Adding Angle of Arrival Modality to Basic RSS Location Management Techniques

Eiman Elnahrawy, John Austen-Francisco, Richard P. Martin
{eiman,deymious,rmartin}@cs.rutgers.edu
Department of Computer Science, Rutgers University
110 Frelinghuysen Rd, Piscataway, NJ 08854

Abstract:

In this paper we describe a radio-based localization approach that is based on the use of rotating directional antennas and a Bayesian network that combines both angle-of-arrival (AoA) and Received Signal Strength (RSS). After describing our network, we extensively characterize the accuracy of our approach under a variety of measured signal distortion types. Next, using a combination of synthetic and trace-driven experiments, we show the impact of different signal distortions on localization performance. We found the use of directional antennas was effective at averaging out multi-path effects in indoor environments, which helped reduce the amount of training data required compared to previous approaches.

I. Introduction

Wireless networks have gained much popularity as a localization infrastructure in addition to being used for communication purposes. Researchers have been motivated by the cost and deployment savings when using such infrastructure for locating objects as opposed to investing in specific localization systems, such as ceiling-based ultrasound. Approaches that use such an infrastructure and measured Received Signal Strength (RSS) for localization, however, have been proven to provide only a room-level localization accuracy; with an average localization error of 10-15ft [3].

In this work we investigate incorporating the Angle of Arrival (AoA) of the received signal information into basic localization techniques that rely on RSS. Our motivation is that one would expect to improve the accuracy of the localization result when providing extra information about the mobile object to the localization system, in this case the angle from a landmark.

We describe a statistical localization approach for wireless LANs and sensor networks that combines both AoA and RSS. In particular, We extend our prior work using Bayesian networks and RSS-to-distance estimation [9]. Our approach is to add a rotating directional antenna to the access points which allows the incorporation of AoA information into the basic RSS-to-distance Bayesian network. The crux of our approach rests in incorporating an abstraction of the *AoA curve* into a Bayesian network. This curve describes the relative changes in signal strength as a function of the angle of the antenna to the object to be localized.

We extensively characterize the accuracy of our approach under a variety of measured signal distortion types using different collected data sets as well as simulated data. We show our approach is applicable to a variety of technologies by showing that the propagation models work well for both 802.11 as well as 802.15.4 networks. Our simple propagation model fits the observed AoA curves in outdoor settings, where the biases generated by indoor environments do not exist. It can give highly accurate results – a mean error of less than 1ft and a maximum error of 3ft.

We found that indoor environments, however, have a high degree of error in both the AoA measurement as well as the distance to signal measurement which results in performance that is only marginally better than not using angle information at all. We also found that variations in distance to signal estimation was the major cause for this degraded performance, as opposed to angle estimation errors. In order to understand the differences between the indoor and outdoor environments, we characterize measured indoor AoA curves and show they deviate significantly from what free-space models would predict. An interesting result we found is that using a directional antenna in combination with a Bayesian network did help in averaging out some noise and multi-path effects. The directional antenna was thus useful for reducing the size of labeled training data needed for learning-based localization approaches that use RSS for position estimation.

In summary, our contributions include: (1) A new Bayesian network for location estimation that uses both AoA and RSS-to-distance, (2) A measured characterization of both open-space and indoor signal propagation as a function of angle using directional antennas, (3) A quantification of the effectiveness of our Bayesian approach, and (4) an investigation of the sensitivity of our approach to errors in the angle and distance estimation.

The rest of this paper is organized as follows. We first describe related work in Section II. We briefly describe our hardware and Bayesian network in Section III. We also show the results of open-space experiments.

Section IV describes our experimental data sets. In Section V we evaluate our approach using both synthetic and trace-driven data from 2 different kinds of access points. Finally, we conclude in Section VI.

II. Related Work

There is a rich history of localization approaches that rely on existing 802.11 communication infrastructure for positioning as we propose [1], [17], [3], [7], [13], [15]. Recent years have also seen tremendous efforts at building small- and medium-scale localization systems for sensor networks, 802.11, custom radios, and ones that use ultrasound or infrared e.g., [10], [4], [16], [6], [8], [5], [12], [14]. Unlike our approach, many of them either require large training sets, rely on properties other than those of the radio, or are intended to be used outdoors where radio propagation is less prone to noise and bias.

A recent theoretical work incorporated AoA into a localization algorithm that is based on semidefinite programming [2]. However, the work most related to ours is VORBA [11]. There, the authors presented a positioning architecture based on 802.11 and AoA. We also implemented and investigated the use of AoA for indoor localization. However we extensively analyze and quantify the various types of errors that distort the signal and angle measurements and show their effect on accuracy. Our approach is superior to VORBA in that we get the same performance using a fewer number of base stations, in our case 4 vs. 7.

Bayesian networks for indoor localization using RSS only were introduced in [9], [3], and shown to be robust and competitive with state of the art positioning approaches. In this paper we generalize this approach to incorporate AoA information. We show that such an approach helps reduce the size of the training examples needed to reach the same performance of Bayesian networks that rely on RSS-to-distance estimation alone.

III. Incorporating The Angle Information

In order to collect the AoA data we constructed a simple base station landmark incorporating a rotating directional 19dB parabolic antenna as shown in Figure 1. The base station can construct a curve of the measured RSS as a function of the angle of the antenna. This function of RSS to angle, or the, *AoA curve*, provides the base radio property of our approach. Specifically, this curve describes the relative changes in signal strength as a function of the angle of the antenna to the object to be localized.

We statistically model an abstraction of the *AoA curve* into our localization system using a Bayesian Graphical Model (BGM). Due to space limitation we briefly describe our Bayesian network here, for more information about BGMs and a description of our prior networks for localization using RSS-to-distance alone please refer to [9].

Fig. 1. Our rotating directional antenna landmark.

Our goal is to arrive at a simple composite model that predicts the RSS given the distance between the transmitter and receiver and relative angle of the antenna. Our AoA propagation model is based on two parts. First, the mean of the curve should follow a standard path-loss model, such as a log-linear model (e.g., $b_1 + b_2 \log(D)$). Second, we study the variation of the standard path loss model as a function of the angle of the directional antenna and we quantize the angle over the whole rotation. We found that a cosine function is a reasonable trade-off in model simplicity to predictive performance for our localization application. Specifically, we model the signal strength received from landmark i at angle quantization j, $S_i[j]$ as being normally distributed around the model described above, as follows, where the bs are the propagation parameters, σ is the standard deviation of the signal, a_i is the true angle from the mobile to landmark i and θ is the quantized angle which varies from 0 to 2π with a granularity that depends on the number of signal strength readings the mobile receives from the landmark in one rotation of the directional antenna.

$$\begin{aligned} S_i[j] \quad &\sim \quad N((b_{i0} + b_{i1}D_i)cos(\frac{\pi}{3}(a_i - \theta[j] - \pi)) + \\ &\qquad (b_{i2} + b_{i3}D_i), {\sigma_i}^2), i = 1, \ldots, n, j = 0 \ldots 2\pi \end{aligned}$$

Figure 2 shows our proposed AoA Bayesian network. We assume a two-dimensional location estimation problem with n access points. The vertices X and Y represent location. The vertex D_i represents the Euclidean distance between the location (X, Y) and the ith landmark, which is assumed to be known. The model assumes that X and Y are marginally independent and are independent of the angle measurement.

We used our hardware to collect RSS and AoA data indoors and outdoors for our experiments, for both 802.11 as well as 802.15.4 wireless networks. Our outdoor data showed that our technique is feasible. Our proposed Bayesian model fits the observed AoA curves in outdoor settings, where the biases generated by indoor environments do not exist. Figure 3 shows the outdoor data. For the angle plots, three sample curves are shown: the raw data collected at 10 degree intervals, a smoothed curve, and a cosine function

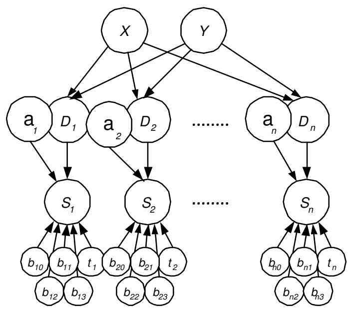

Fig. 2. A Bayesian network for location estimation that combines both AoA and signal to distance in one graphical model.

constructed from the smoothed curve. The applied smoothing function is simply the average of each point along with its six closest neighbors, three on each side (with wraparound). The third curve is one half of a cosine period fitted to smoothed curve. In most cases the fit was relatively good, resulting in an average root mean square difference between the cosine and the actual measured data of just over 4dBm, with some even better. We noticed that the overall performance of the motes and the Lucent cards (WLAN) is functionally identical with regard to signal strength versus AoA.

IV. Experimental Data

We experimented with 3 data sets. The first set, "set 1", and its results, are generated from our indoor experiments. For the second set, "set 2", we used data from previous work [11]. The indoor experiments were similar for both sets. A base station with a rotating directional antenna (landmark) was placed on the floor, and a series of mobiles (laptops) sent packets as the antenna was rotated, generating an AoA curve. Figure 4(a) shows the ground truth placement of the base stations and mobiles for set 1 (top) and set 2 (bottom). Hallways are shaded in grey, offices and laboratories are white. We did not collect data in all rooms because many of the side rooms are private offices.

The relevant differences between the data sets are the type of antenna used, the number of base stations, the orientation of the localized radio and the resulting smoothing function. For set 1, the 19dB antenna, shown in Section III was used in 4 locations and measurements using the mobiles were performed at 20 locations. The antenna was rotated 10^o degrees at a time, starting at 0^o while the mobile's radio remained at 0^o for the duration of the experiment. Data set 2 is the one used in [11]. The same Lucent cards were used, but with an 8dB phase array antenna instead. The base station was placed at 7 positions, and the mobiles at 27 positions. For set 2, the orientation of the mobile's radio was traversed by 90^o intervals, and the resulting AoA curve was smoothed using the 4 different orientations. In this work, we used the resulting smoothed curves.

In order to investigate the cause of our observed localization performance, we also created a set of synthetic AoA curves, which we call the "synthetic" set. It matches the characteristics (i.e., propagation parameters) of our collected outdoor data, from Section III. We generated it assuming same base station and measurement positions as set 1, however.

Unless otherwise explicitly specified, we use a leave-one-out approach to characterize the localization accuracy. Specifically we use a training set of some size m where we repeat the experiment m times for a training-testing split using $m-1$ observations for training and one observation for testing.

V. Localization Performance

In this Section, we detail our experiments using AoA with RSS, the Bayesian model of Section III, and the data sets described above.

We show the localization accuracy of our AoA Bayesian network applied to data with little distortion (e.g., outdoors) can achieve good accuracy, with average errors on the order of 1m. Moving indoors, next we show the performance of our approach using data sets 1 and 2. The performance was not better than approaches that do not incorporate angle information. To explore why, we characterize the errors of AoA curves in indoor environments and show that resulting inaccuracies are about what should be expected given the noise and biases compared to the cleaner outdoor data. Finally, we show how using the directional antenna, however, reduces the amount of training samples needed to localize for certain Bayesian networks.

A. Outdoor Localization Accuracy

Figure 3 showed our measured outdoors data closely follows propagation models introduced in literature. Hence, using propagation parameters estimated from this data we generated "synthetic" data to simulate a data set with little distortion, e.g., outdoors. The top curve of Figure 6(a) shows the performance of our approach on this data. Recall the error CDF is the cumulative distribution function of the difference in the estimate of a mobile's position from its true location. We will describe the figure in detail later. Here, however, we show that the performance is excellent with mean error of less than 1ft and a maximum error of 3ft. Even with random shift errors in the angle up to 45^o, the maximum error is still only 7ft.

B. Indoor Localization Accuracy

Figure 4(b), top, presents the error CDF for set 1, and bottom for set 2. The figures also include error CDFs from a Bayesian network, M1, from prior work, that does not use angle information [9]. The results show the angle information incorporated into the network seems to make no appreciable difference in performance. Indeed, the CDFs are nearly identical in the two cases.

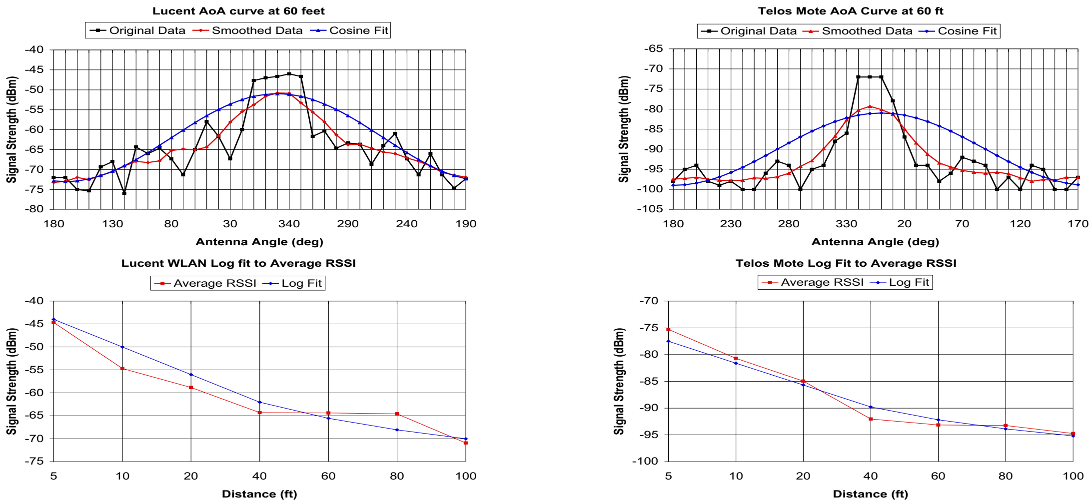

(a) (b)

Fig. 3. Sample outdoor AoA curves (top) and RSS to distance curves (bottom) for the 802.11 (a), and 802.15.4 (b) platforms.

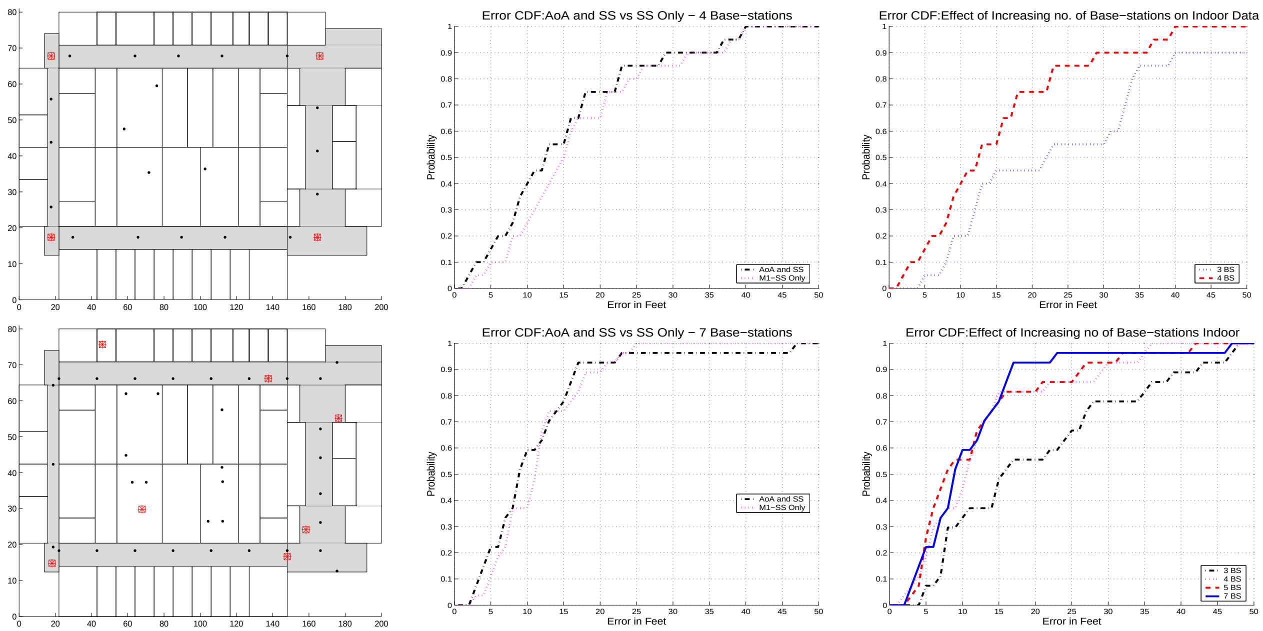

(a) (b) (c)

Fig. 4. (a) Base station positions and measurement points, antennas are red squares and mobiles are black dots, (b) CDF of the localization results, (c) results using different number of base stations. Top graphs for data set 1 and bottom for data set 2. CDFs from Bayesian networks using only RSS to distance (M1) are presented for comparison.

Figure 4(c) top, is useful for comparing our work to VORBA [11], bottom. For both sets, performance tops out with 4 corner base stations. In VORBA, comparable performance (a very similar median error of 18ft and max error of 40ft) was not obtained until the system had 7 landmarks. In other words, using our approach we were able to decrease the number of access points by almost half without sacrificing accuracy.

C. *Indoor AoA Curves*

Here we categorize and analyze the different sources of errors in AoA curves obtained indoors. We found three primary sources of distortions:

a) Shift errors: when the peak center is shifted by a number of degrees from the true angle. These might occur when the signal goes around a corner.

b) Distance errors: when the mean of a curve differs from what an RSS to distance estimate would predict. For example, if the radio is behind many walls, the curve will look lower than what is predicted due to power lost as the signal passes through the walls.

c) Lobe errors: when the main peak is difficult to discriminate because additional peaks due to side lobes are a significant fraction of the main peak. A radio-reflective object, e.g., a large metallic bookcase or elevators, can cause artificial side lobe-like effects.

Our data generally included many low quality AoA curves with multiple peaks. Figure 5 shows the distribution of shift and distance errors for data sets 1 (top) and 2 (bottom). Figure 5(a) is a histogram of angle shifts. We computed this as difference of the angle at the maximum of the curve from the true angle. To create a histogram, we sorted the errors into bins with 10 degree granularity.

The figure shows the shift error for set 1 is much better than set 2. However, even for set 1, the bulk of the curve still varies by 60^o from the true angle. Most of the difference between the two sets is likely to be caused by the differences in the beam width of the different antennas. Figure 5(b) shows a scatter plot of the distribution of distance errors. For both curves, we see a tight distribution at low distances. However, above 50 ft, both curves widen considerably, making predicting the distance given the RSS quite a challenge.

D. Results Using Modified AoA Curves

Here we alter our data sets to reflect the classes and magnitudes of the distortions we observed. Our goal is to quantify how well our model works as we eliminate different kinds of errors in the curves. Our approach works in both directions, namely we modified the "synthetic" data curves to make them appear more like measured indoors data and modified the real indoors data, specifically "set 1", to look more like outdoor curves.

We performed 3 correctional transformations on the real AoA curves that correspond to (1) errors in the signal-to-distance function, (3) errors in the angle of arrival function and (3) magnified side lobes. Removing errors of type (1) requires shifting the entire curve up or down by a percentage, S, of the error from the true mean. For type (2), we shifted it left or right by a percentage, A, of the angular distance to the true peak, and for (3) we smoothed out lobes in the peak by performing a sliding window average over the points of width W^o. We also performed the inverse transformations (distortions) on the synthetic data. That is, we took a high quality centered cosine curve based on the outdoor Lucent cards and adjusted the entire curve up or down by a random amount, shifted left or right by a random angle, and added side lobes by adding additional 2 cosine curves with an amplitude a fraction of the height of the main curve at random points in the main curve.

Figure 6(a) shows the impact of distortions on the synthetic data. Absolute distortion values are presented. We see that with clean data localization performance is excellent. However, a small random signal loss between 0-5dBm greatly impacts the performance, increasing the mean error to 10ft and the max error to 20ft. We also see that adding 2 random side lobes that are 80% of the peak also impacts the performance significantly. Finally, for every increase in the range of random dBm loss, we see a very large impact on the accuracy.

Figure 6(b) shows the impact of correcting the real AoA curves for data set 1. Here, distances and shifts are expressed as percentages of the error corrected. E.g., A=100, S=100 means that 100% of the shift and distance errors are corrected. The W variable is the angle over which a point in the curve is averaged out. Thus, W=60 means that each point in the curve is averaged over its +/- 30 degree neighbors. Finally, the L term describes the angle of the peak that is preserved; i.e. not averaged out. We can see that no single correction does much. In fact, we need to correct both the distance and shift errors to get good performance. The figure also shows it is possible to average out the curves too much — 60 degrees seems to be the best amount. Finally, excluding the peak from averaging is not a good idea; it is better to average out the entire curve rather than try to keep the peak undisturbed.

E. Antenna Smoothing

A final crucial effect we observed is that using a rotating directional antenna greatly reduces the number of training samples needed by M1 and other RSS-based localization approaches. Averaging the signal while rotating the directional antenna might improve the performance by smoothing out bias and noise in indoor environments. Omni-directional antennas are not as effective in this regard. Figure 6(c) shows that M1 performs similarly to our angle network with only 20 training samples if we use the average RSS over the entire rotation. If we use data collected in the same setting from omni-directional antennas, 55 training samples were required to obtain similar performance, and adding additional sampling did not improve the performance much.

VI. Conclusion

We presented a localization approach based on the use of rotational directional antennas and Bayesian networks. We showed that our Bayesian network can give highly accurate results (1-3ft) and is robust to significant errors in the angle estimation. We extensively characterized our approach using data collected outdoors, from both WiFi and ZigBee radios, as well as indoor and synthetic data. We showed that AoA does not help much indoors because of fundamental errors in RSS and angle information, but appears to be a viable technique for outdoor localization. We characterized the different sources of errors indoors. We also showed that a rotating directional antenna is useful for averaging out variances in signal strength, which helps reduce the amount of profiling needed for RSS-based localization.

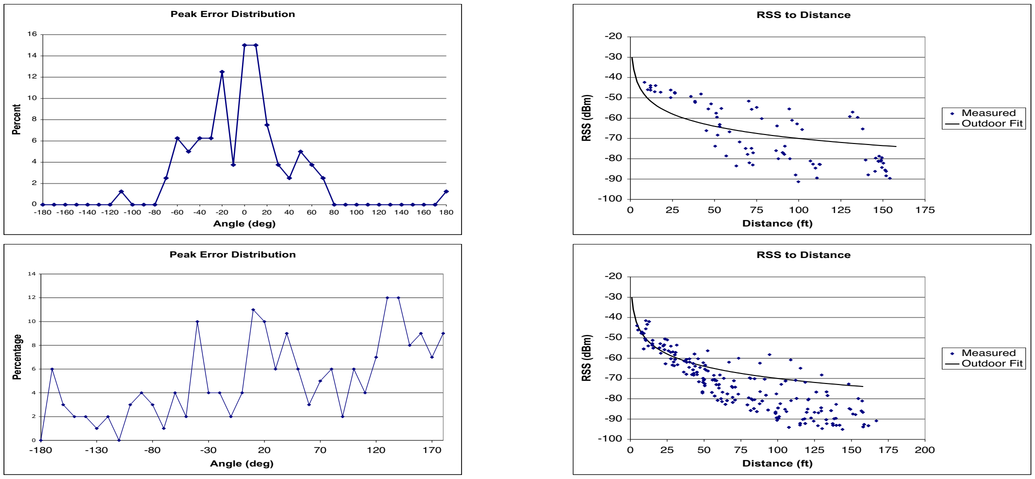

(a) (b)

Fig. 5. Angle error histogram in 10 degree bins (a) and distance errors scatter plot (b) for data set 1 (top) and for data set 2 (bottom).

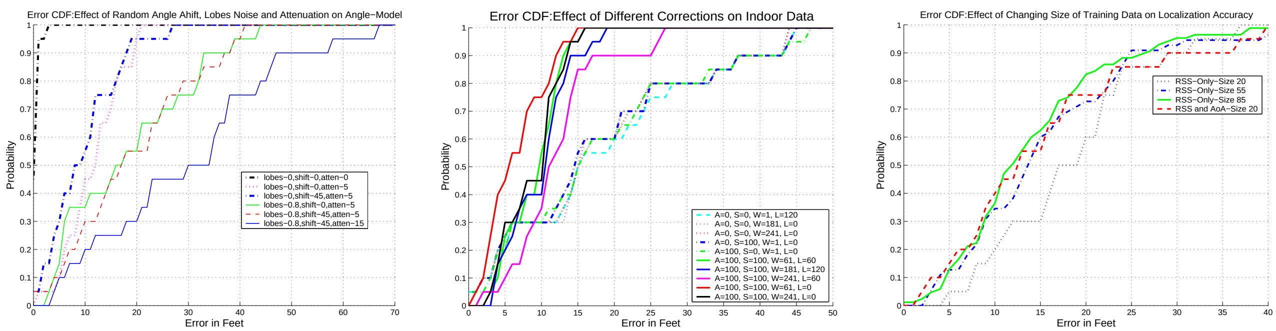

(a) (b) (c)

Fig. 6. (a) presents error CDFs of the synthetic data with varying degrees of distortion, while (b) presents data set 1 with varying amounts of correction. (c) The impact of using rotating antenna to average out noise and bias. It reduces the amount of training needed by M1 from 55 to 20 points.

REFERENCES

[1] P. Bahl and V. N. Padmanabhan. RADAR: An In-Building RF-Based User Location and Tracking System. In *INFOCOM*, March 2000.

[2] P. Biswas, H. Aghajan, and Y. Ye. Integration of angle of arrival information for multimodal sensor network localization using semidefinite programming. In *Proceedings of 39th Asilomar Conference on Signals, Systems and Computers*, Oct. 2005.

[3] E. Elnahrawy, X. Li, and R. P. Martin. The limits of Localization Using Signal Strength: A Comparative Study. In *SECON*, 2004.

[4] L. Fang, W. Du, and P. Ning. A beacon-less location discovery scheme for wireless sensor networks. In *INFOCOM*, 2005.

[5] M. Hazas and A. Ward. A high performance privacy-oriented location system. In *Proceedings of the First IEEE International Conference on Pervasive Computing and Communications (PerCom)*, Mar. 2003.

[6] T. He, C. Huang, B. Blum, J. A. Stankovic, and T. Abdelzaher. Range-free localization schemes in large scale sensor networks. In *Proceedings of the Ninth Annual ACM International Conference on Mobile Computing and Networking (MobiCom'03)*, Sept. 2003.

[7] P. Krishnan, A. S. Krishnakumar, W.-H. Ju, C. Mallows, and S. Ganu. A System for LEASE: Location Estimation Assisted by Stationary Emitters for Indoor RF Wireless Networks. In *INFOCOM*, Oct. 2004.

[8] K. Lorincz and M. Welsh. Motetrack: A robust, decentralized approach to rf-based location tracking. To Appear in Springer Personal and Ubiquitous Computing, Special Issue on Location and Context-Awareness,, 2006.

[9] D. Madigan, E. Elnahrawy, R. P. Martin, W.-H. Ju, P. Krishnan, and A. Krishnakumar. Bayesian Indoor Positioning Systems. In *Infocom*, 2005.

[10] D. Moore, J. Leonard, D. Rus, and S. Teller. Robust distributed network localization with range measurements. In *SENSYS*, 2004.

[11] D. Niculescu and B. Nath. VOR Base Stations for Indoor 802.11 Positioning. In *MOBICOM*, pages 2926–2931, 2004.

[12] N. Priyantha, A. Chakraborty, and H. Balakrishnan. The Cricket Location-Support system. In *ACM International Conference on Mobile Computing and Networking (MobiCom)*, Boston, MA, Aug. 2000.

[13] T. Roos, P. Myllymaki, and H.Tirri. A Statistical Modeling Approach to Location Estimation. *IEEE Transactions on Mobile Computing*, 1(1), Jan-March 2002.

[14] A. Savvides, C.-C. Han, and M. Srivastava. Dynamic Fine-Grained Localization in Ad-Hoc Networks of Sensors.

[15] A. Smailagic and D. Kogan. Location sensing and privacy in a context aware computing environment. *IEEE Wireless Communications*, 9(5), Oct. 2002.

[16] R. Stoleru and J. Stankovic. Probability grid: A location estimation scheme for wireless sensor networks. In *SECON*, 2004.

[17] M. Youssef, A. Agrawal, and A. U. Shankar. WLAN location determination via clustering and probability distributions. In *Proceedings of IEEE PerCom'03*, Fort Worth, TX, Mar. 2003.