Computing Location from Ambient FM Radio Signals

Adel Youssef
Department of Computer Science
University of Maryland
A.V. Williams Building
College Park, MD 20742
adel@cs.umd.edu

John Krumm, Ed Miller, Gerry Cermak, Eric Horvitz
Microsoft Research
Microsoft Corporation
One Microsoft Way
Redmond, WA USA 98052
{jckrumm, edmiller, gcermak, horvitz}@microsoft.com

Abstract- We present a method for computing the location of a device down to a radius of several miles within a greater metropolitan area by analyzing the signal strengths observed from commercial FM radio stations. The use of ambient commercial radio signals allows for wide coverage, both indoor and outdoor reception, client-side computing for privacy, and the feasibility of employing inexpensive, low-power measurement hardware. Our technique is based on a model for computing the likelihood of locations using both received signal strengths and information from a simulated signal strength map. Using simulated signal strengths relieves the burden of manually measuring signal strength as a function of location. We account for the inevitable measurement variations among devices by comparing rankings of radio stations by signal strength. Our experiments show we can measure location down to a median error of about 8 kilometers (5 miles) in the greater Seattle area by listening to seven different radio stations.

Keywords: networks, location determination, smart personal object technology, pervasive computing, Bayesian inference.

Computing Location from Ambient FM Radio Signals

Adel Youssef
Department of Computer Science
University of Maryland
A.V. Williams Building
College Park, MD 20742

John Krumm, Ed Miller, Gerry Cermak, Eric Horvitz
Microsoft Research
Microsoft Corporation
One Microsoft Way
Redmond, WA USA 98052

***Abstract-* We present a method for computing the location of a device down to a radius of several miles within a greater metropolitan area by analyzing the signal strengths observed from commercial FM radio stations. The use of ambient commercial radio signals allows for wide coverage, both indoor and outdoor reception, client-side computing for privacy, and the feasibility of employing inexpensive, low-power measurement hardware. Our technique is based on a model for computing the likelihood of locations using both received signal strengths and information from a simulated signal strength map. Using simulated signal strengths relieves the burden of manually measuring signal strength as a function of location. We account for the inevitable measurement variations among devices by comparing rankings of radio stations by signal strength. Our experiments show we can measure location down to a median error of about 8 kilometers (5 miles) in the greater Seattle area by listening to seven different radio stations.**

I. INTRODUCTION

Accessing information about a user's location has been a central challenge in ubiquitous computing research, as many context-sensitive services depend critically on location. We have pursued methods for determining a user's location, both indoors and outdoors, that can be supported by inexpensive, low-power hardware. We explore in this paper our refinements of methods for identifying location from the analysis of ambient commercial radio broadcasts. We focus, in particular, on developing location services that can be supported by a small receiver, configured in a wristwatch format. The device, called the Smart Personal Object Technology (SPOT) wristwatch, has been developed by Microsoft into a product that is targeted at serving as both a time piece and receiver of information. Fig. 1 shows one of the prototype SPOT devices and sample displays of real time data which can include news, weather, traffic, sports, stock quotes, and instant messages. The information is transmitted on the sidebands of existing, commercial FM radio stations.

We set out to develop methods that could make this device location-aware without requiring extra hardware, fundamental new designs, nor additional requirements on battery power. Information about a user's current location would enable a number of enhancements for the SPOT information service, including the provision of more targeted, location-centric services, such as the times of movies playing in nearby theaters, specials at proximal restaurants, and street maps.

Unfortunately there is no inexpensive, tiny, low-power, add-on circuit that can measure the position of the watch. As examples of available wristwatch-sized location systems,

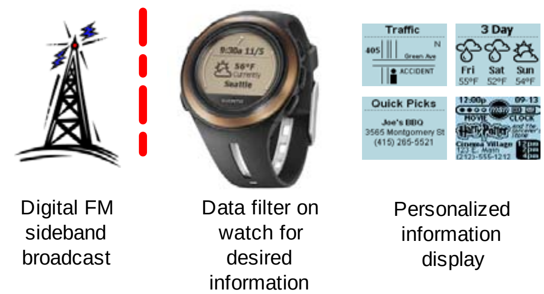

Fig. 1. The SPOT watch receives digital data encoded on the sideband of commercial FM radio stations. It filters the received data into customized channels for the user.

there are at least three commercially available watches with GPS functionality, all of which are much bulkier than normal watches, as shown in Fig. 2(a)–(c). Cell phones can also give location, but are bulky even when attempts are made to minimize their size and providing means for clipping such devices to a wrist, as shown in Fig. 2(d). Although we cannot measure location as accurately on a SPOT device as is common with cell phones, the wristwatch has the advantages of easy glancebility and indoor functionality.

Given that the SPOT watch can already measure signal strengths from commercial FM radio stations, our goal was to develop a technique that uses signal strengths to provide approximate locations for the devices. The advantage of this approach is that it uses the device's existing hardware, requiring only the addition of software. The idea of using signal strengths for location identification has precedent. For instance, the RADAR system [1] uses Wi-Fi signal strengths to locate an 802.11 device down to an accuracy of a few meters. The SmartMoveX active badge [2] uses short range radio signal strengths to localize an active badge inside a building. These efforts share with ours the fact that the radio signals in question do not contain timing information, leaving them to depend on signal strength as the only indication of location.

We reported on our early investigation of the feasibility of localization from the analysis of ambient FM signal strengths in [3]. That effort showed that we could correctly classify a SPOT device into one of six suburbs with a classification accuracy of about 80%. The work described in this paper extends the earlier results in the following ways:

- We classify the device into a uniform grid of locations rather than seeking to identify one of a few, spatially

(a) This Casio watch has a built-in GPS receiver for self-location.

(b) The Wherify watch is designed for tracking children. The GPS-derived position is broadcast wirelessly and can be accessed from a web site.

(c) This Timex watch is the hub of a wireless network of devices carried on the body, including the larger GPS receiver.

(d) NTT DoCoMo's Wirstomo unclips from a wrist to become a handheld cell phone.

Fig. 2. (a) – (c) are three examples of GPS-equipped watches and (d) is an example of a wrist-mounted cell phone. Their extra bulk limits their appeal as watches.

separated suburbs. With the finer-grained resolution, we measure error in units of distance rather than classification accuracy for regions.

- Instead of manually training locations as a function of signal strength as before, we now use simulated signal strength maps that eliminate the need to physically visit locations and measure signal strengths for training.
- We present a principled approach to choosing a good set of radio stations to measure, based on analysis of the simulated signal strength maps.

The remainder of this paper is organized as follows: in the next section, we discuss the signal strength features and describe methods employed in the SPOT device to measure signal strength of FM transmitters. Section 3 explains how the signal strength maps are generated. Our approach is presented in Section 4. In Section 5, we report on a set of experiments in the greater Seattle area and summarize with results, highlighting the ability to localize position down to a few miles. Finally, Section 6 concludes the paper and discusses future work.

II. Signal Strength Features

We use signal strenths measured from commercial FM radio stations for localization. In normal operation, the device listens to selected stations that are broadcasting SPOT data on their sidebands, given prior agreements between the stations and Microsoft on the leasing of these channels. In order to find the best SPOT-enabled station, the device scans through a preprogrammed list of stations and measures their relative signal strengths in order to pick the strongest one. Our technique exploits this ability by adding extra radio stations to the scan list. These extra stations do not carry SPOT data, but instead are used solely to help localize the device. We explain in Section 3 how we select the list of radio stations to scan.

The device's *received signal strength indicator* (rssi) comes from an analogue-to-digital converter (ADC) in the device. The raw digital measurements from each frequency are scaled and then averaged over 20 readings for 13 milliseconds. The ADC and associated circuitry are not carefully calibrated to measure rssi in any particular units nor to be consistent from device to device. This expected inconsistency poses a problem for location measurement, because we cannot expect that a given location will be characterized by a consistent set of signal strengths across multiple devices. One solution to such variation is to calibrate each device at the factory and add a function on the device to compute a canonical signal strength from every measurement. We deemed this approach as being too expensive and pursued methods that could provide robust training and inference without such detailed calibration.

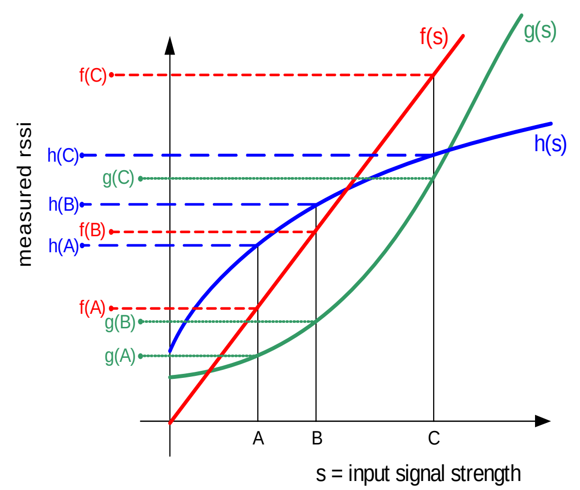

Fig. 3. If the mobile device's measured rssi is a monotonically increasing function of input signal strength, then the signal strength order of inputs will be preserved. For instance, if $A < B < C$, then these values transformed through the three different rssi-measuring functions will stay in the same signal strength order.

As we cannot assume that different locations are characterized by consistent signal strengths, we make the weaker assumption that the relative signal strengths will be consistent, *i.e.* that the ordered list of radio stations sorted by signal strength will be stable for a given location. More precisely, we assume that the relationship between input signal strength and measured rssi is monotonically increasing, as illustrated in Fig. 3. If this is the case, then the transformation between input signal strength and measured rssi will preserve the signal strength order of the input. This helps our algorithm work for a wide variety of devices in spite of device-to-device variations in how they measure signal strengths.

We compute our rank feature by first scanning though a list of n radio stations, each identified with an index $[1\ldots n]$. The signal strength of station i is denoted s_i. Measuring the rssi of each station results in a set of ordered pairs giving the station index and signal strength of each radio station: $\{(1,s_1),(2,s_2),\ldots,(n,s_n)\}$. This set of radio station–strength tuples is then sorted by signal strength to get a rank vector. For example, suppose $n=3$ and the scan results in $\{(1,40),(2,30),(3,35)\}$. Sorting this set of ordered pairs on signal strength gives $\{(2,30),(3,35),(1,40)\}$. The rank vector is then the radio station indices taken in signal strength order, *i.e.* $\underline{r}=(2,3,1)$, indicating that $s_2 < s_3 < s_1$. Equal signal strengths are sorted arbitrarily. For n stations there are $n!$ different rank vectors. We assign a unique integer hash code to each rank vector for more efficient storage. The hash code is computed from an algorithm presented by Knuth [4], which maps each rank vector to an integer $r \in [0\ldots n!-1]$ using a mixed-radix representation of the integers. The next section explains how we predict what the rank hash codes will be as a function of location, which is used ultimately to localize the device.

III. Signal Strength Maps

To infer a location from a rank hash code, we need to have a relationship between location and hash codes. One way to ascertain this relationship is to physically visit a grid of locations on the ground and record signal strengths. For the general case, this is tedious and it requires ongoing maintenance if a radio station changes its transmitter's power or location.

To minimize the need to make a large number of measurements by physically visiting locations, we explored the value of using high-fidelity simulations of radio signal propagation. These simulations are typically used in the commercial radio industry for estimating the reach of stations. In particular, we experimented with RadioSoft's ComStudy software to generate simulated FM radio signal strength maps. Fig. 4 displays such a map. In Section A, we explain how the FM radio maps were generated. Then, in Section B, we discuss how the radio maps were validated against actual signal strength measurements. Finally, in Section C, we show how we selected a good subset of radio stations for efficient computation of location.

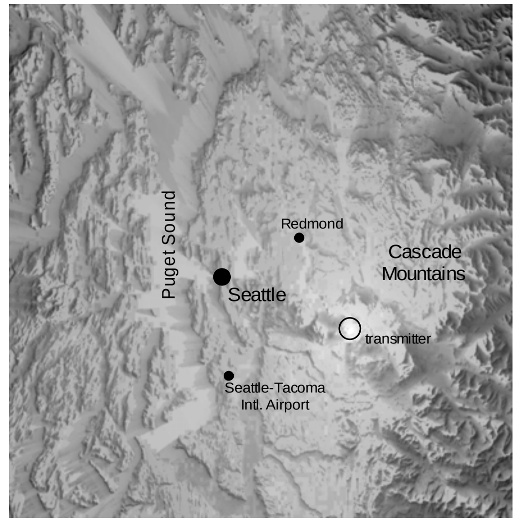

Fig. 4. Display of the simulated signal strengths of a radio station in the Seattle, WA USA area generated by the ComStudy program. The effect of terrain is apparent.

A. *Generating Simulated FM Signal Strength Maps*

The ComStudy software supports many radio propagation models that can be used to predict FM radio maps. We chose the Longley-Rice model, also known as NBS 101[1], for its known accuracy. It is generally the most accurate of the choices since it incorporates reflection, refraction (bending of the rays as they rise through the atmosphere), and several types of diffraction (spilling of signal over hills). The model uses a terrain map to simulate the effect of hills and valleys. The terrain effects are obvious in the simulated rssi map shown in Fig. 4.

ComStudy parameterizes each radio station transmitter by its frequency, transmitting power, and location as (latitude, longitude). For each transmitter, a field strength matrix is generated. The matrix is a grid of points spread over the chosen area on the ground, with each point containing the signal level information from the transmitter in question. The points are 6 arc seconds apart (about 185 meters north-south and 124 meters east-west). This is equivalent to having 40 points per square kilometer. ComStudy then applies the chosen propagation model to calculate field strength at each point. We generated maps for 28 local FM radio stations.

[1] National Bureau of Standards (NBS) Tech Note 101.

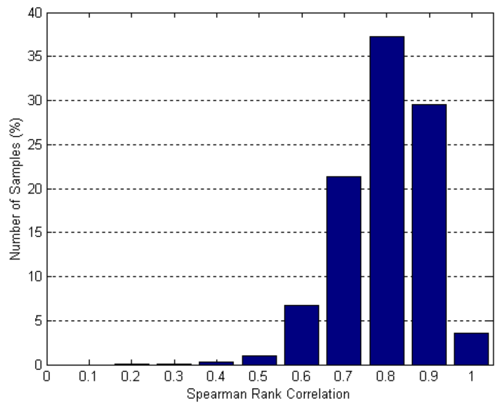

Fig. 5. Histogram of Spearman's rank correlation between simulated and actual signal strength rankings shows that the simulated signal strength maps do a good job of predicting the actual rankings.

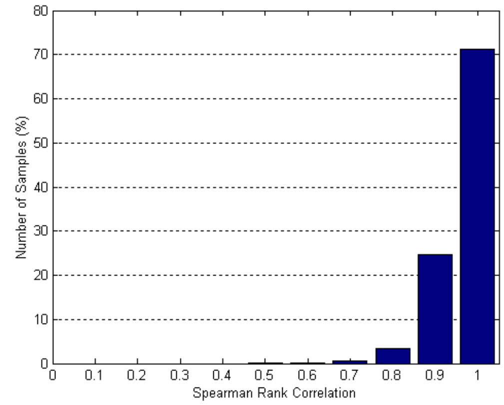

Fig. 6. Compared to the histogram in 5, the reduced set of radio stations shows better rank correlation between the simulated and measured rankings.

B. Validating the Simulated Signal Strength Maps

Before we used the simulated signal strength maps for location inference, we had to be confident that the maps accurately predicted the rank vectors. To address this concern, a SPOT watch was programmed to measure signal strengths of the same 28 FM radio stations for which we had generated maps. We drove to locations within this area with the device, logging the signal strengths of all 28 stations, taking one 28-station scan per second, resulting in about 6290 readings for each station. We simultaneously logged (latitude, longitude) from a GPS receiver in the same vehicle. To assess the accuracy of the simulation, we calculated the Spearman rank correlation [5] between the ranked radio stations from the measured signal strengths and from the simulated signal strengths, using the GPS readings to determine which simulated strengths to use from the maps. We chose the Spearman rank correlation instead of Kendall's because differences between data values ranked further apart are given more weight with Spearman. Both correlations range over [-1 … 1], with "1" indicating equivalent rankings and "-1" indicating opposite rankings.

Fig. 5 shows a histogram of Spearman's correlations between the measured signal strength and the simulated signal strength. The histogram shows that more than 95% of the measured rankings were correlated with the simulated rankings with a factor of at least 0.6 on a scale of [-1 … 1]. We considered this as a sign that the simulated results are potentially accurate enough for determining the location of a SPOT watch. In the next section, we show that even better correlation exists if the number of stations is reduced.

C. Reducing the Number of Radio Stations

With 28 available radio stations, we have $28! \approx 3.05\text{x}10^{29}$ different rank vectors. Considering all of these signals as independent would clearly demand infeasible computation and storage. However, many of the stations are broadcast from the same tower with some having similar transmission power. Measuring signal strengths from co-located, similarly powered stations is redundant. To eliminate these redundancies, we divided the 28 stations into groups of approximately similar stations. We used the following criteria to assess similarity:

1. Calculate Pearson's correlation coefficient [5] between all pairs of stations based on spatially corresponding points on the simulated signal strength maps. Pearson's correlation coefficient is the standard linear correlation coefficient, not a rank correlation.
2. Find groups of stations that are correlated with each other greater than a certain threshold ρ. For our implementation, we used $\rho = 0.95$.
3. From each correlated group, select the station with the highest average signal strength to represent the group.

Applying the above criteria to the simulated signal strength maps of the Seattle area, we found that there are seven groups of correlated stations with Pearson's correlation factor of at least 0.95.

Applying the above criteria, the number of stations was reduced from 28 to 7. This means instead of having 28! different ranking vectors, we had just 7! = 5040, a reduction of 25 orders of magnitude. This has a significantly positive effect on the computational and memory performance of the technique used for location determination. This is particularly important for the resource-limited devices we are targeting. Fig. 6 shows a histogram of Spearman's correlations between the measured signal strength and the simulated signal strength of the seven chosen stations. The figure shows a significant improvement over the similar figure with 28 stations (Fig. 5), with an even higher correlation between the measured and simulated signal strengths. More than 95% of the measured signal strength rank vectors are correlated with the simulated signal strength with a factor of 0.8 or above. This further emphasizes that the simulated results are valid to use for determining the location of a SPOT watch.

IV. Location Estimation From Rank Histograms

Each point in the simulated signal strength maps was converted to a rank hash code, which was ultimately used for inferring location from measured rank hash codes. This section describes our inference process, starting with our technique for producing a smoothed likelihood function giving the probability of a rank hash code as a function of location. We then show how we use a Bayesian inference technique to infer the location of the device.

A. *Rank Hash Code Likelihood as a Function of Location*

Our simulated signal strength maps were generated on a grid with points about 124 meters apart east-west and 185 meters apart north-south. This is near the limit of the ComStudy's maximum simulation resolution, and it is much finer than we need for our device's intended purpose. Furthermore, this fine resolution means we have 442,806 points (811 east-west by 546 north-south) to represent our test area around Seattle. The rank hash code ranges over [0 … 7!-1], meaning that each fine point needs a 13-bit integer to represent its hash code. The radio map would then need 811 x 546 x 13/8 ≈ 0.7 MB to represent the fine grid of hash codes for our test area. This is too much for our device, and we would likely never achieve the location resolution of these small cells anyway.

To alleviate this storage problem and to represent the hash codes at a reasonable resolution, we created a coarse grid of cells over the fine grid of points and represented each cell as a histogram of rank hash codes from its underlying points. We varied the size of the cells, with a typical size is 8 km x 8 km which covers 43 x 64 = 2752 points.

Each cell's histogram was normalized to give an estimate of the probability distribution of rank hash codes for the cell. Mathematically, this likelihood estimate is $P'(r|c_i)$, where $r \in [0 \ldots n!-1]$ represents the rank hash code of the n radio stations and $c_i, i \in [1 \ldots m]$ represents one of the m cells.

In practice, the histograms are sparsely populated because of the large number of possible rank hash codes. Due to noise, orientation of the device, and radio propagation effects that are not well simulated, we often measured rank hash codes in a cell for which the simulated likelihood $P'(r|c_i)$ was zero. We thus chose to smooth our likelihood estimate to fill in the gaps. Unfortunately simple smoothing over the rank hash codes r was not reasonable, since adjacent hash codes do not necessarily represent similar rank vectors. Instead, we smoothed by replacing the value in each histogram bin by the maximum value over all bins whose Spearman correlation coefficient with the bin in question was above a chosen threshold ρ_s. In equation form, the (non-normalized) smoothed likelihood was computed as

$$P_u(r|c) = \max_{r':S(r,r')\geq\rho_s} P'(r'|c)$$

Here $S(r,r')$ is the Spearman correlation coefficient between the rank vectors represented by the hash codes r and r'. $P_u(r|c_i)$ was then normalized over r into $P(r|c_i)$ to give the smoothed likelihood function of a rank hash code given a cell. Intuitively, ρ_s serves as a smoothing parameter, with lower values giving more smoothing. We evaluate the effect of this parameter in the results section.

B. *Location via Bayesian Inferenece*

Given a measured rank hash code r, the probability of being in a cell c_k is given by Bayes rule:

$$P(c_k|r) = \frac{P(r|c_k)P(c_k)}{\sum_{i=1}^{m} P(r|c_i)P(c_i)}$$

Given r, this denominator is constant, and we make no *a priori* assumptions on which of the m cells the device is in, meaning $P(c_k) = 1/m$. Thus we take the maximum likelihood estimate as the location of the device:

$$c = \arg\max_{c_k:k\in[1\ldots m]} P(r|c_k)$$

In actual practice, we smooth these results by taking the most frequently inferred cell over the last K signal strength scans. This helps keeps the location inference from jumping unreasonably from place to place, although an excessive value of K can cause the inference to be too "sticky". The next section shows the results of this algorithm on real data, including the effect of different values of K.

For deploying to an actual device, the watch would monitor radio stations delivering digital SPOT data. (There are usually two such stations for every U.S. metropolitan area.) This data already includes coarse location information to enable local weather reports. If the coarse location changed, then the watch would use the same digital feed to

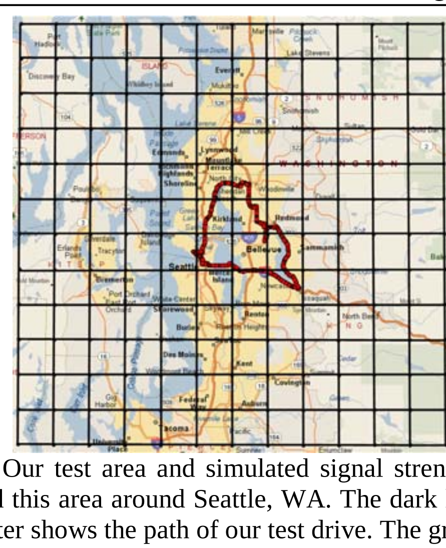

Fig. 7. Our test area and simulated signal strength maps covered this area around Seattle, WA. The dark loop near the center shows the path of our test drive. The grid size in this image has square cells with width 8 km.

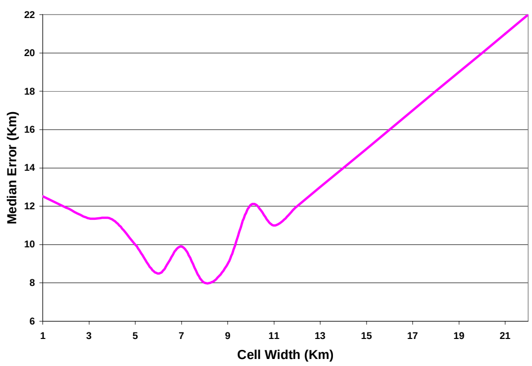

Fig. 8. The relation between grid cell width (w) and median error for $\rho_s = 0.97$ and $K = 30$. We used $w = 8$ kilometers.

download discretized versions of the smoothed likelihood functions $P(r|c_i)$, which would be transmitted periodically along with other SPOT data.

V. Results

In this section, we summarize our experiment and discuss the effect of different parameters on the accuracy of the inference algorithm. We performed the experiment in the greater Seattle area. A SPOT watch was programmed to measure the signal strengths of the reduced set of 7 local FM radio stations, as discussed in Section 3.C. We drove around the area with the watch, logging the signal strengths of all 7 stations, taking one 7-station scan per second, resulting in about 6290 readings for each station. We simultaneously logged (latitude, longitude) from a GPS receiver in the same vehicle. Meanwhile, a simulated radio map for the 7 stations was generated, as discussed in Section 3.A. Fig. 7 shows the drive path along with the grid with cell width of 8 km that was used for inferring the location. The drive path includes both highways and local roads.

We explored the parameter space of correlation threshold ρ_s, temporal window size K, and grid cell width w, testing the median error resulting from different settings. We took the inferred location as the center of the most likely cell. Our experiment showed that we can determine a device's location down to a median error of about 8 kilometers (5 miles) using the technique described in Section 4. The distribution of cumulative error probabilities is shown in Fig. 9 for the optimal settings of the parameters. We found that the best parameter settings were a cell width of $w = 8$ kilometers, a histogram-smoothing correlation threshold of $\rho_s = 0.97$, and a temporal window of $K = 30$.

Our experiment shows that the grid cell width w is the most important parameter in terms of affecting the accuracy of the location inference. Increasing the cell width has the advantage of having more simulated points in the histogram estimate of $P'(r|c)$ and hence more information about the ranking probability distributions. Also, increasing the cell

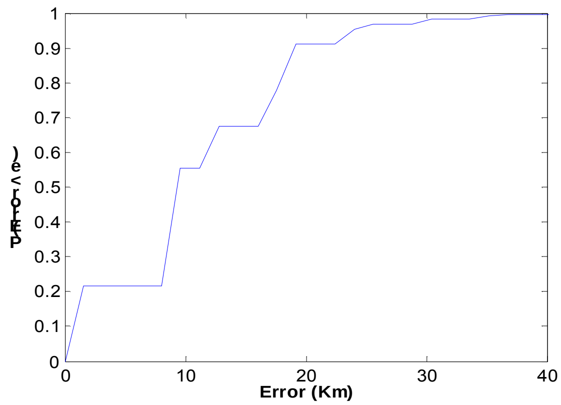

Fig. 9. The error cumulative distribution for $w = 8$ kilometers, $\rho_s = 0.97$ and $K = 30$.

width decreases the overall number cells in the grid and thus reduces the computational overhead. However, as the cell width increases, the location resolution naturally decreases. The relation between cell width and median error is shown in Fig. 8.

The second factor that affects the accuracy of the inference algorithm is the correlation smoothing threshold ρ_s. As discussed in Section 4.A, the correlation threshold was introduced to smooth our likelihood estimate. Therefore, decreasing ρ_s will result in filling more gaps in the histogram-based estimate of $P'(r|c_i)$. However, decreasing ρ_s below a certain value will result in too much smoothing, decreasing the distinguishability of different locations. The plot in Fig. 10 shows the relation between ρ_s and median error. The plot shows that there a minimum of the median error at $\rho_s = 0.97$.

The last factor that affects the accuracy is the temporal window size K. Because we scanned radio stations at 1 Hz, K corresponds to the number of seconds of data we used to infer position. Clearly, increasing K will enhance the accuracy;

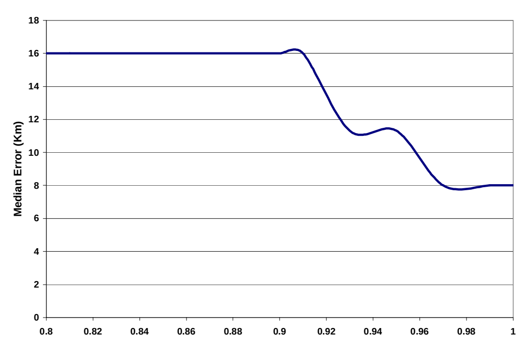

Fig. 10. The relation between correlation threshold ρ_s and median error for $w = 8$ kilometers and $K = 30$. We used $\rho_s = 0.97$.

however, it will increase memory storage and computational power. Fig. 11 shows the relation between the window size and median error. It is clear that we can get a median error of 8 kilometers with a window size $K \geq 30$.

VI. Conclusions And Future Work

We have described a method for inferring the location of a device based on FM radio signal strengths. The advantages of the method include the wide coverage of FM radio, spanning indoor and outdoor locations, and the readiness of the target SPOT device for measuring radio signal strengths. For our applications, an accuracy of several miles is adequate. We made the technique robust to measurement differences between devices by basing our inferences on rankings of radio stations rather than on their absolute signal strengths. Our method requires no manual survey of signal strength as a function of location, because we use simulated signal strengths whose validity we verified against actual measurements. Using smoothed histograms of rank hash codes, we can infer a device's location down to accuracy of about 8 kilometers (5 miles). We believe this method has potential for widespread use in small, wearable devices.

Our future work in this area will include using natural constraints on where and how fast people can move, with an HMM being a likely candidate for imposing such constraints. Computerized route-planning software will be a good source of these constraints. This may lead to an alternate representation instead of a grid of cells. For instance, a natural representation of peoples' location might be islands of high population connected by networks of roads. Other constraints could be derived from querying users about where they normally go or do not go. For instance, a user could indicate that he or she rarely spends time on water and hardly ever visits certain suburbs.

References

[1] P. Bahl, and V.N. Padmanabhan, "RADAR: an in-building RF-based location and tracking system," in IEEE INFOCOM 2000.

[2] J. Krumm, L. Williams, and G. Smith, "SmartMoveX on a graph -- an inexpensive active badge tracker," in Ubicomp 2002.

[3] J. Krumm, G. Cermak, and E. Horvitz, "RightSPOT: a novel sense of location for a smart personal object," in UbiComp (Ubiquitous Computing) 2003.

[4] D. Knuth, The Art of Computer Programming, Addison-Wesley, Massachusetts, 1981, pp. 64 (Algorithm P).

[5] W. Press, B. Flannery,S. Teukolsky, and W. Vetterlinq, Numerical Recipes in C, Cambridge: Cambridge University Press, 1992.