

Robotics-Based Location Sensing using Wireless Ethernet

Andrew M. Ladd
Rice University
aladd@cs.rice.edu

Kostas E. Bekris
Rice University
bekris@cs.rice.edu

Algis Rudys
Rice University
arudys@cs.rice.edu

Guillaume Marceau^{*}
Brown University
gmarceau@cim.mcgill.ca

Lydia E. Kavraki
Rice University
kavraki@cs.rice.edu

Dan S. Wallach
Rice University
dwallach@cs.rice.edu

Abstract

A key subproblem in the construction of location-aware systems is the determination of the position of a mobile device. This paper describes the design, implementation and analysis of a system for determining position inside a building from measured RF signal strengths of packets on an IEEE 802.11b wireless Ethernet network. Previous approaches to location-awareness with RF signals have been severely hampered by non-linearity, noise and complex correlations due to multi-path effects, interference and absorption. The design of our system begins with the observation that determining position from complex, noisy and non-linear signals is a well-studied problem in the field of robotics. Using only off-the-shelf hardware, we achieve robust position estimation to within a meter in our experimental context and after adequate training of our system. We can also coarsely determine our orientation and can track our position as we move. By applying recent advances in probabilistic inference of position and sensor fusion from noisy signals, we show that the RF emissions from base stations as measured by off-the-shelf wireless Ethernet cards are sufficiently rich in information to permit a mobile device to reliably track its location.

Categories and Subject Descriptors

C.2.1 [Computer Systems Organization]: Network Architecture and Design—*Wireless communication*; G.3 [Mathematics of Computing]: Probability and Statistics—*Markov processes, Probabilistic algorithms*; I.2.9 [Computing Methodologies]: Robotics—*Sensors*; I.5.1 [Pattern Recognition]: Models—*Statistical*

General Terms

Algorithms, Design, Experimentation

^{*}Work was completed while visiting Rice University.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MOBICOM'02, September 23–26, 2002, Atlanta, Georgia, USA.
Copyright 2002 ACM 1-58113-486-X/02/0009 ...\$5.00.

Keywords

802.11, wireless networks, mobile systems, localization, probabilistic analysis

1. Introduction

There has been great progress in wireless communications over the last decade, causing the available mobile tools and the emerging mobile applications to become more sophisticated. At the same time, wireless networking is becoming a critical component of networking infrastructure. Wireless technology enables mobility which, in turn, creates a need for location-aware applications. The recent interest in location sensing for network applications and the growing need for large-scale commercial deployment of such systems has brought network researchers up against a fundamental and well-studied problem in the field of robotics: determination of physical position using uncertain sensors (localization).

Many mobile devices and many buildings, both commercial and residential, are already equipped with off-the-shelf IEEE 802.11b wireless Ethernet, a popular and inexpensive technology. Furthermore, most wireless Ethernet devices already measure signal strength of received packets as part of their standard operation and the signal strength varies noticeably as the distance and obstacles between wireless nodes change. If a reliable localization system could be developed using only this technology, then many existing systems could be retrofitted in software and new systems could be deployed using readily available parts.

The development of efficient and accurate location-support systems for indoor environments, which would also have the potential of being widely available, is a challenging task. The limitations usually stem from the harsh nature of the signal and the sensors one has to work with. Indoor environments affect the propagation of wave signals in non-trivial ways, causing severe multi-path effects, dead-spots, noise and interference [5]. These effects make it infeasible to construct a simple and accurate model of the signal's propagation in the space. A location support system has to overcome the high uncertainty due to the behavior of the indoor wireless channels but at the same time it should keep the cost and the complexity of large-scale deployment as small as possible.

1.1 Motivation

Location-Awareness. In the wireless world many desirable applications require context-awareness. The context of an application refers to the information that is part of the application's operating environment. Typically this includes information such as location, activity of people, and the state of other devices [18]. Algorithms

and techniques that allow an application to be aware of the location of a device on a map of the environment are a prerequisite for many of these applications.

The growing need for location support systems underscores the importance of addressing location-awareness problem. For example, government initiatives require that cellular phone providers should develop a way to locate any phone that makes an emergency call [12]. In outdoor settings, GPS [29] has been used in many commercial applications, as in the case of locating automobiles. Despite the extraordinary advances in GPS technology, though, many indoor spaces cannot reliably receive GPS signals.

An indoor system must use different sensors, such as infrared (IR), sonar, vision, or radio (RF), to infer position of a mobile device. Location-aware applications based on these sensors could enable users to discover resources in their physical proximity, such as active maps of their surroundings and adaptive interfaces to the user's location [18]. Specific applications of such a system vary from tracking a guard's position in a penitentiary institution [7] to hospitals where equipment and people must be efficiently located [40]. These applications can also be useful in large office environments, where the loss of valuable equipment such as laptop computers has become a serious problem and locating resources such as printers takes time and disrupts other activities.

Wireless Security. We are also interested in the utility of a location support system over an existing wireless network related to security applications. A principal difference between wired networks and wireless networks is that physical security is no longer sufficient to ensure the security of the network. In addition, in a wireless network, the location of an intruder is considerably more difficult to determine versus a traditional wired network where cables can be traced to their source. Notably, a mobile device which is transmitting on a wireless Ethernet network is leaking its position. This information can be used to locate the intruders who make no deliberate effort to decorrelate their signal from their position. Although this can already be achieved using expensive directional antennas, off-the-shelf hardware is less conspicuous and more readily-available.

Mobile Robotics. Many mobile robot platforms make extensive use of wireless networking to communicate with off-line computing resources, other robots, and various user-interface devices. Since the advent of inexpensive wireless networking, many mobile robots have been equipped with 802.11b wireless Ethernet. In many applications, a sensor from which position can be inferred directly without the computational overhead of image processing or the material expense of laser range-finders is of great use. Many robotics applications would benefit from being able to use wireless Ethernet for both sensing and communication. For example, exploration, map-building and navigation with low-cost wheeled robots could be readily achieved using wireless Ethernet and sonar.

1.2 Our Approach

In this paper, we describe a system that achieves robust indoor localization using only RF signal strength as measured by an IEEE 802.11b wireless Ethernet card communicating with standard base stations. Since the required equipment for a wireless Ethernet network is usually already present in the workspace, serving communication purposes, this reduces the cost of providing localization services in an indoor environment. This also reduces the complexity for the user of a mobile device who wishes to take advantage of this localization service. To achieve our goal, we have adapted standard approaches from robotics-based localization, notably the

explicit manipulation of noise distributions and the modeling of position as a probability distribution.

Our method for localizing a mobile station is divided in two phases. Initially, there is a training phase, where a sensor map of the environment is built by sampling the space and gathering data at various predefined checkpoints of the indoor environment. Later, the operator of a mobile computer walks in the same workspace and the system locates and tracks the operator's position. Our system currently assumes that the environment remains consistent from training to localization. In particular, we assume that people are minimally present when we attempt to localize.

Section 2 presents the algorithms and methodology for our localization system. The results of our experiments are reported in Section 3 and a discussion of our work is presented in Section 4. In Section 5, we discuss related work in the fields of location-aware computing and robot localization.

2. Methodology

In this Section, we discuss our methodology for determining a user's location using wireless network signal strength. We begin by discussing the platform and environment we considered. We then discuss RF signal propagation and describe some problems with devising a signal attenuation model for wireless Ethernet. Finally, we discuss our algorithms for determining the user's location.

2.1 Experimental Setup

Hardware. Our experiments were conducted by a human operator carrying a HP OmniBook 6000 laptop with a PCMCIA LinkSys wireless Ethernet card. This particular card uses the Intersil Prism2 chipset. We modified the standard Linux kernel driver for this card to support a number of new functionalities, including the scanning and recording of hardware MAC addresses and signal strengths of packets, using promiscuous mode, and the automatic scanning of base stations.

We needed a constant source of signal from all base stations for optimum results. Unfortunately, this meant we could not simply be a passive observer. While we could simply put the network interface adapter into promiscuous mode and listen to all packets being transmitted by base stations, this can only guarantee a stream of packets from one base station: the one that the card is currently associated with. While base stations do send out beacon packets several times a second, we could not get access to this signal using our hardware.

Instead, we were forced to use the base station probe facility of 802.11 [23]. Client nodes can broadcast a probe request packet on a wireless network. Base stations that receive such a request respond with a probe response packet. The client then collects these packets and, judging by the strengths of the incoming signals, can determine the closest base station to connect to. We analyze these signal strengths to determine our location relative to the base stations.

A given base station can appear anywhere between zero and four times in the packets the firmware returned to us. For each packet, we get an eight-bit reading representing the signal strength. This value is computed by the network card, and we have no way of determining or affecting how it is calculated. Unless the sender is very close to the receiver, signals in the top half of this range rarely occur. Certain other signal strengths simply never occur. The lowest order bit tends to be very noisy. When compared to other sensors, such as sonar, this signal is very thin: at most 5 usable bits of signal per packet.




Figure 1: Map of the region of the Duncan Hall where we conducted our tests. Base stations are indicated by circles on the map. Note that additional base stations outside of this region (including on other floors) were used in our experiments.

Building Geometry. We operated on the third floor of Duncan Hall at Rice University, in the four hallways shown in Figure 1. The two longer hallways (hallways 1 and 2) measure 32.5 meters, and the two shorter hallways (hallways 3 and 4) measure 17 meters. Hallway 1 has a base station near one end, and hallway 2 has a base station really close to the middle. Hallways 3 and 4 are notable in that they are open above and either partially (in the case of hallway 4) or totally (in the case of hallway 3) open on the sides.

There were nine base stations distributed on this floor. Those within the area described by the map in Figure 1 are marked with circles. The base stations were Apple AirPort base stations and were mounted between two and three meters off the ground. We had a fairly precise map of the building that we had processed to mark off free space and obstacles. The pixel resolution was roughly six centimeters in this map.

2.2 RF Signal Propagation in Wireless Ethernet

The IEEE 802.11b High-Rate standard uses radio frequencies in the 2.4 GHz band, which is attractive as it is license-free in most places around the world. The available adapters are based on spread spectrum radio technology, where the information signal is spread over several frequencies [9], so interference on a single frequency does not block the signal.

The main problem with this sensor is that an accurate prediction of the signal's strength in every position of the environment is a very complex and difficult task because the signal propagates in many unpredictable ways [31]. The received signal is further corrupted by unwanted random effects such as noise, interference from other sources and interference between channels.

As waves propagate through an environment, the environment scatters the waves in a variety of different ways. Reflection, ab-

sorption, and diffraction occur when the waves encounter opaque obstacles; refraction occurs when the waves encounter translucent obstacles. Scattered waves can either decrease or increase the signal strength at the reception point. Changes in atmospheric conditions like air temperature can also affect the propagation of waves and the resulting signal strengths. Unfortunately, 2.4 GHz is a resonant frequency of water, so people absorb radio waves in the 2.4 GHz frequency band that we are using.

Interference occurs when another radio frequency source generates a signal at the same frequency that is of comparable or higher strength than the transmitted signal, as measured by the recipient. The interfering device does not need to be a radio based transmission device [9]. In the 2.4 GHz frequency band, microwave ovens, Bluetooth devices, 2.4 GHz cordless phones and welding equipment can be sources of interference.

Due to reflection, refraction, diffraction, and absorption of radio waves by structures and people inside a building, the transmitted signal often reaches the receiver by more than one path, resulting in a phenomenon known as multi-path fading [20]. The signal components arriving from indirect paths and the direct path, if this exists, combine and produce a distorted version of the transmitted signal.

These difficulties are particularly acute when operating indoors. Since there is rarely a line of sight between the transmitter and the receiver, the received signal is a sum of components that are often caused by some combination of the previously described phenomena.

The received signal varies with respect to time and especially with respect to the relative position of the receiver and the transmitter. However, signal profiles corresponding to spatially coupled locations are expected to be roughly similar as the various external variables remain approximately the same over short distances [20].




Figure 2: Samples of signal strength taken at the same positions facing opposite directions




Figure 3: Examples of signal strength distributions of two different base stations, measured simultaneously from one location

The local average of the signal varies slowly with the displacement. These slow fluctuations depend mostly on environmental characteristics and are known as long-term fading.

While much effort has been made to model radio signal propagation and attenuation in indoor environments, no single consistent model is available. During our initial experiments we took numerous measurements at various positions in our environment. Our objective was to try to see if the variables in the system could be captured with a simple theoretical model to minimize the training phase. We observed a number of interesting properties of RF signals in our environment.

Orientation Matters. The authors of RADAR established a correlation between orientation and measured signal strength [3, 2]. We also observed this. The laptop and the operator affect the signal in a measurable way. It is interesting to note that the presence of the operator affects signal strength and gives the omnidirectional signals some weakly directional properties. Typically the mean signal strength varies less than the statistical distribution of signal strengths. In Figure 2, we give an example of two distributions sampled at the same points while facing in opposite directions.

Noise Distribution Non-Gaussian. The noise distributions at a fixed position were very heterogeneous as we varied the pose and base station that we sampled. In Figure 3, we show two typical examples of the signal strength at different base stations measured simultaneously at the same physical position. Several hundred samples were taken in about forty five seconds for these particular histograms. Notice that the first-order properties of these distributions differ greatly from each other. In general, we have observed that distributions were asymmetric and had multiple modes. There is usually a dominant mode which often differs from the mean. We

concluded that distributions were essentially non-Gaussian. Since the noise behavior is an extremely complex physical phenomenon and explicit histograms are fairly compact, we decided that it would be better to work directly with these distributions rather than reduce the data to average values.

We found it useful to postprocess the sampled distributions by applying a small window smoothing convolution, adding a very small uniform baseline distribution and then normalizing. This is done to try to artificially compensate for sampling errors and allow for a small probability of unexpected measurements in the Bayesian inference calculations that follow. These corrections produced minor but noticeable improvements in the precision of the calculations.

2.3 A Bayesian Inference Algorithm

We model the world as a finite space $S = \{s_1, \dots, s_n\}$ of states with a finite observation space $O = \{o_1, \dots, o_m\}$. The sensor model is some learned or predicted model of the conditional probabilities of seeing some observation o_j at state s_i , in other words $Pr(o_j|s_i)$. A state vector π is a probability vector (distribution) over the various states.

Position is represented as a probability distribution over the states. The inference calculation consists of conditioning on the observations and then selecting a representative point from the resulting distribution.

Given a prior estimate of our state, π , we can construct a new estimate of our state, π' , after observing o_j by calculating the individual conditional probabilities π'_i for each $1 \leq i \leq n$ using Bayes' rule,

$$\pi'_i = \frac{\pi_i \cdot Pr(o_j|s_i)}{\sum_{i=1}^n \pi_i \cdot Pr(o_j|s_i)}.$$

This is a simple principal on which probabilistic inference schemes are built. Of course, the devil is in the details. To implement our system we made several design decisions. We first chose appropriate state and observation spaces. This involved deciding on a sampling granularity for both spaces. We then learned the conditional probability distributions for plugging into the formula above.

2.3.1 Our Model

Our initial experiments and literature search indicated that a priori models of RF signal propagation would be difficult to set up without some on-site training. After verifying that simple assumptions such as fitting analytic curves and surfaces to the means and Gaussians or other simple distributions to the variances provide poor fits to sampled data, we opted for the simpler, more robust scheme of sampling the conditional probabilities directly. The reasoning for this is discussed further in Section 2.1.

We began by defining our state space. We chose a set of points on the map, each tuple (x, y, θ) a location and orientation on the floor of Duncan Hall where our experiments took place. There is no indication that adding an additional parameter for three-dimensional localization would be any harder, although we did no experiments to verify this. Our state space S consisted of a set of n points

$$S = \{s_1 = (x_1, y_1, \theta_1), \dots, s_n = (x_n, y_n, \theta_n)\}.$$

Each observation in our observation space consisted of the measurements that occurred in a single scan from our base station scanner. Each base station scan returned a set of k base station signal strength measurements. Each base station could appear in the scan up to four times. We represent each observation as a vector

$$o = \langle k, f_1, \dots, f_N, (b_1, \lambda_1), \dots, (b_k, \lambda_k) \rangle$$

where k is the total number of base station signal strength measurements, N is the total number of unique base stations represented, f_i is the frequency count for the i th base station, b_j represents the base station in the j th measurement and λ_j is the signal strength of that measurement.

In the training phase, at each point s_i , we take an observation. For each base station we build two histograms for that point. The first is the distribution of frequency counts over the sampled observations. The second is a distribution of observed signal strengths. Based on this training, we can calculate two conditional probabilities. $Pr(f_j = a|s_i)$ is the probability that the frequency count for the j th base station is a when we are at state s_i . $Pr(\lambda_j|b_j, s_i)$ is the probability that base station b_j has signal strength λ_j at state s_i . By multiplying these conditional probabilities we obtain the conditional probability of receiving a particular observation. For $o = \langle k, f_1, \dots, f_n, (b_1, \lambda_1), \dots, (b_k, \lambda_k) \rangle$, we compute

$$Pr(o|s_i) = \left(\prod_{j=1}^N Pr(f_j|s_i) \right) \cdot \left(\prod_{j=1}^k Pr(\lambda_j|b_j, s_i) \right).$$

We note that one observation is typically enough information to decide on one's position. However, errors in the training phase can lead to inaccuracy during localization. Significant causes of such error are subsampling and time-dependent phenomena. Subsampling can create a posteriori model of the noise as measured at that point. Certain measurements that occur rarely may never occur in the subsample. When the measurement occurs online, the hypothesis can be rejected entirely based on a conditional probability of zero for that position. We describe heuristics compensating for this difficulty in Section 4.

After trying several possible schemes, we decided to solve a global localization problem for each observation rather than keep a

running estimate because each observation usually contains enough information to get a good guess of our position. The resulting stream of guesses can be combined in a post-processing step to create a more refined estimate of position. One such mechanism is described in Section 2.4.

The exact calculation proceeds as follows: before each observation we choose our prior state distribution π as the uniform distribution. This is a common Bayesian assumption; we assume we are lost so every position is equally likely. This provides a conservative estimate of our location; any attempt to bias this initial estimate may inhibit accurate localization right from the start. When we make the observation, we simply use Bayes' rule to compute π' , the probability distribution over the states. Then it is simply a matter of choosing appropriate candidate locations.

2.4 Sensor Fusion

We used a post-processing technique called sensor fusion to refine our initial location estimate. Sensor fusion is the process of combining multiple independent observations to obtain a more robust and precise estimate of the measured variables. We implemented a filter which takes the output of the inference engine as a stream of timed observations and tries to stabilize the distribution by noting that a person carrying a laptop typically does not move very quickly. It also takes into account some probability of error on the part of the inference engine.

We model a moving operator trying to track her position as a hidden Markov model (HMM). We use a more finely discretized state space than the Bayesian inference engine and try to interpolate our position out of the stream of measurements coming from the inference engine. This design decision was made after noticing that naive averaging of the inference engine's output produced results with twice the precision we expected for points where we had not taken any training samples.

For our purposes, an HMM is a set of states $S = \{s_1, \dots, s_n\}$, a set of observations $O = \{o_1, \dots, o_m\}$, a conditional probability function $\lambda : S \times O \rightarrow [0, 1]$, and a transition probability matrix A . Each state and each observation is a point (x, y, θ) .

The transition probability matrix semantics describe how the system being modeled evolves with time. In this case, it describes how a person travels through the state space. If π is a probability distribution over S , then $\pi' = A\pi$ is the probability distribution after some discrete time step. The idea is that the random state change occurs "hidden" from the observer. We generate the transitional probability matrix A using a relatively simple heuristic, that people don't travel too fast or change directions too frequently.

The observation function λ has semantics identical to observation in the Bayesian inference of position. $\lambda(s, o) = Pr(o|s)$, the probability of observing o while at s . The conditional probability function λ is also defined using a relatively simple heuristic; smaller distances from an observation to a given state lead to higher probabilities of making that observation at that state. As each observation arrives, λ is used to update the probability of being in a given state in S , and then A is used to transition states. If λ accurately models the behavior of the inference engine and A accurately models the behavior of a person transitioning from state to state, the sensor fusion will have superior results to Bayesian inference alone.

3. Results

In this Section we describe several experiments which try to objectively measure the precision and reliability of our system. We




Figure 4: Bulk cumulative error distribution for 1307 packets over 22 poses in a hallway localized using the position of maximum probability as calculated by direct application of Bayes' rule.




Figure 5: Bulk cumulative error distribution for 1465 packets over 22 poses in a hallway localized using the position of maximum probability as calculated by merging distributions over a one second window.

first present the results for static localization. We then describe the results for user tracking using sensor fusion.

Our system was trained by taking samples at various points in the world, as discussed in Section 2.3.1. The amount of data taken at each point is varied adaptively according to a simple heuristic which measures the rate of convergence to a stable distribution. Once the sampled distribution at each visible base station had converged beyond a threshold, we halt the process. This allowed us to adaptively determine how much sampling is necessary as a function of variation in the signal. In our case, usual sampling times ranged from ten seconds to about a minute.

3.1 Static Localization in a Hallway

This subsection describes experiments executed in hallway 1 of our test area (see Figure 1), which was sampled in two different orientations at every 5 feet. The purpose of this was to test the precision of the Bayesian inference localizer. Timed tests occurred at various

positions and at both orientations in the hallway and bulk statistics were calculated.

The training data was taken by two different operators, with each operator training the localizer in one of the two directions. All experiments were executed by a third operator. The purpose of this was to demonstrate a degree of operator-independence.

Basic Bayesian Inference. Using the algorithm described in Section 2.3, we measured a total of 1307 packets over both orientations on 11 different positions. The positions were spread every 10 feet to be exhaustive. The algorithm reported positions back discretized to 5 feet. In Figure 4, we show the cumulative probability of obtaining error less than a given distance. We have observed that error is within 1.5 meters with probability 0.77.

Simple Averaging Improves Results. In the second experiment, we post-processed the probability distributions computed by Bayesian inference with the following simplistic sensor fusion transform: for each $1 \leq i \leq n$, where n is the number of states,

$$\pi'_i = (\pi_i + u_1) \cdot (\eta_i + u_2).$$

π is the prior distribution on position, π' is the revised distribution, η is the probability distribution computed with the algorithm of Section 2.3 and u_1, u_2 are small constants representing artificial uniform distributions. The resulting distribution π' needs to be normalized after this calculation.

This simple calculation improved our results significantly and is usable as a tracker. Our results are summarized in Figure 5. The measured error was within 1.5 meters with probability 0.83. This is an 8 percent improvement over the raw filter. As a tracker, we observed that it lagged behind the actual position and we attempted to improve our results by using more sophisticated methods described in the next section.

Operator Bias. The above results, with training and experimenting done by different people, tend to suggest that operator bias tends to be less significant than sampling error and time dependent effects. In particular, operator bias is not so significant as to cause the results to be unstable.

3.2 Experiments with Tracking

We attempted to improve these results by implementing a more sophisticated sensor fusion based on a hidden Markov model (HMM), as described in Section 2.4. We then walked round-trips of the four hallways in our test area, shown on the map in Figure 1, tracking our current position and recording the output of both the static localization as described in Section 3.1 and the sensor fusion. The results are shown in Figures 6 through 9.

For hallways 1 and 2, sensor fusion increased by 44% and 40%, respectively, the probability of error less than one meter. The traces show that while static localization is good at tracking, sensor fusion improves the results by effectively ignoring outliers. See Figures 6 and 7 for the results on these hallways.

The results for hallway 4 was somewhat more disappointing. The probability of error less than one meter was increased by a scant 8%. Sensor fusion loosely tracked actual movement, but the signal from the static localizer was too noisy to allow for the level of accuracy achieved on hallways 1 and 2. We attributed this noise to the fact that this hallway was open. See Figure 8 for the results on this hallway.

The worst result was on hallway 3, which is entirely open on one side. The probability of error less than one meter actually went down by 10%. As seen in Figure 9, the static localizer for the most

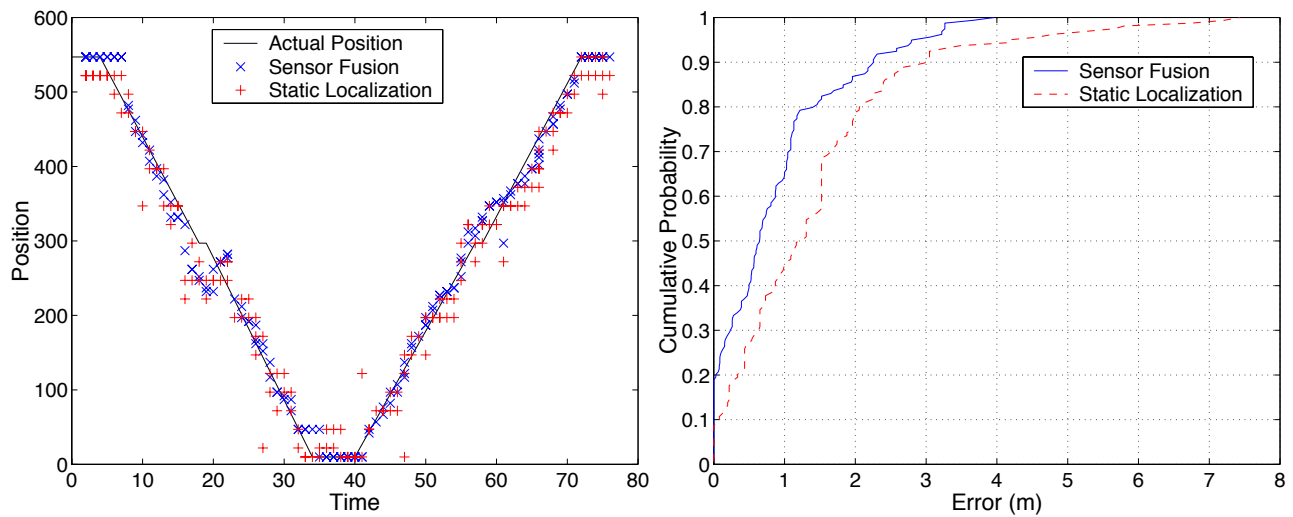


Figure 6: Tracking a round-trip walk of hallway 1 in our test area (see Figure 1 the building map). Measured error for the track, shown on the right graph, is within one meter with probability 0.64, an improvement of 45% over static localization. This improvement is illustrated in the actual tracking performance, shown in the left graph. Position in the left graph is measured in pixels on our map; 50 pixels is approximately equal to 3 meters.

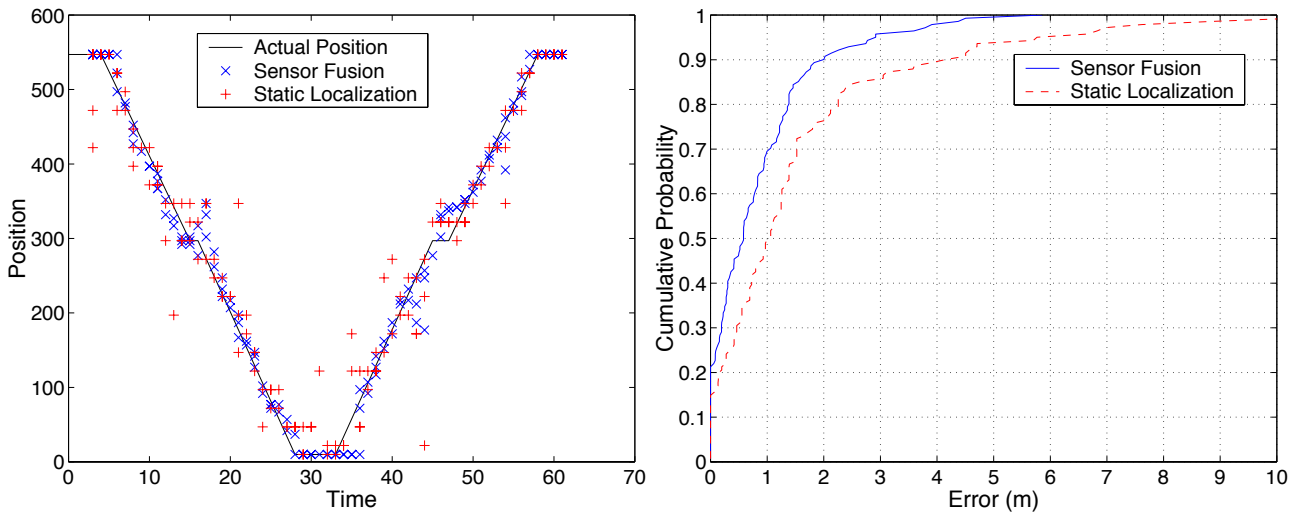


Figure 7: Tracking a round-trip walk of hallway 2 in our test area (see Figure 1 the building map). Measured error for the track, shown on the right graph, is within one meter with probability 0.7, an improvement of 40% over static localization. This improvement is illustrated in the actual tracking performance, shown in the left graph. Position in the left graph is measured in pixels on our map; 50 pixels is approximately equal to 3 meters.

part tended to choose either an endpoint or one of two particular points in the middle of the hallway. This was caused in part by the fact that this hallway is exposed to a large open area, diluting the signal. In addition, all base stations are some distance off to a side, which means our distance (and thus the signal strength) to these base stations does not vary much as we walk the hallway.

Note that the conditional probability function and transition probability matrix we used to initialize the hidden Markov model were generated based on Gaussian distributions. While these were good fits for hallways 1 and 2, they failed to model the noisiness of the static localizer on hallways 3 and 4. A conditional probability function trained to the actual points would likely provide better results.

4. Discussion

The probabilistic robotics-based location-support method with RF-signals that has been described in this paper efficiently reports and tracks the two dimensional position and orientation of a mobile wireless device in an indoor environment. While this is not the first application of probabilistic techniques to the field of location-aware computing, it is one of the first application of such techniques for wireless computing in an indoor environment with commodity hardware. This and the rigorous application of state-of-the-art techniques borrowed from robot localization are the main contributions of this paper. Our work provides a strong indication that localization can be achieved with widely available and inexpensive 802.11b wireless Ethernet hardware. This section will discuss some advantages and disadvantages of our techniques.

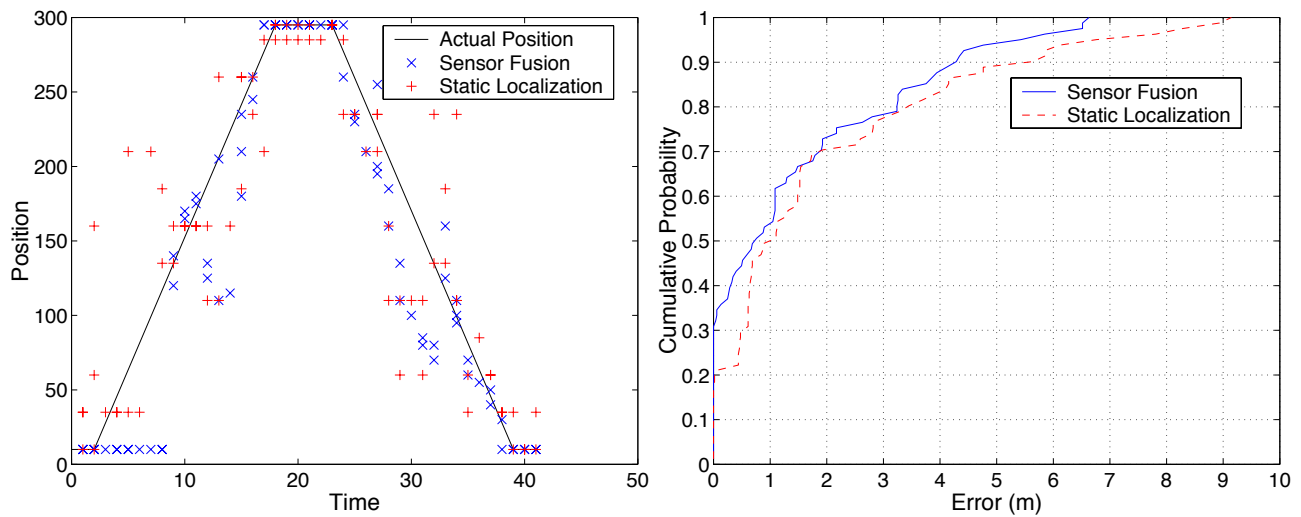


Figure 8: Tracking a round-trip walk of hallway 4 in our test area (see Figure 1 the building map). While sensor fusion provided some improvement, it was not significant. As shown in the left graph, when static localization was significantly off, so was sensor fusion, but when static localization appears to track actual movement, sensor fusion is surprisingly accurate despite the noise. Position in the left graph is measured in pixels on our map; 50 pixels is approximately equal to 3 meters.

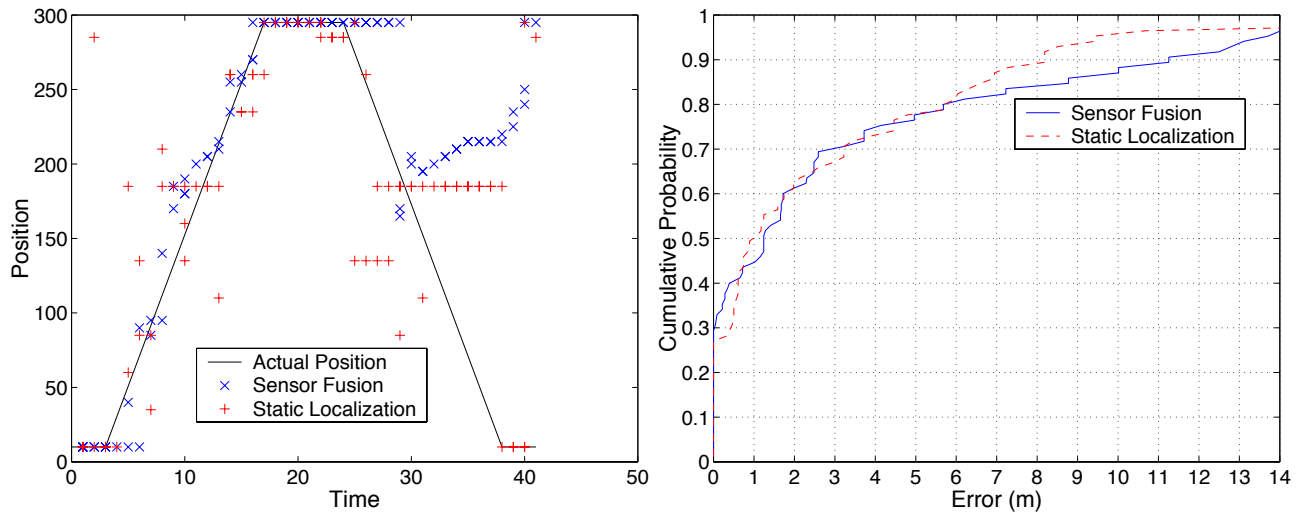


Figure 9: Tracking a round-trip walk of hallway 3 in our test area (see Figure 1 the building map). Sensor fusion did not provide a significant improvement in error, and at times increased error, as shown in the right graph. However, as shown in the left graph, the raw data was already extremely noisy in this case. Position in the left graph is measured in pixels on our map; 50 pixels is approximately equal to 3 meters.

4.1 Advantages

Accuracy. The accuracy of RF based localization is substantially improved in our experimental setup over the reported resolution and accuracy of similar previous efforts. RADAR [3, 2] exhibited a median resolution in the range of 2 to 3 meters. Our results indicate that we can get a resolution of less than 1.5 meters with an accuracy 83% given suitable base station layout. At a coarse resolution, we are very reliable. This is because noise texture varies significantly over relatively large distances, especially when there are intervening obstacles. Inside a room, there are ambiguities in sensing that lead to error. In all of our experiments, we never observed coarse granularity errors except at corners and doorways where the operator is transitioning from one area to another. Our sensor fusion can

improve precision while tracking a moving object by interpolating between sampled points and taking advantage of spatial continuity assumptions to probabilistically reject outliers.

Orientation. Our method explicitly tries to solve for orientation. This is necessary since as we and others [3, 2] have observed orientation is a factor in observed signal strength. In fact, our experiments show that orientation can be coarsely determined by signal strength variations which shows the correlation is often highly non-trivial. By explicitly modeling position and direction, we greatly improve static localization and sensor fusion although orientation determination tended to be much noisier than position. This allows us to overcome difficulties that weakened the applicability of the results of RADAR. On the other hand, we strongly believe the vari-

ations in signal due to orientation are not sufficiently large to ever obtain more than a coarse estimate of direction without employing differential methods with a moving observer.

Cost and Complexity. The advantage of using wireless Ethernet RF signals for localization is that the sensor doubles as a communication device. The infrastructure for such networks already exists in many real-world environments and consequently, for many mobile devices, this sort of localization can be implemented as a software-only solution. This is an attractive option for a number of real-world applications.

Extensibility and Scalability. The methods we use are very general and experiments with a variety of robot localization applications have proven the approach very adaptable. In particular, the framework can be used with other sensors. For example, by using ultrasound sensors such as those used in Cricket [32], we estimate that we could increase our precision to the order of twenty centimeters. This increase in precision is alluded to by the authors of Cricket as a point of future work when they suggest employing Kalman filters [32, 33].

We believe that localization with wireless Ethernet signal strengths scales well into much larger arenas than our experimental test-bed with the caveat that the layout of base stations should be non-pathological. Our evidence for this comes from robot localization and the experimental observation that, at room granularity, signal strength distributions differ greatly.

The particular algorithms we present do not scale if used verbatim. The computational cost of localization in the algorithms we present grows as a linear (Bayesian) or quadratic (sensor fusion) function of the number of possible poses. The vectors and matrices involved however are almost always very sparse. The typical approach in larger cases is to proceed by Monte Carlo (MC) integration of the conditional probability distributions [37]. The computational efficiency of MC is validated by the successful implementation of these algorithms for mobile robots with severely restricted computational power such as the Sony AIBO robot [30].

Privacy and Security. It has been claimed in previous works, such as Cricket [32, 33], that a location support system can be implemented in such a way as to localize a user only if she is willing to be localized. This assertion, though, breaks down if the mobile device is not passive, for example if it is using an active localization scheme or is using wireless networking to communicate. This raises issues of anonymity, privacy, and security. Third-party observers using conventional hardware could conceivably determine the position of a mobile device broadcasting on a wireless Ethernet network without the device's knowledge or permission. Likewise, a network administrator could use the network to track users by having the base stations monitor observed signal strengths.

4.2 Disadvantages

Environment Dependence. Every localization system is hampered by a dependence on the environment it is executed in. In our case, we noticed that some of the areas we tested, notably hallway 3, provided lower accuracy than other areas. The placement of the base stations, the materials in the building, and the building's geometry can affect the difficulty of localizing at a given point. A more worrisome challenge is the variation induced by people absorbing RF signals and other dynamic effects. When working with 2.4 GHz RF signals both static and dynamic environmental conditions can be difficult to predict and have complex behaviors. We believe that continued research on heuristics for coping with these

problems either by judicious placement of base stations or by improvements in the localization algorithm can produce usable results for many applications even in the face of such environmental flux.

Training. The complexity of indoor RF signal propagation is avoided by building a sensor map. The time spent training these maps is a limitation of all localization approaches using a sampling technique for generating maps. As it is, maps were built by marking the workspace and taking measurements at each point. Further automation might be necessary to facilitate deployment of an approach in this spirit. In mobile robotics, map building and exploration for such localization approaches is an important area of research. By augmenting the operator with some extra sensors, for example an accelerometer and magnetic compass to use for dead reckoning, a walk around the building could be used together with a mapping algorithm [36] to automate training further.

4.3 Future Work

This work can be extended in a number of different directions. Most directly, we could expand the experimental area, possibly considering multiple floors and significant amounts of area within rooms. There are also a number of algorithmic aspects of mobile node location tracking that could be explored.

Compensating for dynamic occlusion in robotics localization is a studied problem but is also quite difficult. Many approaches try to predict some variables describing dynamic state. For example, a tour-guide museum robot needs to model the motion of people in the museum to avoid collisions [4]. Multi-robot, collaborative localization is another branch of localization research [14]. Much of the work in this area is relevant to collaborative localization in an ad hoc wireless network. This is a fascinating problem which mixes issues in protocol design and communication with uncertainty and localization. Relative and differential techniques may be of use in combating variations that occur due to environmental effects. For example, landmark based navigation operates using only the angle of deflection to the base station [1]. Pursuit-evasion robotics studies the problem of capturing an active evader under various sensing and environmental constraints. In location-aware security for wireless networks, studying how to intercept a moving intruder under various assumptions about sensing could be an interesting and challenging problem.

5. Related Work

5.1 Location Aware Computing

Many other systems have been built to support indoor localization. These systems vary in many parameters, such as the sensors, the cost, the required hardware, the infrastructure and the resolution in time and space [21].

The AT&T Cambridge Laboratory's Active Badge location system [38] and the more recent Active Bat system [39] are two of the first systems in the field. Active Badge uses diffuse IR technology while Active Bat uses an ultrasound time-of-flight technique to provide accurate physical positioning. Users and objects have to carry Active Bat tags, emitting an ultrasonic pulse to a grid of ceiling-mounted receivers and a simultaneous "reset" signal over a radio link. Each ceiling sensor measures the time interval from reset to ultrasonic pulse arrival and computes its distance from the Bat.

The Cricket Location Support System [32, 33] also uses ultrasound emitters and embeds low-cost receivers in the object being located. Cricket uses additional radio frequency signals to synchronize time measurements and to distinguish ultrasound signals that

are a result of multi-path effects. The main localization techniques that are employed in Cricket are based on triangulation relative to the beacons. Cricket trades accuracy for simpler hardware and infrastructure. It does not require a grid of ceiling sensors with fixed locations as in the Active Bat system but returns an estimation of the user's position with a possible error of a four foot by four foot region, while the Active Bat has an accuracy of nine centimeters. Both of these systems provide excellent localization primitives by employing specialized hardware.

Computer vision has also been used in location support systems. Microsoft Research's Easy Living uses stereo-vision cameras to measure three-dimensional position in a home environment [25]. Camera-based approaches are expensive in terms of hardware infrastructure due to the cost of the camera and the computational overhead of image processing.

RF-Based Systems. The RADAR system [3, 2] uses only a wireless networking signal, employing nearest neighbor heuristics and other pattern recognition techniques for localization. The authors report localization accuracy of about 3 meters of their actual position with about fifty percent probability. They also discuss the problems of localizing in the face of multiple floors and changing environmental conditions, as well as tracking of moving users. While our work has similar design goals to RADAR, we take a very different algorithmic approach, using a probabilistic technique popular in many robotics applications.

The PinPoint location system [40] is similar to RADAR, but uses expensive, proprietary base station and tag hardware to measure radio time of flight. PinPoint's accuracy is roughly 1 to 3 meters.

In the SpotOn system [22], special tags use radio signal attenuation to estimate distance between tags. The aim in SpotOn is to localize wireless devices relative to one another, rather than to fixed base stations, allowing for *ad-hoc localization*. The probabilistic framework we are proposing could also be applied in the case of ad-hoc location sensing.

A number of systems have been built using probabilistic techniques to determine location based on RF signal strength for cellular telephone systems. Liu et al. [28] use Markov modeling and Kalman filtering to predict when a mobile node will cross cell boundaries. Yamamoto et al. [41] use Bayesian analysis to determine the absolute location of a mobile node.

RF Signal Attenuation. Much effort has been made to model radio signal propagation in an indoor environment [17, 31]. Different experiments in the literature have arrived at different distributions. Although each result may be justifiable for a certain set of conditions that govern a certain set of measurements, a consistent model that would give a signal strength distribution under a diversified set of conditions is unavailable. However, experiments with 12000 impulse response profiles in two office buildings have shown good log-normal fit [19]. A general empirical model [17] for indoor propagation of radio signals can be expressed as

$$PL(d) = PL(d_0) + 10 \cdot n \cdot \log\left(\frac{d}{d_0}\right) + X_\sigma$$

where $PL(d)$ is the path loss in dB at distance d , $PL(d_0)$ is the known path loss at the reference distance d_0 , n denotes the exponent depending on the propagation environment and X_σ is the variable representing uncertainty of the model. We note that decibels are a log-scale.

Based on this general formulation, many empirical models have been derived in the field of indoor propagation modeling in the

wireless community. Parameter n is very sensitive to the propagation environment, like the type of the construction material and type of the interior [31], limiting the value of these models.

5.2 Robot Localization

Robot localization is a well-studied problem in robotics. Robot localization is the process of maintaining an ongoing estimate of a robot's location with respect to its environment, given a representation of this environment and some sensing ability within the environment. The importance of this problem in the context of building reliable robot systems cannot be overstated; determining the pose (position and orientation) of the robot from physical sensors is often referred to as "the most fundamental problem to providing a mobile robot with autonomous capabilities" [8]. In our case, we can consider any wireless device as a mobile robot.

If there is no a priori estimate of the robot's location, the problem is referred to as *global localization*, which is a particularly challenging case of localization. This is the type of problem we discussed. We have no information where the wireless device is before it starts communicating with the network's base stations. Furthermore, there is the need to refine the estimate of the device's pose continuously. This task is known as *pose maintenance*.

Sensor based localization is based on the premise that we use sensor data in conjunction with the representation of the environment to produce a refined position estimate, such that this estimate is more likely to predict the true positions. By sensor, we mean any device which can measure attributes of the environment in a way that can be correlated to position. Typical sensors that are deployed in robotics are IR transmitters, ultrasound or laser proximity sensors and camera images.

Sensor fusion is another important notion in robot localization. A broad definition of sensor fusion is the combination of multiple independent observations to obtain a more robust and precise estimate of the measured variables. This can be implemented in terms of integrating sensor readings over time or in the synthesis of measurements from multiple sensors. Most of the recent work in robot localization has been in improving and implementing sensor fusion for many systems.

Much progress has been made in developing localization techniques since the problem first appeared in the literature. Dead reckoning can be used for pose maintenance, but requires some initial knowledge of location. Some of the simplest methods for global localization include landmark-based localization and triangulation. Probabilistic techniques, such as Kalman filtering, and later, Bayesian analysis, were developed to address flaws in these systems. Finally, for when a grid-based map is inappropriate to the application or environment, topological approaches have been developed.

Dead Reckoning. Perhaps the simplest approach to the pose maintenance task is to keep track of how far the robot moves in each direction and then to sum these motions to produce a net displacement that can be added to an initial position estimate. Keeping track of how much one moves by observing internal parameters without reference to the external world is known as dead reckoning and is usually implemented with an odometer. If only *dead reckoning* is used for position estimation, these errors are added to the absolute pose estimate and errors are accumulated. Long-term localization must make reference to the external world for position correction. This involves the use of sensory data for recalibrating a robot's sense of its own location with the environment. In some circumstances, such as the case of a wireless device that a person is moving around in space, we have no analogue of odometry.

Triangulation. Distance to known landmarks is frequently used to determine pose as this can be computed with cameras, laser range-finders, IR transmitters, sonar and other commonly used sensors. A naïve approach is to take three distance measurements and triangulate position. This works when the sensors are reliable and relatively noise-free but leaves several problems unaddressed. When the sensors are noisy, the calculations for triangulation become unstable for many positions and landmark arrangements and lead to significant loss of precision. Typically, multiple measurements are merged over time to try to compensate for this, however some care must be taken in choosing the method of merging or poor results will be obtained [11]. In some cases where the sensors are fairly reliable and have simple noise distributions, direct triangulation or triangulation with differential windowing can produce excellent results. Noisy sensors, however, complicate triangulation adding uncertainty to the results. GPS [29] is perhaps the most-used sensor based on triangulation.

Kalman Filter. In 1987, Smith and Cheeseman introduced the use of Kalman filters to the problem of determining position [34]. Many systems in robot localization, since then, have been based on Kalman filtering [16, 27, 10]. The robot's pose estimation is maintained as a Gaussian distribution in $\mathbb{R}^2 \times S$ and sensor data from dead reckoning and landmark observations is fused to obtain a new position distribution. This method is provably optimal when all distributions are linear but typically fails when these assumptions break down. Extended Kalman filters address this problem by linearizing the system. In practice, obtaining linearizations for many sensing systems is difficult and errors can propagate very quickly through the system.

Bayesian Approaches. Possibly the most powerful family of global localization algorithms to date is based on Bayesian inference, in particular Markov localization [24, 15] and Monte Carlo localization [13, 37]. These are generalizations of the Kalman filter. These algorithms estimate posterior distributions over robot poses which are approximated by piecewise constant functions instead of Gaussians, enabling them to represent highly multi-modal distributions. In this way, they can be applied in the case of sensors that are non-linear and have non-Gaussian noise distributions. The accuracy of the results, however, is limited by the resolution of the approximation. Due to the very complex nature of some sensors and usually also of the environment, many systems have difficulties modeling outliers and other artifacts. These difficulties can be addressed by sampling the distributions of the sensor signals in the target environment and using this directly as a model, as in the case of the sensor map we built in the first phase of our method. By explicitly integrating the conditional probability distributions, we can obtain precise approximations of the robot's positional distribution. This approach is both computationally tractable and effective [37]. Many excellent examples of this method exist in the literature [35]. This the approach we took in implementing localization using wireless Ethernet, as described in Section 2.

Topological Approaches. Typically the Bayesian approach is applied in the case when we have a grid-based representation of the environment. Another alternative for modeling the environment is with a topological map, represented as a generalized Voronoi graph [6]. Localization on the topological map is based on the fact that the robot automatically identifies nodes in the graph from geometric environmental information [26].

6. Conclusions

In this paper, we provide strong evidence that reliable localization with wireless Ethernet can be achieved. In our experiments, we can measure and track position robustly with the first meter of error distributed within a standard deviation. We used the Intersil Prism2 chipset for our wireless Ethernet cards and Apple AirPort base stations, both readily available and inexpensive hardware. The building we operated in had fairly complicated geometry and the base stations were laid out more than a year before we began our work. The methods we employed were general methods from robotics and followed the Bayesian approach to localization. These methods were readily adaptable to the problem at hand and can be applied to other location problems that might arise in mobile computing.

7. Acknowledgements

The authors would like to thank Moez Abdel-Gawad and Skye Schell for their help with taking measurements. They would also like to thank Dave Johnson for his advice and comments. Thanks also to Scott Crosby and to the anonymous MOBICOM'02 reviewers for their comments.

Andrew Ladd is partially supported by FCAR 70577, by NSF-IRI-970228 and a Whitaker grant. Kostas Bekris is partially supported by NSF-IRI-970228. Algis Rudys is supported by NSF-CCR-9985332. Dan Wallach is supported by an NSF Career Award CCR-9985332, an NSF Special Projects Award ANI-9979465, and a Texas ATP award. Lydia Kavradi is supported by NSF Career Award IRI-970228, a Whitaker grant, a Texas ATP award and a Sloan Fellowship.

8. References

- [1] A. A. Argyros, K. E. Bekris, and S. C. Orphanoudakis. Robot homing based on corner tracking in a sequence of a panoramic images. In *Proc. of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2001)*, volume 2, pages 3–10, Kauai, Hawaii, Dec. 2001.
- [2] P. Bahl and V. N. Padmanabhan. Enhancements to the RADAR user location and tracking system. Technical Report MSR-TR-2000-12, Microsoft Research, Feb. 2000.
- [3] P. Bahl and V. N. Padmanabhan. RADAR: An in-building RF-based user location and tracking system. In *Proc. of IEEE Infocom 2000*, volume 2, pages 775–784, Tel Aviv, Israel, Mar. 2000.
- [4] W. Burgard, A. Cremers, D. Fox, D. Hahnel, G. Lakemeyer, D. Schulz, W. Steiner, and S. Thrun. The interactive museum tour-guide robot. In *Proc. of the Fifteenth National Conference on Artificial Intelligence (AAAI-98)*, Madison, Wisconsin, July 1998. Outstanding paper award.
- [5] A. Chakraborty. A distributed architecture for mobile, location-dependent applications. Master's thesis, Massachusetts Institute of Technology, May 2000.
- [6] H. Choset and K. Nagatani. Topological simultaneous localization and mapping (SLAM): Toward exact localization without explicit localization. *IEEE Transactions on Robotics and Automation*, 17(2):125–137, Apr. 2001.
- [7] T. W. Christ and P. A. Godwin. A prison guard duress alarm location system. In *Proc. IEEE International Carnahan Conference on Security Technology*, Oct. 1993.
- [8] I. Cox. Blanche - an experiment in guidance and navigation of an autonomous robot vehicle. *IEEE Transactions on Robotics and Automation*, 7(2):193–204, 1991.

- [9] T. Cutler. Wireless Ethernet and how to use it. *The Online Industrial Ethernet Book*, Issue 5, 1999.
- [10] A. J. Davison and N. Kita. 3D simultaneous localization and map-building using active vision for a robot moving on undulating terrain. In *Proc. IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2001)*, volume 1, pages 384–391, Kauai, Hawaii, Dec. 2001.
- [11] G. Dudek and M. Jenkins. *Computational Principles of Mobile Robotics*. Cambridge University Press, 2000.
- [12] Federal Communications Commission Report and Order 96-264: Revision of the commission's rules to ensure compatibility with Enhanced 911 emergency calling systems, July 1996. <http://www.fcc.gov/Bureaus/Wireless/Orders/1996/fcc96264.txt>.
- [13] D. Fox, W. Burgard, F. Dellaert, and S. Thrun. Monte Carlo localization: Efficient position estimation for mobile robots. In *Proc. of the Sixteenth National Conference on Artificial Intelligence (AAAI-99)*, pages 343–349, Orlando, Florida, 1999.
- [14] D. Fox, W. Burgard, H. Kruppa, and S. Thrun. A probabilistic approach to collaborative multi-robot localization. *Autonomous Robots*, 8(3):325–344, June 2000.
- [15] D. Fox, W. Burgard, and S. Thrun. Markov localization for mobile robots in dynamic environments. *Journal of Artificial Intelligence Research (JAIR)*, 11:391–427, Nov. 1999.
- [16] J. Guivant and E. Nebot. Optimization of the simultaneous localization and map building algorithm for real time implementation. *Journal of Robotics Research*, 17(10):565–583, 2000.
- [17] P. Harley. Short distance attenuation measurements at 900MHz and 1.8GHz using low antenna heights for microcells. *IEEE Journal on Selected Areas in Communications (JSAC)*, 7(1):5–11, Jan. 1989.
- [18] A. Harter, A. Hopper, P. Steggle, A. Ward, and P. Webster. The anatomy of a context-aware application. In *Proc. of the 5th Annual ACM/IEEE International Conference on Mobile Computing and Networking (MOBICOM 1999)*, pages 59–68, Seattle, WA, Aug. 1999.
- [19] H. Hashemi. Impulse response modeling of indoor radio propagation channels. *IEEE Journal on Selected Areas in Communications (JSAC)*, 11:967–978, Sept. 1993.
- [20] H. Hashemi. The indoor radio propagation channel. *Proc. of the IEEE*, 81(7):943–968, July 1993.
- [21] J. Hightower and G. Borriello. Location systems for ubiquitous computing. *IEEE Computer*, 34(8):57–66, Aug. 2001.
- [22] J. Hightower, R. Want, and G. Borriello. SpotON: An indoor 3D location sensing technology based on RF signal strength. Technical Report UW CSE 00-02-02, University of Washington, Department of Computer Science and Engineering, Seattle, WA, Feb. 2000.
- [23] Institute of Electrical and Electronics Engineers, Inc. *ANSI/IEEE Standard 802.11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, 1999.
- [24] K. Konolige and K. Chou. Markov localization using correlation. In *Proc. of the Seventeenth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1154–1159, Seattle, Washington, Aug. 1999.
- [25] J. Krumm, S. Harris, B. Meyers, B. Brumitt, M. Hale, and S. Shafer. Multi-camera multi-person tracking for EasyLiving. In *Third IEEE International Workshop on Visual Surveillance*, Dublin, July 2000.
- [26] B. Kuipers and Y. T. Byan. A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations. *Journal on Robotics and Automatic Systems*, 8:47–63, 1991.
- [27] J. F. Leonard and H. Durrant-Whyte. Mobile robot localization by tracking geometric beacons. *IEEE Transactions Robotics and Automations*, 7(3):376–382, June 1991.
- [28] T. Liu, P. Bahl, and I. Chlamtac. Mobility modeling, location tracking, and trajectory prediction in wireless ATM networks. *IEEE Journal on Selected Areas in Communications*, 16(6):922–936, Aug. 1998.
- [29] T. Logsdon. *Understanding the Navstar: GPS, GIS and IVHS. Second edition*. Van Nostrand Reinhold, New York, 1995.
- [30] G. Marceau. The McGill's RedDogs legged league system. *Robocup 2000*, pages 627–630, 2000.
- [31] A. Neskovic, N. Neskovic, and G. Paunovic. Modern approaches in modeling of mobile radio systems propagation environment. *IEEE Communications Surveys*, Third Quarter 2000.
- [32] N. Priyantha, A. Chakraborty, and H. Balakrishman. The Cricket location support system. In *Proc. of the 6th Annual ACM/IEEE International Conference on Mobile Computing and Networking (MOBICOM 2000)*, pages 32–43, Boston, MA, Aug. 2000.
- [33] N. Priyantha, A. Miu, H. Balakrishman, and S. Teller. The Cricket compass for context-aware mobile applications. In *Proc. of the 7th Annual ACM/IEEE International Conference on Mobile Computing and Networking (MOBICOM 2000)*, pages 1–14, Rome, Italy, July 2001.
- [34] R. Smith and P. Cheeseman. On the representation of spatial uncertainty. *Journal of Robotics Research*, 5(4):56–68, 1987.
- [35] S. Thrun. Probabilistic algorithms in robotics. *AI Magazine*, 21(4):93–109, 2000.
- [36] S. Thrun, W. Burgard, and D. Fox. A probabilistic approach to concurrent mapping and localization for mobile robots. *Machine Learning*, 31(1-3):29–53, 1998.
- [37] S. Thrun, D. Fox, W. Burgard, and F. Dellaert. Robust Monte Carlo localization for mobile robots. *Artificial Intelligence*, 101:99–141, 2000.
- [38] R. Want, A. Hopper, V. Falco, and J. Gibbons. The Active Badge location system. *ACM Transactions on Information Systems*, 10:91–102, Jan. 1992.
- [39] A. Ward, A. Jones, and A. Hopper. A new location technique for the active office. *IEEE Personal Communications*, 4(5):42–47, Oct. 1997.
- [40] J. Werb and C. Lanzl. Designing a positioning system for finding things and people indoors. *IEEE Spectrum*, 35(9):71–78, Sept. 1998.
- [41] R. Yamamoto, H. Matsutani, H. Matsuki, T. Oono, and H. Ohtsuka. Position location technologies using signal strength in cellular systems. In *Proc. of the 53rd IEEE Vehicular Technology Conference*, May 2001.