Energy-Aware Localization Using Mobile Phones

I. Constandache, Duke University
M. Sayler, Duke University
S. Gaonkar, Univ. of Illinois
R. Roy Choudhury, Duke University
L. Cox, Duke University

1. INTRODUCTION

Mobile computing applications are increasingly relying on location information. Examples include geotagging, geocasting, asset tracking, context-aware search, and visualization. Most applications assume GPS capabilities, and propose to fall back upon less accurate WiFi/GSM triangulation, when GPS capabilities are unavailable. This poster identifies a critical tradeoff between localization accuracy and energy consumption that has not been studied in the past. Our experiments show that while GPS can be accurate to within several meters, it can drain a mobile phone's battery within nine hours. WiFi and GSM based localization sustain higher battery life, but suffer from a proportionally greater localization error, as shown in Table 1. To address this tradeoff, we propose an energy-aware localization solution that multiplexes between alternate localization schemes, with the aim to maximize location accuracy under a given energy-budget. Our scheme exploits a variety of opportunities, including in-built accelerometers, user's mobility profile, and the knowledge of an application's needs. We briefly sketch our approach next.

	GPS	WiFi	GSM
Approx. Life	10 hours	40 hours	60 hours
Approx. Error	8m	25m	500m

Table 1: The energy-accuracy tradeoff

2. SOLUTION DESIGN

Our objective is to find an optimal schedule for sampling the phone's location, such that the average localization error is minimized while ensuring that the energy consumed by the schedule is less than the budget. We approach this through a dynamic programming formulation. Without loss of generality, we sketch the dynamic program for the case of GPS alone, with discrete time. Fig 1 illustrates this approach.

The staircase (solid) line represents the rate of increase in location error as a phone moves away from its starting point. In the absence of any localization (i.e., if the phone only knew its original location), the average error until t_n would be the area under this curve. However, when a GPS sample is taken at some point, t_i, the location error at that instant becomes zero. Of course, the error starts rising once again (shown by the dotted staircase line). Observe that the benefit of sampling at t_i can be expressed as the area of the shaded rectangle $(E_i \times (t_n - t_i))$. A second reading at time

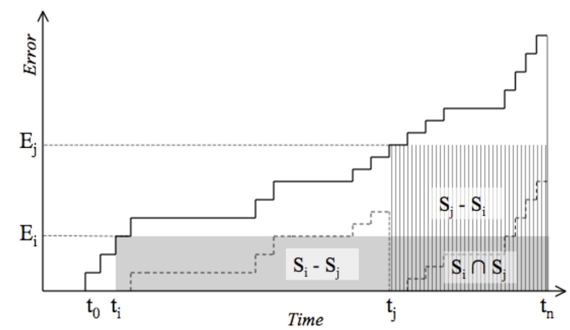

Figure 1: Cumulative error over time

t_j would offer an additional benefit of $(E_j - E_i) \times (t_n - t_j)$, the striped box minus the shaded box. Our dynamic program essentially determines the time points at which location samples must be taken, such that the union of all these areas is maximized. The equation is as follows.

$$P(j, K) = max_i\{P(i, K-1) + S(i, j)\},\ i < j, \quad (1)$$

$$P(i, 1) = area\ of\ rectangle\ i \quad (2)$$

where $P(j, N)$ is the maximum improvement with N readings given that the next rectangle we choose is j, and $S(i, j)$ is the additional benefit from taking a reading at time t_j, given a prior reading at time t_i.

3. ONGOING WORK

The above framework provides us upper and lower bounds on localization error, for a given energy budget. The upper bound is obtained in the absence of any mobility information (i.e., the phone is assumed to always move away from the starting location with velocity V_{max}). The lower bound is under complete information of a node's mobility pattern. Our ongoing work is utilizing this framework to dynamically determine near-optimal localization schedules, based on partial mobility traces (from t_0 to t). Towards this, a user's mobility profile is being exploited. For example, by correlating recent accelerometer readings with the distribution of the user's pause time, our scheme may be able to predict that the user will remain stationary for a long duration. At this point, a GPS sample can be cost-effective. However, as the user moves, the phone can switch to less accurate measures, since the samples will anyway become stale. When combined with application-specific requirements, further optimizations are possible. We are evaluating these ideas on a testbed of Nokia N95 mobile phones.