Practical Metropolitan-Scale Positioning
for GSM Phones

Mike Y. Chen[1], Timothy Sohn[2],
Dmitri Chmelev[3], Dirk Haehnel[1], Jeffrey Hightower[1], Jeff Hughes[3],
Anthony LaMarca[1], Fred Potter[3], Ian Smith[1], Alex Varshavsky[4]

[1] Intel Research Seattle, USA, [2] University of California at San Diego, USA
[3] University of Washington, USA, [4] University of Toronto, Canada

Abstract. This paper examines the positioning accuracy of a GSM beacon-based location system in a metropolitan environment. We explore five factors effecting positioning accuracy: location algorithm choice, scan set size, simultaneous use of cells from different providers, training and testing on different devices, and calibration data density. We collected a 208-hour, 4350Km driving trace of three different GSM networks covering the Seattle metropolitan area. We show a median error of 94m in downtown and 196m in residential areas using a single GSM network and the best algorithm for each area. Estimating location using multiple providers' cells reduces median error to 65-134 meters and 95% error to 163m in the downtown area, which meets the accuracy requirements for E911. We also show that a small 60-hour calibration drive is sufficient for enabling a metropolitan area similar to Seattle.

1   Introduction

While several research and commercial efforts now exist for WiFi beacon-location [6, 9, 25, 27], little research exists on how beacon-location extends to the most ubiquitous mobile computing platform today: the mobile phone. Mobile phones are an attractive platform for emerging location-aware applications [19, 20] with an estimated 2 billion subscribers world-wide as of 2005–about ten times as many as the total number of mobile PCs in-use [22, 24]. Specifically, the GSM family of technologies, which includes W-CDMA, has more than 1.5 billion subscribers, which is more than the total number of Internet users world-wide [23, 24]. In addition, mobile phones have long battery life, constant connectivity, and are usually at hand and powered on.

This paper examines the feasibility of a client-side, beacon-based GSM location system and whether the methods from the WiFi literature can be retargeted to GSM phones in a metropolitan setting. Our approach differs from operator-provided network-based GSM location solutions in that the phone can position itself in a *privacy-observant* manner and can use cell towers from *all network operators* to compute location without requiring assistance from the network operators. The approach also uses the *existing hardware* in mobile phones without requiring any additional hardware.

There are several important differences between WiFi and GSM. First, the range of a GSM cell can be up to 35Km, which is 70 times larger than WiFi's maximum range of around 500m. Second, the deployment of GSM networks is stable and planned compared to the more ad hoc deployment of WiFi access points. Third, GSM operates in a licensed frequency band and is thus less prone to interference caused by other electronic devices such as cordless phones and microwaves. To study the effect these differences have on accuracy, coverage, and calibration overhead, we collected an extensive 208-hour, 4350Km driving trace of three major GSM networks covering the Seattle metropolitan area: AT&T, Cingular, and T-Mobile.[1] Collecting such a large trace was necessary to assess the algorithms' accuracies without introducing significant experimental error, as we will describe in Section 2.2. For each GSM network, we sampled the radio environment at 1Hz throughout our drive. Each sample from each device on each network contained up to 7 different cell IDs with their associated signal strengths, network provider ID, and area ID.

Our contributions are the following:

- We characterize the accuracy of three published positioning algorithms: a simple centroid algorithm that does not model radio propagation, fingerprinting, and Monte Carlo localization with a Gaussian Processes signal propagation model. Our experiments show that we can achieve a median accuracy of 94m in downtown and 196m in residential areas. For GSM, choosing a good algorithm can result in up to 388% improvement in position accuracy, which is in stark contrast to past wide-area analyses of WiFi positioning algorithms where the choice of algorithm was comparatively irrelevant since accuracy only varied 20% across different algorithms [4].

- We show that using *cross-provider* GSM beacons can significantly improve positioning accuracy compared to only using cells from a single provider, achieving median accuracy of 65-134 meters. In the downtown Seattle area, this technique meets the E911 positioning requirements for network-based solutions.

- We show that *cross-device* operation is possible with only 6-8% degradation in accuracy for the centroid algorithm and 57%-63% for Gaussian Processes on devices with completely different radios and antennas.

- We also show that a small 60-hour calibration drive is sufficient for enabling a metropolitan area similar to Seattle.

- We have publicly released our GSM positioning toolkit [13] with the algorithmic implementations described in the paper along with tower location traces, allowing researchers at other institutions to experiment with new positioning algorithms. This data contains information about all three major GSM providers in the USA as well as WiFi scans over the 208 hours of data collection.

The rest of the paper is organized as follows: Section 2 presents our data collection methodology and the positioning algorithms, Section 3 characterizes the accuracy of the algorithms under various conditions, Section 4 discusses the implications of our findings, Section 5 presents related work, and Section 6 discusses future work and concludes.

---

[1] Cingular recently acquired AT&T Wireless, but many AT&T-identified towers still existed at the time of this study.

2 Methodology

This section describes our data collection methodology, the trace characteristics, and the positioning algorithms.

2.1 Data Collection

Our data collection hardware consisted of an IBM Thinkpad T30 laptop with a WiFi card, two GPS units, three Sony Ericsson GM28 GSM modems and three Audiovox SMT5600 phones (also known as the HTC Typhoon phones), shown in Figure 1. The GSM phones and GSM modems were fitted with SIM cards from each of the three GSM network providers in the area: AT&T, Cingular, and T-Mobile. A second, identical setup provides redundancy in the event of equipment failure. In all, we used 12 GSM devices to collect GSM traces and 4 redundant GPS units to provide ground truth for location.

Our data collection software is implemented in C#. It records the attached cell ID as well as observed signal strength in dBm for up to 7 GSM cells for each of the modems and the phones, independently once every second for the phones and the maximum scan rate of once every three seconds for the modems. Readings from each GPS unit are recorded once a second. For ground truth of location, we use the latitude and longitude values from the readings that meet the following criteria: valid GPS lock with 5 or more satellite and a low horizontal dilution of precision (HDOP).

A dense calibration trace allows us to characterize the best-case positioning accuracy of a GSM-based location system and perform sensitivity analyses. To collect such a trace, we put our data collection setup in a car and drove every publicly accessible street in the Seattle metropolitan area as shown in Figure 2. External antennas for the GSM modems and GPS units were placed on the car roof to improve signal reception. We drove over 208hrs, or 4350Km, and collected over 24GB of traces over a period of three months. The complete trace contains 6756 unique cells across the three network providers. During this effort, we had one laptop failure due to rain and one car accident.

Figure 1. *Our GSM/WiFi data collection device: one WiFi card, two GPS units (left), and three Sony Ericsson GM28 GSM modems (center) are connected to an IBM Thinkpad T30 laptop. The modems require external antennas (bottom) and we have modified them to be powered via USB hubs. Three Audiovox SMT5600 phones, one for each network provider, are shown on the right,*
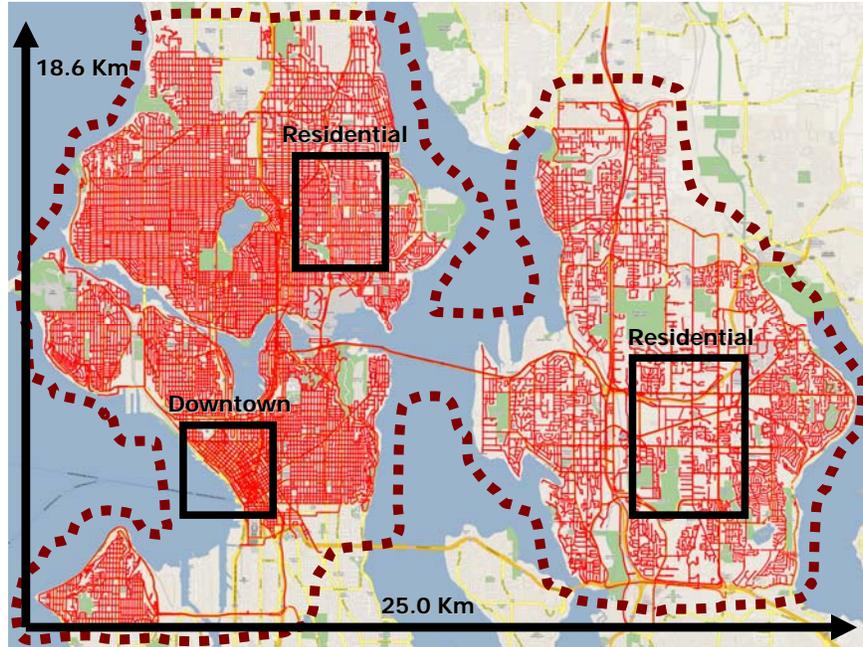
Figure 2. *Dotted-line area and street highlights show the roads driven during our GSM trace collection spanning an 18Km x 25Km region of the Seattle metropolitan area. The three solid rectangles outline the test areas:* Downtown *with high cell tower density and* Residential, *two neighborhoods with lower cell density.*

The calibration trace is used to train the three positioning algorithms. To measure the accuracy of these algorithms, we selected three test neighborhoods and collected a second complete test trace in each neighborhood two weeks after the training trace was collected. Downtown Seattle was chosen as one test neighborhood as it had the highest tower density at 66 cells/Km^2. In addition, we selected two residential neighborhoods with lower cell tower densities. On average, these neighborhoods had 26 cells/Km^2 or 39% of Downtown's cell density. In our results, the data for the two low-density neighborhoods have been combined. Table 1 summarizes the properties of the training and the test traces.

Table 1. *Properties of the collected training trace and two testing traces*

	Training Trace	*Testing Traces*	
		Downtown	**Residential**
Name			
Duration	208hr	70min	169min
Distance	4350Km	24Km	89Km
Dimension	25.0 x 18.6Km	2.7 x 2.3Km	2.6 x 4.1Km + 4.6 x 5.5Km
Area	Greater Seattle	Downtown Seattle	Ravenna + East Bellevue
Avg. Cell Density	28 cells/Km^2	66 cells/Km^2	26 cells/Km^2

2.2   Training Area Size

A valid question to ask is why the algorithmic evaluation requires collecting training traces from a much larger area than the regions tested. The answer is that the training data must be large enough to cover the complete cells of all towers seen in the testing traces. If the training data does not meet this constraint then we risk artificially inflating the measured accuracy of the algorithms by not letting them make mistakes they might otherwise make. For example, an algorithm sensing a tower located on the border of the testing area should be allowed to err by estimating the phone's position to be in the part of that tower's cell that is outside the testing area. Indeed, by shrinking the training area to match the size of the testing area we found that we could artificially inflate the accuracy up to 41% for some algorithms. Therefore, we were careful to collect a wide-scale training trace of the entire greater Seattle area extending well beyond the boundaries of all our testing areas.

2.3   Positioning Algorithms

In this paper, we measure the performance of three positioning algorithms from the research literature: a centroid algorithm that does not model radio propagation [9], a radio fingerprinting algorithm [3], and Monte Carlo localization with a learned Gaussian Processes signal propagation model [5, 14]. While a wide variety of radio-based location algorithms have been published, we chose these three as they are representative of the spectrum of positioning algorithms and vary in complexity and expected accuracy.

All our algorithms have been implements in a C# location toolkit that runs on Microsoft Windows Mobile Smartphones, PDAs, and PCs running Windows. On an HTC Typhoon phone, our toolkit can poll GSM readings and calculate its location four times per second using the centroid algorithm and the cell tower maps size for the 6756 cells we observed in Seattle occupy only 44KB of compressed data.

There are two phases to analyzing the positioning accuracy of an algorithm. First, a training trace containing time-stamped GSM and GPS measurements is used to build a model that is specific to that algorithm. In the second phase, the algorithm uses the GSM measurements in an independent testing trace to estimate its position, and outputs latitude/longitude values in its position estimate. Positioning error is computed by calculating the distance between the positions estimated by the algorithm and the ground truth positions provided by GPS. We used the Haversine Formula of distance between two points over the earth as the distance metric [21].

2.3.1   Centroid Family
The centroid algorithm [9] is very fast to compute and, in its basic form, does not employ a radio propagation model. Given a lookup table of <Cell ID, Latitude, Longitude> entries, the centroid algorithm estimates the phone's position to be the geometric center of all the cells that are seen in a measurement. Weighting by the received signal strength observed in the scans is an extension that can offer a small improvement in accuracy. Our experiments used the centroid algorithm in its basic form without any modeling of radio propagation.

Downtown	**Residential (cropped)**

Figure 3. *Cell location map for the three network providers; each dot represents the estimated location of a cell. The left map shows Downtown with an average density of 66 cells/Km². The right map shows a cropped Residential region with an average cell density of 26 cells/Km².*

Centroid is the only algorithm we consider that depends on having an estimate of the true tower positions. Because the true cell tower positions are not publicly available in the USA, the training phase estimates tower positions by averaging the places where the highest signal strengths in each cell was observed. Figure 3 shows map of the estimated tower positions for two of our test areas. To evaluate our tower placement accuracy, we randomly selected six cell towers and physically visited these towers to precisely verify their true location. In this test, we found an average error of 56m and a maximum error of 76m. These values are reasonable estimated lower bounds for position error in the centroid algorithm.

2.3.2   Fingerprinting

Radio fingerprinting is a positioning method that assumes the radio beacons and associated signal strengths observed at a particular location is stable over time. The training phase constructs a search index mapping radio fingerprints to locations. In testing, upon seeing a similar radio fingerprint the algorithm can use the index to deduce the phone's position. The RADAR system [3] is an example of this technique applied to WiFi positioning in laptops. Fingerprinting has also more recently been used with GSM radio signatures for server-side, wide-area positioning [10, 28] and indoor location and floor estimation [12].

The fingerprinting algorithm does not create a map of estimated tower positions nor does it model radio propagation. Instead, it creates a search index of radio fingerprints to latitude/longitude coordinates. To position a device, the algorithm uses the constructed index and calculates the Euclidean distance in signal strength space between the current fingerprint and all available fingerprints in the index [4]. It then selects k fingerprints with the smallest Euclidean distance as potential indicators of the current location. The location of the device is estimated as an average of the latitude and longitude coordinates of the best k matches. The accuracy of the location estimate is highly dependent on the density of the set of collected fingerprints. The indoor WiFi and GSM localization papers cited above collected fingerprints at a density of around one fingerprint per square meter. This paper investigates how well

fingerprinting works with sparser calibration and less uniformly distributed set of GSM fingerprints at a metropolitan scale. We also characterize the effects of several practical factors on positioning accuracy.

2.3.3   Monte Carlo Localization with Gaussian Processes Signal Models

Gaussian Processes-based Monte Carlo localization uses a radio propagation model and Markov localization to predict the phone's position. The idea is to build a sensor model to predict the signal strength at each location and then use this information to compute the likelihood of measurements. The phone's position is estimated using a Bayesian particle filter.

   This approach is like fingerprinting, except it uses an abstract parametric model of the signal environment instead of building a direct search index of the calibration data itself as is done with fingerprinting. To model the signal propagation, we use Gaussian Processes, which are nonparametric models that estimate Gaussian distributions over functions based on the training data [14]. In order to achieve fast execution we pre-process the signal propagation function to a grid with 15m grid-cells. The computation of the signal propagation can then be implemented simply by a look-up function in the maps of the cell towers. With the predicted signal strength, we can compute the likelihood of an observation given the phone is at a particular location.

   To represent the posterior probability distribution about the position of the phone, we apply standard Monte-Carlo localization called particle filtering [5]. In Monte-Carlo localization, the belief about the phone's position is represented by a set of random samples. Each sample consists of a state vector of the underlying system, which is the position of the mobile phone, and a weighting factor. The weight is the likelihood of the measurement at the particle's location. The posterior is represented by the distribution of the samples and their importance factors. This particle filter algorithm used by our system is also known as sequential importance sampling [1].

3   Results

This section presents analyses that explore the effects of five factors on positioning accuracy: algorithm selection, scan set size, simultaneous use of cells from different providers, training and testing on different devices, and calibration drive density. Because our goal is to characterize the positioning accuracy of GSM, we have anonymized the names of the three network providers and we will refer to them as Provider A, B, and C. For single-provider results, we report the median values among the three providers.

3.1   Effect of Algorithm Selection on Positioning Accuracy

We evaluated the positioning error for each of the three algorithms described in Section 2.3 using the test traces collected from the test areas. Table 2 shows the median and 90[th] percentile error for each algorithm grouped by the test areas. All

Table 2. *Median and 90th-percentile positioning errors (in meters) for Centroid, Fingerprinting, and Gaussian Processes algorithms in the two test areas*

	Downtown (higher density)		Residential (lower density)	
	50%	**90%**	**50%**	**90%**
Centroid	232	574	760	2479
Fingerprinting	94	291	277	984
Gaussian Processes	126	358	196	552

three algorithms performed better in the higher tower density area, Downtown, than the lower tower density area, with the median error ranging from 94-232 meters compared to 196-760 meters.

Since the centroid algorithm does not model signal strength and assigns equal weight to each cell, the position estimate can be greatly affected by the density and the placement of the cells resulting in the worst position estimates. The fingerprinting algorithm performs much better than centroid, achieving a 94m median error in Downtown and 277m in Residential. One reason for this decrease in error is that Downtown has more obstructions due to large buildings than the Residential area. These obstructions actually prove advantageous because they help form unique fingerprints to allow the algorithm to differentiate between nearby locations. The more open space coupled with the lower cell density in the Residential area results in less unique fingerprints, thereby increasing the positioning error.

The Gaussian Processes algorithm models signal strengths with continuous functions, and therefore is not able to capture the sharp changes in signal strengths due to obstructions as well as fingerprinting, resulting in slightly worse accuracy in Downtown. In the more open environments, however, it models the sparse training data sufficiently to produce the best accuracy in the residential areas.

Positioning error from WiFi localization techniques has been shown to be highly dependent on access point density. Moreover, the use of complex positioning algorithms only provide a 20% improvement in accuracy compared to the simple algorithms such as centroid [4]. With the much larger cell sizes of GSM, we found that algorithmic improvements can improve positioning accuracy by 247%- 388%.

3.2 Effects of Scan Set Size

At any given time, a GSM device may be within range of a large number of GSM cells. A client-side location system, however, may be limited in the number of nearby cells it can sense. For many models of mobile phones, normal user-level application are not allowed any information about which cell tower the phones are associated with or the observed signal strength. The Series 60 phones [26] (e.g. Nokia 6600 and N92) allow user programs to find out the ID and observed signal strength of the single cell with which the phone is currently associated. Other devices, such as the GSM modems and phones we used in our study, provide information about the cell the phone is associated with, as well as six other nearby cells (making a total of seven). In this experiment, we investigate what effect the size of this set has on position
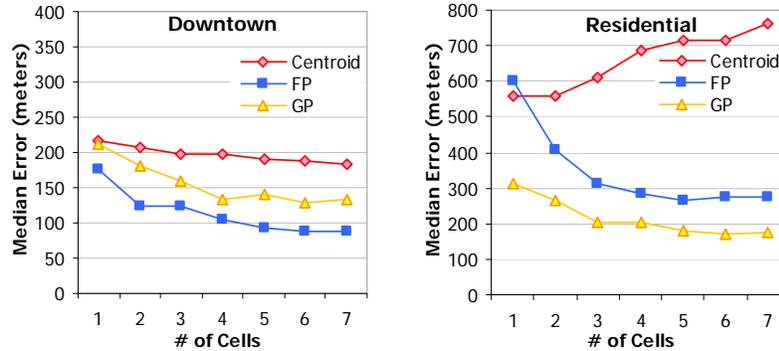
Figure 4. *Median positioning error as a function of the N cells with the strongest signal strengths for the Downtown and Residential test areas. FP stands for fingerprinting and GP stands for Gaussian Processes.*

accuracy for all three algorithms. We vary the number of cells between one and seven by sorting the observed cells by signal strength and using only the *n* strongest.

Figure 4 shows sensitivity analysis of positioning error versus the number of cells available for both Downtown and Residential. Both the fingerprinting and Gaussian Processes algorithms improve in accuracy as more cells are used for positioning, for a 50% and 37% improvement in Downtown, respectively. The additional cells, regardless of signals strengths, provide information that help differentiation among similar measurements and help improve the position estimates.

Centroid also exhibits the same trend in areas of high cell density, showing improved accuracy with more cells in Downtown. However, as cells become sparser and distant in the Residential area, the cells with weak signal strengths contribute to increased positioning error. This trend is evident in Figure 4, where surprisingly, using only 1-2 cells gives the best accuracy, and using 5 or more cells produce the worst positioning accuracy.

3.3   Effects of Using Towers From Multiple Providers

Once associated with a cell tower, GSM devices only monitor cells from that tower's network, even though cell towers from other providers may be closer and have stronger signals. Up to this point, all of the results have reflected this limitation: our AT&T phone estimates its location using the AT&T training data, the T-Mobile phone estimates its location using the T-Mobile data, etc. We now consider what would happen if a GSM phone could scan for cells from all available networks. Ideally, it would increase both the number of nearby observable cell towers and the number of strongly observable towers, thus providing an opportunity to improve positioning accuracy.

Table 3. *Cross-provider median positioning error in meters when 1) using only cells from a single provider, 2) using the top 7 strongest cells across providers, and 3) using all available cells across providers*

	Downtown			Residential		
	Single Provider (7)	Cross-Provider (7)	Cross-Provider (all)	Single Provider (7)	Cross-Provider (7)	Cross-Provider (all)
Centroid	187	166	170	647	456	574
Fingerprinting	94	153	245	277	313	297
Gaussian Processes	126	87	65	196	147	134

To evaluate its effectiveness, we simulate a cross-provider device by combining measurements from the three network providers: AT&T, Cingular, and T-Mobile that were collected within 5 seconds of each other, and use them to estimate location. Because each GSM device gets information for up to 7 nearby cells on its network, our simulated device will hear up to 21 nearby cells. This raised a concern that accuracy improvements may only be due to the magnitude of the scan set. To measure this effect, we first limit our simulated cross-provider device to only use the strongest 7 cells it hears across the three networks. We then evaluate the positioning accuracy when all available cells from all providers are used.

Table 3 shows the effect on positioning accuracy when each algorithm uses only cells from a single provider, the top 7 strongest cells across providers, and all possible cells. For Centroid, the increased ratio of nearby towers is the most significant with the strongest 7 cells, which shows an improvement of 11-30% compared to a single provider. Increasing the number of cells reduces the improvement.

Surprisingly, fingerprinting performs much worse when cells from multiple providers are used. A closer inspection shows that one of the network providers, Provider A, added new cells to the downtown area (or renamed existing cells) during our 3-month calibration drive. Because we drove the downtown area during the first month, we did not observe these new cells with strong signal strengths at their true location. Rather, we observed them with weak signal strengths kilometers away, resulting in skewed position estimates produced by fingerprinting and centroid. Gaussian Processes is more resistant to this effect because the models only contain information about these cells at weaks signal strengths, limiting their impact during the tests drive when these cells were observed at strong signal strengths.

For Gaussian Processes, using the strongest 7 cells across providers improves the median error by 27-40% compared to a single provider. Using all the available cells, it further improves the median error for a total improvement of 45-55%, achieving a median error of 65m for Downtown and 134m for Residential.

The cross-provider accuracy improvement is more significant towards the tail end of the distribution, as evident in the more vertical curves in the Cumulative Distribution Function (CDF) shown in Figure 5. When all available cells are being used, Gaussian Processes achieves 136m at the 90th percentile in Downtown compared to 308-907m for the three individual providers. Specifically, the 67th percentile error is 88m and the 95th percentile is 163m, which is within the E911 requirements of 100m and 300m, respectively, for network-based solutions.
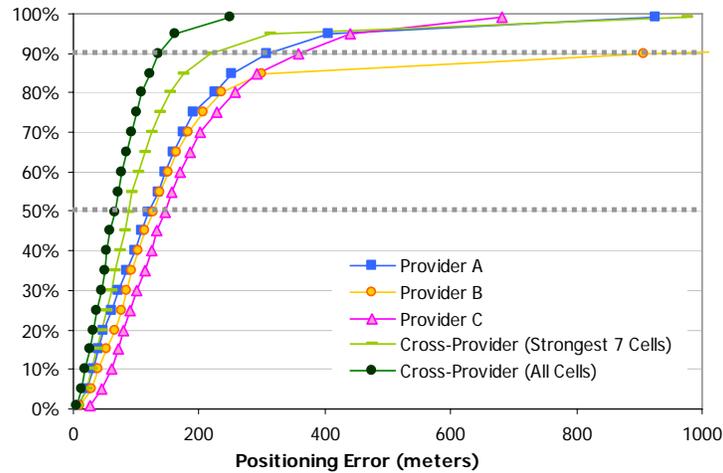
Figure 5. *CDF of positioning error for the Gaussian Processes algorithm in Downtown using 1) cells from a single provider, 2) the top 7 strongest cells across multiple network providers, and 3) all available cells from multiple network providers.*

3.4 Effects of Training on One Device and Testing on Another

Our results thus far have all been presented based on a single GSM device: the HTC Typhoon phone. Moreover, we used the same device to gather data for both testing and training. A practical question is if other GSM devices, including another HTC Typhoon phone, can be used to achieve comparable positioning results using the same calibration data. For simple algorithms such as centroid, the only algorithmic requirement is that the device observes a similar set of cell towers as the HTC Typhoon. For other algorithms, the device must observe similar towers, and its signal strength values would have to correlate well with the HTC Typhoon. Given a strong correlation, a transformation function could convert the device's signal strength values to those reported by the HTC Typhoon.

We compared the common cells seen, and the signal strength values for those cells, among three GSM devices: a duplicate HTC Typhoon phone, a HTC Tornado phone, and a Sony Ericsson GM28 modem. These devices represent GSM units that have different radio and different antenna designs. All devices report a signal strength value with each observed cell.

Table 4 shows the average number of common towers seen per scan between each device and the reference device, and the Pearson correlation coefficient for the signal strength values of these common cells. Even though the HTC Tornado has a different antenna and the modem has both a different radio and antenna, all devices exhibit strong signal strength correlation that are statistically significant to the 0.001 level.

Table 4. *Similarity between different GSM devices and the reference HTC Typhoon phone, showing the number of cells that are the common when two different devices scan at the same time. Pearson correlation coefficient and significance are shown for the signal strengths of these common cells between each device and the reference phone.*

Devices	Radio	Antenna	Average # of Common Cells	Signal Strength Correlation	Correlation Significance
HTC Typhoon (reference)	*Same*	*Same*	*7.000*	*1.000*	*.000*
HTC Typhoon (duplicate)	Same	Same	6.484	0.828	<.001
HTC Tornado	Same	Diff.	5.018	0.789	<.001
Sony Ericsson GM28 Modem	Diff.	Diff.	4.283	0.874	<.001

Given that a linear transformation function can map the signal strength from one device to another, we now explore the effect of training and testing on different devices. Because the handoff behavior are different on different GSM devices due to radio and antenna design, techniques such as fingerprinting that do exact pattern matching of signal strengths should see the most degradation.

Table 5 shows the median error when we train on the data from the HTC Typhoon phone and test on the GM28 modem after applying a linear transformation function for the observed signal strengths. These accuracy numbers are conservative because the phone and the modems had completely different radio and antenna designs, with the modems using roof-mounted antennas.

The Gaussian Processes algorithm had the best accuracy in both Downtown and Residential, with median accuracy of 206m and 307m, and degradation of 63% and 57%, respectively. Fingerprinting performed the worst in Downtown at 366m, and has the largest increase in error in both Downtown and Residential compared to training and testing on the same device. This is likely due to the modems switching over to different cells at slightly different times than the phones when they moved across cell boundaries. Centroid was the most robust with the least degradation at 6-8%.

Table 5. *Cross-device median positioning error and % change when training with the trace collected on one device (the HTC Typhoon phones) and testing on another device (the Sony Ericsson modems)*

	Downtown		Residential	
	50% (meters)	**% change**	**50% (meters)**	**% change**
Centroid	245	5.6%	818	7.6%
Fingerprinting	366	289%	803	190%
Gaussian Processes	206	63%	307	57%

Figure 6. *Example of a generated, virtual street grid that simulates a drive density equivalent to 10% of the full training trace.*

3.5 Effects of Reducing Calibration Drive Density

We characterized the tradeoff between calibration drive density and positioning error by simulating a sparser driving pattern from our comprehensive data set. In contrast to prior experiments on WiFi density that effectively simulate devices with slower scanning rates [4], we are interested in spatial thinning that simulates the effect of driving fewer streets. Understanding this effect is useful to estimate the resource and cost necessary to calibrate a GSM-based positioning system to support accuracy requirements of the intended applications. Simple algorithms that do not model signal propagation characteristics only require the location of the cell towers, which can potentially be supplied by the GSM network providers. More sophisticated algorithms, however, rely on calibration data to improve positioning accuracy, but it is unclear how much calibration is necessary.

To simulate these sparse drives, we super-impose a virtual street grid pattern on our dense calibration trace and filter measurements that do not fall on the virtual street grid. By varying the width between the virtual streets, varying levels of driving density can be simulated. In order to reduce systematic error due to the interaction of the virtual and real grids, we use five random offsets for each grid width. We then average them for each width to estimate the positioning error given the level of density being simulated.

Figure 7 shows the median error versus the percentage of data dropped using our simulated street grids for Downtown. With centroid, using only 20% of the density produces statistically equivalent positioning accuracy as our full drive density. Both fingerprinting and Gaussian Proccesses algorithms show slight, but gradual
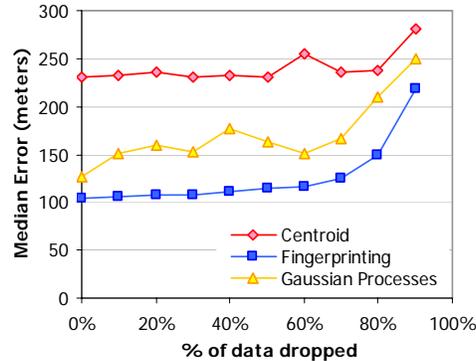
Figure 7. *Median positioning error as a function of the calibration drive density, simulating the effect of driving fewer streets. Dropping 70% of the data still provides comparable median error to the full calibration drive.*

degradation as the amount of data decreases, because the quality of the radio models degrades with less calibration. All three algorithms had median positioning error above 200m when only 10% of the density is used. Although a comprehensive drive still produces the best positioning error, less dense drives can still produce similar accuracy without the added expense. For the Seattle metropolitan area, 30%, or 60 hours of driving is sufficient to calibrate a GSM-based positioning system without a significant loss in accuracy. More sophisticated street selection can further reduce the amount of calibration required.

4   Discussion

We have presented a variety of results to show how three location estimation algorithms performed using GSM traces from three neighborhoods. By using data from different devices and by simulating varying scan-set sizes and training trace densities, we modeled the algorithms performance in a variety of situations. Table 6 show a high-level summary of these results. From this table, we can draw a number of conclusions about when and where these algorithms are the most appropriate.

A number of mobile phones, including the popular Series 60 phones [26], only provide information about a single tower at a time. For these devices, as well as phones with constrained storage and computation capability, the best algorithm is the centroid. The centroid algorithm uses very few resources, and is extremely robust to the quality of the training set. Our data showed the centroid to be the most resistant to sparse training data as well as training data collected on a different device. This simplicity comes at the cost of being the least accurate in areas with both high and low tower densities.

Table 6. *Summary of the characteristics of the three positioning algorithms*

	Storage Required (for Seattle, compressed)	CPU Usage	Accuracy (Dense Towers)	Accuracy (Sparse Towers)	Required Density of Training Data	Requires Same-Device Training Set	Benefits from Cross-Provider Scanning	Tolerant of Phones Exposing Single Cell
Centroid	Low (44KB)	Low	232m	760m	Low	No	Yes	Yes
Finger-printing	High (188MB)	Med.	94m	277m	High	Yes	No	No
Gaussian Processes	Med. (80MB)	High	126m	196m	Med.	No	Yes	Yes

For mobile phones with more storage and processing power and the ability to scan for multiple cell towers at a time, the Monte Carlo localization with Gaussian Processes model probably offers the best overall mix of accuracy and practicality. It was the most accurate in the residential neighborhoods and was close to the best algorithm in downtown. It was also robust to single cells positioning as well as cross-device training data.

Fingerprinting was the most accurate in the high-density urban area and performed well in the low-density area as well. The fingerprint training sets consume large amounts of storage, but with optimized indexing the relevant fingerprints can be accessed quickly. However, fingerprinting was the most fragile of the algorithms, requiring dense training data collected on similar, if not identical, mobile phones. It was also the most fragile to the addition of new cells during the calibration drive.

The most worrisome results in this study are the sensitivities we measured to cross-device accuracy degradation. Although we tested a particularly dramatic cross-device scenario (phone versus modem) with completely different antenna configuration (internal versus vehicle roof-mounted), the error increase was higher than we would like to see. We believe more research, both at the hardware and algorithmic levels is warranted to understand and mitigate the errors in this area.

Our data also shows that if future mobile phones support the ability to scan for towers across service providers, the Gaussian Process algorithm becomes the clear winner. While the Centroid and Fingerprinting showed no benefit from additional tower data, the Gaussian Process algorithm saw up to a 50% reduction in median error. Perhaps more importantly, the 95[th]-percentile error showed an even more pronounced drop, improving by over 70% in our downtown test area.

5   Related Work

The most common location technology today is the Global Positioning System (GPS). Although GPS is effective in open environments, it does not work well when the GPS receiver is indoors or in dense urban areas. The research community generally agrees that many location-enhanced mobile applications require another technology to

augment or replace GPS to meet coverage and accuracy needs. GSM and WiFi-based location techniques are common ways to overcome the shortcomings in GPS.

The four papers most related to this work are follows: LaMarca et al [9] presents an overview of Place Lab, with experiments that are appropriate to study WiFi positioning accuracy, but not sufficient for characterizing GSM. Trevisani and Vitaletti [15] studies single cell-ID location in mobile phone networks that compares the accuracy and coverage of operators' location systems in different cities, Laitinen et al [10] analyzes the positioning accuracy of a server-side fingerprinting approach in both urban and suburban environments. Otsason et al [12] explores client-side indoor GSM location and floor differentiation. Our work characterizes the wide-area positioning accuracy for three classes of algorithms and investigates the effects of several practical issues such as cross-device positioning and calibration drive density. We also present a novel cross-provider positioning technique that significantly improves positioning accuracy.

The E911/E112 initiatives in the US and Europe specify requirements on localization accuracy for mobile phones placing emergency calls. These initiatives have catalyzed a market for network operator-provided location capabilities and services like AT&T Wireless' friend-finder and Sprint-NexTel's fleet management tools. Operators calculate mobile phone positions using hybrid network-client techniques like Assisted GPS (AGPS) where the network data links provide aiding information to a limited in-phone GPS chip, or network-only techniques like enhanced observed time difference (EOTD), angle of arrival (AOA), and time difference of arrival (TDOA) [15]. AGPS is more accurate than network-only techniques and extends the coverage of standalone GPS to operate in areas with limited GPS reception, such as urban canyons and some indoor environments

Beacon-based location with WiFi positioning is also a well-studied problem by the Place Lab project and other researchers [4, 6, 9, 10, 11] and has been commercialized by Microsoft Virtual Earth [25] and SkyHook Wireless [28]. They have demonstrated that the WiFi beacon approach is viable and can result in a good indoor-outdoor location system with high coverage and sufficient accuracy for many mobile applications. Metrics including coverage, accuracy, beacon density, mapping and calibration drive overhead, and performance have all been evaluated. Median accuracy of the approach is 15m-60m with nearly 100% coverage in urban areas. Wide-area beacon-based approaches complement the many indoor positioning systems that provide high precision in indoor environments but require specialized hardware or have high installation costs. Examples of these systems include Cricket [15], Active Badge [16], and Active Floor [17].

A related but distinct class of research to the work in this paper is the problem of learning places using mobile devices carried by the user. These place-learning systems do not provide real-time navigation or fulfill the "dot-on-a-map" application scenarios. Instead, they provide the ability to recognize previously visited destinations by using metrics like GPS dropout [2], cell handoff patterns [8], or radio signatures [7] to match previously learned places.

6   Conclusions

This paper examined the positioning accuracy and the practical challenges in deploying a GSM beacon-based location system in a metropolitan environment. To do this, we collected data over a larger area than past WiFi beacon-based location projects because the range of GSM cells are up to 70 times larger than WiFi access points. We collected an extensive 208hr, 4350Km trace covering the Seattle metropolitan area.

Our results show that existing GSM devices can achieve a positioning accuracy with a median error of 94-196 meters using cells from a single provider. We have observed that the positioning accuracy varies significantly across algorithms, by a factor of almost 4x, compared to past characterizations for WiFi that have a maximum variation of 20%.

Our analysis on calibration drive density suggests that 30% of our dataset was sufficient to provide comparable positioning accuracy – suggesting that 60hrs of driving can cover a metropolitan area similar to the size of Seattle. Our results show that cross-device positioning is possible with only 6-8% degradation in accuracy for the centroid algorithm and 57%-63% for Gaussian Processes on devices with completely different radios and antennas. We believe more research, both at the hardware and algorithmic levels is warranted to understand and mitigate the errors in this area.

Finally, we have identified an opportunity to significantly improve accuracy by scanning cells across all available providers, for a median error of 65-134m, which is a factor of 3-4x of the published accuracy for WiFi.

Acknowledgements

References

1. S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, "A Tutorial on Particle Filters for Online Non-Linear/Non-Gaussian Bayesian Tracking", *IEEE Transactions on Signal Processing*, 50(2):174–188, 2002.
2. D. Ashbrook, T. Starner, "Using GPS to Learn Significant Locations and Predict Movement across Multiple Users", *Personal and Ubiquitous Computing* 7 (2003) 275–286
3. P. Bahl and V. N. Padmanabhan, "RADAR: An In-Building RF-Based User Location and Tracking System", in *Proceedings of IEEE INFOCOM* 2000, Vol. 2: 775-784
4. Y. Cheng, Y. Chawathe, A. LaMarca and J. Krumm, "Accuracy Characterization for Metropolitan-scale WiFi Localization", in *Proceedings of Mobisys* 2005.
5. D. Fox, W. Burgard, F. Dellaert, and S. Thrun, "Monte Carlo localization: Efficient Position Estimation for Mobile Robots", in *Proceedings of AAAI*, 1999.

6. A. Haeberlen, E. Flannery, A. M. Ladd, A. Rudys, D. S. Wallach, and L. E. Kavraki, "Practical Robust Localization over Large-scale 802.11 Wireless Networks," in *Proceedings of Mobicom*, 2004.

7. J. Hightower, et al, "Learning and Recognizing the Places We Go", in *Proceedings of Ubicomp*, pp. 159-176, Sep. 2005.

8. K. Laasonen, M. Raento, and H. Toivonen. "Adaptive On-device Location Recognition", in *Proceedings of the Second International Conference on Pervasive Computin*g. Volume 3001 of Lecture Notes in Computer Science., Springer-Verlag (2004) 287–304

9. A. LaMarca, et al, "Place Lab: Device Positioning Using Radio Beacons in the Wild," in *Proceedings of the Third International Conference on Pervasive Computing*, May 2005.

10. H. Laitinen, J. Lahteenmaki, T. Nordstrom, "Database correlation method for GSM location", *IEEE 53rd Vehicular Technology Conference*, 2001.

11. J. Letchner, D. Fox, and A. LaMarca, "Large-Scale Localization from Wireless Signal Strength", In *Proceedings of the National Conference on Artificial Intelligence (AAAI-05).*

12. V. Otsason, A. Varshavsky, A. LaMarca, E. de Lara: "Accurate GSM Indoor Localization," in *Proceedings of Ubicomp 2005*: 141-158.

13. Privacy-Observant Location System, http://pols.sourceforge.net/

14. A. Schwaighofer, M. Grigoras, V. Tresp, and C. Hoffmann. "GPPS: A Gaussian Process Positioning System for Cellular Networks", in *Proceedings of NIPS 2003*.

15. E. Trevisani and A. Vitaletti. "Cell-ID Location Technique, Limits and Benefits: An Experimental Study." In *Proceedings of WMCSA 2004*. 51—60.

16. N. B. Priyantha, A. Chakraborty, H. Balakrishnan. "The cricket location-support system." In *Proceedings of Mobicom 2000*, pp. 32-43

17. R. Want, A. Hopper, V. Falco, J. Gibbons. "The Active Badge Location System." *ACM Transactions on Information Systems* 10, 1 (1992), 91-102

18. M. D. Addlesee, A. Jones, F. Livesey, and F. Samaria. "The ORL Active Floor." *IEEE Personal Communications* 4, 5 (1997), 35-41.

19. T. Sohn, et al, "Place-Its: A Study of Location-Based Reminders on Mobile Phones." In *Proceedings of Ubicomp 2005*.

20. I. Smith, et al, "Social Disclosure of Place: From Location Technology to Communication Practice." In *Proceedings of Pervasive 2005*.

21. R.W. Sinnott, "Virtues of the Haversine", Sky and Telescope, vol. 68, no. 2, 1984, p. 159

22. Computer Industry Almanac Press Release. "Mobile PCs In-Use Surpass 200M." June, 2005. http://www.c-i-a.com/pr0605.htm

23. Computer Industry Almanac Press Release. "Worldwide Internet Users will Top 1 Billion in 2005." Sept 2004. http://www.c-i-a.com/pr0904.htm

24. GSM Association Press Release. "Worldwide cellular connections exceeds 2 billion." Sept 2005. http://www.gsmworld.com/news/press_2005/press05_21.shtml

25. Microsoft Virtual Earth. http://virtualearth.msn.com

26. Series 60 Phone Platform. http://s60.com

27. Skyhook Wireless. http://www.skyhookwireless.com

28. US Wireless. http://web.archive.org/web/20031124182802/http://uswcorp.com