

Accurate GSM Indoor Localization

Veljo Otsason¹, Alex Varshavsky², Anthony LaMarca³, and Eyal de Lara²

¹ Tartu University, veljo.otsason@utoronto.ca

² University of Toronto, walex.delara@cs.toronto.edu

³ Intel Research Seattle, anthony.lamarca@intel.com

Abstract. Accurate indoor localization has long been an objective of the ubiquitous computing research community, and numerous indoor localization solutions based on 802.11, Bluetooth, ultrasound and infrared technologies have been proposed. This paper presents the first accurate GSM indoor localization system that achieves median accuracy of 5 meters in large multi-floor buildings. The key idea that makes accurate GSM-based indoor localization possible is the use of *wide* signal-strength fingerprints. In addition to the 6-strongest cells traditionally used in the GSM standard, the wide fingerprint includes readings from additional cells that are strong enough to be detected, but too weak to be used for efficient communication. Experiments conducted on three multi-floor buildings show that our system achieves accuracy comparable to an 802.11-based implementation, and can accurately differentiate between floors in both wooden and steel-reinforced concrete structures.

1 Introduction

The accurate localization of objects and people in indoor environments has long been considered an important building block for ubiquitous computing applications [7,8]. Most research on indoor localization systems has been based on the use of short-range signals, such as WiFi [3,5,11], Bluetooth [1], ultra sound [15], or infrared [16]. This paper shows that contrary to popular belief an indoor localization system based on wide-area GSM fingerprints can achieve high accuracy, and is in fact comparable to an 802.11-based implementation.

GSM-based indoor localization has several benefits: (i) GSM coverage is all but pervasive, far outreaching the coverage of 802.11 networks; (ii) the wide acceptance of cellular phones makes them ideal conduits for the delivery of ubiquitous computing applications. A localization system based on cellular signals, such as GSM, leverages the phone's existing hardware and removes the need for additional radio interfaces; (iii) because cellular towers are dispersed across the covered area, a cellular-based localization system would still work in situations where a building's electrical infrastructure has failed. Moreover, cellular systems are designed to tolerate power failures. For example, the cellular network kept working during the massive power outage that left most of the Northeastern United States and Canada in the dark in the Summer of 2003; (iv) GSM, unlike 802.11 networks, operates in a licensed band, and therefore does not suffer

from interference from nearby devices transmitting on the same frequency (e.g., microwaves, cordless phones); and (v) the significant expense and complexity of cellular base stations⁴ results in a network that evolves slowly and is only reconfigured infrequently. While this lack of flexibility (and high configuration cost) is certainly a drawback for the cellular system operator, it results in a stable environment that allows the localization system to operate for a long period before having to be recalibrated.

This paper presents the first accurate GSM-based indoor localization system. The key idea that makes accurate GSM-based indoor localization possible is the use of *wide* signal-strength fingerprints. The wide fingerprint includes the 6-strongest GSM cells and readings of up to 29 additional GSM channels, most of which are strong enough to be detected, but too weak to be used for efficient communication. The higher dimensionality introduced by the additional channel dramatically increases localization accuracy.

We present results for experiments conducted on signal-strength fingerprints collected from three multi-floor buildings located in Toronto and Seattle. These structures span a wide spectrum of urban densities, ranging from a busy downtown core to a quiet residential neighborhood. The results show that our GSM-based indoor localization system can effectively differentiate between floors and achieves median within-floor accuracy as low as 2.5 meters.

We make the following contributions: (i) we present the first accurate GSM-based indoor localization system and show that it achieves accuracy comparable to an 802.11-based implementation; (ii) we show that a GSM-based localization system can effectively differentiate between floors for both wooden and steel-reinforced concrete structures; (iii) we show that there is significant signal diversity across metropolitan environments⁵ and that this diversity enables the GSM-based system to achieve high localization accuracy; and (iv) we show that the availability of signal strength readings from cells other than the 6-strongest cells traditionally used in GSM increases localization accuracy by up to 50%.

The rest of this paper is organized as follows. Section 2 describes related work. Section 3 describes our methodology, and Section 4 presents results for our experimental evaluation. Finally, Section 5 concludes the paper and discusses directions for future work.

2 Related Work

This paper examines the effectiveness of GSM fingerprinting as an indoor localization technique. While this combination is new, indoor localization, radio fingerprinting and use of GSM for localization have all been explored before. We describe these efforts and key distinctions between these efforts and ours.

⁴ A macro-cell costs \$500,000 to \$1 million. Micro-cells cost about a third as much, but a larger number is needed to cover the same area [14].

⁵ In all three indoor environments (including the private residence) we were able to detect at least 24 different GSM signals.

2.1 Indoor Localization

While outdoor localization is almost exclusively performed using the Global Positioning System (GPS), indoor location systems have successfully employed a variety of technologies. The original Active Badge system [7] and follow on commercial systems like Versus [20] use infrared emitters and detectors to achieve 5-10m accuracy. Both the Cricket [15] and the Bat [16] systems use ultrasonic ranging to estimate location. Depending on the density of infrastructure and degree of calibration, ultrasonic systems have accuracies between a few meters and a few centimeters. Most recently, ultra-wideband emitters and receivers have been used to achieve highly accurate indoor localization [19]. The common drawback of all of these systems is that they require custom infrastructure for every area in which localization is to be performed. As a result, these systems have not seen significant deployment outside of high-value applications like hospital process management. In contrast, GSM fingerprinting makes use of the existing GSM infrastructure, obviating the need for infrastructure investment and greatly increasing the possible area in which the system will work. This increases the likelihood of GSM fingerprinting achieving popular adoption.

2.2 Indoor Localization Using 802.11 Fingerprinting

Bahl *et al.* observed that the strength of the signal from an 802.11 access point does not vary significantly in a given location. They used this observation to build RADAR [3], a system that performed localization based on which access points would be heard where, and how strongly. This was the first *fingerprinting* system that showed that it is possible to localize a laptop in the hallways of a small office building within 2-3 meters of its true location, using fingerprints from four 802.11 access points. There have been improvements to RADAR's fingerprint matching algorithm that have improved accuracy [2,11,17] and differentiated floors of a building with a high degree of precision [6]. In addition, commercial localization products have been built using 802.11 fingerprinting [18]. The differences between our work and 802.11 fingerprinting systems are primarily due to the differences between 802.11 and GSM that were outlined in Section 1: Due to higher coverage, GSM fingerprinting works in more places than 802.11 fingerprinting. Due to more stable infrastructure, 802.11 radio maps will degrade more quickly than GSM radio maps. Due to the larger range of GSM cells, 802.11 fingerprinting will be more accurate than GSM fingerprinting given the same number of radio sources.

2.3 Localizing Using GSM

A number of systems have used GSM to estimate the location of mobile clients. The Place Lab system employed a map built using war-driving software and a simple radio model to estimate a cell phone's location with 100-150 meter accuracy in a city environment [13]. The goal of Place Lab was to provide coarse-grained accuracy with minimal mapping effort. This is different, and complementary to our goal of doing accurate indoor localization given a detailed radio

survey. Another distinction is that Place Lab used a cell phone platform that only programmatically exported the single associated cell tower.

Laitinen *et al.* [12] used GSM-based fingerprinting for outdoor localization. They have collected sparse fingerprints from the 6-strongest cells, achieving 67th percentile accuracy of 44m. Finally, Laasonen *et al.* used the transition between GSM cell towers to build a graph representing the places a user goes [10]. Like Place Lab, Laasonen’s system used cell phones that only exported the single cell-tower the phone was associated with. In contrast to the other systems we have mentioned, Laasonen’s system did not attempt to estimate absolute location, but rather assigned locations symbolic names like *Home* and *Grocery Store*.

These previous efforts to use GSM for localization differ from the work reported in this paper in that they are based on sparse fingerprints collected tens to hundreds of meters apart from each other. Moreover, these efforts used *narrow* fingerprints obtained from commercial GSM phones that report the signal strength for the current cell [10,13] or the 6-strongest cells [12]. In contrast, we collected GSM fingerprints in a dense grid with 1.5 meters granularity. Moreover, we collected *wide* fingerprints that include up to 29 different GSM channels in addition to the 6-strongest GSM cells.

3 Methodology

This section first gives an overview of GSM and wireless signal fingerprinting. We then describe our data collection process and the localization algorithms that we use in our evaluation.

3.1 GSM Primer

GSM is the most widespread cellular telephony standard in the world, with deployments in more than 100 countries by over 220 network operators [4]. In North-America, GSM operates on the 850 MHz and 1900 MHz frequency bands. Each band is subdivided into 200 KHz wide physical channels using Frequency Division Multiple Access (FDMA). Each physical channel is then subdivided into 8 logical channels based on Time Division Multiple Access (TDMA). There are 299 non-interfering physical channels available in the 1900 MHz band, and 124 in the 850 MHz band, totaling 423 physical channels in North-America.

A GSM *base station* is typically equipped with a number of directional antennas that define sectors of coverage or *cells*. Each cell is allocated a number of physical channels based on the expected traffic load and the operator’s requirements. Typically, the channels are allocated in a way that both increases coverage and reduces interference between cells. Thus, for example, two neighboring cells will never be assigned the same channel. Channels are, however, reused across cells that are far-enough away from each other so that inter-cell interference is minimized while channel reuse is maximized. The channel to cell allocation is a complex and costly process that requires careful planning and typically involves field measurements and extensive computer-based simulations of radio

signal propagation. Therefore, once the mapping between cells and frequencies has been established, it rarely changes.

Every GSM cell has a special broadcast control channel (BCCH) used to transmit, among other things, the identities of neighboring cells to be monitored by mobile stations for handover purposes. While GSM employs transmission power control both at the base station and the mobile device, the data on the BCCH is transmitted at a full constant power. This allows mobile stations to compare signal strength of neighboring cells in a meaningful manner and choose the best one for further communication. It is these BCCH channels that we use for localization. In the rest of this paper, we refer to the BCCH channels simply as *channels*.

3.2 Fingerprinting

Two factors lead to the good performance of radio fingerprinting in the wireless band used by GSM and 802.11 networks. The first is that the signal strengths observed by mobile devices exhibit considerable spatial variability at the 1-10M level. That is to say, a given radio source may be heard stronger or not at all a few meters away. The second factor is that these same signal strengths are consistent in time; the signal strength from a given source at a given location is likely to be similar tomorrow and next week. In combination, this means that there is a radio profile that is feature-rich in space and reasonably consistent in time. Fingerprinting-based location techniques take advantage of this by capturing this radio profile for later reference.

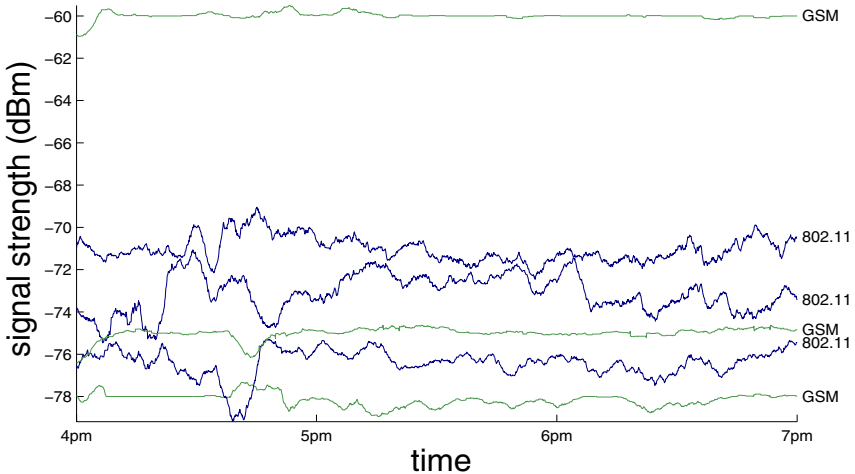


Fig. 1. 802.11 and GSM signal stability over time.

To compare the stability of GSM and 802.11 signals, we recorded the signal strength of nearby 802.11 access points (AP) and 6-strongest GSM cells at several

locations in one of the buildings that houses the Department of Computer Science at the University of Toronto. Figure 1 shows a 3-hour segment of the signal strength measurements at a location on the fifth floor of the building during a workday afternoon. The plot shows the three-strongest GSM cells and the three-strongest 802.11 APs. GSM signals appear to be more stable than 802.11 signals. We believe that this is because 802.11 uses unlicensed overcrowded 2.4 GHz band, and therefore suffers from interference from nearby appliances such as microwaves and cordless phones. An analysis of GSM signal stability under different weather conditions (e.g., rain, snow, fog) is left for future work.

Fingerprinting relies on a “training phase” in which a mobile device moves through the environment recording the strength of signals emanating from a group of radio sources (e.g., 802.11 access points, GSM base stations, FM radio [9] or TV stations). We refer to the physical position where the measurement is performed as a *location*, to the radio scan as a *measurement* and to the recording of the signal strength of a single source as a *reading*. That is, to build a radio map of the building, a mobile device takes a series of measurements in multiple locations of the building. Each measurement is composed of several readings; one for each radio source in range. The set of data recorded in a single location is also referred to as a *training point*. Since fingerprinting systems do not model radio propagation, a fairly dense collection of radio scans needs to be collected to achieve good accuracy. The original RADAR experiments, for example, collected measurements every square meter on average[3]. To achieve their advertised accuracy, the commercial 802.11 fingerprinting product from Ekahau [18] recommends a similar density.

Once the training phase is complete, a client can estimate its location by performing a radio scan (or equivalently collecting a *testing point*) and feeding it to a *localization algorithm*, which estimates the client’s location based on the similarity of the signal strength signatures between the testing and the training points. The similarity of signatures can be computed in a variety of ways, but it typically involves finding measurements in the training points that have the same radio sources with similar signal strengths. The easiest technique for estimating location is to choose the location of the training point with the closest Euclidean distance in a signal space. Better accuracy can be achieved by averaging the location of the K closest neighbors (or training points) in the radio map, where K is some small constant. It is also beneficial to use weighted averaging, so that neighbors closer in signal space are given higher weights.

In this paper, we compare the accuracy of localization based on 802.11 and GSM fingerprinting using the popular weighted K nearest neighbors algorithm. Investigating the applicability of other localization algorithms to GSM fingerprinting is a topic for future work.

3.3 Data Collection

We collected multi-floor measurements in two office buildings and one private detached house. The office buildings are the home to the Intel Research Seattle Lab and part of the Department of Computer Science of the University of

Toronto. In the rest of this paper, we refer to these buildings as: University, Research Lab, and House. University is located in Toronto’s busy downtown core, and Research Lab and House are located in Seattle’s commercial midtown and a quiet residential neighborhood, respectively.

University is a large ($88m \times 113m$) 8-storey building with lecture rooms, offices and research labs. Since we had no access to the offices, we collected training points in the hallways of the 5th and 7th floors of the building⁶. Research Lab is a medium size ($30m \times 30m$) 6-storey building. Space inside the building is partitioned with semi-permanent cubicles. Due to access restrictions, we collected readings from the whole 6th floor, but only a half of the 5th floor. House is a 3-storey wooden structure ($18m \times 6m$) that includes a basement and two floors above ground. We collected measurements on all 3 floors. The distance between floors is about 6 meters for University and Research Lab, and about 3 meters for House.

We collected 802.11 and GSM fingerprints using a laptop running Windows XP. To collect 802.11 fingerprints, we used an Orinoco Gold wireless card configured in active scanning mode, where the laptop periodically transmits probe requests and listens to probe responses from nearby 802.11 APs.

We collected GSM fingerprints using a Sony/Ericsson GM28 GSM modem, which operates as an ordinary GSM cell phone, but exports a richer programming interface. The GSM modem provides two interfaces for accessing signal strength information: *cellsAPI* and *channelsAPI*. The *cellsAPI* interface reports the cell ID, signal strength, and associated channel for the n strongest cells. While the modem’s specifications does not set a hard bound on the value of n , in practice in the 3 environments we measured n was equal to 6. The *channelsAPI* interface simultaneously provides the signal strength for up to 35 channels, 13 of which can be specified by the programmer, with up to 22 additional channels picked by the modem itself. In practice, 6 of the 35 channels typically corresponds to the 6-strongest cells. Unfortunately, *channelsAPI* reports signal strength but does not report cell IDs. We speculate that the cell ID information for other than the 6-strongest cells cannot be determined because the signals of those cells are strong enough to be detected, but too weak to be used for efficient communication.

	University (downtown)	Research Lab (midtown)	House (residential)
Cells	-87.69	-76.74	-88.35
Channels	-96.41	-102.19	-105.27

Table 1. Average signal strength (dBm) for cells and channels.

Table 1 shows the average signal strength returned by the *cellsAPI* and *channelsAPI* interfaces. As expected, the average signal strength reported by *cellsAPI*

⁶ We did not take measurements on the 6th floor because at the time of this study it was under going extensive renovations.

is significantly higher than the average reported by *channelsAPI*. Note that the average signal strength reported by the *channelsAPI* interface is close to modem's stated receiver sensitivity⁷ of -102 dBm. Efficient GSM communication requires an SNR higher than -90 dB.

The lack of cell ID information for some channels raises the possibility of *aliasing*, i.e., a situation when two or more cells transmitting simultaneously on the same channel appear to be a single radio source and therefore cannot be differentiated. In the extreme case, a fingerprinting system that relies exclusively on channel-based data may suffer from worldwide aliasing. Because channels are reused throughout the world, fingerprints taken in two far-away locations may produce similar fingerprints. To alleviate the aliasing problem, we combine the information returned by the *cellsAPI* and *channelsAPI* interfaces into a single fingerprint. We then restrict the set of fingerprints to which we compare a *testing point* to fingerprints that have at least one cell ID in common with the *testing point*. This practice effectively differentiates between fingerprints from our three indoor environments.

As we show in Section 4, even with the potential for aliasing, our localization system based on wide GSM fingerprinting significantly outperforms GSM fingerprinting based on the 6-strongest cells, and is comparable to 802.11 based fingerprinting. This is because our fingerprints are wide (have many readings), and therefore, in order for the aliasing to reduce accuracy, many readings in the fingerprints of distant locations need to match, which is highly unlikely in practice.



Fig. 2. Experimental setup



Fig. 3. Measuring signal strength and identifying location by clicking on the map

⁷ In practice, the modem reports signal strength as low as -115 dBm.

	University (downtown)		Research Lab (midtown)		House (residential)		
	5th	7th	5th	6th	basement	1st	2nd
per floor	130	154	53	181	17	44	50
per building	284		234		111		

Table 2. Training points collected on each floor for the three buildings.

We developed a simple Java-based application to assist us in the process of gathering fingerprints. To record a fingerprint, we first identify the current position by clicking on a map of the building. The application then records the signal strengths reported by the 802.11 card and the *cellsAPI* and *channelsAPI* interfaces of the GSM modem. To collect the measurements, we placed the laptop on an office chair and moved the chair around the building. While primitive, this setup assures measurements collected at a constant height. Figures 2 and 3 show our experimental setup and a screen shot of the Java-based application, whereas Table 2 summarizes the number of training points collected on each of the floors of the three buildings. In all three indoor environments, we collected 802.11 and GSM fingerprints for points located 1 to 1.5 meters apart. We collected 2 measurements per location, waiting 5 seconds between the scans (the default value according to the modem specification).

Practical Considerations We collected our wide fingerprints using a programmable Sony/Ericsson GSM modem, which operates as an ordinary GSM cell phone, but exports a richer programming interface that provides access to readings from up to 35 GSM channels. In contrast, commercial phones limit access to signal strength information to the 6-strongest cells or even just the current cell. However, we speculate that the software on commercial phones could be easily enhanced to provide signal strength measurements for a richer set of channels. Once extended, those phones could take advantage of the wide-fingerprinting technique introduced in this paper. We base this speculation on the observation that the Sony/Ericsson GSM modem is implemented using standard GSM electronics, and that the GSM standard requires phones to be able to scan all channels in the GSM band.

3.4 Localization Algorithms

We implemented four localization algorithms which differ in the structure of their fingerprints: (i) *802.11*, uses only readings from 802.11 access points; (ii) *onecell*, uses the reading of the single strongest GSM cell; (iii) *cell*, uses readings of the 6-strongest GSM cells; and (iv) *chann*, uses readings from up to 35 GSM channels.

All our localization algorithms use the K-nearest neighbors algorithm described in Section 3.2. For each algorithm, we varied the number of nearest neighbors to average over, and selected the value of K that gave the best results.

In most cases, the best K was a small constant (2 or 4). We also experimented with assigning higher weights to neighbors that appear closer in the signal space. The weight is the reciprocal of the distance in signal strength space between the testing point and the specific nearest neighbor.

An initial evaluation of `chann` uncovered cases in which the algorithm selected points that are neighbors in the signal space, but are actually located far away from the true location of the testing point in the physical space. To ameliorate the effect of these false positives, we applied the K -mean clustering algorithm to split the set of nearest neighbors into two geographical clusters⁸. We then removed the points that belong to the smaller cluster from the final location calculation. In the rest of this paper, we refer to the version of `chann` that uses geographical clustering as `channcl`.

4 Evaluation

In this section, we first analyze the collected data and then evaluate localization accuracy obtained by 802.11 and GSM fingerprinting.

4.1 Data Analysis

Table 3 shows the total number of 802.11 APs, GSM cells, and GSM channels recorded during the data collection phase for each of the 3 buildings. The University building has a much denser 802.11 deployment than the Research Lab building both because the University building is much larger and because while the APs at the Research Lab building are maintained by IT personnel, numerous APs at the University building are owned and maintained by independent research groups.

The total number of GSM cells seen at the University building is larger than in other buildings because of the higher coverage and the larger building size. The lower number of cells seen at the Research Lab is the consequence of both the much smaller building size and the much stronger signal received from nearby cells. Because of the proximity of a few base stations, the strongest cells reported by the modem in the Research Lab benefit from less variations than in other buildings (i.e., the same group of cells appears in most of the cell measurements). The total number of channels seen in the residential area is somewhat lower than in other areas due to lower coverage.

Figure 4 plots the cumulative distribution function (CDF) of the number of transmitters per location detected at the University building. Shown are fingerprints based on 802.11 AP, GSM cells, and GSM channels. The figures for the Research Lab and the House show similar patterns and are therefore not included. The median width of 802.11 AP and GSM cells fingerprints is 5 and 6, respectively. In contrast, the median width of GSM channel fingerprints is 25. We will show in the next section that the larger fingerprint has a dramatic effect on localization performance.

⁸ We experimented with different numbers of clusters, but 2 clusters produced the best results.

	University (downtown)	Research Lab (midtown)	House (residential)
802.11 APs	44	10	5
Cells	58	14	18
Channels	34	33	24

Table 3. Total number of 802.11 APs, GSM cells and GSM channels recorded.

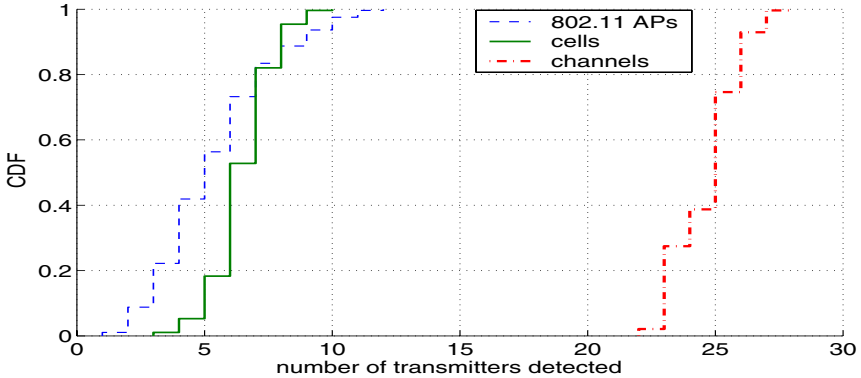


Fig. 4. Cumulative distribution of the number of transmitters per location detected at the University building. Shown are fingerprints based on 802.11 AP, GSM cells, and GSM channels.

4.2 Performance

The results reported in this section were obtained by taking one point at a time out of the training set and using it as the testing point. This technique is similar to that used by Bahl [3], and is a somewhat pessimistic approach since it makes sure that there is no training point in the radio map that has exactly the same fingerprint as the testing point. On the other hand, all our measurements were performed with a single modem in a course of two days. In the future, we plan to investigate the effects of using different hardware for training and testing, as well as the effect of separating the training and testing of the system by larger time intervals.

Floor Classification Table 4 summarizes the effectiveness with which the localization algorithms introduced in Section 3.4 differentiate between floors in the three indoor environments. All algorithms predict the current floor as the one where most of the K -nearest neighbors are located. chann_{cl} achieves similar performance to chann and is therefore not shown.

As expected, 802.11 does an excellent job differentiating between floors in the University and Research Lab buildings. The reinforced concrete floors in these structures effectively block the propagation of 802.11 signals between floors,

	University (downtown)	Research Lab (midtown)	House (residential)
802.11	100	100	62.16
chann	89.08	97.01	93.69
cell	89.08	81.2	51.35
onecell	74.65	77.35	57.66

Table 4. % of succesful floor classifications.

significantly simplifying the task of floor prediction. These results are consistent with previous findings [6].

In the house environment, however, 802.11 achieves low classification accuracy as the house’s wood structure presents little obstacle to radio propagation, making it harder to differentiate between signal fingerprints on different floors. Not surprisingly, all but 3 of the 42 misclassifications happen at locations on the first and second floors of the house. In the house scenario, 4 out of 5 of the available 802.11 signals emanate from neighboring residences. These signals propagate easily through the wooden frame of the first and second floors, but suffer significant attenuation propagating through dirt and the house’s foundations to reach the basement. The low power at which neighboring access points are heard (if at all) in the basement helps to identify basement locations. On the other hand, the 802.11 signals from neighboring households contribute little to improving the accuracy of predictions for the above-ground floors.

In contrast, the GSM-based chann algorithm shows strong performance across all three buildings, and significantly outperforms 802.11 for the House environment. Overall, chann achieves up to 42% better accuracy than cell and onecell. This is strong evidence that extending fingerprints to include signal strength information from channels other than the 6-strongest cells, even when the identity of the transmitter cannot be determined, can dramatically increase localization accuracy.

Within-Floor Localization Error Table 5 summarizes the localization errors within specific floors for the 5 algorithms introduced in Section 3.4 for the three indoor environments. For each floor, the table shows the 50-percentile localization error, calculated as the Euclidean distance between the *actual* and *predicted* location of the point within the specific floor. All calculations assume a training set restricted to include only points that are on the same floor as the point whose position is being determined.

Table 5 also presents results for random, an algorithm that determines localization by picking an arbitrary position in the building. Therefore, random provides a lower bound on the performance of localization systems for a given floor and building. The localization error in random depends on the size of the floor, which accounts for difference in its localization error across floor and building.

	University (downtown)		Research Lab (midtown)		House (residential)		
	5th	7th	5th	6th	basement	1st	2nd
	802.11	4.22	4.78	2.20	2.59	3.49	3.43
chann _{cl}	5.44	3.98	2.48	4.77	3.28	2.95	3.96
chann	6.47	4.07	3.40	4.82	3.28	3.36	4.55
cell	11.06	8.02	4.82	6.99	3.41	3.40	5.27
onecell	15.05	14.64	8.39	7.93	3.42	4.85	6.13
random	33.87	30.43	10.40	13.35	4.68	6.21	7.07

Table 5. Single-floor median localization error (meters).

Across the three buildings, 802.11 achieves median accuracy between 2.2 and 4.8 meters. These results are consistent with results previously reported in the literature. Differences in accuracy between building reflect discrepancies in the granularity of the measurement grid which varied between 1 and 1.5 meters.

There are large differences in the performance of the various GSM-based algorithms. `chann` and `channcl` outperform `cell` and `onecell` in all cases. Moreover, `channcl` achieves between 25% to 50% better performance than `cell` for at least one floor in each of the three buildings. Across the three buildings, `channcl` achieves median accuracy between 2.5 and 5.4 meters, and in 3 out of the 7 floors, `channcl` even achieves better accuracy than 802.11 (e.g., 7th floor of University building).

The strong performance of `channcl` demonstrates the advantage of wide fingerprints including measurements from a large number of channels rather than just the 6-strongest cells. Moreover, the significant accuracy improvement of `channcl` over `chann` shows that geographical clustering manages to reduce the effect of false-positives introduced by channel aliasing. Geographical clustering, on the other hand, did not have a significant effect on the performance of 802.11 as channel aliasing does not occur in this case.

Figure 5 shows the cumulative distribution (CDF) of the localization error of all algorithms for the 7th floor of the University building. Most remarkable is the closeness with which `channcl` approximates 802.11, and the large difference in performance between `channcl` and `cell`.

Effects of Multi-floor Fingerprints In the previous section, we evaluated within-floor localization accuracy assuming that the training set was limited to fingerprints in the same floor, i.e., we predicted the floor first, and then predicted position within that floor. In contrast, in this section, we evaluate the effects on within-floor localization accuracy of including in the training set fingerprints taken on different floors. For this purpose, we project the training points collected on different floors of a building onto a single X, Y plane, therefore removing all floor information. We then ran the K-nearest neighbors on the extended training set. Table 6 shows the results of this experiment.

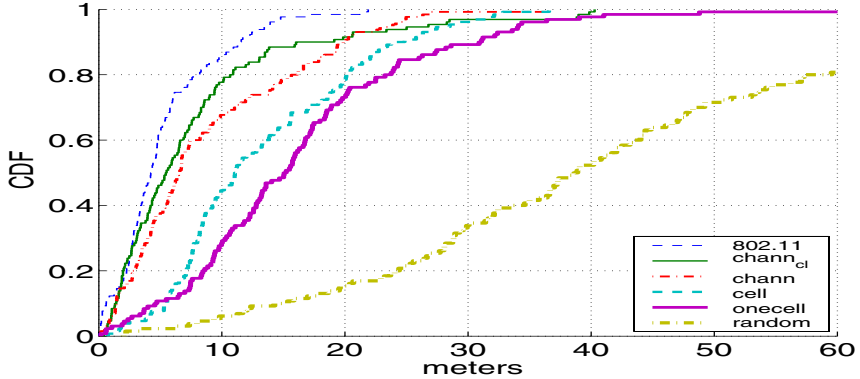


Fig. 5. CDF of the localization error for 7th floor of the University building.

	University (downtown)		Research Lab (midtown)		House (residential)	
	50%-ile	90%-ile	50%-ile	90%-ile	50%-ile	90%-ile
802.11	4.40	10.27	2.49	4.94	3.11	5.80
chann _{cl}	4.98	18.74	4.41	9.43	3.66	7.02
chann	5.76	21.75	4.72	9.44	4.10	7.18
cell	9.86	22.31	6.41	11.64	4.35	8.05
onecell	14.92	29.80	8.55	14.31	4.67	8.95
random	35.61	59.36	13.85	21.33	6.46	15.18

Table 6. Localization error with multi-floor fingerprints (meters).

Projecting the points collected on different floors onto a single plane has several effects. On the one hand, this practice may reduce the localization accuracy as the training points of other floors add “noise” (e.g., potential aliasing), which may result in larger localization errors. On the other hand, if the training points at a specific $\langle X, Y \rangle$ location on all floors have similar signal strength signatures, combining the training data from multiple floors will increase the density of the measurement’s grid, which may result in higher accuracy.

The multi-floor performance of 802.11 in the House is better than in any of the single-floor experiments. We found that the signal strength from the APs outside the building varies more with distance within a floor than within similar position on different floor. As a result, the training data from multiple floors overlaps, tightening the grid and increasing localization accuracy. The performance of 802.11 in a multi-floor setting in the University and Research Lab buildings is close to the average of the single-floor experiments, which is further indication that 802.11 can effectively differentiate between floors in office buildings with heavy concrete and steel frames.

The multi-floor localization error for GSM-based algorithms is also close to the average of the single-floor experiments. Therefore, for most cases, first identi-

ifying the floor and then performing localization using single-floor training data results in higher accuracy than performing the localization using multi-floor data. However, when the number of readings per location is low or differences in signal strength across floors are small, combining the training sets of multiple floors may produce higher localization accuracy.

4.3 Sensitivity Analysis

In this section, we analyze the best GSM performer, chann_{cl} , in more detail. Specifically, we test the localization accuracy of chann_{cl} as a function of the number of channels used, the number of measurements collected per location and the training grid size.

Number of Channels Figure 6 plots the median localization error for the multi-floor experiment as a function of the number of channels used. Increasing the number of channels results in a larger fingerprint, which allows a better comparison between neighboring points and therefore for increased localization accuracy. The channels picked are sorted by popularity (i.e., the number of fingerprints on which a specific channel appears). For example, the median localization error for 6 channels, corresponds to an algorithm where the 6 (fixed) most popular channels are picked from the training set. Notice that the accuracy of the algorithm that picks the 6 most popular channels is lower than of the cell algorithm. This is because the cell algorithm picks the 6 strongest cells for each measurement, which may result in much larger fingerprint vector (e.g., completely different 6 cells may be picked in two distant locations, increasing the fingerprint vector to 12 entries).

Figure 7 plots the percentage of incorrect floor classifications as a function of the number of channels. As expected, picking more channels decreases classification error. Interestingly, in all cases, picking about 20 channels is sufficient for achieving good localization accuracy.

Number of Measurements Per Location Although all the results reported so far were based on the average of 2 measurements per location, we actually obtained 10 measurements per location for the University building dataset. However, experiments varying the number of measurements per location between 2 and 10 showed virtually no difference in the accuracy of the algorithms. This is because our readings are stable and therefore adding more measurements per location does not improve localization accuracy.

Data Collection Grid Size Figure 8 and Figure 9 show the effects of reducing measurement grid size on the median multi-floor localization error and the floor classification error, respectively. We simulate the effect of increasing the measurement grid size by uniformly removing points from the training set. In most cases, reducing the measurement grid size results in lower localization accuracy, but occasionally we do see anomalies. As it turns out, decreasing the size of the

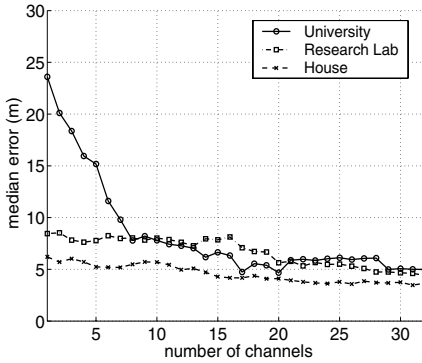


Fig. 6. Localization error as a function of fingerprint size.




Fig. 7. % of erroneous floor classifications as a function of fingerprint size.

grid may eliminate (in some cases) “problematic” or “aliased” points, which in turn increases localization accuracy.




Fig. 8. Localization error as a function of the measurement grid size.

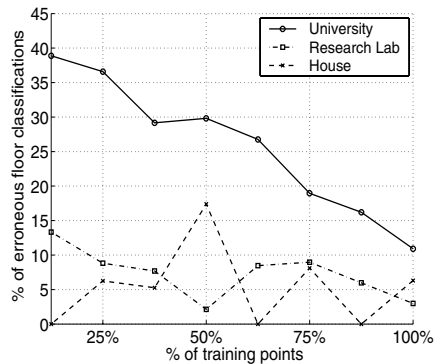


Fig. 9. % of erroneous floor classifications as a function of the measurement grid size.

4.4 Combined 802.11 and GSM Localization

In this section, we present an initial attempt to combine 802.11 and GSM fingerprinting. Since we collected both 802.11 and GSM channels information simultaneously, we have been able to combine the readings of both into one large fingerprint. The results are summarized in Table 7. The combined algorithm achieves moderately better accuracy in the University building, underperforms 802.11 in the Research Lab, and achieves similar performance in the House. An

explanation for the lackluster performance of the combined algorithm may be found in the way in which we combine the fingerprint data. By simply concatenating fingerprint vectors we implicitly give more weight to the more numerous and less accurate GSM readings. In the future, we plan to investigate better ways of combining the two fingerprints (e.g., give higher weight to 802.11 readings).

	University (downtown)		Research Lab (midtown)		House (residential)	
	50%-ile	90%-ile	50%-ile	90%-ile	50%-ile	90%-ile
802.11+chann _{cl}	4.03	8.65	3.35	6.39	3.24	4.29
802.11	4.40	10.27	2.49	4.94	3.11	5.80
chann _{cl}	4.98	18.74	4.41	9.43	3.66	7.02

Table 7. Multi-floor localization error (meters).

5 Conclusions and Future Work

We presented the first GSM-based indoor localization system that achieves median accuracy comparable to an 802.11-based implementation. We showed that accurate indoor GSM-based localization is possible thanks to the use of *wide* signal-strength fingerprints that include readings of up to 29 GSM channels in addition to the 6-strongest cells.

While the lack of cell ID information for some channels raises the possibility of *world wide aliasing*, we showed that filtering fingerprints based on the subsets of the cell IDs of the 6-strongest cells is sufficient for differentiating between locations in our three indoor environments.

We presented evaluation results of our system in three multi-floor buildings located in the Toronto and Seattle metropolitan areas, covering a wide range of urban densities. Our GSM-based indoor localization system achieves a median accuracy ranging from 2.48m to 5.44m in large multi-floor buildings. Moreover, our GSM-based system effectively differentiates between floors in both wooden and steel-reinforced concrete structures, achieving correct floor classifications between 89% and 97% of the time. In contrast, in the wooden building, the 802.11-based fingerprinting system achieved correct classifications only 62% of the time due to a limited fingerprint size.

In the future, we plan to test applicability of additional positioning algorithms to GSM-based fingerprinting. We also plan to extend our system to a working prototype that will allow for accurate localization both indoors and outdoors.

Acknowledgments

We would like to thank our anonymous UbiComp reviewers and especially our shepherd John Krumm for their detailed and sharp comments. We would also like to thank Neil Ernst for his contributions to earlier versions of this work.

References

1. L. Aalto, N. Gothlin, J. Korhonen, and T. Ojala. Bluetooth and WAP push based location-aware mobile advertising system. In *MobiSYS '04: Proceedings of the 2nd international conference on Mobile systems, applications, and services*, pages 49–58. ACM Press, 2004.
2. P. Bahl, A. Balachandran, and V. Padmanabhan. Enhancements to the RADAR User Location and Tracking System, Microsoft Research, Technical Report, Feb. 2000.
3. P. Bahl and V. N. Padmanabhan. RADAR: An in-building RF-based user location and tracking system. In *INFOCOM*, pages 775–784, 2000.
4. J. Eberspacher, H.-J. Vogel, and C. Bettstetter. GSM switching, services and protocols, 2001.
5. E. Elnahrawy, X. Li, and R. Martin. The limits of localization using signal strength: A comparative study. In *Proceedings of the 1st IEEE International Conference on Sensor and Ad Hoc Communications and Networks*, Santa Clara, CA, Oct. 2004.
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 the Tenth ACM International Conference on Mobile Computing and Networking (MOBICOM)*, Philadelphia, PA, Sept. 2004.
7. A. Hopper, A. Harter, and T. Blackie. The active badge system. In *Proc. of INTERCHI-93*, pages 533–534, Amsterdam, The Netherlands, 1993.
8. T. Kindberg and A. Fox. System software for ubiquitous computing. *IEEE Pervasive Computing*, 1(1):26–35, Jan. 2002.
9. J. Krumm, G. Cermak, and E. Horvitz. RightSPOT: A Novel Sense of Location for Smart Personal Objects. In *Proceedings of Ubicomp*, 2003.
10. K. Laasonen, M. Raento, and H. Toivonen. Adaptive on-device location recognition. In *Proceedings of the Second International Conference on Pervasive Computing*, pages 287–304. Springer-Verlag, 2004.
11. A. Ladd, K. Bekris, G. Marceau, A. Rudys, L. Kavraki, and D. Wallach. Robotics-based location sensing using wireless ethernet. In *Proceedings of the Tenth ACM International Conference on Mobile Computing and Networking (MOBICOM)*, 2002.
12. H. Laitinen, J. Lahteenmaki, and T. Nordstrom. Database correlation method for GSM location. In *Proceedings of the 53rd IEEE Vehicular Technology Conference*, Rhodes, Greece, May 2001.
13. A. LaMarca, Y. Chawathe, S. Consolvo, J. Hightower, I. Smith, J. Scott, T. Sohn, J. Howard, J. Hughes, F. Potter, J. Tabert, P. Powledge, G. Borriello, and B. Schilit. Place Lab: Device Positioning Using Radio Beacons in the Wild. *3rd Annual Conference on Pervasive Computing*, 2005. To appear.
14. L. Luxner. The Manhattan Project: AT&T Wireless invades the Big Apple with microcells. *Telephony*, 8(20), Feb. 1997.
15. N. B. Priyantha, A. Chakraborty, and H. Balakrishnan. The cricket location-support system. In *Mobile Computing and Networking*, pages 32–43, 2000.
16. A. Ward, A. Jones, and A. Hopper. A new location technique for the active office. In *IEEE Personnel Communications*, 4(5), pages 42–47, Oct. 1997.
17. M. Youssef, A. Agrawala, and U. Shankar. WLAN Location Determination via Clustering and Probability Distributions. In *IEEE PerCom 2003*, March 2003.
18. Ekahau, <http://www.ekahau.com>.
19. Ubisense, <http://www.ubisense.net>.
20. Versus Technologies, <http://www.versustech.com>.