seamlessly between the three positioning modes. When a reliable A-GPS position fix is available, the location service publishes latitude, longitude, altitude and an estimate of positional error. On the default mapping application (Google Maps), this position is indicated by a bright blue dot with a pulsating blue location circle around it (Figure 3a). A transparent blue disc is shown to indicate the estimated positional error, although this is often invisible behind the bright blue dot depending on the display scale. When a reliable A-GPS position fix is not available the positioning mode is switched to WiFi, and when no reliable WiFi position fix can be obtained the positioning mode is switched to cellular. For both WiFi and cellular positioning the location service publishes latitude, longitude and an estimate of positional error – but no altitude. On the default mapping application the WiFi and cellular position fix is shown as a blue location circle with tick marks in the four cardinal directions. The circle is centered on the estimated position, but this location itself is not highlighted directly (Figures 3b and 3c). The diameter of the circle corresponds to the estimated positional error. Without additional knowledge a typical user cannot determine whether a WiFi or cellular position fix is used, although the size of the circle for cellular positioning is typically much larger (at the same scale).

While the iPhone was selected for this particular research effort, other smart phones have adopted a similar approach to using a hybrid positioning system, including the G1 phone released in October 2008 by T-Mobile and Google.

One of the strengths of the iPhone platform is that it is relatively easy to develop applications which can be distributed through the App Store in iTunes. Apple has released a Software Development Kit (SDK) for developers and a number of third-party manuals have appeared to support application development. Developing applications




Figure 3 Positioning modes of the 3G iPhone: (a) GPS position; (b) WiFi position; and (c) cellular position

that use the hybrid location service is relatively easy. Since the release of the 3G iPhone in July 2008 there has been a dramatic increase in the number of applications that use the location service. As of April 2009 there were more than 1,900 such applications in the App Store including local searches, navigation/routing, social networking and many others. The Android platform on the G1 phone provides similar possibilities. As of April 2009 nearly 300 applications had been developed using the location service of the G1 phone.

The hybrid positioning system of the 3G iPhone was featured prominently on the cover story of the February 2009 issue of *Wired* magazine. In a brief explanation of the three positioning modes the positional accuracy was characterized as 10 m for A-GPS, 30 m for WiFi positioning and 500 m for cellular positioning, although no sources or test results were reported. In fact, to date there have been no published studies on the performance of the hybrid positioning system of the iPhone (or the G1 phone). Apple, Skyhook and Google have not published any test results of the iPhone's performance. Several iPhone blogs have postings on user experiences, but do not include any specific accuracy evaluations. The blogs and other published material also reveal a substantial degree of confusion about how the hybrid positioning systems works. Many users for example did not realize that getting an A-GPS position fix indoors is often impossible. One published manual in fact explained (incorrectly) how WiFi positioning on the iPhone is based on the IP address and therefore inaccurate by several miles. Some users also expressed frustration with the accuracy of the locations. Specifically, when cellular positioning is employed the area covered by the blue location circle can be as large as several square kilometers, which may not be very informative based on a user's needs.

The current study is designed to determine the performance of the hybrid positioning system on the 3G iPhone. The emphasis is on the positional accuracy of each of the three positioning modes under static conditions. Other aspects of the iPhone positioning

system, including switching between positioning modes, time-to-first-fix, coverage, reliability under variable field conditions, battery usage, etc. are not addressed in the current study.

2 Data and Methods

A 3G iPhone was used to collect locations using three different modes: Assisted GPS, WiFi and cellular positioning. The 3G iPhone has a built-in location service which can be utilized by third-party applications. The location service provides coordinates in the format of latitude/longitude and also an altitude when the A-GPS mode is employed. Since only a single positioning mode is provided at any particular time, this presents some challenges for comparative analysis. Therefore, the following strategy was adopted. A-GPS locations were collected at outdoor sites under ideal conditions, i.e. excellent satellite visibility. WiFi and cellular positions were collected at indoor sites where A-GPS position fixes were not available. Switching between these two positioning modes was accomplished by turning the iPhone's WiFi receiver on and off – when no A-GPS or WiFi is available, the iPhone's location service defaults to cellular positioning. A third-party application was used to record the locations as waypoints and these were transferred to GIS software for processing. What follows are the specific details on the collection and processing of the data for each positioning mode.

2.1 A-GPS Positions

Locations in A-GPS mode were collected outdoors at 10 different surveyed benchmarks within the Albuquerque, NM metropolitan area. Appropriate benchmarks were selected from the Albuquerque Geodetic Reference System (AGRS). All benchmarks ($n = 853$) were plotted and overlaid on 6-inch color orthophotos from 2006. Datasheets for all benchmarks were also obtained and reviewed. Benchmarks with poor visibility as determined from the orthophotos and the obstruction diagrams in the datasheets were removed from the sample, as well as all those benchmarks in unsafe locations (e.g. located in the median of a road or very close to a road). From the remaining benchmarks a random sample of 10 was selected with the additional condition that no two benchmarks could be closer together than 1 km. For each benchmark the datasheet included the Northing and Easting in U.S. survey feet (in the State Plane coordinate system, New Mexico Central, NAD83) as well as ellipsoidal and orthometric heights.

At each of the 10 benchmarks, position fixes were recorded using both the 3G iPhone in A-GPS mode and a Garmin GPSMAP 60Cx unit (with a SiRF III chipset) in autonomous mode. Both units were placed vertically in a mount attached to the top of a survey tripod. The tripod was placed level directly over the benchmark and the height from the survey disk to the top of the antennae was measured. The units were placed at the same height, at opposite sides of the survey pole, approximately 5 inches apart horizontally. The horizontal displacement of the antennae relative to the center of the survey pole was not considered in the analysis since it contributed very little to the overall positional error, expected to be in the order of several meters.

Position fixes were logged with both units every 5 seconds for 20 minutes, resulting in 240 positions for each unit at each benchmark. The default WGS84 datum was used for both units. Locations were plotted in ArcGIS 9.3 and projected in the State Plane

coordinate system of the benchmark datasheets using the appropriate datum transformation. Northing, Easting and altitude for each position fix were compared to the values for these parameters reported in the datasheets. Corrections were made for the antenna height, as well as for the fact that the iPhone's A-GPS records altitude in ellipsoidal height while Garmin units records orthometric height. The distributions of horizontal and vertical error values were characterized using percentiles and Root Mean Square Error (RMSE) for both units at each of the 10 benchmarks.

2.2 *WiFi and Cellular Positions*

WiFi and cellular positions were collected at indoor sites where no A-GPS position fix could be adopted. To determine suitable locations for indoor sites, the following sampling strategy was obtained. First, a data file of address points was obtained from the City of Albuquerque – this includes the location of every occupied building within the city limits. Using the zoning information in the parcel data all commercial and institutional properties were selected since access to those locations would be easiest. A random sample of 65 properties was obtained, with the additional conditions that no two locations could be closer together than 300 m. These 65 locations were visited in the field. If access to the particular building was restricted or impractical, a new random location was selected. If an A-GPS position fix could be obtained inside the building, preventing the use of WiFi and cellular positioning on the iPhone, a new random location was chosen. In total 87 buildings were visited with 22 resulting in an A-GPS position fix. Only the results for WiFi and cellular positions at the remaining 65 locations were used in the analysis.

Within each building a location was selected that could easily be recognized on the 6-inch color orthophotos from 2006, for example, near an entrance, window, or corner. A laptop preloaded with the orthophotos was used in the field to digitize the estimated location of the indoor sites. At each location a single WiFi and cellular position fix was recorded. In the initial testing phase multiple position fixes were recorded at 5 second intervals, but these turned out to be identical (to 6-decimals in lat/long format) for each positioning mode. As a result, only a single position fix was recorded in the final field data collection effort. The field laptop was also used to confirm the availability of WiFi positioning at each indoor site. To access the WiFi positioning system without having to rely on connections to potentially weak or encrypted WiFi networks, the laptop was equipped with a cellular broadband Internet connection. The WiFi positions obtained using the laptop were not used in the analysis, but only served to confirm the availability of the WiFi positioning system. When A-GPS and WiFi positioning are not available, the iPhone defaults to cellular positioning and the logged positions of WiFi and cellular positioning on the iPhone are indistinguishable (in contrast to A-GPS positions). When no A-GPS fix can be obtained and the WiFi receiver is turned off and cellular positioning is not available, the iPhone's location service reports an error. So at each indoor site the availability of both the WiFi and cellular positioning was recorded as well as the position fix for each positioning mode when available.

Processing of the position fixes of WiFi and cellular positioning was similar to the procedure described for A-GPS position fixes, with the exception of the altitude information since this is not reported for these two modes. Since only a single position is recorded at each of the 65 indoor sites, accuracy statistics were determined for the combined set of positions.

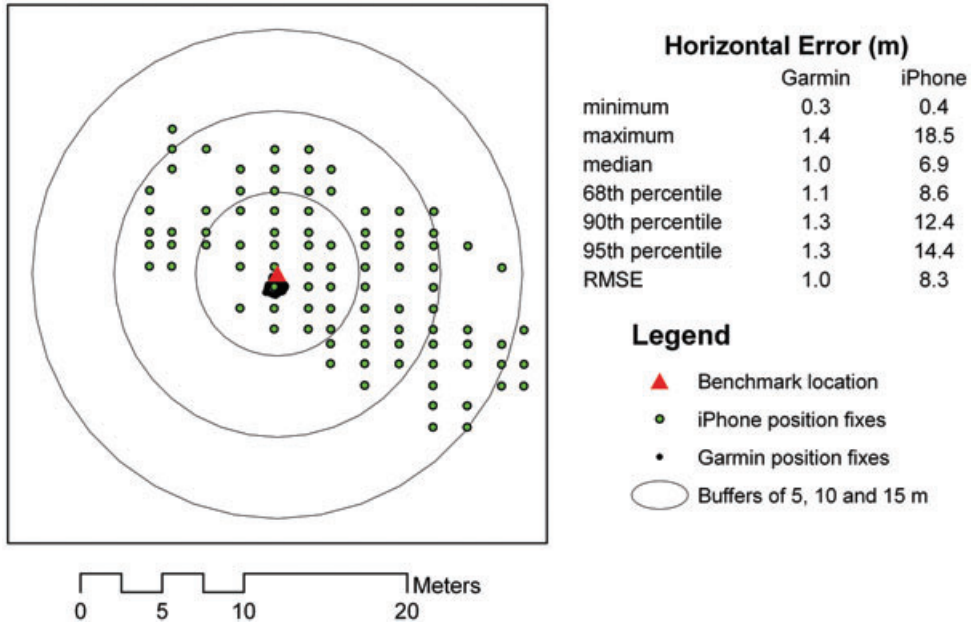


Figure 4 Horizontal accuracy of iPhone A-GPS and Garmin autonomous GPS locations

3 Results

3.1 A-GPS Positioning

For the ten 20-minute field tests under ideal conditions both the iPhone in A-GPS mode and the Garmin unit in autonomous mode were able to log valid position fixes 100% of the time, confirming excellent availability. Figure 4 shows an example scatterplot of a single 20-minute field test. The Garmin position fixes are strongly clustered close to the benchmark location with a maximum horizontal error of approximately 1.4 m and a RMSE of 1.0 m. The A-GPS position fixes reveal a much greater spread with a maximum error of 18.5 m and a RMSE of 8.3 m. The A-GPS positions also reveal a “gridded” pattern. While the original coordinates in decimal degrees are logged with six decimals, the *de facto* precision of the A-GPS coordinates is only 1 to 2 m which corresponds to five decimals. This is relatively common in older models of consumer grade GPS receivers, but newer models record with a precision of six or seven decimals. This gridding also explains that visually there appear to be much fewer than 240 positions since many have identical coordinates. While this results in a somewhat peculiar pattern, the typical error for the A-GPS positions is much greater than this precision and hence the truncation to five decimals does not appear to influence positional accuracy very much.

Results such as those reported for a single test in Figure 4 were obtained for every one of the 10 outdoor tests, in addition to the results for altitude. Values for the 68th percentile and RMSE were consistently similar across the tests for the two units, and therefore the error distributions were considered to be fairly close to normal (see Zandbergen 2008). Table 1 summarizes the horizontal and vertical errors reported as the median and RMSE values for the 10 tests. Results reveal that both the horizontal and vertical errors for the iPhone’s A-GPS positions are consistently much larger compared to

Table 1 Horizontal and vertical positional accuracy for iPhone locations in A-GPS mode and Garmin locations in autonomous mode under ideal outdoor conditions

Site ID	Horizontal Error (m)				Vertical Error (m)			
	Garmin		iPhone		Garmin		iPhone	
	Median	RMSE	Median	RMSE	Median	RMSE	Median	RMSE
#1	1.1	1.1	5.2	6.2	1.3	1.4	4.4	5.6
#2	0.8	1.1	10.1	12.4	3.3	3.1	6.4	9.6
#3	0.6	0.7	5.9	7.3	1.3	1.3	5.2	8.1
#4	2.5	2.6	8.1	9.0	8.8	9.0	9.7	11.7
#5	0.4	0.5	7.7	7.6	2.5	2.7	8.7	11.1
#6	1.0	1.6	12.6	15.5	4.5	4.3	10.1	17.3
#7	2.1	2.2	5.2	6.1	1.6	2.0	10.6	11.6
#8	3.4	3.4	11.2	11.4	4.5	4.4	7.4	10.0
#9	0.9	1.7	4.3	5.8	1.4	1.3	4.9	7.5
#10	1.0	1.0	6.9	8.3	0.9	1.4	12.1	13.6
Average	1.4	1.6	7.7	9.0	3.0	3.1	8.0	10.6

those for the Garmin's autonomous GPS positions. The average horizontal RMSE of 9.0 m for the iPhone is more than five times the value of 1.6 m for the Garmin unit and the average vertical RMSE of 10.6 m for the iPhone is more than three times the value of 3.1 m for the Garmin unit. Statistical testing of the difference between the RMSE values using a two-sample t-test showed a statistically significant difference for both horizontal ($t = 7.032$, $p < 0.001$) and vertical ($t = 5.964$, $p < 0.001$) accuracy.

Despite the relatively poor positional accuracy compared to the Garmin unit, the performance of the iPhone's A-GPS positioning was fairly consistent. For example, the largest error observed in all tests combined (2,400 positions total) was 27.7 m horizontal and 48.4 m vertical, which easily meets the FCC positioning requirements. For most LBS purposes a positional error of up to 20 m or so is also quite acceptable, but it clearly does not come close to the performance of a dedicated GPS receiver under ideal outdoor conditions.


3.2 WiFi and cellular positioning

The results for WiFi and cellular positioning are summarized in Table 2. WiFi positioning was not able to determine a position fix at eight locations, resulting in an availability of 87.7%. At five locations no WiFi signals were available as determined by using a laptop. At three locations WiFi signals were available, but no WiFi position fix could be achieved after repeated attempts. Cellular positioning failed at only one location, resulting in an availability of 98.5%.

In terms of positional accuracy only horizontal error was determined, since no elevation data is provided as part of WiFi or cellular positioning. Table 2 reports the

Table 2 Positional accuracy of iPhone locations by using WiFi and cellular positioning

	WiFi Positions	Cellular Positions
Number of observations	65	65
Number of valid position fixes	57	64
Percent valid fixes	87.7%	98.5%
Horizontal error (m)		
Minimum	16	30
Maximum	562	2,731
Median	74	599
68th percentile	88	827
RMSE	128	962
# Observations with error <20 m	3	0
# Observations with error <50 m	15	1
# Observations with error <100 m	41	3

**Figure 5** Cumulative distribution function of the positional error in WiFi and cellular positioning

error statistics and Figure 5 summarizes the results in a cumulative frequency distribution. WiFi positioning has a median error of 74 m and is clearly much more accurate than cellular positioning with a median error of 599 m. Figure 5 also shows that the error distribution is not normal as a result of a relatively small number of very larger error values. The RMSE value is very sensitive to these outliers as illustrated by the fact the

RMSE values are much larger than the 68th percentile. The RMSE value therefore becomes somewhat unreliable as a metric to summarize the error distribution (Zandbergen 2008) and the median error and 68th percentile should be used instead. Despite the fact that WiFi positioning is generally much more accurate than cellular positioning, at three locations cellular positioning was more accurate.

The results for the accuracy of the WiFi positioning system are in sharp contrast to those reported by Skyhook Wireless (2008). While Skyhook's WiFi positioning system is stated to provide 20 m accuracy, only 3 of 57 observations were found to be within 20 m of the reference location. The median error of 74 m determined in this study is several times larger than the 20 to 30 m reported by Skyhook.

To examine the spatial pattern in the performance of WiFi positioning, Figure 6 shows the field locations employed in the study. The eight locations where no WiFi position could be determined are located throughout the study area and do not appear to be concentrated in outlying lower density areas. Similarly, there is no clear spatial pattern in the positional error of WiFi positioning, with both high and low values spread throughout the study area. It should be noted that the 57 valid WiFi position fixes resulted in 57 unique coordinate pairs. This is in contrast to cellular positioning where 64 valid position fixes resulted in only 51 unique coordinate pairs: several position fixes at locations in relatively close proximity resulted in identical coordinates (to six decimals in decimal degrees). Only one cellular position corresponded exactly to the location of a known cellular tower, which suggests that the positioning algorithms may rely on cell tower locations of relatively poor positional accuracy or at least in part on enhanced cell ID techniques using cell tower sectors as opposed to strictly cell ID positioning.

Given the large observed positional errors of WiFi positioning relative to published accuracy metrics, the nature of this positional error was examined in more detail. Figure 7 shows four close-up examples of the results of WiFi positioning, including a range of different error values. In the case of Figures 7a and 7d, the estimated position lies exactly at the road median and this is the case for nearly half of all observed WiFi positions. It appears as if the location estimates are "snapped" to the road network. In the case of Figures 7b and 7c, the estimated location does not lie exactly on the road, but does appear to be influenced by a similar effect.

4 Discussion and Conclusions

The performance of A-GPS on the iPhone at outdoor locations was substantially less than that achieved using dedicated consumer-grade GPS receivers. While availability of A-GPS was 100%, the positional error for A-GPS was quite a bit larger compared to autonomous GPS. The average RMSE value for ten 20-minute tests was 9.0 m horizontal and 10.6 m vertical, several times larger than those for the consumer-grade GPS receiver. This can likely be attributed to the concessions that are made in the design of the A-GPS hardware on the iPhone, including antenna, power and other considerations. The observed gridding of the data is likely a result of a truncating or rounding step in the processing of the position fixes by the GPS location server. While this creates a somewhat odd-looking scatterplot, this lack of precision in the coordinates does not have a substantial impact on the positional accuracy and is similar to the pattern observed in older and low-cost GPS receivers.

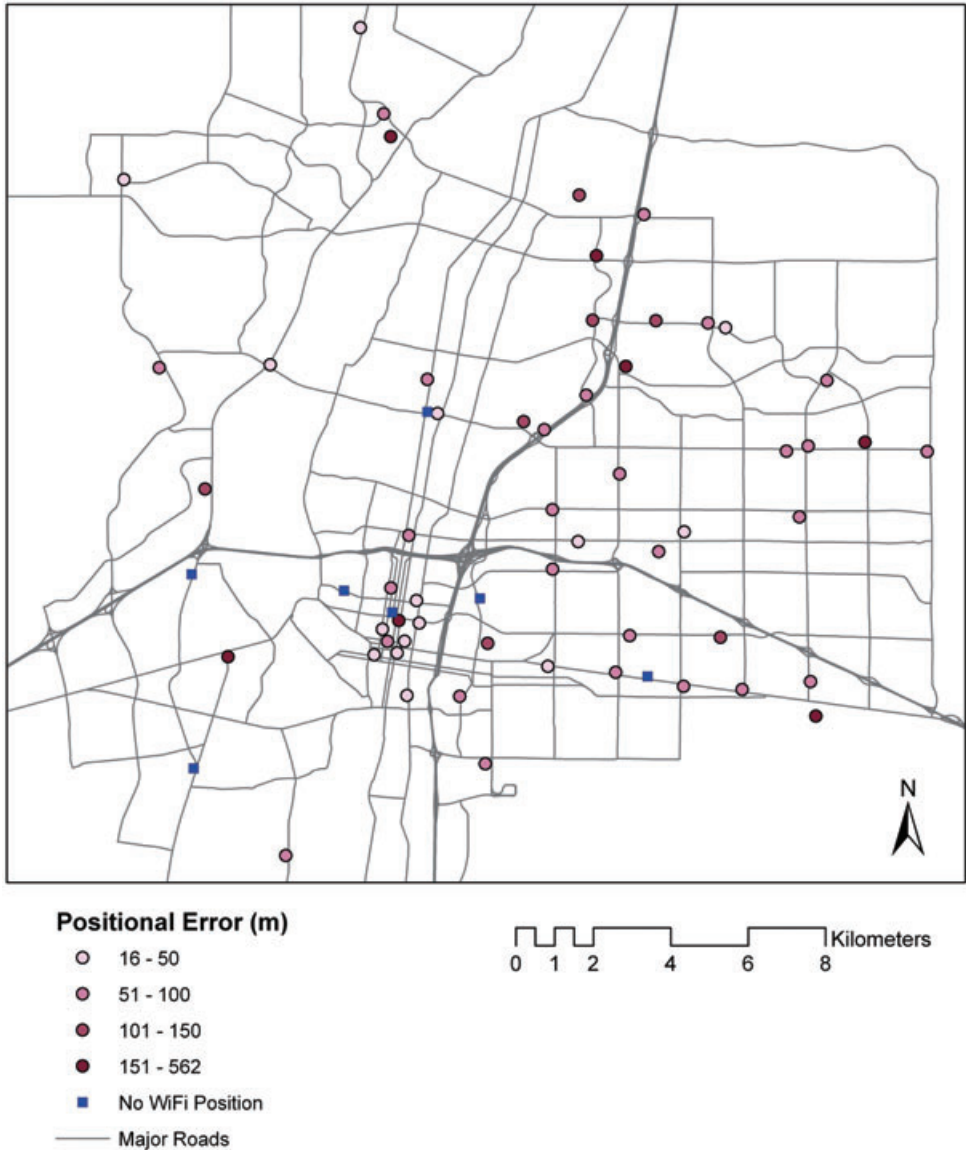


Figure 6 Spatial distribution of positional error of WiFi positioning

The performance of WiFi positioning on the iPhone at indoor locations was substantially less than the performance of A-GPS outdoors and in fact was far below expectations based on published performance measures by Skyhook Wireless (2008). Claims that WiFi positioning is able to acquire a location nearly 100% of the time in urban settings could not be confirmed and only 87.7% availability was achieved. Claims that WiFi positioning is able to achieve a median horizontal accuracy of 20 to 30 m could not be confirmed and a much larger median error of 74 m was found based on 57 observations. The 68th percentile was 88 m and the RMSE value was 128 m. No vertical




Figure 7 Examples of positional error in WiFi positioning

position can be estimated using WiFi positioning. The spatial pattern in the positional error suggests that the nature of the calibration effort greatly influences positional accuracy. Calibration data is collected by driving a vehicle on public roads. A substantial number of WiFi positions appear to be “snapped” to the road network, indicating that the positioning algorithms employed (a proprietary version of fingerprinting) are not able to reliably estimate locations at some distance from roads with the calibration data available.

Availability of cellular positioning on the iPhone at indoor locations was 98.5%, substantially higher compared to WiFi positioning at the same locations. Positional accuracy was much lower than for WiFi positioning, but similar to the results published by previous studies on cellular positioning (e.g. Mohr et al. 2008). Median horizontal error was 599 m, the 68th percentile was 827 m and the RMSE was 962 m based on 64 observations. The spatial pattern of cellular positions suggest that relatively simple positioning algorithms are employed (i.e. cell ID), which do not perform as well as other more complex algorithms.

There are several limitations to the current study. First, only a single metropolitan area was used in the field data collection. While this is not so much a concern for the evaluation of A-GPS, the accuracy of WiFi and cellular positioning is expected to be related to the density of APs and cell towers, respectively, as well as the rigor of the calibration effort and the time since the last calibration. However, there are no indications that the Albuquerque area has a low density of APs or cell towers, or suffers from poor calibration compared to other metropolitan areas. Based on ongoing research using WiFi positioning on a laptop throughout locations in Albuquerque there are very few locations reporting fewer than five APs (unpublished data). Based on records from the FCC, the City of Albuquerque (488 km²) contains a total of 82 antenna structures (mostly in the urban core) and an additional seven standalone cell towers (mostly in outlining areas). These densities are comparable to those in other urban areas. As far as the calibration effort goes, that information is not available. However, the coverage maps provided by Skyhook Wireless include the entire City of Albuquerque and surrounding urbanized areas, with the exception of a few large areas of open space where no indoor locations were selected as part of this study. The selection of locations where WiFi signals are not very likely was also prevented by using only commercial and institutional buildings in the sampling. A second limitation is the fact that only static conditions were evaluated. It would be meaningful to compare the performance of the positioning methods under dynamic conditions, such as driving through an urban canyon, or walking from outdoors to indoors and back. A third limitation is that A-GPS was only evaluated under ideal conditions. Since it is not possible to disable A-GPS positioning on the iPhone when an A-GPS position fix is available, a direct comparison of the three positioning systems under more challenging conditions (urban canyons, indoors) is not possible using only the 3G iPhone. Modifications to the iPhones's firmware or the use of multiple devices with different capabilities would be required for such a comparison. Finally, only a single type of device was employed and the results for other A-GPS enabled devices or other WiFi and cellular software platforms may be somewhat different.

WiFi and cellular positioning were able to obtain a position fix at most of the indoor sites where A-GPS failed. This result confirms the promise of the hybrid positioning approach. However, the positional error for both techniques is substantial and, for the WiFi positions in particular, much larger than published performance metrics indicate. Both developers and users of LBS on mobile services should be cautious when these services are based on non-GPS positions.

The somewhat disappointing results for WiFi positioning warrant further discussion. The first potential reason for the relatively poor accuracy is that calibration data is only collected along public roads, which prevents the fingerprinting algorithms to produce reliable position estimates at some distance from a road. When APs are located on either side of a road, war driving along the road does not allow for determining on what side of the road the AP is located on unless it is also picked up from another road at a different

angle. This ambiguity evidently introduces substantial error in the fingerprinting algorithm. This is consistent with the findings of Kim et al. (2006) who found a substantial improvement in the accuracy of outdoor WiFi positioning when regular wardriving was supplemented with data collection on foot around and between buildings, and a further improvement when the known locations of APs were used instead of the estimated locations based on the calibration data. The second potential reason is that the fingerprinting algorithm is based on the assumption that the observed WiFi signals correspond closely to those recorded in the calibration phase. APs change over time through additions, removals and changed locations. The reliability of fingerprinting therefore is expected to decrease as the calibration data ages. Skyhook does not reveal when a particular area was last covered in field calibration efforts, but this could definitely be a factor. A third potential reason is that radio signal propagation suffers from time-correlated fading effects, resulting from the interference from other devices, multi-path effects, changes to buildings and natural features, and the presence of moving objects. This means that observed signal strengths may deviate significantly from those recorded during the calibration phase (e.g. see Yin et al. 2008). This is an inherent limitation of static fingerprinting techniques and it is not well known how much of a factor this is for outdoor WiFi positioning. Finally, the only two published studies on the performance of metropolitan-scale positioning (Cheng et al. 2005, Skyhook Wireless 2008) were potentially biased in the sense that their field testing followed the exact same roads where calibration data had been collected. Results from these previous studies were therefore likely much too optimistic about the positional accuracy that can be achieved at locations at some distance from roads.

Determining the relative importance of the factors that contribute to the performance of WiFi positioning would be a worthwhile effort. For example, at present it is not well known which strategy would be most effective in improving the positional accuracy of WiFi positioning. More frequent wardriving? More intensive wardriving? Collecting data on foot in and around buildings? More refined positioning algorithms? While all of these efforts should contribute to improved positional accuracy, their relative importance is not well understood.

There have been substantial efforts in recent years to refine WiFi positioning for controlled indoor locations, including the development of novel approaches to fingerprinting (e.g. Kushki et al. 2007), development of time-of-arrival techniques (e.g. Golden and Bateman 2007), and modeling the performance of WiFi positioning systems (e.g. Swangmuang and Krishnamurthy 2008). Metropolitan-scale WiFi positioning, however, has received very limited attention from the research community. In fact, since the Place Lab project was terminated in 2005 there have been virtually no peer reviewed publications on the subject. Given the widespread adoption of WiFi positioning on a range of mobile devices, future research in this area should try to extend the recent progress made in controlled indoor environments to larger indoor/outdoor environments.

Acknowledgements

This research was funded by the Research Allocation Committee of the University of New Mexico and the Geography and Regional Science program of the National Science Foundation.

References

- AeroScout 2008 Aeroscout Technical Documentation. WWW document, <http://www.aeroscout.com/content/technology>
- Bahl P and V N Padmanabhan 2000 *Enhancements to the RADAR User Location and Tracking System*. Redmond, WA, Microsoft Research Technical Report No MSR-TR-2000-12
- Bensky A 2008 *Wireless Positioning: Technologies and Applications*. Boston, MA, Artech House
- Chen M Y, Sohn T, Chmelev D, Haehnel D, Hightower J, Hughes J, LaMarca A, Potter F, Smith I, and Varshavsky A 2006 Practical metropolitan-scale positioning for GSM Phones. In *Proceedings of Ubicomp 2006*, Newport Beach, California: 225–42
- Cheng Y, Chawathe Y, LaMarca A, and Krumm J 2005 *Accuracy Characterization for Metropolitan-scale Wi-Fi Localization*. Seattle, WA, Intel Research Report No IRS-TR-05-003
- Ekahau 2008 *Ekahau Real-Time Location System*. Helsinki, Ekahau
- Gezici S 2008 A survey on wireless positioning estimation. *Wireless Personal Communication* 44: 263–82
- Golden S A and Bateman S S 2007 Sensor measurements of Wi-Fi location with emphasis on time-of-arrival ranging. *IEEE Transactions on Mobile Computing* 6: 1185–98
- Gustafsson F and Gunnarsson F 2005 Mobile positioning using wireless networks: Possibilities and fundamental limitations based on available wireless network measurements. *IEEE Signal Processing Magazine* 22(4): 41–53
- Harter A, Hopper A, Steggle P, Ward A, and Webster P 1999 The anatomy of a context-aware application. In *Proceedings of Mobicom 1999*, Seattle, Washington: 59–68
- Hightower J and Borriello G 2004 Particle filters for location estimation in ubiquitous computing: A case study. In *Proceedings of Ubicomp 2004*, Tokyo, Japan: 88–106
- Hightower J, LaMarca A, and Smith I 2006 Practical lessons from Place Lab. *IEEE Pervasive Computing* 5(3): 32–9
- Juurakko S and Backman W 2004 Database correlation method with error corrections for emergency location. *Wireless Personal Communications* 30(2–4): 183–94
- Kaemarungsi K and Krishnamurthy P 2004 Modeling of indoor positioning systems based on location fingerprinting. In *Proceedings of IEEE Infocom 2004*, Hong Kong: 1012–22
- Kim M, Fielding J, and Kotz D 2006 Risks of using AP locations discovered through war driving. In *Proceedings of the Fourth International Conference on Pervasive Computing*, Dublin, Ireland: 67–82
- Kolodziej K W and Hjelm J 2006 *Local Positioning Systems: LBS Applications and Services*. Boca Raton, FL, Taylor and Francis
- Kushki A, Plataniotis K N, and Venetsanopoulos A N 2007 Kernel-based positioning in wireless local area networks. *IEEE Transaction on Mobile Computing* 6: 689–705
- LaMarca A, Chawathe Y, Consolvo S, Hightower J, Smith I E, Scott J, Sohn T, Howard J, Hughes J, Potter F, Tabert J, Powledge P, Borriello G, Schilit B N 2005 Place Lab: Device positioning using radio beacons in the wild. In *Proceedings of the Third International Conference on Pervasive Computing*, Munich, Germany: 116–33
- Liao I E and Kao K F 2008 Enhancing the accuracy of WLAN-based location determination systems using predicted orientation information. *Information Science* 178: 1049–68
- Lin D B and Juang R T 2005 Mobile location estimation based on differences of signal attenuations for GSM systems. *IEEE Transactions on Vehicular Technology* 54: 1447–54
- Lin T S and Lin P C 2005 Performance comparison of indoor positioning techniques based on fingerprinting. In *Proceedings of the Sixth IEEE International Symposium on a World of Wireless Mobile and Multimedia Networks*, Taormina, Italy: 382–7
- Lorincz K and Welsh M 2007 MoteTrack: A robust, decentralized approach to RF-based location tracking. *Personal Ubiquitous Computing* 11: 489–503
- Manodham T, Loyola L, Miki T 2008 A novel wireless positioning system for seamless Internet connectivity based on the WLAN infrastructure. *Wireless Personal Communication* 44: 295–309
- Mohr M, Edwards C, and McCarthy B 2008 A study of LBS accuracy in the UK and a novel approach to inferring the positioning technology employed. *Computer Communications* 31: 1148–59

- Mok E and Retscher G 2007 Location determination using WiFi fingerprinting versus WiFi trilateration. *Journal of Location Based Services* 1: 145–59
- Navizon 2007 *Bringing WiFi and Cellular Positioning to Mobile Devices and to the Web*. New York, Mexens Technology, Inc., Navizon Technical Paper
- Otsason V, Varshavsky A, LaMarca A, and de Lara E 2005 Accurate GSM indoor localization. In *Proceedings of Ubicomp 2005*, Tokyo, Japan: 141–58
- Priyantha N B, Chakraborty A, Hari Balakrishnan H 2000 The Cricket location-support system. In *Proceedings of the Sixth ACM MOBICOM*, Boston, Massachusetts
- Roos T, Myllymaki P, Tirri H, Misikangas P, and Sievanen J 2002 A probabilistic approach to WLAN use location estimation. *International Journal of Wireless Information Networks* 9(3): 155–64
- Rosum 2008 Rosum TV Technology. WWW document, http://www.rosun.com/rosun_technology.html
- Schwaighofer A, Grigoras M, Tresp V, and Hoffmann C 2003 GPPS: A Gaussian Process Positioning System for cellular networks. In *Proceedings of Neural Information Processing Systems 2003*, Vancouver and Whistler, British Columbia
- Skyhook Wireless 2008 *WiFi Positioning System: Accuracy, Availability and Time to Fix Performance*. Boston, MA, Skyhook Wireless Technical White Paper
- Swangmuang N and Krishnamurthy P 2008 An effective location fingerprint model for wireless indoor localization. *Pervasive and Mobile Computing* 4: 836–50
- Tippenhauer N O, Rasmussen K B, Pöpper C, and Capkun S 2009 Attacks on public WLAN-based positioning systems. In *Proceedings of the ACM/USENIX International Conference on Mobile Systems, Applications and Services (MobiSys '09)*, Kraków, Poland
- Wallbaum M 2007 A priori error estimates for wireless local area network positioning systems. *Pervasive and Mobile Computing* 3: 560–80
- Want R, Hopper A, Falcao V, and Gibbons J 1992 The active badge location system. *ACM Transaction in Information Systems* 40: 91–102
- Wayn C J, Li B, Dempster A, and Rizos C 2009 GPS/WiFi real-time positioning device: An initial outcome. In Gartner G, Cartwright W, and Peterson M P (eds) *Location Based Services and TeleCartography II*. Berlin, Springer Lecture Notes in Geoinformation and Cartography Vol. 22: 439–56
- Weiss A J 2003 On the accuracy of a cellular location system based on RSS measurements. *IEEE Transactions on Vehicular Technology* 52: 1508–18
- Yim J 2008 Introducing a decision tree-based indoor positioning technique. *Expert Systems with Applications* 34: 1296–302
- Yin J, Yang Q, and Ni L M 2008 Learning adaptive temporal radio maps for signal-strength based location estimation. *IEEE Transactions on Mobile Computing* 7: 869–83
- Youssef M, Agrawala A, and Shankar U 2003 WLAN location determination via clustering and probability distributions. In *Proceedings of PerCom 2003*, Dallas-Forth Worth, Texas: 143–50
- Zandbergen P A 2008 Positional accuracy of spatial data: Non-normal distributions and a critique of the National Standard for Spatial Data Accuracy. *Transactions in GIS* 12: 103–30