This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE INFOCOM 2009 proceedings.

EnLoc: Energy-Efficient Localization for Mobile Phones

Ionut Constandache (Duke), Shravan Gaonkar (UIUC), Matt Sayler (Duke),
Romit Roy Choudhury (Duke), Landon Cox (Duke)

***Abstract*—A growing number of mobile phone applications utilize physical location to express the context of information. Most of these location-based applications assume GPS capabilities. Unfortunately, GPS incurs an unacceptable energy cost that can reduce the phone's battery life to less than nine hours. Alternate localization technologies, based on WiFi or GSM, improve battery life at the expense of localization accuracy. This paper quantifies this important tradeoff that underlies a range of emerging services. Driven by measurements from Nokia N95 phones, we develop an energy-efficient localization framework called *EnLoc*. The framework characterizes the optimal localization accuracy for a given energy budget, and develops prediction-based heuristics for real-time use. Evaluation on traces from real users demonstrates the possibility of achieving good localization accuracy for a realistic energy budget.**

I. Introduction

Mobile phones are powerful platforms for sensing, sharing, and querying people-centric information. A variety of applications are on the rise, many of which utilize *location* to express the context of information. Most of these location based applications assume GPS capabilities. While GPS offers good location accuracy of around 10m, it incurs a serious energy cost that can drain a fully charged phone battery in 8.5 hours [1]. We make a few observations in light of this energy-accuracy profile.

(1) In real life, the phone battery must be shared with voice calls, SMS, emails and pictures. The energy budget for localization alone is a small fraction of the battery capacity. If this fraction is assumed to be 25%, *continuous* GPS localization is available for less than 2.5 hours.

(2) One may argue that continuous GPS is unnecessary, and can be activated only on demand. While this may suffice for some services (e.g. geo-tagging a photo), many emerging applications rely on the feasibility of continuous localization over reasonably long time scales. Examples include GeoLife [2], Micro-Blog [1], TrafficSense [3] and Pothole Patrol [4].

(3) Continuous localization over long time scales results in higher *average* error. For an energy budget of K GPS readings and a duration of $T > K$ time units, at $(T-K)$ time units location can only be estimated, and hence, it is more erroneous than an actual GPS reading. When averaged over all, actual and estimated readings, the *average localization error* is higher than the GPS instantaneous error ($\sim$10m).

(4) WiFi and GSM-based localizations are not obvious replacements to GPS. While these schemes are less energy-hungry, they incur higher (instantaneous) localization error (around 40m and 400m respectively). This permits a greater number of location readings, each of which is less accurate. Whether this results in lower *average error* than few, but accurate, GPS readings, is an open question.

This paper investigates the space of energy-efficient localization for mobile phones and expands on the following contributions:

- *Identify the space of energy-accuracy tradeoff.* Measurements on Nokia N95 phones quantify this tradeoff.
- *Analysis of the optimal localization accuracy for a given energy budget.* For a given mobility trace, an offline dynamic program (DP) computes the maximum location accuracy achievable using GPS, WiFi, GSM and combinations thereof. When results show that the theoretical optimal may not suffice for high-accuracy applications, we explore the usefulness of prediction.
- *Exploit habitual activity of individuals and behavior of populations to predict location.* Predictions incorporated into the DP offer offline optimal solutions. Online heuristics permit energy-efficient localization in real time.
- *Evaluate heuristics in real life situations.* Performance is compared with the theoretical optimum using a custom trace-based simulator and mobility traces collected in the UIUC campus. Results confirm the feasibility of achieving good localization over a day's energy budget.

II. Energy Measurements

We used a software monitor to measure fine-grained power consumption in Nokia N95 phones. This section reports the accuracy and energy measurements that motivate EnLoc.

A. Localization Using WiFi and GSM

As an alternative to GPS hardware and its unavailability indoors, project Place Lab [5] proposed using WiFi and GSM sensors for localization. Specifically, authors create a wireless map of a region by war-driving in the area. The wireless map is composed of sampled GPS locations, WiFi access points and GSM towers audible at these locations. This wireless map is then distributed to phones. When a phone travels through the mapped area, it estimates its own location by matching its list of audible WiFi APs/GSM towers to the wireless map. Place Lab experiments in downtown Seattle exhibit a median positioning error of 13 to 40m with WiFi, and around 94 to 196m with GSM. When performed in Champaign, IL, and Durham, NC, WiFi accuracy ranged between 25 to 40m, while GSM ranged between 300 to 400m.

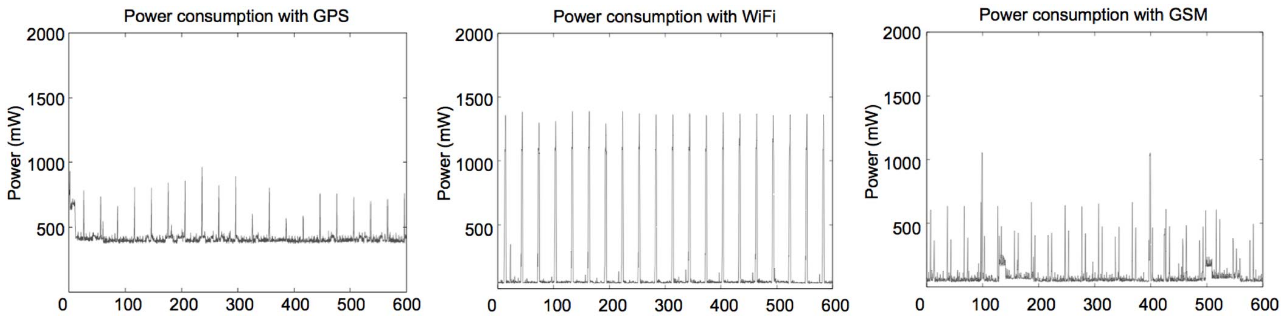

Fig. 1. Power consumption in mW for Nokia N95 phones, sampled at 30s intervals: (a) GPS measurement (b) WiFi measurement (c) GSM measurement.

B. Energy Measurements for GPS, WiFi, and GSM

We measured the energy consumption on Nokia N95 phones for each localization sensor. We charged the phone battery to full capacity and turned on only the location sensor we intended to measure. An energy monitoring program probed the location sensor at a chosen interval T_{probe}.

Figure 1(a) shows GPS power consumptions as a function of time, for T_{probe} = 30 seconds. We see periodic spikes on top of a baseline energy consumption at approximately 400 mW. The spikes correspond to a GPS sensor read operation and a write operation into the phone's file system for logging the location data. The baseline corresponds to the power consumed by the GPS receiver.

Similar measurements are reported in Figure 1(b) and (c) for WiFi and GSM. Observe that while the baseline power consumption for WiFi is low (55 mW), it exhibits a high spike (of around 1100) for every probe. GSM based samples exhibit similar characteristics, however, their power consumption is less. When power consumption is translated to net battery life, we found that GPS allowed for 9 hours, while WiFi and GSM sustained for 40 and 60 hours, respectively. Viewed against corresponding localization accuracies of 10m, 40m, and 400m, the energy-accuracy tradeoff is evident.

III. EnLoc: Framework Design

A. Average Location Error

The energy-efficient localization problem can be defined as follows. Given an energy budget B and time duration T, design a strategy that will minimize the *average localization error* (ALE). Formally, denote $L_{reported}(t_i)$ and $L_{actual}(t_i)$ to be the reported and actual locations of the phone at time t_i. Assuming T discrete time-points, the ALE is:

$$\delta_{avg} = \sum_{i=1}^{T} \Big(\frac{L_{reported}(t_i) - L_{actual}(t_i)}{T} \Big)$$

Assuming GPS to be the ground truth, a GPS reading at time t_j implies that $L_{reported}(t_j)$ is same as $L_{actual}(t_j)$. Similarly, a WiFi reading at t_j implies that $L_{reported}(t_j) - L_{actual}(t_j)$ is in the order of 40m. The problem then is to minimize δ_{avg} for a given energy budget. We model this energy-accuracy tradeoff as an optimization problem.

B. Problem Formulation

Our goal is to determine a schedule with which the location sensors should be triggered such that the *average localization error* (ALE) is minimized for a given energy budget. The schedule is a set of time instants, $\{t_1, t_2, t_3, ..\}$ and the corresponding sensors $\{s_1, s_2, s_3, ..\}$, where $s_i \in [GPS, WiFi, GSM]$. The optimal schedule should trigger a reading of sensor s_i at time t_i to minimize ALE.

We developed a dynamic programming (DP) solution to the problem above. The DP takes as input an entire user trace (i.e., GPS, WiFi and GSM readings at all time-points along the user's path) and outputs a sensor reading schedule that achieves the minimum ALE. In the interest of space we omit the details of the dynamic program and only present our main findings.

C. Optimal Localization Error

To obtain the best localization accuracy, we executed the DP on mobility traces collected on the UIUC campus. We war-drove the campus [5] and generated a wireless map of the area. Then, we distributed phones to students to gather mobility traces. A custom simulator integrated the traces with the wireless map, and executed the dynamic program. We specified an energy budget of 25% of the phone battery and the duration of operation was 24 hours. For intermediary time points, at which sensor readings are not performed, we report "the last known location". We assumed that sensors can be sampled every 30 seconds.

	OptGPS	OptWiFi	OptGSM	OptComb
Trace 1	164.999	78.52	352.909	78.52
Trace 2	105.35	75.16	327.116	58.66
Trace 3	125.848	62.134	370.621	62.134

TABLE I
Optimal performance for different traces

We evaluate four optimal schemes, namely, *Optimal GPS/WiFi/GSM/Combined*. As the name suggests, *Optimal GPS* corresponds to the minimum ALE achieved when only GPS readings are used. Table I reports results from three mobility traces. Observe that *Opt WiFi* outperforms *Opt GPS* indicating that greater number of less accurate readings is better for localization. Also, *Opt Combined* outperforms the others, and is close to *Opt WiFi* in many of the traces. However, it is surprising that the offline optimal error (with knowledge of the entire trace) was typically more than 60m. Online versions of these schemes (that do not have the entire trace) will naturally perform worse. Reporting the last known location between consecutive location readings is a source of this inefficiency; we address this through mobility prediction.

D. Prediction Opportunity

In reality, human behavior/mobility is amenable to prediction [6], [7], [8]. Driving on straight highways, turns on one way streets, habitual office hours, are examples of prediction opportunities. EnLoc attempts to exploit them.

Simple Linear Predictor: We begin by considering a basic linear predictor. The location of a phone at time t_k, denoted $L(t_k)$, can be a linear extrapolation of the two previously sampled locations, $L(t_i)$ and $L(t_j)$. This can be effective when a phone moves on a straight road. However, if the phone's movement is not straight, or if $L(t_i)$ and $L(t_j)$ were highly erroneous, linear prediction may be unsuitable. We have modified the DP to incorporate linear prediction (LP) and compute the minimum average error values.

Human Mobility Patterns: While linear prediction is a general approach, recognizing individual human behavior may facilitate better predictions. The intuition is that humans have habitual activity patterns, and sampling the activity at a few uncertainty points may be sufficient for predicting the rest. For instance, given that a person goes to lunch at either 12:00pm, 12:50pm, or 1:00pm, the phone may trigger GPS readings just after these times. Learning that the person has started out for lunch, her subsequent locations can be predicted (i.e., locations along the habitual path from office to the cafe). Similarly, GPS readings between 12:30am to 7:00am can be obviated if the person habitually sleeps in this time window.

Deviations: To cope with deviations from habitual paths, we hypothesize that statistical behavior of large populations provide useful hints. Knowing that most of the vehicles take a left turn at a traffic intersection can be valuable for prediction. In general, if a "probability map" can be generated for a given area, an individual's mobility in that area can be predicted. We extended our DP to incorporate probability maps and extract the optimal localization schedule for a given trace. Intuitively, the DP is expected to schedule location readings at points where the individual's behavior differs from the population's statistical behavior.

IV. EnLoc: System Design

This section attempts to translate the above ideas into a working system, called EnLoc. The system exploits both habitual mobility patterns and population-driven probability maps. EnLoc is an online solution, and unlike the DP, does not assume knowledge of the user's entire trace.

Exploiting Habitual Mobility

A study with $100,000$ people has shown that individuals exhibit habitual space-time movements, with reasonably small variation [6]. To visualize this, we collected GPS-based mobility traces of several people and plotted them over Google Maps. Figure 2(a) shows a simplified example.

One may envision the Google maps plot as a tree, with branches at certain points – we call this the *logical mobility tree* (LMT). The vertices of this tree are the branching points on the person's actual mobility paths. Uncertainty arises at these branching points (e.g., at a traffic intersection where a person may go straight or turn right), and hence, the vertices of the LMT are also called *uncertainty points.* The edges of the LMT represent physical paths that connect consecutive uncertainty points. Each edge is associated with (1) the starting time of that physical movement, (2) the average velocity on that path, and (3) the duration of travel on that physical path. Figure 2(b) shows the LMT corresponding to the physical mobility in Figure 2(a).

Our key idea is to schedule location readings right *after* the uncertainty points on the LMT. Such a location reading will resolve the uncertainty since the phone will be placed in one of the paths emanating from that uncertainty point. Thus, the phone's location can be reasonably predicted until it encounters the next uncertainty point, at which time, another location reading will be necessary.

Observe that the LMT in Figure 2(b) is a spatial representation of a person's mobility profile. In reality, a person traverses the same edge on the LMT at many different times. Each of these possibilities translates into a distinct edge in the LMT. Figure 2(c) shows a hand-constructed example of such a space-time LMT representation. To accurately know when the phone leaves a particular node of the LMT, EnLoc will need to sample all the edges emanating from that node. However, energy-budget limitations will allow only a fraction of the emanating edges to be sampled. EnLoc designs a heuristic to sample a subset of the edges branching out of a node. We explain this with the example of Figure 2(c) which is not derived from an actual mobility trace.

Assume that current time is 8:00am, and the phone is located at home, H. Also assume that the remaining energy budget is $B_{remaining}$. The heuristic begins by identifying all the paths from H to the leaves of the tree, i.e., $P_1 = home \Rightarrow coffee$, and $P_2 = home \Rightarrow walmart$. Then, the number of location readings N_i, necessary to track the phone with certainty, is computed for each path P_i. Thus $N_1 = 4 + 6 = 10$ readings. Observe that 4 readings are necessary to track the phone leaving Home, and 6 readings for going from Office to Coffee. These 6 readings include the 5 edges from Office to Coffee (the latest being 6:10pm), as well as the 6:00pm edge from Office to Gym. If the 6:00pm edge is not included, EnLoc may not know if the phone has started moving towards the Gym. Similarly, $N_2 = 4 + 8 + 3 = 15$ readings. Now, the heuristic computes $M = max(N_i)$, a pessimistic estimate of the number of readings necessary in the future; $M = 15$ for this example. Then, the heuristic computes $F = \frac{e_H}{M}$, where e_H is the number of emanating edges from Home. In this example, $F = \frac{4}{15}$. The phone is allocated $F \times B_{remaining}$ amount of energy for detecting its departure from home. Assuming $B_{remaining}$ is 10, there

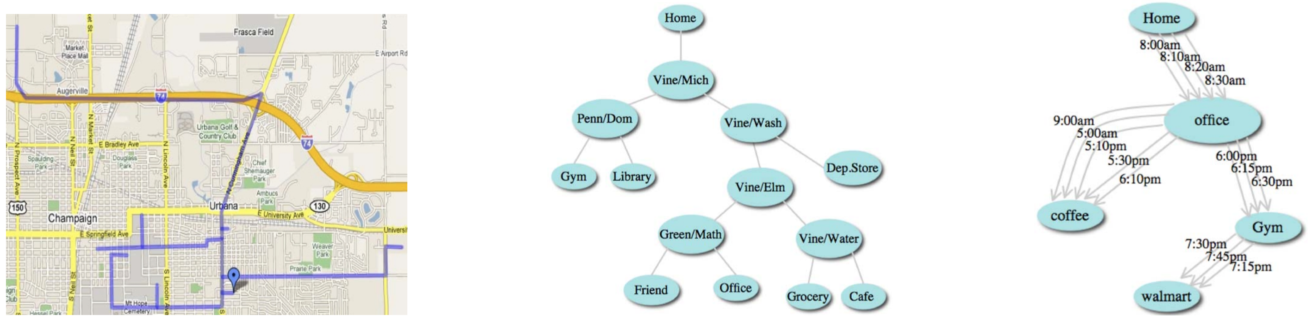

Fig. 2. Personal mobility profile: (a) An anonymous user's movement over two weeks (b) A spatial logical mobility tree (LMT) (c) A spatio-temporal LMT.

are approximately 2 GPS readings available. The heuristic randomly chooses 2 time-points out of the 4, and samples the phone's location. Once that phone is found to be on one of the paths going out of Home, the heuristic predicts the phone's location based on the habitual velocity on that edge. At the next uncertainty point, the phone recomputes F using the scheme above.

Addressing Deviation from Habits

Users may deviate from their habitual paths. Even though deviations are not the common case, they are important because several applications may be triggered due to deviation. Micro-blogging [1] may be more active when people go for vacations; location-specific information may be necessary when people are driving down unfamiliar paths. To addresses the case of deviation, EnLoc exploits mobility of large populations as a potential indicator of the individual's mobility. The basic idea is as follows: consider a person approaching a traffic intersection from Street A. Since the person has not visited this street in the past, it is difficult to predict how she will behave at the imminent intersection. Now, if a large fraction of the population is known to take a left turn onto Street B, then the person's movement can be guessed accordingly. EnLoc develops mobility maps of large populations and exploits them for prediction.

EnLoc detects a deviation when a scheduled location reading discovers the phone in an unexpected location (i.e., not on the LMT). At this time, EnLoc switches to the *Deviation Mode* of operation. In this mode, the residual energy budget is divided into equally time-spaced WiFi readings across the remainder of the day. Now, once the first location sample has been obtained, EnLoc uses the population activity map to predict the phone's movement. The velocity and turns at different intersections are estimated from the activity map. Incorrect predictions obviously incur location error. The error accumulates until the next reading, when EnLoc makes a new prediction using the new location as the starting point.

Without loss of generalization, let us consider 4-way traffic intersections. EnLoc computes 4 probabilities for each intersection, i.e., an user entering the intersection from Street A, either turns left, turns right, continues straight, or takes a U-turn. One may envision this as a 4×4 matrix, where element ij denotes the probability that the user entering street i exits through street j.

We generated the probability maps for UIUC campus using Google maps. First, we identified all roads that border the campus. Further, we identified roads that intersect the bordering roads, and enter the campus. We call these feeder roads. We also identified all parking lots within the campus and their capacities. Now, we simulated vehicles that enter the campus through a feeder road, and drive to a pre-specified parking lot. The pre-specified parking lot is randomly chosen from the distribution of parking lot sizes. For each vehicle, we obtained its driving direction through Google Map APIs, and parsed it to extract the vehicle's movement through each traffic intersection (i.e., left/right/straight/U-turn). Simulation of thousands of vehicles produces the probability matrix for each intersection. A phone that is installed with this matrix should be able to predict/localize itself in the UIUC campus.

V. Performance Evaluation

We evaluate EnLoc using traces collected on UIUC campus. The energy budget was set to 25% of the battery capacity and the duration of operation was 24 hours. An ideal evaluation of EnLoc should characterize the average localization error (ALE) over a person's complete mobility pattern (i.e., habitual and deviant paths). However, we found that the actual deviations extended far beyond the war-driven/probability-mapped UIUC campus. Therefore, we evaluated population-based prediction using portions of traces that were within the campus. Then, we evaluated the *habitual mobility profile*-based prediction by pruning the deviations from a user's LMT. We believe that the localization error will be close to the average of these two cases.

Deviant Paths

At each traffic intersection, the mobility of the phone was predicted based on the maximum probability at the intersection. The error was computed whenever the prediction was inconsistent with the actual user's movement. Section III-C and Table II describe the schemes examined in this section.

Opt-LP-GPS/WiFi/GSM	Optimal GPS/WiFi/GSM + Lin. Pred (LP)
Opt-LP-Comb	Optimal+Combined GPS,WiFi,GSM + LP
OptMap	Optimal using map predictor
Heu-Eq-GPS/WiFi/GSM	Heuristic+Equally-spaced GPS/WiFi/GSM+LP
EnLoc-Eq-Map	Heuristic+Equally-spaced GPS on Map

TABLE II
Optimal and heuristic schemes

Figure 3(a) reports the optimal localization error averaged over all mobility traces. Figure 3(b) present the performance

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE INFOCOM 2009 proceedings.

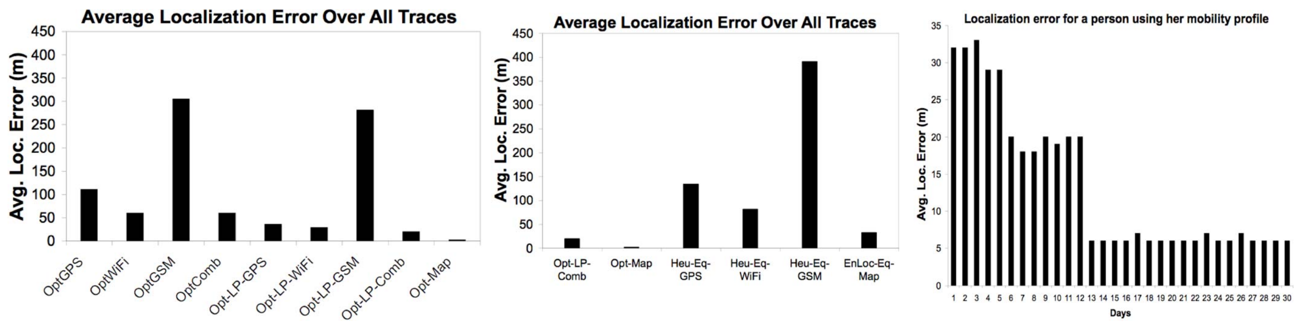

Fig. 3. ALE (a) Deviant Path Optimal Schemes, (b) Deviant Path Heuristics, (c) Individual Mobility Profile.

of online heuristics. We make the following observations.

As mentioned earlier, OptWiFi consistently outperforms OptGPS. This trend holds for linear prediction as well. We conclude that, if scheduled carefully, WiFi offers better energy-efficient localization than GPS/GSM.

Linear prediction performs well even for mobility traces that take frequent turns. We examined the optimal schedule for "staircase like" movements. When the distances between consecutive turns were short, the linear predictor approximated the movement with a straight line cutting diagonally through the staircase. If the trace has long stretches of straight lines, the Opt-LP schemes naturally predict well, scheduling 2 consecutive readings at the beginning of each straight line.

When using probability maps, the optimal ALE is small. This is because the number of mis-predictions (at the intersections) are typically fewer than the number of location readings permitted by the budget. As a result, Opt-Map schedules a location reading wherever there is a mis-prediction. We assumed that the velocity of the phone can be perfectly predicted, and hence, errors arise only after mis-predictions.

Consistent with our earlier observation, Heu-Eq-WiFi outperforms both Heu-Eq-GPS and Heu-Eq-GSM. However, note that EnLoc-Eq-Map outperforms Heu-Eq-WiFi, indicating that heuristics based on probability maps are effective for achieving energy-efficient localization.

Habitual Mobility

We present the localization error when a person's mobility profile is utilized for prediction. We use an anonymous student's mobility profile derived from 30 days of traces. We processed her mobility traces and manually generated the logical mobility tree (LMT). For each day, we executed the *Mobility Profile Heuristic*. Figure 3(c) shows that for the allocated energy budget of 25% for 24-hours the average localization error averages around 12m.

VI. Limitations and Future Work

We discuss the key limitations of our current approach:

We assumed that while moving along a predicted path, the location of the phone is accurately tracked. In reality, varying speeds or pauses cause this prediction to be imprecise. Our evaluation results do not account for these errors. Nonetheless, with accelerometers available on modern phones, speed variations may be estimated and used for accurate prediction [9]. Moreover, compasses may be able to dynamically sense the movement/orientation of a person, and obviate EnLoc-specified GPS readings. For instance, a GPS reading scheduled *after* an intersection can be eliminated if the compass provides the orientation of the user, indicating that the user took a left turn. Since compasses/accelerometers can be less energy-hungry, they may present opportunities for dynamic, adaptive localization (as opposed to the static schedule in EnLoc).

EnLoc does not proactively identify deviations from habitual paths. Techniques are necessary to quickly detect departures without investing excessive energy. Phone sensors may again be effective here. If the phone behaves differently from its habitual behavior at that time, EnLoc may suspect deviation and schedule a sensor reading. We are investigating these possibilities in our ongoing work.

Probability maps may be harder to generate for places unlike university campuses (a town or city). Location updates gathered over time from many mobile phones and statistics from transportation departments may be useful.

Lastly, the mobility profile used in EnLoc are those of graduate students, and may be less diverse (more predictable) than that of a traveling salesman. Thus, the reported errors may be optimistic. However, we believe that the algorithms/results presented in this paper validate the intuition that individual mobility profiling and large population statistics are an effective tool for energy-efficient localization.

References

[1] S. Gaonkar, J. Li, R. R. Choudhury, L. Cox, and A. Schmidt, "Micro-blog: Sharing and querying content through mobile phones and social participation," in *ACM MobiSys*, 2008.

[2] T. Sohn, K. A. Li, G. Lee, I. E. Smith, J. Scott, and W. G. Griswold, "Place-its: A study of location-based reminders on mobile phones," in *Ubicomp*, 2005.

[3] J. Yoon, B. Noble, and M. Liu, "Surface street traffic estimation." in *MobiSys '07*, 2007, pp. 220–232.

[4] J. Eriksson, L. Girod, B. Hull, R. Newton, H. Balakrishnan, and S. Madden, "The pothole patrol: Using a mobile sensor network for road surface monitoring," in *ACM MobiSys*, 2008.

[5] Y.-C. Cheng, Y. Chawathe, A. LaMarca, and J. Krumm, "Accuracy characterization for metropolitan-scale wi-fi localization," in *MobiSys*, 2005.

[6] M. C. Gonzalez, C. A. Hidalgo, and A.-L. Barabasi, "Understanding individual human mobility patterns," *Nature*, vol. 453, 2008.

[7] H. Lee, M. Wicke, B. Kusy, and L. Guibas, "Localization of mobile users using trajectory matching," in *ACM MELT*, 2008.

[8] I. Burbey and T. Martin, "Predicting future locations using prediction-by-partial-match," in *ACM MELT*, 08.

[9] A. Ofstad, E. Nicholas, R. Szcodronski, and R. R. Choudhury, "Aampl: Accelerometer augmented mobile phone localization," in *ACM MELT Workshop (in conjunction with Mobicom)*, 2008.