RADAR: An In-Building RF-based User Location and Tracking System

Paramvir Bahl and Venkata N. Padmanabhan
Microsoft Research
{bahl, padmanab}@microsoft.com

Abstract

The proliferation of mobile computing devices and local-area wireless networks has fostered a growing interest in location-aware systems and services. In this paper we present RADAR, a radio-frequency (RF) based system for locating and tracking users inside buildings. RADAR operates by recording and processing signal strength information at multiple base stations positioned to provide overlapping coverage in the area of interest. It combines empirical measurements with signal propagation modeling to determine user location and thereby enable location-aware services and applications. We present experimental results that demonstrate the ability of RADAR to estimate user location with a high degree of accuracy.

Keywords: *location-aware services, user location and tracking, wireless LAN, radio-frequency wireless network.*

1 Introduction

The proliferation of mobile computing devices and local-area wireless networks has fostered a growing interest in location-aware systems and services. A key distinguishing feature of such systems is that the application information and/or interface presented to the user is, in general, a function of his or her physical location. The granularity of location information needed could vary from one application to another. For example, locating a nearby printer requires fairly coarse-grained location information whereas locating a book in a library would require fine-grained information.

While much research has focussed on developing services architectures for location-aware systems (e.g., [Maa97,Nel98]), less attention has been paid to the fundamental and challenging problem of locating and tracking mobile users, especially in in-building environments. The few efforts that have addressed this problem have typically done so in the context of infrared (IR) wireless networks. The limited range of an IR network, which facilitates user location, is a handicap in providing ubiquitous coverage. Also, the IR network is often deployed for the sole purpose of locating people and does not provide traditional data networking services. To avoid these limitations, we focus on RF wireless networks in our research. Our goal is to complement the data networking capabilities of RF wireless LANs with accurate user location and tracking capabilities, thereby enhancing the value of such networks.

In this paper, we present RADAR, an RF-based system for locating and tracking users inside buildings. RADAR uses signal strength information gathered at multiple receiver locations to *triangulate* the user's coordinates. Triangulation is done using both empirically-determined and theoretically-computed signal strength information.

Our experimental results are quite encouraging. With high probability, RADAR is able to estimate a user's location to within a few meters of his/her actual location. This suggests that a large class of location-aware services can be built over an RF local-area wireless data network.

The remainder of this paper is organized as follows. In Section 2, we survey related work in location determination technologies. In Section 3, we discuss our research methodology. Section 4 contains the core of the paper where we present and analyze the empirical and the signal propagation modeling methods. A discussion of extensions to the base RADAR system appears in Section 5. Finally, we present our conclusions in Section 6.

2 Related Work

Related work in the area of user location and tracking falls into the following broad categories: (1) in-building IR networks, (2) wide-area cellular networks (based on RF), and (3) Global Positioning System (GPS).

The *Active Badge* system [Wan92,Har94] was an early and significant contribution to the field of location-aware systems. In this system, a badge worn by a person emits a unique IR signal every 10 seconds. Sensors placed at known positions within a building pick up the unique identifiers and relay these to the location manager software. While this system provides accurate location information, it suffers from several drawbacks: (a) it scales poorly due to the limited range of IR, (b) it incurs significant installation and maintenance costs, and (c) it performs poorly in the presence of direct sunlight, which is likely to be a problem in rooms with windows.

Another system based on IR technology is described in [Azu93]. IR transmitters are attached to the ceiling at known positions in the building. An optical sensor on a head-mounted unit senses the IR beacons, which enables the system software to determine the user's location. This system suffers from similar drawbacks as the Active Badge system.

The system described in [ATC97] is based on pulsed DC magnetic fields. Multiple sensors are placed on body-mounted peripherals, such as data gloves, and their output is processed to determine a person's location and orientation with a high degree of precision. This technology is used

IPR2024-00397
Apple EX1014 Page 1

extensively in the computer animation industry. It is, however, quite expensive and, like IR, severely range limited, hence unsuitable for large-scale deployment.

Recently, several location systems have been proposed for wide-area cellular systems [Tek98]. The technological alternatives for locating cellular telephones involve measuring the signal attenuation, the angle of arrival (AOA), and/or the time difference of arrival (TDOA). While these systems have been found to be promising in outdoor environments, their effectiveness in indoor environments is limited by the multiple reflections suffered by the RF signal, and the inability of off-the-shelf and inexpensive hardware to provide fine-grain time synchronization.

Systems based on the Global Positioning System [GPS99], while very useful outdoors, are ineffective indoors because buildings block GPS transmissions.

The Daedalus project [Hod97] developed a system for coarse-grained user location. Base stations transmit beacons augmented with their physical coordinates. A mobile host estimates its location to be the same as that of the base station to which it is attached. Consequently, the accuracy of the system is limited by the (possibly large) cell size.

Our work differs from previous work in that we tackle the problem of user location and tracking on a widely available *radio frequency* based wireless network in an *in-building* environment. RF networks offer a significant advantage over IR networks in terms of range, scalability, deployment, and maintenance. With speeds of up to 11 Mbps, these systems have gained rapid acceptance and are being widely deployed in offices, schools, homes, etc.

We recently became aware of the Duress Alarm Location System (DALS) described in [CG93]. While their work and ours are similar in some ways, they also differ in significant ways. Briefly, their system (1) is dependent on specialized hardware, (2) does not use propagation modeling to build a radio map of the building, (3) does not factor in user body orientation, and (4) requires infrastructural deployment over and above a wireless data network. These points are clarified in the following sections.

3 Research Methodology

We begin with a description of our experimental testbed. We then discuss the data collection process, including tools we developed for this purpose. Finally, we describe the processing we performed on the data as a precursor to the analysis described in Section 4.

3.1 Experimental Testbed

Our experimental testbed is located on the second floor of a 3-storey building. The layout of the floor is shown in Figure 1. The floor has dimension of 43.5 m by 22.5 m, an area of 980 sq. m (10500 sq. ft.), and includes more than 50 rooms.

We placed three base stations, BS_1, BS_2 and BS_3, at the locations indicated in Figure 1. Each base station was a Pentium-based PC running FreeBSD 3.0 equipped with a wireless adapter. Our mobile host, carried by the user being tracked, was a Pentium-based laptop computer running Microsoft Windows 95.

Each base station and the mobile host was equipped with a Digital RoamAbout™ network interface card (NIC), based on Lucent's popular WaveLAN™ RF LAN technology. The network operates in the 2.4 GHz license-free ISM (Industrial, Scientific and Medical) band. It has a raw data rate of 2 Mbps and a one-way delay of 1-2 ms. The range of the network, as specified in [Roa96], is 200 m, 50 m, and 25 m, respectively, for open, semi-open, and closed office environments. This classification is based on the type and density of obstructions between the transmitter and the receiver. With this nomenclature, our testbed environment would be classified as being open along the hallways where the base stations are located and closed elsewhere. The base stations provide overlapping coverage in portions of the floor, and together cover the entire floor.

3.2 Data Collection

A key step in our research methodology is the data collection phase. We record information about the radio signal as a function of the user's location. As discussed in Section 4, we use the signal information to construct and validate models for signal propagation during off-line analysis as well as to infer the location of a user in real time. We refer to the former as the *off-line phase* and the latter as the *real-time phase*.

Among other information, the WaveLAN NIC makes available the *signal strength* (SS) and *the signal-to-noise ratio* (SNR). SS is reported in units of dBm and SNR is expressed in dB. A signal strength of *s* Watts is equivalent to $10*log_{10}(s/0.001)$ dBm. A signal strength of *s* Watts and a noise power of *n* Watts yields an SNR of $10*log_{10}(s/n)$ dB. For example, a signal strength of 1 Watt is equivalent to 30 dBm. Furthermore, if the noise power is 0.1 Watt, the SNR would be 10 dB.

The FreeBSD 3.0 WaveLAN driver extracts the SS and the SNR information from the WaveLAN firmware each time a broadcast packet is received[1]. It then makes the information available to user-level applications via *ioctl* system calls. We used the *wlconfig* utility, which provides a wrapper around the *ioctl* calls, to extract the signal information.

In our testbed, we chose to have the Windows-based mobile host broadcast packets (*beacons*) periodically and have the FreeBSD base stations record signal strength information. However, in a production system with many more mobiles than base stations, it may be desirable to have the latter transmit the beacons and the former measure the signal strength. Nevertheless, the accuracy of the user location and tracking is not impacted by this choice[2].

We wrote a simple application using Tcl/Tk [Ous94] and Perl [Wal96] to control the entire data collection process

[1] It is quite easy to modify the driver to record information for other packets as well, but we found no reason to do so.

[2] While our analysis does not assume symmetry of signal strength, the few instances where we measured signal strength at both ends indicate little asymmetry.

IPR2024-00397
Apple EX1014 Page 2

from the mobile host. The process operates as follows. First, the clocks on the mobile host and the base stations are synchronized (to within the round-trip latency of the wireless link, essentially less than 5 ms). The mobile host then starts broadcasting UDP packets, each with a 6-byte payload and spaced apart uniformly, at a default rate of 4 per second. Each base station (*bs*) records the signal strength (*ss*) measurement[3] together with a synchronized timestamp *t*, i.e., it records tuples of the form (*t*, *bs*, *ss*). This information is collected both during the off-line phase and the real-time phase.

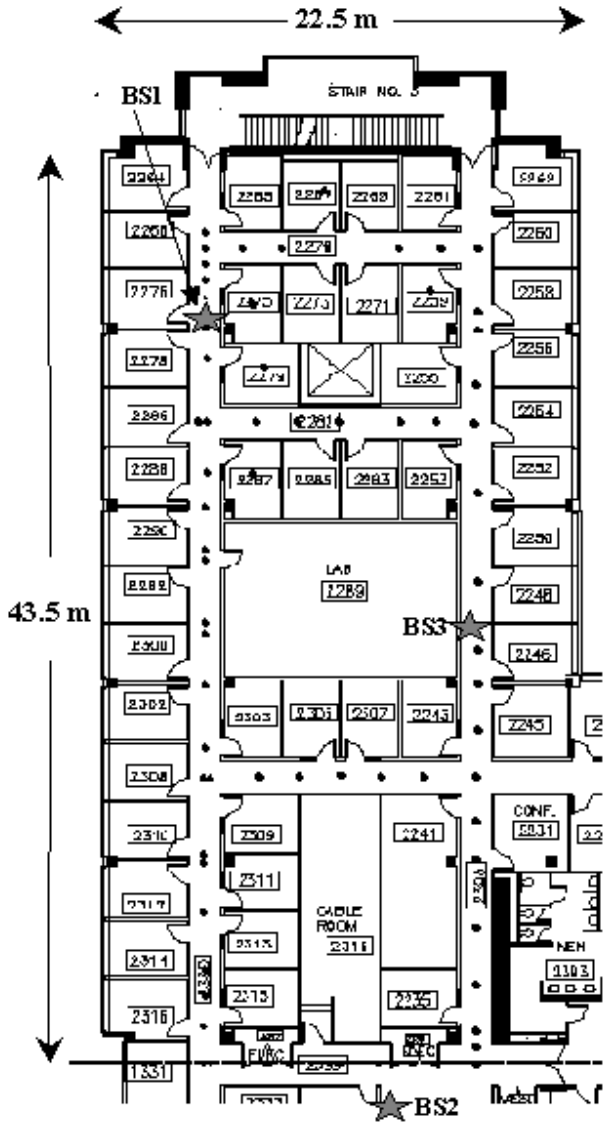

Figure 1 Map of the floor where the experiments were conducted. The black dots denote locations were empirical signal strength information was collected. The large stars show the locations of the 3 base stations. The orientation is North (up) and East (right).

In addition, during the off-line phase (but not the real-time phase), the user indicates his/her current location by clicking on a map of the floor. The user's coordinates (*x*,*y*) and timestamp *t* are recorded.

During our experiments, we discovered that signal strength at a given location varies quite significantly (by up to 5 dBm) depending on the user's orientation, i.e., the direction he/she is facing. In one orientation, the mobile host's antenna may have line-of-sight (LoS) connectivity to a base station's antenna while in the opposite orientation, the user's body may form an obstruction. So, in addition to user's location (*x*,*y*), we also recorded the direction (*d*) (one of north, south, east, or west) that he/she is facing at the time the measurement is made[4]. Thus, the mobile host records tuples of the form (*t*,*x*,*y*,*d*) during the off-line phase. We discuss the implications of the user's orientation in more detail in Section 4.

In all, during the off-line phase, we collected signal strength information in each of the 4 directions at 70 distinct physical locations on our floor. For each combination of location and orientation (i.e., (*x*,*y*,*d*) tuple), we collected at least 20 signal strength samples.

3.3 Data Processing

We outline the data processing that we performed as a precursor to the analyses discussed in Section 4.

3.3.1 Signal Strength Information

Using the synchronized timestamps, we merged all of the traces collected during the off-line phase into a single, unified table containing tuples of the form (x,y,d,ss_i,snr_i), where $i \in \{1,2,3\}$ corresponding to the three base stations. For each *(x,y,d)* tuple, we computed the mean, the standard deviation, and the median of the corresponding signal strength values for each of the base stations. For much of our analysis, we use this processed data set (primarily the mean) rather than the original, raw data set.

We wrote routines to search through the processed data set to determine exact as well as closest matches. There is a fair amount of database research literature that describes efficient data structures and algorithms for such multi-dimensional searches (e.g., *R-Tree* [Gut84], *X-Tree* [Ber96], *optimal k-nearest neighbor search* [Sei98], etc.) However, we chose a simple linear-time search algorithm because our relatively small data set and dimensionality (at most 3, as explained in Section 4) did not warrant the complexity of the aforementioned algorithms. Moreover, the focus of our research is on the analysis rather than on developing an optimal closest match implementation.

3.3.2 Building Floor Layout Information

We obtained the layout information for our floor, which specified the coordinates of each room. We also obtained the coordinates of the three base stations. Using these and

[3] During the course of our experiments, we discovered that the signal strength is a stronger function of location than the signal-to-noise ratio. The latter is impacted by random fluctuations in the noise process. So we only use signal strength information in our analysis.

[4] While there are other sources of fluctuation, such as the movement of other people and objects, these tend to be random. In contrast, the body of the person carrying the mobile host introduces a systematic source of error.

IPR2024-00397
Apple EX1014 Page 3

the *Cohen-Sutherland line-clipping algorithm* [Fol90], we computed the number of walls that obstructed the direct line between the base stations and the locations where we had collected the empirical signal strength data. We use this to build an accurate signal propagation model (Section 4.2).

4 Algorithms and Experimental Analysis

In this section, we discuss several algorithms for user location and tracking, and present an analysis of how well these perform using our experimental data.

A premise of our work is that signal strength (SS) information provides a means of inferring user location. To demonstrate that this is a reasonable premise, we show in Figure 2 shows how the SS measured at each of the 3 base stations varies as the user walks along the outer hallway of the floor in a counter-clockwise direction (Figure 1). The walk begins and terminates at the north-west corner (close to BS_1). There is a definite trend in the SS recorded at the three base stations as the user walks around the loop. Not surprisingly, the signal received at a base station is the strongest when the user is close to it and weakest when the user is far away. This strong trend is an indication that using SS to infer location is a promising approach.

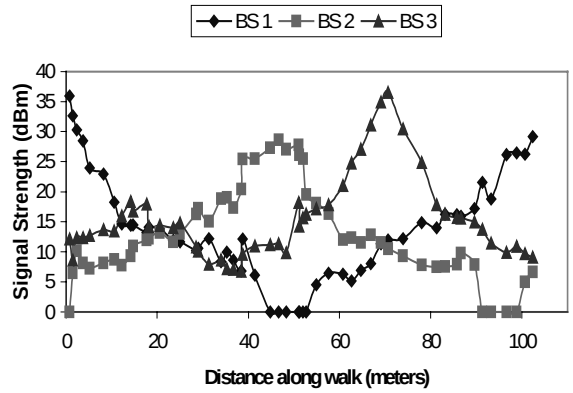

Figure 2 Signal strength recorded at the three base stations as the user walks around the floor.

Our basic approach is *triangulation*[5]. Given a set of signal strength measurements at each of the base stations, we determine the location that best matches the observed signal strength data. We then "guess" that to be the location of the user. There are multiple variations of this basic idea that arise because of several choices for each of the following:

- Ways of summarizing the signal strength samples measured at the base stations.
- Basis for determining the best match.
- Metric for determining the best match.

We discuss each of these in turn.

[5] It is just coincidental that we have three base station in our testbed. Triangulation can also be done with fewer or more base stations.

First, we summarize multiple signal strength samples from a base station using the sample mean. In the case of a static user whose location and orientation are fixed (the *user location* problem), it is clear which signal strength measurements should be included in the sample set. In the case of a mobile user (the *user tracking* problem), it is less clear what the sample set should be. In the latter case, we define the sample set to be all samples that lie within a sliding time window.

Second, to determine the location and orientation that best match a given (summarized) set of SS measurements, we first need to determine what the SS (at each base station) should be used for a particular combination of user location and orientation. We consider a couple of alternative approaches. The first is the *empirical method* where we use the location and SS data gathered during the off-line phase (Section 3.2). The second approach is *signal propagation modeling*. As discussed in Section 4.2, we have developed a model that accounts for both free-space loss and loss due to obstructions in computing the SS at each base station corresponding to given a user location.

Third, we need a metric and a search methodology to compare multiple locations and pick the one that best matches the observed signal strength. We term our general technique *nearest neighbor(s) in signal space (NNSS)*. The idea is to compute the *distance* (in signal space) between the observed set of SS measurements, (ss_1,ss_2,ss_3), and the recorded SS, (ss'_1,ss'_2,ss'_3), at a fixed set of locations, and then pick the location that minimizes the distance. In our analysis, we use the *Euclidean distance* measure, i.e., $sqrt((ss_1-ss'_1)^2+(ss_2-ss'_2)^2+(ss_3-ss'_3)^2)$. It is possible to use other distance metrics, for example, the sum of the absolute differences for each base station (the "*Manhattan" distance* [Cor90]) or a metric weighted by the signal strength level at each base station. We experimented briefly with these alternatives, but do not present the results here due to space limitations.

In all of our analyses, we characterize the goodness of our estimate of the user's location using the *error distance*, which is the Euclidean distance between the actual (physical) location of the user and the estimated location.

4.1 Empirical Method

In this case, we use the empirical data obtained in the off-line phase (Section 3.2) to construct the search space for the NNSS algorithm. We present results of the various analyses that we performed. Unless otherwise indicated, we assume the user to be stationary.

4.1.1 Basic Analysis

For the basic analysis, we use all of the (more than 20) signal strength samples collected for each of the 70*4 = 280 combinations of user location and orientation. In the analysis, we pick one of the locations and orientations at random, and then conduct an NNSS search for the corresponding signal strength tuple in the space defined by the remaining 69 points times 4 orientations. This emulates the process of locating a (stationary) user during the real-time phase. Note that the exclusion of one of the 70 physical points from the search space would result in a somewhat

IPR2024-00397
Apple EX1014 Page 4

worse measured accuracy than would be obtained with the RADAR system in real use. In this sense, the accuracy results presented here are somewhat conservative.

We compare the empirical method with two other methods: *random selection* and *strongest base station selection* [Hod97]. With random selection, we estimate the user's location by picking one of the 70 points at random, regardless of the SS information. With strongest base station selection, we guess the user's location to be the same as the location of the base station that records the strongest signal. A comparison with these simple methods enables us to evaluate how worthwhile the increased sophistication of our techniques is.

Figure 3 shows the cumulative distribution function (CDF) of the error distance for the empirical, strongest base station, and random methods. The empirical method performs significantly better than both of the other methods. Table 1 summarizes the information in the figure in terms of the 25th, 50th (median), and 75th percentile values of the error distance for each method.

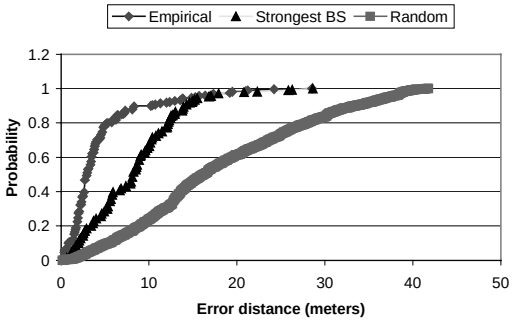

Figure 3 CDF of the error in location estimation.

Considering the median (50th percentile), for instance, the empirical method has a resolution of under 3 meters, which is about the size of an office room in our building. In terms of linear resolution, it is 2.8 times better than the strongest base station method and 5.5 times better than the random method. In terms of spatial resolution, the improvement is even greater: 7.7 and 30.6 times, respectively. We use the percentile values for the empirical method in Table 1 as a basis for comparison in the rest of the analysis.

Method	**25th (meter)**	**50th (meter)**	**75th (meter)**
Empirical	1.92	2.94	4.69
Strongest	4.54 (2.4x)	8.16 (2.8x)	11.5 (2.5x)
Random	10.37 (5.4x)	16.26 (5.5x)	25.63 (5.5x)

Table 1 The 25th, 50th, and 75th percentile values of the error distance. The numbers in parentheses indicate the degradation of the strongest BS and random methods compared to the empirical method.

In summary, the empirical method performs quite well. We now discuss ways of making it perform even better.

4.1.2 Multiple Nearest Neighbors

Unlike the basic analysis where we only considered the single nearest neighbor in signal space, we now consider *k* nearest neighbors, for various values of *k*. The intuition is that often there are multiple neighbors that are at roughly the same distance from the point of interest (in signal space). Given the inherent variability in the measured signal strength at a point, there is fundamentally no reason to pick only the closest neighbor (in signal space) and reject others that are almost as close.

A second, and equally important, reason for considering additional neighbors is that it is likely that the *error vector* (in physical space) corresponding to each neighbor is oriented in a different direction. So averaging the coordinates of the neighbors may yield an estimate that is closer to the user's true location than any individual neighbor is. Figure 4 illustrates this for *k=3* nearest neighbors.

Our experimental analysis of averaging over *k* nearest neighbors shows that for small *k*, averaging has some benefit though not very significant. For instance, for *k=5*, the 25th percentile of error distance is 1.5 m (22% better than the 1.92 m in Table 1) and the 50th percentile is 2.75 m (9% better). For large *k*, accuracy degrades rapidly because points far removed from the true location also are included in the averaging procedure, thereby corrupting the estimate.

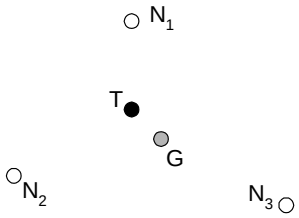

Figure 4 An illustration of how averaging multiple nearest neighbors (N_1, N_2, N_3) can lead to a guess (G) that is closer to the user's true location (T) than any of the neighbors is individually.

The reason why the benefits of averaging are not very significant even for small *k* is that often the *k* nearest neighbors in signal space are *not k* physically distinct points. In many instances, multiple nearest neighbors in signal space correspond to different orientations at the same point in physical space. So averaging in physical space does not improve the location estimate by very much.

4.1.3 Max Signal Strength Across Orientations

Since the dependence of signal strength on orientation creates a challenge for location estimation, we analyze how well the empirical method would perform if orientation were not an issue. For each user location in the off-line data set, we compute the maximum SS at each base station across the four possible orientations at that location[6]. Note that the maximum for each base station may correspond to a different orientation of the user. The goal is to emulate the case where

[6] For each base station and user location, we first compute the mean SS for each of the four orientations at that location, and then pick the maximum among the four means.

the signal generated by the mobile host is *not* obstructed by the user's body. While this may not be realistic given the antenna design and positioning for existing wireless LANs, it may be possible to approximate this "ideal case" with new antenna designs (e.g., omnidirectional wearable antenna)

We repeat the analysis of the previous sections with the smaller "maximum signal strength" data set of 70 data points (instead of 70*4=280 data points in the original data set). In Figure 5, we plot the 25th and the 50th percentile values of the error distance with averaging over neighbor sets of various sizes.

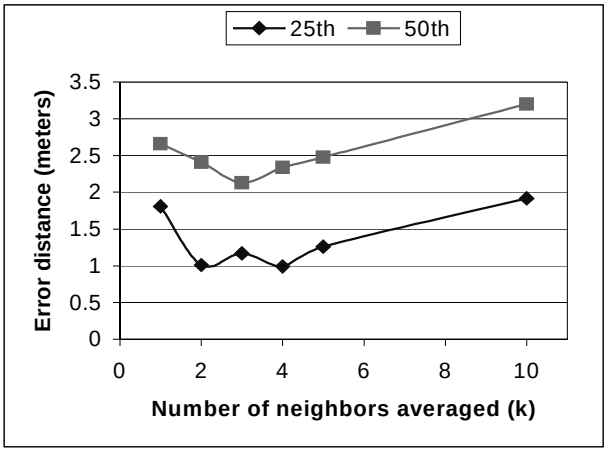

Figure 5 The error distance for the empirical method with averaging on the data set containing the max signal strength measurement for each location.

We make a couple of observations. First, just as expected, the use of the maximum SS data set improves the accuracy of location estimation slightly even in the absence of averaging (*k=1*). The 25th percentile value of the error distance is 1.8 m and the 50th percentile 2.67 m, 6% and 9% better, respectively, compared to Table 1. Second, averaging over 2-4 nearest neighbors improves accuracy significantly; the 25th percentile is about 1 m (48% better) and the 50th percentile is 2.13 m (28% better). Averaging is more effective here than in Section 4.1.2 because the set of *k* nearest neighbors in signal space necessarily correspond to *k* physically distinct locations.

4.1.4 Impact of the Number of Data Points

We now investigate how the accuracy of location estimation would be impacted if we had data from fewer than the 70 distinct physical locations considered thus far.

For each value of *n*, the number of physical locations (ranging between 2 and 70), we conducted 20 runs of our analysis program. In each run, we picked *n* points at random from the entire data set collected during the off-line phase and used this subset to construct the search space for the NNSS algorithm. We collated the error distance data from all the runs corresponding to the same value of *n* (Figure 6).

For small *n* (5 or less), the error distance is a factor of 2 to 4 worse than when the entire empirical set containing 70 physical points is used. But the error distance diminishes rapidly as *n* increases. For *n=20*, the median error distance is less than 33% worse and for *n*=40, it is less than 10% worse. The diminishing returns as *n* becomes large is due to the inherent variability in the measured SS. This translates into inaccuracy in the estimation of physical location. So there is little benefit in obtaining empirical data at physical points spaced closer than a threshold.

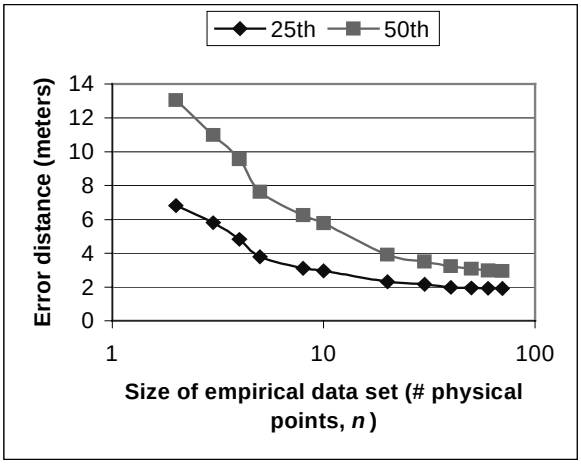

Figure 6 The error distance versus the size of the empirical data set (on a log scale).

In summary, for our floor, the empirical method would perform almost as well with a data set of 40 physical points as with a set of 70 points. In practice, we could make do with even fewer points by picking physical locations that are distributed uniformly over the area of the floor rather than at random.

4.1.5 Impact of the Number of Samples

In the analysis presented so far, we have worked with the mean of *all* of the samples recorded during the off-line phase for each combination of location and orientation. While it may be reasonable to construct the empirical data set with a large number of samples (since it is a one-time task), there may be constraints on the number of samples that can be obtained in real-time to determine a user's location. So we investigate the impact of a limited number of *real-time* samples (while retaining the entire *off-line* data set for the NNSS search space) on the accuracy of location estimation. Our analysis shows that only a small number of real-time samples are needed to approach the accuracy obtained using all of the samples (Table 1). With just 1 real-time sample, the median error distance is about 30% worse than when all samples were considered. With 2 samples, it is about 11% worse and with 3 samples it is under 4% worse.

4.1.6 Impact of User Orientation

As we have already discussed, the user's orientation has a significant impact on the SS measured at the base stations. In Section 4.1.3, we did a best-case analysis using the maximum SS across all four orientations. We now consider, in some sense, the worst case where the off-line data set only has points corresponding to a particular orientation (say north) while the real-time samples correspond to the opposite orientation (i.e., south). We compute the error distance for all four combinations of opposing directions: north-south, south-north, east-west, and west-east.

We observe a fairly significant degradation in the accuracy of location estimation. For instance, for the north-south case, the 25th percentile of the error distance is 2.95 meters (54% worse than in Table 1) while the 50th percentile (median) is 4.90 meters (67% worse). This degradation underscores the importance of obtaining empirical data for multiple orientations to construct the NNSS search space. However, even in this worst case, the empirical method outperforms the strongest base station and random methods by a factor of 2 to 4.

4.1.7 Tracking a Mobile User

In this sub-section, we analyze the problem of *tracking a mobile user* rather than *locating a stationary user*, as we have done so far. For this analysis, we collected a new SS data set corresponding to random walks by the user along the hallways of our floor. We collected 4 signal strength samples per second at each of the base stations. Assuming that the user walked at a uniform pace, we are able to determine the true location of the user at each time instant.

We reduce the problem of tracking the mobile user to a sequence of location determination problems for a (nearly stationary) user. We use a sliding window of 10 samples to compute the mean signal strength on a continuous basis. This information is then used with the basic method (Section 4.1.1) to estimate the user's location on a continuous basis.

The error distance for tracking the mobile user is only slightly worse than that for locating a stationary user. The median error distance is 3.5 meters, about 19% worse than that for a stationary user.

4.1.8 Summary of Empirical Method

The empirical method is able to estimate user location with a high degree of accuracy. The median error distance is 2 to 3 meters, about the size of a typical office room. For our experimental environment, much of the accuracy can be achieved with an empirical data set of about 40 physical points and about 3 real-time signal strength samples (at each base station). It is important, however, that the empirical data set contain data corresponding to multiple user orientations at each location.

The main limitation of the empirical method is that significant effort is needed to construct the SS data set for each physical environment of interest (each floor, each building, etc.). Furthermore, the data collection process may need to be repeated in certain circumstances, e.g., when a base station is relocated.

We now discuss a method based on signal propagation modeling, which avoids these limitations.

4.2 Radio Propagation Model

Radio propagation modeling provides an alternative to the empirical method for constructing the search space for the NNSS algorithm.

4.2.1 Motivation

The primary motivation for the radio propagation model is to reduce RADAR's dependence on empirical data. Using a mathematical model of indoor signal propagation, we generate a set of theoretically-computed signal strength data akin to the empirical data set discussed in Section 3.2. The data points correspond to locations spaced uniformly on the floor. The NNSS algorithm can then estimate the location of the mobile user by matching the signal strength measured in real-time to the theoretically-computed signal strengths at these locations. It is clear that the performance of this approach is directly impacted by the "goodness" of the propagation model. In the following subsections, we develop the model and discuss the performance of location determination based on the model.

4.2.2 Determination of the model

For a radio channel, signal propagation in an indoor environment is dominated by reflections, diffraction, and scattering of radio waves caused by structures within the building. The transmitted signal generally reaches the receiver via multiple paths (termed the *multipath* phenomenon). Multipath causes fluctuations in the received signal envelope and phase, and the signal components arriving from indirect and direct paths combine to produce a distorted version of the transmitted signal. Since multipath within buildings is strongly influenced by the layout of the building, the construction material used, and the number and type of objects in the building, characterizing the radio channel in such an environment is challenging.

We considered three different models before settling on one. The first model was the well-accepted *Rayleigh fading model* [Has93], which describes small-scale rapid amplitude fluctuation in the absence of a strong received component. The Rayleigh distribution is widely used to describe multipath fading because of its elegant theoretical explanation and the occasional empirical justification. However, in deriving this distribution, a critical assumption made is that all signals reaching the receivers have equal strength. In general, this is unrealistic. Our empirical data shows that for a number of sample points (along the hallways), there exists a dominant line-of-sight (LoS) component that is not accounted for by this distribution. Hence, we did not use this distribution.

The second model we considered was *the Rician distribution model* [Ric44]. The Rician distribution occurs when a strong path exists in addition to the low level scattered path. This strong component may be the LoS path or a path that encounters much less attenuation than others. The Rayleigh distribution is a special case of the Rician distribution; when the strong path is eliminated, the amplitude distribution becomes Rayleigh. While the model is intuitively appealing, it is very difficult to determine the model parameters (i.e., the local mean of the scattered power and the power of the dominant component) precisely as this requires physically isolating the direct wave from the scattered components. To keep the system simple and easy to deploy, we decided against using this distribution to model the radio channel.

We found a good compromise between simplicity and accuracy in the *Floor Attenuation Factor* propagation model (FAF) suggested by [Sei92]. We like this model because it provides flexibility in accommodating different building layouts while taking into account large-scale path loss. We

IPR2024-00397
Apple EX1014 Page 7

adapted the original model proposed by Seidel and Rappaport, which included an attenuation factor for building floors, to disregard the effects of the floors and instead consider the effects of obstacles (walls) between the transmitter and the receiver. The *Wall Attenuation Factor* (WAF) model is described by

$$P(d)[dBm] = P(d_o)[dBm] - 10n\log\left(\frac{d}{d_o}\right) - \begin{cases} nW * WAF & nW < C \\ C * WAF & nW \geq C \end{cases}$$

where n indicates the rate at which the path loss increases with distance, $P(d_o)$ is the signal power at some reference distance d_o and d is the transmitter-receiver (T-R) separation distance. C is the maximum number of obstructions (walls) up to which the attenuation factor makes a difference, nW is the number of obstructions (walls) between the transmitter and the receiver, and WAF is the wall attenuation factor. In general the values of n and WAF depend on the building layout and construction material, and are derived empirically. The value of $P(d_o)$ can either be derived empirically or obtained from the wireless network hardware specifications.

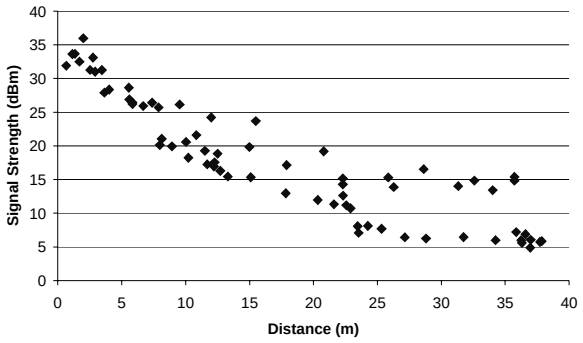

Figure 7 SS as a function of T-R separation derived from the empirical data collected in Section 3.2

Figure 7 illustrates how the signal strength varies with distance between the transmitter and the receiver. The wide difference in signal strengths between points at similar distances is explained as follows: the layout of the rooms in the building, the placement of base stations, and the location of the mobile user all have an effect on the received signal. Signals transmitted from two locations at the same distance from the receiver (base station) may be attenuated by different amounts due to the differences in the number and types of obstructions they encounter. For instance, in Figure 7, we observe that the strength of the signal from two locations approximately 36 meters from the receiver (base station) were approximately 10 dBm apart. This is because there were several walls between one of the locations and the base station, while the other location had line-of-sight to the base station.

Previous work in indoor radio propagation modeling has included extensive characterization of signal loss for various materials and at different frequencies [Rap96]. However, using this information in a practical setting is difficult because the obstructing materials vary considerably in their physical and electrical characteristics. For example, water causes signal attenuation and the human body is made up of water, so the size of a human body and its orientation can result in different amounts of signal loss. It is virtually impossible to characterize such loss precisely since the number and sizes of humans in the building at any time is generally a finite but random number. Thus, despite their complexity, these extensive models fall short in terms of accuracy. So we decided to go with the much simpler WAF propagation model.

We determined the attenuation caused by a wall empirically (described next) and used this in conjunction with the number of intervening walls (determined using the Cohen-Sutherland algorithm mentioned in Section 3.3.2) to compensate for the attenuation caused by the walls.

We conducted the following experiment to determine the Wall Attenuation Factor (*WAF*): we measured the signal strength at the receiver when the receiver and the transmitter had line-of-sight. We then measured the signal strength with varying but known number of walls between the receiver and the transmitter. We computed the average of the difference between these signal strength values to determine the *WAF*. We observed that the amount of additional attenuation dropped off as the number of walls separating the transmitter and the receiver increased. This observation is consistent with [Sei92] where the attenuation between different floors was considered and shown to flatten out as the number of floors between the transmitter and the receiver increased. In general, with a large T-R separation and a large number of intervening walls, free-space path loss dominates the loss due to obstructions. Based on our measurements, we chose *WAF* to be 3.1 dBm and C to be 4 (where C represents number of walls that are factored into the model). Figure 8 shows the result after the measured signal strength has been compensated for signal loss due to the intervening walls between the transmitter and the receiver. We observe that the resulting plot shows a trend similar to the free-space loss trend. This demonstrates that the WAF propagation model compensates effectively for attenuation due to obstructions.

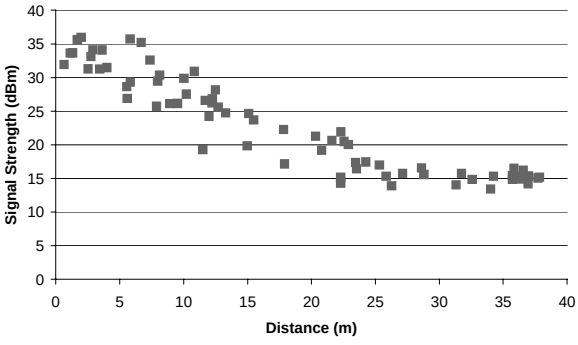

Figure 8 Effect of applying correction for intervening walls between the base station and the mobile user.

After compensating for the effect of walls and creating the "corrected" data for all three base stations, we proceeded to determine the two other parameters, n and P_{do}, of our model. We reduced the propagation model to a form where it exhibits a linear relationship between the theoretically-predicted signal strength and logarithm of the distance

between the transmitter and the receiver, and then applied simple linear regression to determine the parameters of the model [Jai91].

Table 2 contains the numerical values of the model parameters for the three base stations considered separately and when taken together. We note that the values for the path loss exponent (n) and the reference signal strength (P_{do}) for all three base stations are similar despite their different physical locations and surroundings. This result is encouraging since it indicates that the parameter values are not tied to the specific location of the base stations. The values of P_{do} are higher than those published by the manufacturer [Roa96] (for d_0 = 1 meter) because our WAF model does not account for multipath propagation. The values of the path loss exponent are smaller than those reported in previous work on indoor radio propagation modeling [Rap96]. However, they are consistent with our expectations since we compensate the measured signal strength for attenuation due to obstructions, and since we do not consider multipath (which can boost the signal strength at a given location). R^2 represents the *coefficient of determination*, which is a useful measure for indicating the goodness of regression [Jai91]. The high values of R^2 (on a scale of 0 to 1) suggest that there is a good match between the estimated and the measured values of the signal strength. Another value of interest shown in the table is the *mean squared error* (MSE). These numbers reinforce the observation that the WAF propagation model fits the measured data well.

	BS_1	BS_2	BS_3	All
P_{do}	57.58	56.95	64.94	58.48
n	1.53	1.45	1.76	1.523
R^2	0.81	0.65	0.69	0.72
MSE	10.49	13.98	7.34	9.82

Table 2 Parameter estimates using linear regression

The final column in Table 2 shows the values for P_{do} and n when the data from all the transmitter-receiver pairs (i.e., all three base stations) was combined. The motivation for this was to determine a value of P_{do} and n that could be used for all base stations without overly affecting the result. The advantage of using a common value is that it avoids the need for individual measurements of each base station as they are installed in the network, thus greatly reducing the cost of system setup. We can then use these values to estimate the signal strength at various points within the building.

Figure 9 illustrates how the predicted values of the signal strength generated with the propagation model (after compensating for wall attenuation) compares with the actual measurements. We observe a good match between the two. While this plot is for one of the three base stations, plots for the other two base stations exhibit a similar match.

4.2.3 Results using the Propagation Model

To determine the performance of location estimation with the signal propagation modeling method, we used the model to compute the signal strength at a grid of locations on the floor. We then used this data set as the search space for the NNSS algorithm.

Considering the median (50^{th} percentile), the propagation method provides a resolution of about 4.3 m, compared to a resolution of 2.94 m for the empirical method and 8.16 m for the strongest base station method (Table 1). For the 25^{th} percentile the propagation method provides a resolution of 1.86 m compared to 1.92 m for the empirical method and 4.94 m for the strongest base station method.

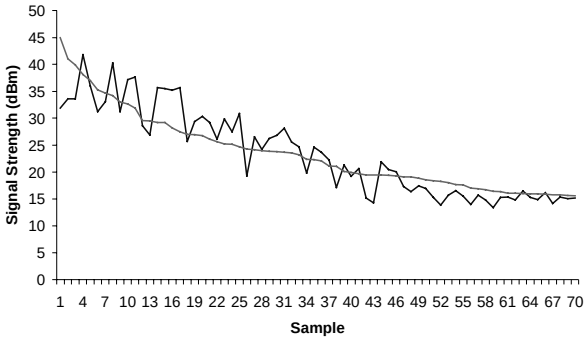

Figure 9 Predicted versus measured signal strength.

While the propagation method is not as accurate as the empirical method, it is significantly better than the strongest BS and random methods. Thus, even without extensive empirical measurements, RADAR based on the propagation model alone would significantly outperform the strongest base station method proposed in [Hod97].

4.2.4 Summary of Radio Propagation Method

The WAF propagation model provides a cost effective means for user location and tracking in an indoor RF wireless network. The model is cost effective in the sense that it does not require detailed empirical measurements to generate a signal strength map and consequently has a low set up cost. A significant result from Section 4.2.2 is that the parameters for the wall attenuation propagation model are similar across base stations despite the latter being in different locations. This suggests that the entire system can be relocated to a different part of the building, but the same parameter values can be used to model propagation and thereby determine a user's location.

5 Discussion and Future Work

We discuss extensions to the RADAR system that would help improve its robustness and accuracy. Due to space constraints, we keep our discussion brief.

We are investigating how user-mobility profiles can supplement signal strength information in locating and tracking users. A profile specifies *a priori* likelihood of user location and/or movement patterns, which can be derived from history [Liu98], calendar information, building layout, etc.

We are also investigating *base station-based environmental profiling* to make RADAR robust in the face

IPR2024-00397
Apple EX1014 Page 9

of large-scale variations in the RF signal propagation environment (caused, for instance, by the varying number of people in a building during the course of a day). Instead of recording just one set of signal strength measurements, we record multiple sets at different times of the day. The base stations probe the channel periodically to determine the current conditions, and accordingly pick the data set that is most appropriate for these conditions.

6 Conclusions

In this paper, we have presented RADAR, a system for locating and tracking users inside a building. RADAR is based on empirical signal strength measurements as well as a simple yet effective signal propagation model. While the empirical method is superior in terms of accuracy, the signal propagation method makes deployment easier.

We have shown the despite the hostile nature of the radio channels, we are able to locate and track users with a high degree of accuracy. The median resolution of the RADAR system is in the range of 2 to 3 meters, about the size of a typical office room.

Our results indicate that it is possible to build an interesting class of location-aware services, such as printing to the nearest printer, navigating through a building, etc., on an RF wireless LAN, thereby adding value to such a network. This, we believe, is a significant contribution of our research.

Our eventual plan is to combine location information services with the RADAR system and deploy this within our organization.

Acknowledgements

We would like to thank Stephen Dahl for his help in setting up our experimental testbed, and the anonymous reviewers for their perspicacious comments.

References

[ATC98] http://www.ascension-tech.com

[Azu93] R. Azuma, "Tracking Requirements for Augmented Reality," *Communications of the ACM*, Vol. 36, No. 7, pp: 50-51, July 1993

[Ber96] S. Berchtold, D. A. Keim, H. P. Kriegel, "The X-tree: An Index Structure for High-Dimensional Data,", Proc. VLDB, 1996

[CG93] T. W. Christ, P. A. Godwin, "A Prison Guard Duress Alarm Location System", Proc. IEEE International Carnahan Conference on Security Technology, October 1993

[Cor90] T. H. Cormen, C. E. Leiserson, R. L. Rivest, "Introduction to Algorithms," The MIT Press, 1990

[Fol90] Foley, van Dam, Feiner, and Hughes, Computer Graphics Principles and Practice, 2nd Edition, Addison-Wesley, 1990

[GPS99] P. Enge, and P. Misra, "Special Issue on GPS: The Global positioning System," *Proc. of the IEEE*, pp. 3-172, January 1999

[Gut84] A. Guttman, "R-Trees: A Dynamic Index Structure for Spatial Searching," ACM, 1984

[Har94] A. Harter and A. Hopper, "A Distributed Location System for the Active Office," *IEEE Network*, January 1994

[Has93] H. Hashemi, "The Indoor Radio Propagation Channel," *Proceedings of the IEEE*, Vol. 81, No. 7, pages 943-968 July 1993

[Hod97] T. D. Hodes, R. H. Katz, E. S. Schreiber, and L. Rowe, "Composable Ad Hoc Mobile Services for Universal Interaction," *MobiCom '97 Proceedings*, pp: 1-12, September 1997

[Jai91] R. Jain, The Art of Computer Systems Performance Analysis, *John Wiley*, 1991

[Liu98] T. Liu, P. Bahl, "Mobility Modeling, Location Tracking, and Trajectory Prediction in Wireless ATM Networks", *IEEE JSAC*, Vol. 16, No. 6, pp. 922-936, August 1998

[Maa97] H. Maass, "Location-Aware Mobile Applications based on Directory Services," *MobiCom '97*, pp. 23-33, September 1997

[Nel98] G. J. Nelson, "Context-Aware and Location Systems," Ph.D. Theses, University of Cambridge, Computer Lab., UK., Jan. 1998

[Ous94] J. K. Ousterhout, "Tcl and the Tk Toolkit," Addison-Wesley Professional Computing Series, 1994

[Ric44] S. O. Rice, "Mathematical analysis of Random Noise," *Bell Systems Technical Journal*, Vol. 23 (1944), & Vol. 24 (1945)

[Rap96] T. S. Rapport, Wireless Communications – Principles and Practice, *IEEE Press*, 1996

[Roa96] "Digital RoamAbout 915/2400 DS/PC Card and ISA Network Adapter: Installation and Configuration," DEC, April 1996

[Sei98] T. Seidl, H. P. Kriegel, "Optimal Multi-Step k-Nearest Neighbor Search," Proc. ACM SIGMOD, 1998

[Sei92] S. Y. Seidel and T. S. Rapport, "914 MHz path loss prediction Model for Indoor Wireless Communications in Multi-floored buildings," *IEEE Trans. on Antennas & Propagation*, Feb. 1992

[Tek98] S. Tekinay, "Wireless Geolocation Systems and Services," Special Issue of the *IEEE Communications Magazine*, April 1998

[Wal96] L. Wall, T. Christiansen, R. L. Schwartz, "Programming Perl," O'Reilly & Associates, Inc., 1996

[Wan92] R. Want, A. Hopper, V. Falcao, J. Gibbons., "The Active Badge Location System," *ACM Transactions on Information Systems*, Vol. 40, No. 1, pp. 91-102, January 1992