

Are GSM phones THE solution for localization?

Alex Varshavsky¹, Mike Y. Chen⁵, Eyal de Lara¹, Jon Froehlich², Dirk Haehnel⁵, Jeffrey Hightower⁵, Anthony LaMarca⁵, Fred Potter², Timothy Sohn⁴, Karen Tang³, and Ian Smith⁵

¹Computer Science
University of Toronto
{walex,delara}@cs.toronto.edu

²Computer Science and Engineering
University of Washington
{jfroehli, fpotter}@cs.washington.edu

³Computer Science
Carnegie Mellon University
kptang@cs.cmu.edu

⁴Computer Science and Engineering
University of California, San Diego
tsohn@cs.ucsd.edu

⁵ Intel Research Seattle
{mike.y.chen, dirk.haehnel, jeffrey.r.hightower,
anthony.lamarca, ian.e.smith}@intel.com

Abstract

In this paper, we argue that localization solution based on cellular phone technology, specifically GSM phones, is a sufficient and attractive option in terms of coverage and accuracy for a wide range of indoor, outdoor, and place-based location-aware applications. We present preliminary results that indicate that GSM-based localization systems have the potential to detect the places that people visit in their everyday lives, and can achieve median localization accuracies of 5 and 75 meters for indoor and outdoor environments, respectively.

1 Introduction

Development of location-aware applications and systems, especially in the social-mobile [19] and health care [6] domains, has been driving the need for more accurate and pervasive localization technologies. We see a taxonomy emerging where location-enhanced applications can be categorized into three types according to their needs: *indoor localization*, *outdoor localization* and *place detection*. Digital homes [13], in-building navigation, and in-building coordination between peers [21] all desire indoor localization with at least room-level accuracy. City-wide tourist guides [7, 2], the US FCC’s E911 mandate, and location-based web search [5] need outdoor wide-area localization where coverage is often paramount to precision. Some applications benefit from the combination of both indoor and outdoor localization on a single easy to carry device. Those include Computer Supported Cooperative Care (CSCC) [6], social-mobile computing [19] and gaming [3]. Finally, new

location-enhanced applications are emerging that use information about places the user visits repeatedly instead of latitude-longitude coordinates [11]. Examples of these *place detection* applications include automatic configuration of wireless network settings based on place, recommendation systems that learn user’ preferences by tracking the places (e.g., restaurants, bars) the user visits, location-enhanced instant messengers, systems that allow setting reminders based on places that are important to the user [20]. In this paper, we argue that localization based on cellular phone technology, specifically GSM phones, can be a sufficient and appropriate solution both in terms of coverage and localization accuracy for this taxonomic spectrum of applications.

There is a conception that client-based GSM localization is inherently inaccurate. GSM’s large cell sizes (GSM macro-cells have a range of 35km, which can be extended if necessary) seem to make it harder to achieve good localization accuracies. In this paper, we argue that this conception is in fact incorrect. We show that GSM localization achieves accuracies that are appropriate for many applications. In indoor environments it is possible to perform room level localization with GSM and achieve median localization accuracies of 2 to 5 meters. We also report our preliminary results for wide-area outdoor GSM localization, achieving up to 75m median error. Finally, we show how having only GSM traces reported by a cell phone allows us to detect places people visit in their everyday lives.

2 Background

Many available localization technologies have low coverage or only work in a specific environment. The most

commonly available location technology today is the Global Positioning System (GPS). Although accurate and very effective in open environments, GPS typically does not work well where people spend their time. For example, GPS does not work well indoors, in urban canyons, or in similar areas with limited view of sky. A recent study showed that GPS coverage is available only 4.5% of the time for a device carried in users' pockets or purse during a typical day [15], although these numbers are admittedly worst-case and they rise if only both mobile and stationary times are considered. GPS-enabled devices are quite valuable and will become more and more widespread, but it is clear that many systems require another technology to meet the coverage and accuracy demands of applications. Infrared [12], ultrasound [18] and Bluetooth localization systems [1] work well indoors, but deploying these technologies to the wide-area is either cost prohibitive or not technically possible, for example, due to infrared interference from the sun.

The wide adoption of WiFi-enabled mobile devices and rapid deployment of WiFi access points make WiFi localization attractive. The RADAR project [4] pioneered indoor WiFi location, achieving 2-3 meter median accuracy and inspiring many follow-up efforts [14, 9]. Place Lab [15] introduced wide-area WiFi location, showing median accuracies ranging between 15 and 60 meters and high coverage. Examples of recent commercial systems using Place Lab's approach include Microsoft's Virtual Earth (<http://virtualearth.msn.com>) and SkyHook Wireless (<http://www.skyhookwireless.com>). Finally, the Beacon Print project [11] showed how the places people go to can be learned and recognized, without relying on a coordinate-based localization system. Unfortunately, because of their high power consumption, current WiFi-enabled devices are not frequently used "on-the-go" and, unless a power line is available nearby, are used intermittently. For example, Henderson *et al.* showed that although people do use their WiFi-enabled devices in several locations, they tend to power them off before moving to a new place and do not power them on unless necessary [10]. As a result, WiFi-equipped devices cannot be used effectively as a platform for location-enhanced applications that rely on continuous network connectivity or spontaneous interactions, for example, social-mobile and monitoring applications.

Fortunately, there are devices that people do carry with them most of the time that have continuous network connectivity: mobile phones. Mobile phones have low power consumption, ubiquitous connectivity, established interface metaphors, wide adoption, and, most important for this paper, research results suggests that they can offer indoor, outdoor, and place detection capabilities. We believe all these characteristics make mobile phones an excellent platform for developing and deploying location-enhanced applications. Phone localization has specific advantages over

WiFi localization: (a) phones operate over a licensed band, meaning no interference from microwave ovens and cordless phones; (b) phones use a managed network, meaning no interference from neighboring access points that happened to be tuned to the same channel; (c) phone networks require significant installation investments, resulting in stable environment that changes less frequently [15]; and (d) phone network coverage is greater than that of WiFi networks.

3 GSM Primer

Global System for Mobile Communication (GSM) is the most widespread cellular telephony standard in the world, with deployments in more than 210 countries by over 676 network operators [8]. 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.

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 and 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.

We collected GSM traces using a Sony Ericsson GM28 GSM modem and an Audiovox SMT 5600 phone, depicted in Figures 1 and 2. The modem operates as an ordinary GSM cell phone, but exports a richer programming interface. Both the modem and the phone provide two inter-



Figure 1. GM28 Sony Ericsson Modem



Figure 2. Audiovox SMT 5600 Smart Phone

faces for accessing signal strength information: *cellsAPI* and *channelsAPI*¹. The *cellsAPI* interface reports the cell ID, signal strength, and associated channel for the n neighboring cells. While the modem’s and the phone’s specification does not set a hard bound on the value of n , in practice we saw the maximum value of n varying from 6-7. The *channelsAPI* interface simultaneously provides the signal strength for up to 35 channels on the modem and 18 channels on the phone. In practice, 6 of the channels typically correspond to the 6-strongest cells. Unfortunately, *channelAPI* reports signal strength but does not report cell IDs.

Although the modem exposes the *cellsAPI* and *channelAPI* explicitly, we are not aware of any GSM phone that makes this information easily available. For example, to access cell and channel information on the Audiovox SMT 5600 phone, we had to write a C tool that reads this data directly from the phone’s memory.

We speculate that the fact that current phones do not expose similar interfaces reflects the unwillingness of network operators to make signal strength information public. Indeed, by not exposing these interfaces, network operators can maintain a monopoly on the provisioning of location-based services. What needs to be done to influence network operators to allow exposing such interfaces is beyond the scope of this paper. Instead, in this paper we explore the opportunities that are made possible by the availability of signal strength information on the phone.

4 Indoor and Outdoor Localization

We begin this section by describing how GSM and WiFi localization works. We then present our experimental GSM localization results and show that although GSM accuracies

¹The terms *cellsAPI* and *channelsAPI* are used to simplify presentation. In practice, the *cellsAPI* correspond to `AT+E2EMM=1` command and the *channelsAPI* correspond to the `AT+E2NBTS?` command on the GM28 GSM modem, respectively

appear slightly lower than those achieved using WiFi localization, they are comparable and sufficient for the same types of location-enhanced applications.

Two approaches to GSM and WiFi localization are *fingerprinting* and *centroid*, both of which require a training phase where given a set of GPS-stamped WiFi or GSM traces the algorithm builds a model of an environment, which it later uses for predicting device’s location. Given the training traces, the *centroid* algorithm learns the positions of radio beacons in the environment (i.e., WiFi APs or GSM cell towers) by positioning the radio beacon in a location where the signal strength for that beacon was observed the strongest. During the testing phase, the *centroid* algorithm predicts a position of a measurement by averaging the positions of the radio beacons that appear in the measurement. Typically, giving a higher weight during averaging to radio beacons with stronger signal strength yields better localization accuracy. Unfortunately, walls, doors and other obstacles attenuate radio signals in an unpredictable way, making the *centroid* algorithm inaccurate in indoor environments. Therefore, we present *centroid* results only for outdoor experiments.

In contrast, the *fingerprinting* algorithm uses the training set to build a mapping from measurements to positions where those measurements were observed. Then, during the testing phase, *fingerprinting* matches every measurement in the testing set to one or more measurements observed during the training phase and then averages the true positions of the best matched measurements. Once again, weighting by the signal strength of the best matched measurements yields better results.

We also show results for a *random* algorithm, which predicts position by randomly picking a measurement from the training set and assigning its position as the predicted position, thus providing a lower bound on the performance of a localization system. The localization error, or the distance between the true and predicted position, of random

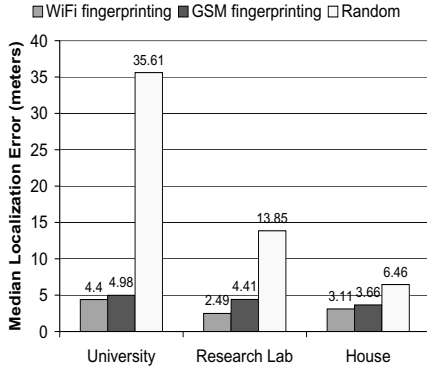


Figure 3. Median indoor localization error.

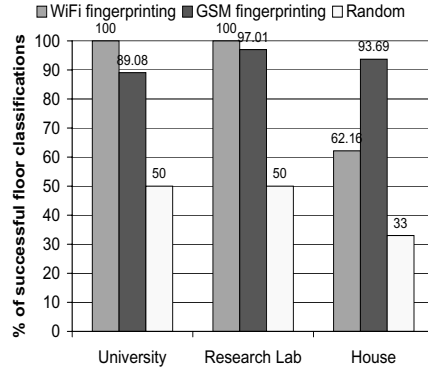


Figure 4. Floor classification accuracy.

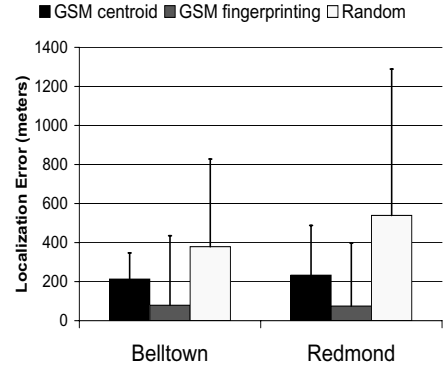


Figure 5. Outdoor localization error (median and 95%).

depends on the size of the area covered by the training set.

4.1 Indoor Localization

In our previous work [17], we presented the first accurate GSM indoor localization system that achieves median accuracy ranging from 2.48m to 5.44m in large multi-floor buildings. We will first briefly summarize this work and then present our new preliminary results for room level localization.

4.1.1 Coordinate level accuracy

To test our system, we collected GSM measurements on the 5th and 6th floors of the Intel Research Seattle Lab building, the 5th and 7th floors of Bahen Center (the home to the Department of Computer Science of the University of Toronto) and at the basement, 1st and 2nd floors of a private house located in a suburban Seattle area. The measurements were collected about 1.5 meters apart. We tested the accuracy with which GSM and WiFi localization systems based on fingerprinting [4] are able to differentiate floors of the buildings and to localize mobile devices within the floor.

Figure 3 and Figure 4 summarize our results. In our experiments, WiFi achieved within-floor localization accuracies consistent with previous findings [4, 9]. Also, because of reinforced concrete floors in the Research Lab and the University buildings, WiFi was able to differentiate floors perfectly. In the house environment, however, WiFi achieves low classification accuracy as the house’s wood structure presents little obstacle to radio propagation, making it harder to differentiate floors.

GSM localization system performs well, achieving within-floor localization results comparable to 802.11 system. Moreover, our GSM system effectively differentiates

between floors in both wooden and steel-reinforced concrete structures, achieving correct floor classifications between 89% and 97% of the time.

4.1.2 Room level accuracy

To test room level localization accuracy, we first collected a training trace by walking around with an Audiovox 5600 SMT cell phone in 8 rooms on the 6th floor of Intel Research Seattle building. One hour later, we gathered an additional similar trace using the same phone. We used fingerprinting to match training and testing measurements.

First, we have repeatedly broken the training set down into two random sets with 90% of points in the training set and 10% in the testing set. In all cases, we obtained 100% classification accuracy, which suggests that two consecutive measurements taken on the phone are typically very similar. We then tried to match the measurements taken within an hour and saw 70% accuracy (87 of 126). More detailed examination revealed that in most incorrect classifications the predicted room was the next closest room.

Although we plan to perform more elaborate testing in the future, these results suggest that room level localization using GSM traces is feasible and GSM phones could support applications like in-building navigation and in-building coordination between peers.

4.2 Outdoor Localization

This section presents our preliminary GSM outdoor localization results. We collected traces from a vehicle driving in two neighborhoods in the Seattle metropolitan area: (a) Belltown, a mix of commercial and residential urban high-rises and (b) Redmond, a medium density residential neighborhood. We collected GPS-stamped GSM traces us-

		Predicted Movement	
		Place	Mobile
Precision	Ground Truth	98.2%	37.5%
		1.8%	62.5%

		Predicted Movement	
		Place	Mobile
Recall	Ground Truth	96.2%	3.8%
		21.3%	78.7%

Figure 6. Place detection accuracy using GSM mobile phones.

ing a laptop connected to a GPS device and the Sony Ericsson GSM modem.

The median and the 95% accuracy results for the centroid, fingerprinting and random algorithms are summarized in Figure 5. As expected, fingerprinting achieves the best accuracy, with median error below 75m in both areas we tested. The centroid algorithm performs worse, achieving a 213m median error. These preliminary results are encouraging, as 75m or even 213m median error is comparable to what is possible with WiFi and likely more than sufficient for many wide-area location-enhanced applications such as social coordination and local web search.

5 Place Detection

In this section, we describe how using a stream of GSM readings, we were able to effectively detect places people visit in their everyday lives. We developed an algorithm that given a stream of time-stamped GSM readings, outputs the times when a person was at a “place”. Here, we consider a place to be a time interval during which our algorithm predicted that the user was stationary for more than 3 minutes. To gather data, we developed an application that runs on an Audiovox SMT 5600 phone, continuously scans nearby GSM cell towers once per second and allows users to use text entry to name and select the places they went. Having the GSM traces labeled with the ground truth data enabled us to test the accuracy of our place detection algorithm.

Our place detection algorithm applies a simple principle: when someone is at a place the stream of GSM readings their phone captures is “stable.” Currently, we measure stability by tracking the Euclidean distance in signal strength space [4] between consecutive GSM readings. The smaller the signal distance between two readings, the more similar these readings are. When the phone is stationary, the distance between consecutive GSM measurements tends to be small, whereas when the cell phone is being carried around, the Euclidean distance between consecutive GSM measurements will oscillate widely. We deployed the system to 5 users in our lab who carried the phones and labeled every

places they went to for a month. Figure 6 shows the effectiveness of our GSM place prediction algorithm in recognizing when the phone is stable or mobile between places. Precision is how often the algorithm’s prediction matches the true state while recall is how many of the true states were correctly identified by the algorithm. With this high accuracy, we argue it should be possible to extend this capability to a full place learning and recognition system for GSM phones, analogous to what BeaconPrint [11] did with WiFi. This capability would allow GSM phones to support applications like visit-driven recommendation systems, place-based device configuration, and context-aware notes and reminders.

6 Conclusions and Future Work

In this paper, we argued that for emerging location-enhanced applications, client-based GSM localization can provide an adequate solution both in terms of coverage and accuracy in a device people already carry. To dispel the notion that location systems using GSM phones are inherently less accurate than systems built for WiFi devices, we presented preliminary results showing that using GSM it is feasible to achieve 2-5 meters median error and room-level localization indoors, 70-200 meters median error outdoors, and to detect places people go in their everyday lives. These results are comparable to what has been demonstrated previously for WiFi.

In this paper, we discussed localization solutions based on GSM cellular network. However, we believe that on-phone localization based on cellular networks is not specific to GSM. Indeed, any cellular technology that transmits stable beacons from the cellular towers (e.g., for the need of hand-off purposes) will make the on-phone localization possible.

The main drawback of many existing localization systems, whether WiFi or GSM, is the non-trivial training required for the system to become usable. For instance, fingerprinting-based solutions require a tedious training data collection, and centroid-based solutions require as-

sembling maps of locations of cellular towers or access points in the target environment. We already made initial steps toward reducing the training needed for wide-area centroid-based location using a technique called beacon self-mapping [16]. Our future plans include investigating novel ways of reducing training requirements for both indoor and outdoor localization.

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] G. D. Abowd, C. G. Atkeson, J. Hong, S. Long, R. Kooper, and M. Pinkerton. Cyberguide: a mobile context-aware tour guide. *Wirel. Netw.*, 3(5):421–433, 1997.
- [3] R. Anastasi, N. Tandavanitj, M. Flintham, A. Crabtree, M. Adams, J. Row-Farr, J. Iddon, S. Benford, T. Hemmings, S. Izadi, and I. Taylor. Can you see me now? a citywide mixed-reality gaming experience. In *Proceedings of the UbiComp*, 2002.
- [4] P. Bahl and V. N. Padmanabhan. RADAR: An in-building RF-based user location and tracking system. In *INFOCOM*, pages 775–784, 2000.
- [5] O. Buyukkocuten, J. Cho, H. Garcia-Molina, L. Gravano, and N. Shivakumar. Exploiting geographical location information of web pages. In *WebDB (Informal Proceedings)*, pages 91–96, 1999.
- [6] S. Consolvo, P. Roessler, and B. Shelton. The carenet display: Lessons learned from an in home evaluation of an ambient display. In *Proceedings of Ubicomp*, 2004.
- [7] N. Davies, K. Cheverst, K. Mitchell, and A. Efrat. Using and determining location in a context-sensitive tour guide. *IEEE Computer*, 33(8), August 2001.
- [8] J. Eberspacher, H.-J. Vogel, and C. Bettstetter. GSM switching, services and protocols, 2001.
- [9] 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.
- [10] T. Henderson, D. Kotz, and I. Abyzov. The changing usage of a mature campus-wide wireless network. In *Proc. of MobiCom*, 2004.
- [11] J. Hightower, S. Consolvo, A. LaMarca, I. Smith, and J. Hughes. Learning and recognizing the places we go. In *Proceedings of the Seventh International Conference on Ubiquitous Computing (UbiComp 2005)*, Lecture Notes in Computer Science, pages 159–176. Springer-Verlag, September 2005.
- [12] A. Hopper, A. Harter, and T. Blackie. The active badge system. In *Proc. of INTERCHI-93*, pages 533–534, Amsterdam, The Netherlands, 1993.
- [13] J. Krumm, S. Harris, B. Meyers, B. Brumitt, M. Hale, and S. Shafer. Multi-camera multi-person tracking for easy living. In *VS '00: Proceedings of the Third IEEE International Workshop on Visual Surveillance (VS'2000)*, Washington, DC, USA, 2000. IEEE Computer Society.
- [14] 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.
- [15] 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. In *Proceedings of the Third International Conference on Pervasive Computing*, Lecture Notes in Computer Science. Springer-Verlag, May 2005.
- [16] A. LaMarca, J. Hightower, I. Smith, and S. Consolvo. Self-mapping in 802.11 location systems. In *Proceedings of the Seventh International Conference on Ubiquitous Computing (UbiComp 2005)*, Lecture Notes in Computer Science, pages 87–104. Springer-Verlag, September 2005.
- [17] V. Otsason, A. Varshavsky, A. LaMarca, and E. de Lara. Accurate gsm indoor localization. In *the Seventh International Conference on Ubiquitous Computing (UbiComp 2005)*, September 2005.
- [18] N. B. Priyantha, A. Chakraborty, and H. Balakrishnan. The cricket location-support system. In *Mobile Computing and Networking*, pages 32–43, 2000.
- [19] I. E. Smith. Social mobile applications. *IEEE Computer*, 38(4):84–85, April 2005.
- [20] T. Sohn, K. Li, G. Lee, I. Smith, J. Scott, and W. G. Griswold. Place-its: Location-based reminders on mobile phones. In *Proceedings of Ubicomp*, Tokyo, Japan, September 2005.
- [21] J. C. Tang, N. Yankelovich, J. Begole, M. V. Kleek, F. Li, and J. Bhalodia. Connexus to awarenex: extending awareness to mobile users. In *CHI '01: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 221–228, New York, NY, USA, 2001. ACM Press.