Journal of Location Based Services
Vol. 6, No. 1, March 2012, 35–50

Comparison of WiFi positioning on two mobile devices

Paul A. Zandbergen*

Department of Geography, University of New Mexico, Bandelier West Room 111, MSC01 1110, 1 University of New Mexico, Albuquerque, NM 87131, USA

(*Received 9 December 2010; final version received 30 September 2011; accepted 4 October 2011*)

Metropolitan WiFi positioning is widely used to complement GPS on mobile devices. WiFi positioning typically has very fast time-to-first-fix and can provide reliable location information when GPS signals are too weak for a position fix. Several commercial WiFi positioning systems have been developed in recent years and most newer model smart phones have the technology embedded. This study empirically determined the performance of WiFi positioning system on two different mobile devices. Skyhook's system, running on an iPhone and a laptop, was selected for this study. Field work was carried out in three cities at a total of 90 sites. The positional accuracy of WiFi positioning was found to be very similar on the two devices with no statistically significant difference between the two error distributions. This suggests that the replicability of WiFi positioning on different devices is high based on aggregate performance metrics. Median values for positional accuracy in the three study areas ranged from 43 to 92 m. These results are similar to earlier independent evaluations of Skyhook's system. The number of access points (APs) observed on the iPhone was consistently lower than that on the laptop. This lower number of APs, however, was not found to reduce positional accuracy. In general, no relationship was found between the number of APs and positional accuracy for either device, counter to earlier findings. The results indicate, however, that WiFi positioning can be achieved with a small number of APs (5–10), but that increased numbers of APs do not contribute to improved positional accuracy. Despite the agreement in aggregate performance metrics between the two devices, the replicability of WiFi positioning using Skyhook's system in terms of getting the same location by using two different devices at approximately the same place and time was relatively poor. Implications for location-based services on mobile devices are discussed.

Keywords: WiFi positioning; positional accuracy; mobile devices

1. Introduction

Location-based services (LBS) rely on the availability of robust location information on mobile devices. Global Positioning Systems (GPS) have emerged as the leading technology to provide this location information. 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

*Email: zandberg@unm.edu

ISSN 1748–9725 print/ISSN 1748–9733 online

http://dx.doi.org/10.1080/17489725.2011.630038
http://www.tandfonline.com

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.

GPS has a number of limitations to support location information on mobile devices. First, it relies on the line of sight to satellites and therefore does not work well indoors and in urban canyons due to limited satellite visibility. In these adverse conditions, positional accuracy is often limited and in many cases no position fix is obtained at all. As a result, the availability of reliable positioning is limited in areas where people spend a substantial portion of their time. Second, the time-to-first-fix can be very substantial, especially under weak signal conditions, from several seconds to several minutes. Third, power consumption can be quite significant and continuous positioning using GPS on a mobile device competes with other features for scarce battery resources. Finally, several types of mobile devices do not have GPS chipsets embedded, including many laptops, netbooks, tablet computers and older model cell phones.

To overcome these limitations, many mobile devices employ other positioning systems in addition or as an alternative to GPS. For example, cell phones can fall back on cellular network positioning techniques, but this tends to result in very large positional errors of several hundred metres. In recent years metropolitan-scale WiFi positioning has become available as a robust alternative in most urban areas in the United States and Europe for WiFi-enabled devices for both indoor and outdoor locations.

The purpose of this study is to determine whether the type of device has an influence on the performance of WiFi positioning on mobile devices. In other words, how replicable is the location information across different devices? The following section first provides a detailed review of WiFi positioning and its relevance to LBS.

2. Background

2.1. *Indoor 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 available (Kolodziej and Hjelm 2006, Bensky 2007). Many of these systems have been very promising 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 specialised 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 and 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.

Cellular positioning results in very limited positional accuracy and the following sections focus on WiFi positioning.

2.2. *WiFi positioning*

WiFi positioning uses terrestrial-based WiFi access points (APs) to determine location. Over the past several years, a very large number of APs using the 802.11 standard have been deployed by individuals, homeowners, businesses, academic institutions, retail stores and public buildings. All these APs repeatedly broadcast a signal announcing their existence to the surrounding area. These signals typically travel several hundred metres 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 the 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 uses 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 does not consistently 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 Media Access Control (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. Recent reviews of these techniques are available (Hightower and Borriello 2004, Gezici 2007). While originally developed for indoor positioning, these techniques 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 too many APs to consider, typically in the range of tens of thousands for a single city and obtaining their exact location would be very impractical. Second, modelling signal strength in a complex and highly variable environment (buildings, structure, 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, the calibration data are collected using a technique known as 'wardriving' – a mobile device with a WiFi receiver (typically a laptop or similar device) 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 neighbour estimation, support vector regression, smallest M-vertex polygon, Bayesian modelling, neural networks and kernelised distance estimation (Youssef *et al.* 2005, Kolodziej and Hjelm 2006, Kushki *et al.* 2007, Yim 2008). *k*-Nearest neighbour estimation has been most widely used, in part due to its computational simplicity and in part because it performs well relative to other techniques.

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 (Manodham *et al.* 2007, Mok and Retscher 2007, Wallbaum 2007, Liao and Kao 2008, Yin *et al.* 2008, Cheong *et al.* 2009, Chiou *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, Swangmuang and Krishnamurthy 2008).

The same approach used for indoor WiFi positioning can be employed for outdoor positioning. While the average AP density (in units per km^2) 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 (Cheng *et al.* 2005, LaMarca *et al.* 2005). For well-calibrated areas, Place Lab was able to achieve median positioning errors between 15 and 40 m (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. Since then, several commercial WiFi positioning systems have been developed.

2.3. *Commercial WiFi positioning systems*

Several commercial WiFi positioning systems are currently in operation. These include Skyhook Wireless, Google, Apple, Microsoft, 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 or operating system). 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. On laptops and similar mobile devices, WiFi positioning is often embedded in internet browsers using a plug-in, whereas on mobile phones WiFi positioning is typically part of the operating system or accessed using a third-party application.

Some WiFi positioning systems such as Navizon and WeFi rely on a user community to populate the database of WiFi signals. Users are encouraged to record their locations (using a GPS signal or other means) together with WiFi signal readings and this information is uploaded to a community database. Coverage is in theory global, but in reality somewhat sporadic, based on the contributions of users. Skyhook Wireless, Google and Microsoft by contrast are fully commercialised systems and employ their own fleet of data collectors, in addition to data collected by consumer devices. Apple, on the other hand, exclusively relies on iPhone devices being used by consumers to collect its database.

Of the various commercial systems currently in operation, Skyhook's system has been the longest in operation and has seen widespread adoption on many different types of devices. Skyhook Wireless was founded in 2003 and started creating a commercial WiFi positioning system, building on the successful efforts of Intel's Place Lab project. Skyhook refined the WiFi positioning technology (it holds 16 patents as of late 2010 and has several dozen pending applications) and in 2008 released its hybrid positioning system (XPS) which combines GPS, WiFi and cellular positioning. Skykook's documentation indicates that its WiFi positioning algorithms are based on fingerprinting, although the specific algorithms are proprietary. 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). Numerous online services have also partnered with Skyhook to provide location-aware web services. As of late 2010, Skyhook's database contains over 250 million APs and covers the most populated areas in North America, Europe, Australia and certain Asian countries.

Despite the widespread adoption by numerous hardware and software manufacturers, WiFi positioning systems have recently been criticised because they are vulnerable to location spoofing and location database manipulation attacks (Tippenhauer *et al.* 2009). The hacker community has also demonstrated that the location services which rely on WiFi positioning systems can also be reverse engineered to identify the approximate location of a particular MAC address (Kamker 2010). These findings clearly highlight potential security and privacy concerns over the collection of WiFi signal patterns and their utilisation in commercial WiFi positioning systems.

2.4. *Performance of WiFi positioning*

Limited information has been published on the performance of the existing metropolitan-scale commercial WiFi positioning systems. 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 utilisation of commercial metropolitan-scale WiFi positioning.

Skyhook states in its online documentation that their WiFi positioning is accurate to within '20 metres, indoors and outdoors'. Results of performance test by Skyhook are provided in the form of a whitepaper (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 20–30 m with availability $>97\%$ and time-to-first-fix <1 s. In several tests, the WiFi positioning system was found to be more accurate than a handheld GPS receiver (Garmin eTrex with SiRF III chip).

Skyhook's results are not confirmed by other published independent research studies. For example, Skyhook's system was evaluated for an urban area in Sydney, Australia (Gallagher *et al*. 2009). Using an indoor survey at 17 locations, the average positional error was 63 m, whereas for an outdoor survey at 10 locations the average positional error was also 63 m. The maximum observed positional error was approximately 400 m. Skyhook's system on Apple's 3G iPhone was also evaluated in Albuquerque, New Mexico (Zandbergen 2009). Using a survey of 65 indoor locations, the median positional accuracy was 74 m with a maximum observed error of almost 500 m. These errors are several times larger than those reported by Skyhook.

This study is designed to determine the variability in the performance of WiFi positioning system in metropolitan areas as a function of the type of device. Despite widespread adoption of WiFi positioning, there has been no published research on the possible influence of the device type. This study employs a laptop and a mobile phone using the same commercial WiFi positioning system. Skyhook's system was selected since it has been the most established system to date and can be accessed on many different types of devices.

3. Methods and data

Two WiFi-enabled devices were used: (1) a Dell Latitude D 630 laptop with an Intel 4965AGN WiFi card; and (2) an Apple 3GS iPhone. WiFi positioning on the laptop was accomplished using Skyhook's Loki service. Access to the WiFi positioning service was established using a cellular broadband access card, making it possible to use an internet connection without relying on a connection to local WiFi networks. WiFi positioning comes standard on the 3GS iPhone as part of its hybrid positioning system. Whenever a GPS position fix cannot be determined, the 3GS iPhone's positioning system automatically switches to WiFi positioning. Apple adopted Skyhook as its source of WiFi positioning on the iPhone in 2008. With the introduction of iOS 3.2 in early 2010, however, Apple started to use its own WiFi positioning system. Devices running older versions of the iPhone's operating system continue to rely on Skyhook. This study used a 3GS iPhone running iOS 3.1 to ensure that Skyhook's systems were employed, not Apple's.

Three study areas were used for the field research: Las Vegas, NV, Miami, FL and San Diego, CA. Field locations were selected from a database of all Starbucks locations in these cities obtained from www.starbucks.com. These locations were considered appropriate field sites because: (1) they are easily accessible; (2) they are spread out throughout urban areas in a density which approximates the density of urban activities in general; and (3) they are typical of locations where users of mobile devices would employ WiFi positioning. Airport locations were removed from the sample due to access limitations. From the remaining locations, a sample of 30 locations in each city was selected at random. During the field work, certain locations proved inaccessible due to closures; additional random locations were selected to achieve a final sample of 30 locations in each city for a total of 90.

The addresses of the field locations were geocoded against address point and parcel databases of business locations provided by local GIS departments (match rate 100%). These geocoded locations were loaded onto the laptop to provide a reference for navigation. Each location was visited and within each building a location was selected that could easily be recognised on high-resolution orthophotos, for example, near an entrance, window or corner. The laptop was preloaded with high-resolution orthophotos for the field sites and this information was used in the field to digitise the estimated location of the indoor sites. Coordinates were recorded in the local UTM coordinate system.

At each of the 90 locations, the same experimental setup was used. First, both the Dell laptop and the 3GS iPhone were turned on and WiFi was activated. The laptop was placed on a table whereas the iPhone was held upright approximately 30 cm above the table, directly above the screen of the laptop. The screens of the two devices were consistently facing in the same direction relative to each other, i.e. facing the observer, to limit the effect of any potential directional bias. After 2 min, the number of WiFi APs detected by each device was recorded. Immediately following, WiFi locations were recorded on each device within a period of 1 min. In the initial testing phase, multiple location estimates were recorded at 5 s intervals, but these turned out to be identical (to 6-decimals in lat/long format), most likely since the location server resends the same result instead of recalculating the position. As a result, only a single location estimate was recorded on each device in the final field data collection effort.

To confirm that the 3GS iPhone employed WiFi positioning (and not cellular positioning), the WiFi card on the iPhone was turned off and a new position fix using cellular position was obtained for comparison. In all cases this new cellular position was substantially different from the position fix with WiFi activated, confirming the availability of WiFi positioning on the 3GS iPhone. While assisted GPS is available on the 3G iPhone, no position fix using assisted GPS was obtained at any of the locations within the testing timeframe. It should be noted that assisted GPS on the 3GS iPhone can work indoors and was regularly observed in selected field locations. However, this occurred only after 5–10 min when WiFi positions had already been recorded.

The WiFi locations were recorded in decimal degrees with six decimals. Coordinates were plotted in ArcGIS and projected into the local UTM coordinate system. The positional error of each WiFi location was determined as the Euclidean distance between the WiFi position and the actual location of the field site. The distance between the WiFi locations obtained using the Dell laptop and the iPhone was used as a measure of replicability. Statistical analysis consisted of a characterisation of the positional error distribution of each device, as well as correlations between the number of APs and positional error. The distributions of positional error for the two devices were compared using a Kolmogorov–Smirnov (K–S) two-sample test. The K–S two-sample test is based on the maximum absolute difference (D) between the distributions for two continuous random variables. Unlike conventional statistical tests, this is a non-parametric test that does not require the variables to be normally distributed. The null hypothesis for the K–S test assumes that there is no difference in the distributions. In the test for this study, the largest observed difference between the two distributions was compared to the critical value of D at the 5% level of significance to determine if there is a statistically significant difference between the distributions.

4. Results

Field evaluation confirmed that WiFi positioning using Skyhook's system was available at all 90 sites using both devices. While previous work had identified gaps in coverage in urban areas (Gallagher *et al*. 2009, Zandbergen 2009), the locations employed in this study were all located in medium to high density urban areas with good WiFi coverage. Locations were also typically located in close proximity to roads, making it more likely for nearby APs to have been picked up by Skyhook's field data collection efforts. It should also be noted that nearly all the visited locations had their own WiFi APs, which likely contributed to the reliability of WiFi positioning in terms of availability.

The positional error of WiFi positioning on the two devices is summarised in Table 1. The differences between the various error metrics within a single city are very small. For example, in San Diego the median error on the laptop is 46.4 m whereas the median error on the iPhone is 42.6 m. Differences for other metrics are also quite small. Median errors in Miami are slightly lower (38.9 m on the laptop and 41.6 m on the iPhone) whereas in Las Vegas they were substantially higher (79.7 m on the laptop and 92.4 m on the iPhone). In general, however, the differences between the results for the two devices for each study area are quite small. Statistical comparison of the distributions using the K–S test revealed no statistically significant difference for any of the three study areas separately (three tests with $n = 30$) and after pooling the results into a single sample (one test with $n = 90$). On the basis of the distribution of the positional errors, therefore, WiFi positioning appears highly replicable across devices. It should also be noted that error distributions for both the iPhone and laptop revealed evidence of non-normal behaviour, specifically the occurrence of a few outliers with a positional error over 200 m. This makes the use of

Table 1. Positional error of WiFi positioning (m) on two different devices in three study areas.

	San Diego		Miami		Las Vegas	
	Laptop	iPhone	Laptop	iPhone	Laptop	iPhone
Sample size	30	30	30	30	30	30
Minimum	4.3	10.2	13.0	8.2	19.9	19.8
Maximum	462.9	344.7	237.5	232.4	443.9	1986.4
Median	46.4	42.6	38.9	41.6	79.7	92.4
68th Percentile	73.5	73.5	61.7	63.8	134.3	133.6
90th Percentile	167.0	153.8	150.1	147.8	287.0	394.6
95th Percentile	208.7	206.8	174.8	179.2	361.6	424.1
RMSE	121.4	102.9	80.2	80.9	167.3	404.6

metrics such as the maximum and root mean squared error (RMSE) somewhat unreliable (Zandbergen 2008).

The correlation between the positional errors on the two devices is examined in Figure 1. With a few notable exceptions, the general agreement between the two sets of observations is strong, as indicated by the R-squared value of 0.6983. While the slope of the regression is not a perfect 1, the very small intercept value and the slope value of 0.94 indicate that there is very little bias, if any, introduced by the use of a different device. In other words, there is no strong evidence to suggest that one device consistently estimates locations that are more accurate than those from another device. Despite the relatively strong correlation in Figure 1, a number of outliers are present. WiFi positioning typically results in a small number of large errors of several hundred metres or more (e.g. Zandbergen 2009) and Figure 1 confirms that these results are not always consistent across different devices. The occurrence of outliers in Figure 1 can most likely be attributed to the different sets of APs observed by each device, as discussed in the next section.

While several studies have found that more APs result in lower positional errors of WiFi positioning (e.g. Cheng *et al*. 2005, King *et al*. 2007, Wallbaum 2007, Chang *et al*. 2010, Feng *et al*. 2010), these have been limited to relatively small numbers of APs. Others have found that for larger number of APs, there is no strong relationship between the number of APs and positional error, and that APs with weak signals may in fact reduce performance (Lin *et al*. 2009, Chan *et al*. 2010,

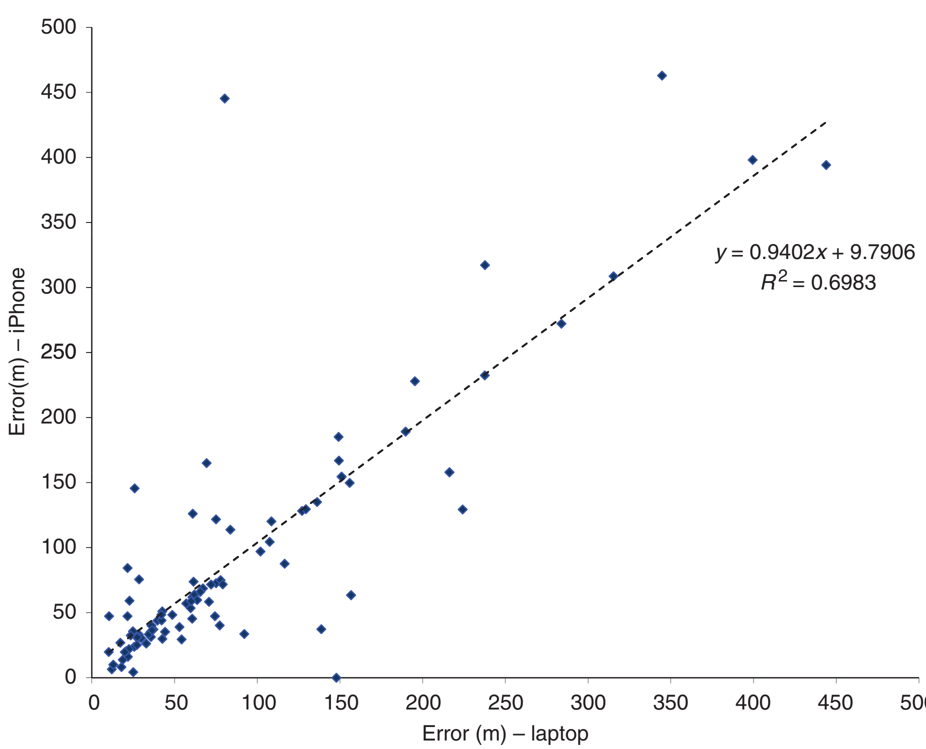

Figure 1. Correlation between the positional error of WiFi positioning on two mobile devices.

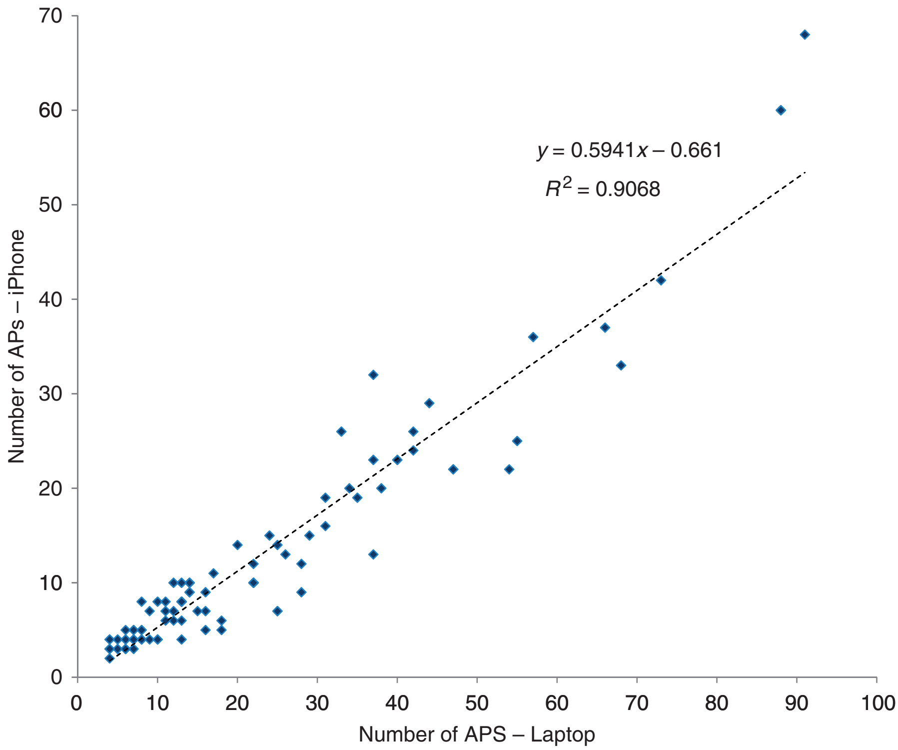

Figure 2. Correlation between the number of WiFi APs on two mobile devices.

Tsui *et al.* 2010). Figure 2 shows the correlation between the number of APs observed on the two devices. Somewhat surprisingly, the number of APs observed on the iPhone is consistently only about half of those observed on the laptop. The R-squared value of the regressions is 0.9068 which indicates very strong agreement, and the slope value of 0.5941 indicates that, on average, only approximately 6 APs are observed on the iPhone for every 10 APs observed on the laptop. Arguably, many of these WiFi signals are far too weak to provide for a reliable connection, but one of the strengths of the WiFi positioning system is that it can utilise even very weak signals. The larger number of APs observed by the laptop can be attributed to its higher quality WiFi antenna. The lower number of APs observed on the iPhone, however, does not appear to have a noticeable effect on positional error. The explanation can be found in Figure 3 that compares the number of APs to the positional accuracies of both devices. Figure 3 shows no relationship for either device. It confirms that WiFi positioning works quite well with only a very limited number of APs (5–10) and that additional APs do not necessarily improve accuracy. This is counter to Wallbaum (2007) which predicted steady improvements in accuracy with an increasing number of APs. The original PlaceLab project (Cheng *et al.* 2005) and other recent studies (King *et al.* 2007, Chang *et al.* 2010, Feng *et al.* 2010) have also found an empirical relationship between an increasing number of APs and decreasing positional error. It should be noted that these earlier studies considered up to 15 unique APs, while in this study more than 15 APs are

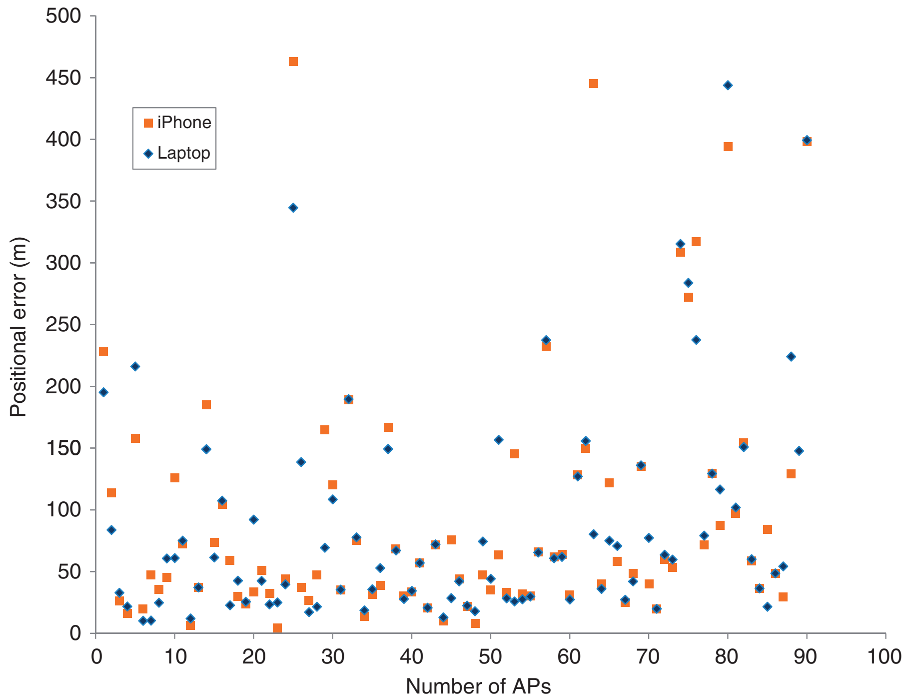

Figure 3. Relationship between the number of WiFi APs and positional error (m) on two mobile devices.

encountered in the majority of locations, reflecting the growth in WiFi infrastructure in urban areas. The results in Figure 3 confirm the findings from other research studies (Chan *et al*. 2010, Tsui *et al*. 2010) that the use of a very large number of APs does not result in the improved performance of WiFi positioning. An optimum of approximately 10 APs has been suggested by Chan *et al*. (2010), although this is expected to vary with the specific positioning algorithm. This study does not indicate that there is an optimum number of APs. One additional observation from Figure 3 is that some of the largest errors of several hundred metres occur for locations with more than 50 APs. It appears that using many APs may in fact present some challenges for WiFi positioning algorithms, as if they are 'overwhelmed' by the number of signals to be potentially used in processing.

The final aspect of the analysis is an examination of the precision of WiFi positioning between the two devices. In other words, how close together are WiFi locations taken at approximately the same place and time but on different devices? Figure 4 shows the WiFi positional error as a scatter plot for the three study areas combined. WiFi location estimates are plotted as the distance and direction from the reference location for each observation. Each pair of WiFi location estimates recorded at the same place and time but on two different devices is connected by a black line. A total of 89 pairs are plotted in Figure 4. One observation can be considered an outlier and extends beyond the range of the scatter plot; the positional error for this particular location estimate on the iPhone is 1986 m.

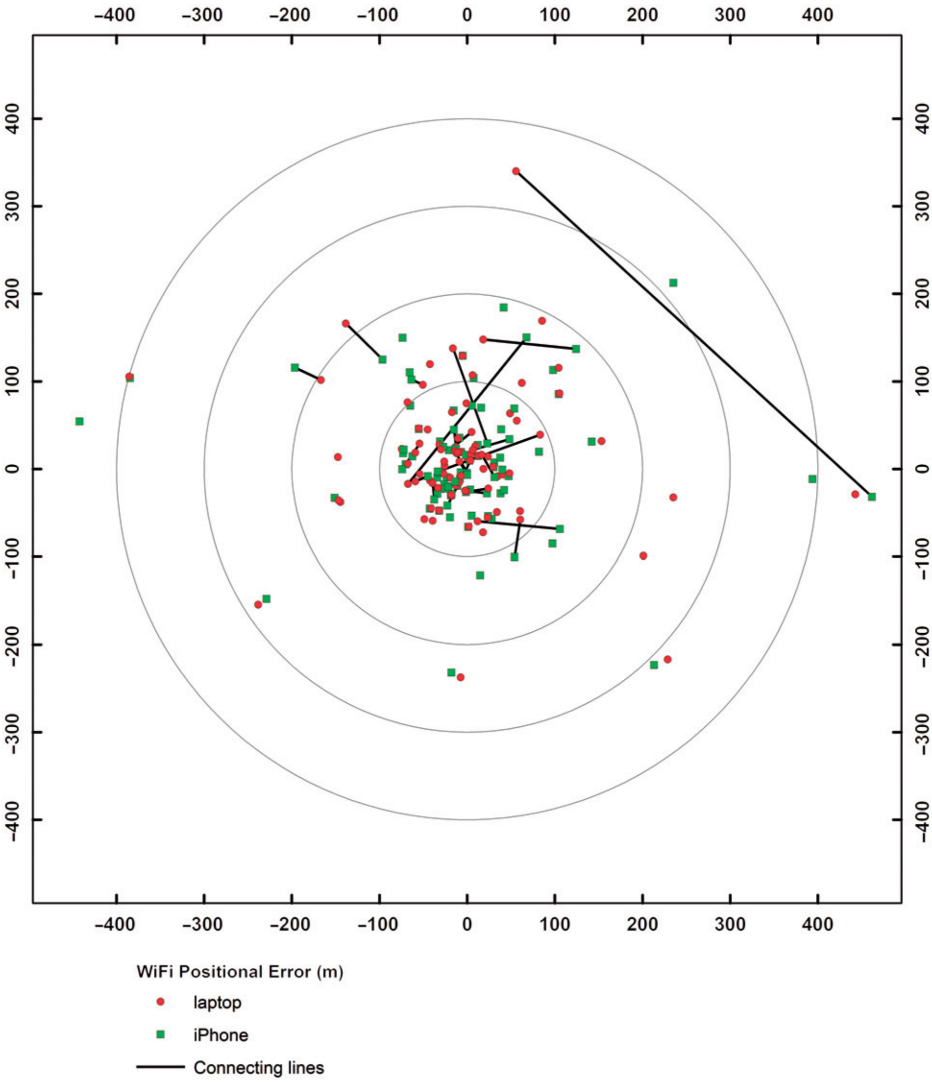

Figure 4. Positional error of WiFi positioning (m) relative to reference location.

Figure 4 reveals some interesting patterns. First, there does not appear to be any dominant direction in the connecting lines, suggesting that no directional bias is introduced by the device type. Second, examples where the laptop position is more accurate are roughly as common as those where the iPhone position is more accurate. This confirms the general agreement in the overall accuracy metrics already discussed. Finally, small and large differences (i.e. the length of the connecting line) occur at both small and large positional errors. There are quite a number of cases where the two observations are very close together, as indicated by a very short connecting line. These cases are indicative of the good replicability of WiFi positioning on the two devices. However, there are a number of cases where the two observations are a substantial distance apart. Out of total of 90 sites, the two observations on different devices are more than 50 m apart at 21 sites and more than 100 m apart at 13 sites. This suggests that the replicability of WiFi positioning between the two devices is not as good as the analysis of aggregate error metrics first indicated. The substantial distances between two observations on two different devices at approximately the same place and time can likely be attributed to the different set of APs observed by the two devices and to the signal strengths reported by the different WiFi chipsets. Previous research has found significant differences in the reported signal strength between WiFi devices (e.g. Tao *et al.* 2003, Haeberlen *et al.* 2004, Lui *et al.* 2011). In addition to the WiFi chipset, reported signal strength has also been found to be dependent on antenna and hardware design, which may present an inherent limitation to the accuracy of WiFi positioning systems (Lui *et al.* 2011).

5. Discussion and conclusions

The positional accuracy metrics of WiFi positioning using Skyhook's system were found to be very similar on the two devices evaluated using a sample of 90 sites. Statistical testing revealed no statistically significant differences between the error distributions. This suggests that, in general, the replicability of WiFi positioning on the two devices is high and that a single device does not consistently produce more accurate location estimates.

Median values for positional error in the three study areas ranged from 43 to 92 m. These results are similar to earlier independent evaluations of Skyhook's WiFi positioning system (Gallagher *et al.* 2009, Zandbergen 2009) but substantially exceed those reported by Skyhook (Skyhook Wireless 2008). In each study area, a small number of observations were recorded with a positional error of several hundred metres, suggesting that great care should be taken when utilising WiFi positioning for LBS that require very accurate locations.

The number of APs observed on the iPhone was consistently lower than that on the laptop. This can be attributed to the higher quality WiFi card on the laptop. The lower number of APs on the iPhone, however, was not found to reduce positional accuracy. In general, no relationship between number of APs and positional accuracy was found for either device, counter to earlier findings. The results indicate, however, that WiFi positioning can be achieved with a small number of APs (5–10), but an increased number of APs does not contribute to improved positional accuracy.

Replicability of WiFi positioning using Skyhook's system in terms of getting the same location by using two different devices at approximately the same place and time was relatively poor. In a number of cases the distance between the two observations was greater than the average positional error of the two observations. Replicability of WiFi location estimates on different devices should therefore be considered quite poor, but the general performance metrics on different devices for a larger sample are nearly identical.

This study has a number of limitations. First, testing considered only Skyhook's WiFi position system and did not include other systems which have been implemented more recently such as those by Google, Apple and Microsoft.

Without additional testing, it is therefore not known how Skyhook's system compares to other systems in terms of positional accuracy, coverage and replicability. Second, testing only included two specific devices and the findings may not be applicable to other types of WiFi devices. Testing across a wider range of devices would be useful to establish a stronger baseline of minimum performance expectations for WiFi positioning. Third, testing was carried out without replication over longer time periods. As the AP landscape changes at a specific location and the areas are recalibrated, location estimates using WiFi positioning are likely to change. This temporal aspect of replicability warrants attention in future research.

Despite these limitations, the results from this study have implications for LBS on mobile devices under the assumption that the performance realised by Skyhook's system on the two devices is applicable to other commercial WiFi positioning systems and other WiFi devices. First, the positional error in metropolitan WiFi positioning can be substantial. Applications that require a positional accuracy in the order of 10–30 m should not employ WiFi positioning. Many consumer applications, however, may not require this level of accuracy and larger errors may not hinder the utility of the service. Second, when two users are both employing WiFi positioning (e.g. in a social networking application), the positional errors are not very likely to cancel each other out, i.e. the errors are not expected to be in the same direction. Instead, the error resulting from the combination of two observations may be compounded. This suggests that WiFi positioning is most appropriate for applications where a positional error in the order of 100 m is acceptable.

References

Bensky, A., 2007. *Wireless positioning technologies and applications*. Boston, MA: Artech House Publishers.

Chan, E.C.L., Baciu, G., and Mak, S.C., 2010. Wi-Fi positioning based on Fourier descriptors. *In*: *International conference on communications and mobile computing*, Shenzhen. Los Alamitos, CA: IEEE Computer Society, 545–551.

Chang, N., Rashidzadeh, R., and Ahmadi, M., 2010. Robust indoor positioning using differential Wi-Fi access point. *IEEE Transactions on Consumer Electronics*, 56 (3), 1860–1867.

Cheng, Y.C., *et al*., 2005. Accuracy characterization for metropolitan-scale Wi-Fi localization. *In*: *Proceedings of the 3rd international conference on mobile systems, applications, and services*, 6–8 June 2005, Seattle, WA. New York: ACM Press, 233–245.

Cheong, J.W., *et al*., 2009. GPS/WiFi real-time positioning device: an initial outcome. *In*: G. Gartner and K. Rehrl, eds. *Location based services and telecartography II*. Section IV. Lecture Notes in Geoinformation and Cartography, Heidelberg, Berlin: Springer-Verlag, 439–456.

Chiou, Y.S., *et al*., 2009. Design of an adaptive positioning system based on WiFi radio signals. *Computer Communications*, 32 (7–10), 1245–1254.

Feng, C., *et al*., 2010. Compressive sensing based positioning using RSS of WLAN access points. *In*: *Proceedings of the 29th conference on information communications*, 15–19 March 2010, San Diego, CA. Piscataway, NJ: IEEE Educational Activities Department, 1631–1639.

Gallagher, T., *et al*., 2009. Trials of commercial Wi-Fi positioning systems for indoor and urban canyons. *In*: *Proceedings of the IGNSS symposium 2009*, 1–3 December, Queensland, Australia.

Gezici, S., 2007. A survey on wireless position estimation. *Wireless Personal Communications*, 44 (3), 263–282.

Haeberlen, A., *et al*., 2004. Practical robust localization over large-scale 802.11 wireless networks. *In*: *Proceedings of the 10th annual international conference on mobile computing and networking*, 26 September–1 October, Philadelphia, 70–84.

Hightower, J. and Borriello, G., 2004. Particle filters for location estimation in ubiquitous computing: a case study. *In: Proceedings of the UbiComp 2004*, 7–10 September, Nottingham, UK, 88–106.

Kamker, S., 2010. How I met your girlfriend. *In*: *Presentation at Black Hat USA conference*, 28–29 July 2010, Las Vegas, NV.

King, T., Haenselmann, T., and Effelsberg, W., 2007. Deployment, calibration, and measurement factors for position errors in 802.11-based indoor positioning systems. *In*: *Proceedings of LoCA 2007*. Vol. 4718, Heidelberg, Berlin: Springer-Verlag, 17–34.

Kolodziej, K.W. and Hjelm, J., 2006. *Local positioning systems: LBS applications and services*. Boca Raton, FL: CRC Press.

Kushki, A., Plataniotis, K.N., and Venetsanopoulos, A.N., 2007. Kernel-based positioning in wireless local area networks. *IEEE Transactions on Mobile Computing*, 6 (6), 689–705.

LaMarca, A., *et al*., 2005. Place lab: device positioning using radio beacons in the wild. *In*: H.W. Gellersen, R. Want and A. Schmidt, eds. *Pervasive computing*. Vol. 3468, Lecture Notes in Computer Science, Heidelberg, Berlin: Springer-Verlag, 116–133.

Liao, I. and Kao, K., 2008. Enhancing the accuracy of WLAN-based location determination systems using predicted orientation information. *Information Sciences*, 178 (4), 1049–1068.

Lin, H., *et al*., 2009. *Enhanced indoor locationing in a congested Wi-Fi environment*. Pittsburgh: Carnegie Institute of Technology.

Lui, G., *et al*., 2011. *Differences in RSSI readings made by different Wi-Fi chipsets: a limitation of WLAN localization*. Sydney: School of Surveying and Spatial Information Systems, University of New South Wales.

Manodham, T., Loyola, L., and Miki, T., 2007. A novel wireless positioning system for seamless internet connectivity based on the WLAN infrastructure. *Wireless Personal Communications*, 44 (3), 295–309.

Mok, E. and Retscher, G., 2007. Location determination using WiFi fingerprinting versus WiFi trlateration. *Journal of Location Based Services*, 1 (2), 145–159.

Skyhook Wireless, 2008. *Wi-Fi positioning system: accuracy, availability, and time to fix performance*. Boston, MA: Skyhook Wireless.

Swangmuang, N. and Krishnamurthy, P., 2008. An effective location fingerprint model for wireless indoor localization. *Pervasive and Mobile Computing*, 4 (6), 836–850.

Tao, P., *et al*., 2003. Wireless LAN location-sensing for security applications. *In*: *Proceedings of the ACM Workshop on Wireless Security (WiSE 2003)*, 10–12 December 2003, Rome, Italy, 11–20.

Tippenhauer, N.O., *et al*., 2009. Attacks on public WLAN-based positioning systems. *In*: *Proceedings of the 7th international conference on mobile systems, applications, and services*, 22–25 June 2009, Kraków, Poland, 29–40.

Tsui, A.W., *et al*., 2010. Accuracy performance analysis between war driving and war walking in metropolitan Wi-Fi localization. *IEEE Transactions on Mobile Computing*, 9 (11), 1551–1562.

Wallbaum, M., 2007. A priori error estimates for wireless local area network positioning systems. *Pervasive and Mobile Computing*, 3 (5), 560–580.

Yim, J., 2008. Introducing a decision tree-based indoor positioning technique. *Expert Systems with Applications*, 44, 1296–1302.

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 (7), 869–883.

Youssef, A., *et al.*, 2005. Computing location from ambient FM radio signals [commercial radio station signals]. *In*: *Proceedings of the wireless communications and networking conference*, Vol. 2, 13–17 March 2005, New Orleans, LA, 824–829.

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 (1), 103–130.

Zandbergen, P.A., 2009. Accuracy of iPhone locations: a comparison of assisted GPS, WiFi and cellular positioning. *Transactions in GIS*, 13 (6), 5–25.