VOR Base Stations for Indoor 802.11 Positioning

Dragoş Niculescu and Badri Nath
Computer Science, Rutgers University
{dnicules,badri}@cs.rutgers.edu

ABSTRACT

Angle of arrival (AOA) has previously been used for outdoor positioning in aircraft navigation and for services like E911. For indoor positioning, the best schemes to date rely either on extensive infrastructure, or on sampling of the signal strength on a dense grid, which is subject to changes in the environment, like furniture, elevators, or people. We present an indoor positioning architecture that does not require a signal strength map, simply requiring the placement of special VOR base stations (VORBA). While our incipient realization of the AOA using 802.11 uses a base station with a revolving directional antenna, a non mechanical implementation would yield comparable performance, even with quantized angles. Performance of positioning with VOR base stations is evaluated though experimentation, simulation, and theoretical analysis.

Categories and Subject Descriptors

C.2.1 [**Network architecture and design**]: Wireless communication; C.2.2 [**Network protocols**]: Applications

Keywords

AOA, ranging, 802.11, indoor positioning, VOR, VORBA

General Terms

Design, Experimentation

1. INTRODUCTION

Indoor positioning is a complex engineering problem that has been approached by many computing communities: networking, robotics, vision, and signal processing. In most of the proposed solutions, certain aspects of the problem domain are so specialized that a solution applicable in one domain does not easily translate into a solution in other domains. For example, many computer vision based techniques require line of sight conditions, and achieve high ac-

curacy. In the context of 802.11 based position aware applications, line of sight is usually not available, but a lower accuracy may be tolerable.

There has been a surge of research activity in the pursuit of finding methods for accurate and robust indoor positioning techniques. Currently, the most convenient solutions for indoor positioning using 802.11 are RADAR [1], and its derivatives [2, 3, 4]. The problem with these approaches is that they require signal strength mapping of the entire area that has to support positioning, mapping which has to be reworked when propagation conditions change. As the human body roughly halves the strength of 2.4GHz signals, a crowded building might render the measurement map unusable. The advantage of these systems however lies in the fact they use off the shelf hardware which is widely available due to the popularity of wireless LANs and 802.11 hotspots. One application in which it is not feasible to build a signal strength map is that of an ad hoc disaster network. In such a setup there is no time to do signal strength mapping, and it may be impossible to deploy such a positioning infrastructure.

Methods that rely on range measurements as a function of signal strength are subjected to variance caused by the environmental surroundings which in previous studies have been shown to be a significant factor [5]. We propose to mitigate the effects of the environment on range measurements by use of a base station (or an access point) that has a rotating directional antenna. Using a 802.11 base station with a revolving directional antenna that can provide angle of arrival (AOA) and range measurements, we can obtain a more robust estimation than that based on signal strength from a standard 802.11 access point. By moving some of the complexity of the positioning support to the base station, we can significantly reduce the work and cost of deploying a range map in order to estimate indoor positions.
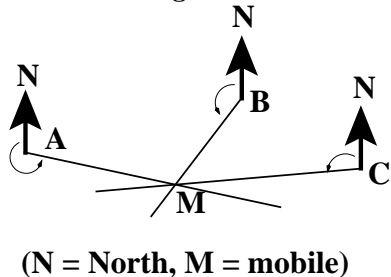
In this paper, we show that there are a number of ways of determining AOA from a 802.11 base station, one of which we have actually built, tested and obtained results with. Some of the newer schemes that improved on the RADAR idea showed better performance, but still retained the requirement of building a dense signal strength map of the building. Accuracy of positions obtained by our system, of 2.1m median error, is comparable to the original RADAR, but works without requiring a map of the signal strength of the area.

We show how to use the idea of VOR (VHF Omnidirectional Ranging) for indoor positioning using 802.11 hardware, with a **customized base station that can measure**

Figure 1: Basic triangulation and trilateration

(N = North, M = mobile)

both angles and ranges. In our incipient realization, the base station is a laptop equipped with a directional antenna that is continually rotating. We are able to derive angles towards mobiles with a median error of 22° and ranges to mobiles with a median error of 2.8m. These enable the use of trilateration procedures (for ranges), triangulation (for angles), and a combination of them.

In summary, here are the main contributions of the paper: first, an indoor positioning architecture that does not require a signal strength map. Second, a more robust positioning system that uses AOA and ranges obtained from a rotating directional antenna. Third, an actual implementation of the base station using off the shelf (garage sale) hardware. Fourth, analysis of position accuracy and robustness based on experiments, simulation and theory.

The rest of the article is structured as follows: the remaining of the section reviews the basics of triangulation and trilateration; section 2 describes our implementation of the VORBA, how it infers angles, and a description of the measurement methodology. Section 3 presents several methods of positioning using VORBAs: based on discrete angles, angle distributions, quantized angles, and ranges. A theoretical analysis of the AOA based positioning is also included in this section. We review related articles in section 4, and conclude with some discussion and summary in section 5.

1.1 Triangulation and trilateration

Trilateration is a positioning procedure in which mobile M knows distances MA, MB, and MC, also known as ranges, in addition to the coordinates of landmarks A, B, C (Figure 1). By solving the nonlinear system :

$$\sqrt{(x_M - x_I)^2 + (y_M - y_I)^2} = MI, I = A, B, C$$

the mobile is able to find an estimate for its own position (x_M, y_M). This procedure is used by GPS with ranges obtained measuring time of flight to precisely synchronized and positioned satellites.

If no distances are known, but the mobile can find the angles under which it is seen by the landmarks, the **triangulation** procedure can be applied. Here the mobile knows the angles $\widehat{NAM}$, $\widehat{NBM}$, and $\widehat{NCM}$. These angles, together with the known positions of the respective landmarks determine half lines whose intersection is at the mobile M. Triangulation has actually been used prior to trilateration, because angles are easier to measure than ranges using simple mechanical / optical methods - for example in topometry, topography, air and sea navigation. One such example is VOR - a ground based navigation aid that still is the pri-

Figure 2: Experimental base station with revolving antenna

mary navigation system for the majority of aircrafts, even after the introduction of GPS. Its principle was the main inspiration for this paper: a landmark sends two signals, one that is periodic and omni-directional, while the another one is directional and rotating about the landmark. The airborne equipment receives both signals, and interprets the difference between the times of the signals as an angle under which the landmark sees the aircraft. The coordinates of the landmark are known, therefore placing the aircraft anywhere on a given line. A second VOR reading provides a second line to be intersected with the first to yield a position.
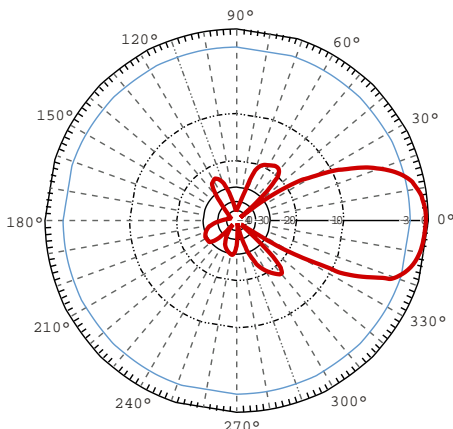
In the next section, we describe VORBA, a prototype base station that measures ranges and angles, enabling mobiles to use triangulation, trilateration, and other positioning procedures.

2. VOR BASE STATION REALIZATION

A way to have AOA functionality on a 802.11 base station is to attach a directional antenna to a wireless access point. When this antenna is rotated, the SS (signal strength) reported by the card is higher in the direction of the mobile, and possibly in other directions as well, due to reflections. To automatize this measurement of the angle, we mounted a small Toshiba Libretto 70ct laptop on a record player (turntable) as shown from a top view in Figure 2. In order to obtain higher difference in the maximums, we chose an antenna that is highly directional. We linked the Lucent 2Mbps 802.11 card to a Hyperlink 14dB gain directional antenna that has the horizontal radiation pattern shown in Figure 3. The vertical pattern is almost identical, the main feature being that the strongest signal is spread only 30° from the center. The antenna is attached to the bottom of the laptop, so that it rotates in the horizontal plane.

One revolution of the turntable corresponds to a complete sweep of all angles in $0..2\pi$. About 500 samples of the SS can be associated with each angle at a 33RPM speed of the turntable (a period of 1.8 seconds). $SS(\alpha)$ is a function describing the strength of the signal of the mobile as seen from the base station. **This function provides all the information needed to derive angles and ranges for triangulation and trilateration.** Most of section 3 shows how a mobile can derive angle information by using $SS(\alpha)$ from several base stations, while subsection 3.4 explains how ranges are derived from the same information.

Figure 3: Hyperlink HG2414 directional antenna gain pattern

Figure 4: A peak in signal strength indicates a possible direction of the mobile. Cartesian (top) and polar (bottom) representation of $SS(\alpha)$.
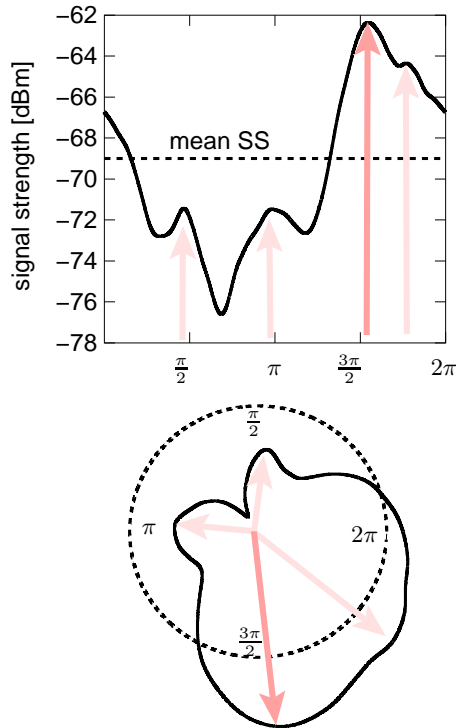

If the revolution speed would be constant, $SS(t)$ measurements in time could be directly translated into angles with respect to North to obtain $SS(\alpha)$. But a problem we encountered early on during experimentation phase was the instability of the turntable period, that varied albeit slowly between 1.7 and 2.3 seconds. This would produce SS samples at variable rates, which wouldn't theoretically be a problem if we could accurately associate each SS reading with an angle. Ideally, it would be convenient to use a digital compass attached to the laptop, but this option was not available due to lack of interfaces on the laptop. Also, compasses might be disturbed near power cables, or large metal objects. We opted instead for synchronization using infrared once every revolution. A continuous IR signal is sent from the device on the left in Figure 2, in this case another laptop. When the revolving antenna on the laptop perfectly aligns its receiver with the fixed IR beam, the base station knows it has hit the horizontal axis of the system, which provides for a good continuous calibration of the system. The fixed IR beam indicates the horizontal axis of our coordinate system (opposite direction), so that all base stations report angles with respect to the same reference.

2.1 Measurements

The mobile requires an $SS(\alpha)$ from each base station, but for the purpose of experimentation, we only realized one VOR base station and took several sets of measurements for different positions of the base station. Figure 5 shows the $56m \times 25m$ building in which 32 measurements points were taken (possible positions of the mobile), indicated with black dots. VOR base stations were placed at locations indicated by stars. Five base stations are sufficient to cover most of the building except the right side, where most points are separated from the base stations by three large elevators. In some experiments, we used two extra base stations indicated with hollow stars only for this side of the building.

In each measurement point, a regular 802.11 equipped laptop took four sets of measurements, one for each pose of the mobile (facing North, East, South, and West). No compass was used for orientation, the user just aiming to have the measurement laptop parallel to the walls. The posi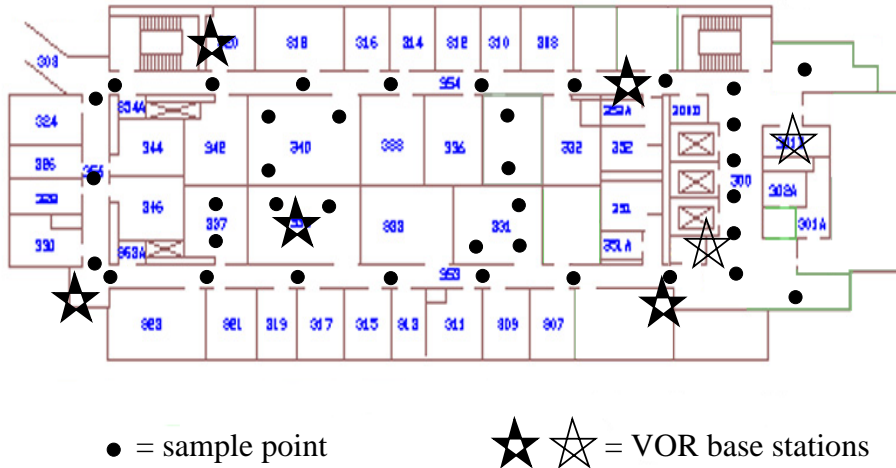tion of the user was randomized in order to include in the measurements the situations in which both a human body and a laptop screen block the shortest path towards the base station. Polarization doesn't seem to matter, but we performed all the measurements keeping the 802.11 card in an horizontal plane, as is standard in most laptops. A measurement for a pose is in fact an average over three or four revolutions of the base station, in order to reduce the effect of temporary factors, such as open doors, or people moving by. Measurements were taken at various times of the day and night, including the busy morning and afternoon hours.

In the initial phases of experimentation, we decided to take as many SS samples per revolution as possible in order to maximize the information in $SS(\alpha)$. The current setup supports 250 pings per second, yielding almost 500 samples per revolution, and any additional traffic would probably introduce jitter in the angle measurements unless some prioritizing for the probe traffic is used. On the other hand, since VORBA is an extension for data access points, user data can be used to sample SS without the need for evenly spaced probes, when a synchronization mechanism such as NTP is available. The necessary accuracy for such a mechanism is on the order of milliseconds ($1° \simeq 5ms$). Therefore, even if in the current experimental setup the positioning support is using all the bandwidth available for evenly spaced SS sampling, it is possible to ameliorate this by a number of methods: 1. employ user traffic for sampling; 2. broadcast the angle periodically from the base station (no RTS, CTS, or ACK overhead); 3. reduce the sampling rate. This last method is probably the most effective, because the high

Figure 5: Map of VOR base stations and sample points

● = sample point   ★ ☆ = VOR base stations

frequencies of $SS(\alpha)$ are not usable anyway. To obtain the signal in Figure 4, we employed a filter of 0.2 seconds to smooth the 2 second periodic signal.

2.2   AOA inference

As will be shown in this section, in most cases the best direction towards the mobile is indicated by one of the strongest peaks in $SS(\alpha)$. To extract the two most powerful peaks of the signals we experimented with several heuristics, including voting between peaks produced by all four poses and averaging the functions from the four poses. Choosing the right peak is a critical part of the system and still a subject of experimentation, but for the results presented in this paper we use the signal pattern obtained through the following procedure: samples from all periods for all poses are each shifted in the interval $[0, 2\pi)$; samples are then sorted based on their corresponding angle and a smoothing Gaussian filter with the size 10% of the period is applied; the two strongest peaks are retained.

An $SS(\alpha)$ measurement is shown in Figure 4. As we expected, there are several peaks, or local maximums, indicated with arrows from the center in the polar representation. The darker arrow indicates the peak towards the actual position of the mobile, within a few degrees. In this case, the best direction is the strongest signal, but that is not always the case. We measured $SS(\alpha)$ in 32 points, each with four orientations, for five base stations, and ranked the peaks based on their SS value. As shown in the right of Figure 6, we found that in 90% of cases the best direction is either the first, or the second peak. If we knew the best of the two, it would have an error of $22°$ standard deviation from the true direction of the mobile. Not knowing which of them is best, we have to use both of them in triangulation.

It is now important to see how the non optimal peaks are distributed, because they have to be somehow integrated in the triangulation procedure. Even when working with the strongest two angles, it is not known which of them points towards the mobile. In the left Figure 6, we see that the total number of maximums obtained is between 2 and 8,

with an average of 4.5 and a seemingly Poisson distribution. In Figure 7, we eliminate the best peak and find that the other maximums lie mostly away from the it. Only 15% of the times other maximums are in a $90°$ sector towards the mobile, and only 33% of the times in a $180°$ sector.

Based on these statistics, it is likely that working only with the strongest two maximums will cover most situations, while the other maximums are mostly grouped in a direction away from the mobile. The next question is how to use the two angles communicated from each base station to infer a position for the moving mobile. In fact, in one of the methods proposed below, we can use not just the two strongest directions, but the entire shape of $SS(\alpha)$.
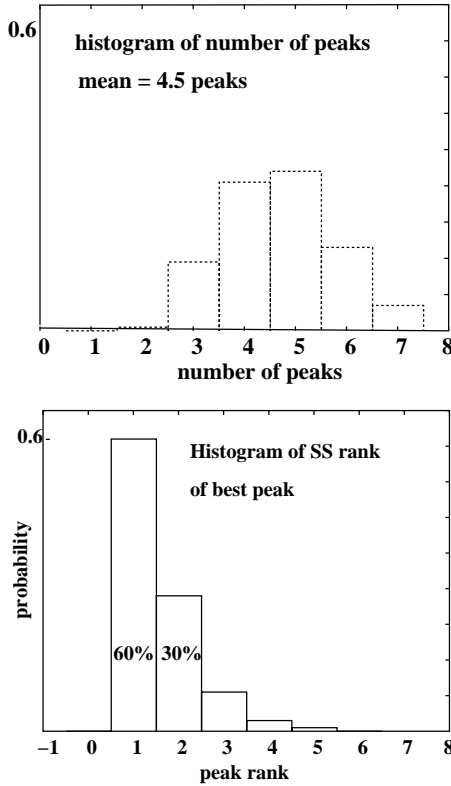
3.   POSITIONING WITH VORBA

In this section, we detail several methods of determining the position based on the information provided by base stations. Two of them are based on angles derived from $SS(\alpha)$, and the last one is based on ranges derived from the average signal strength $\frac{1}{2\pi} \int\limits_{0}^{2\pi} SS(\alpha)d\alpha$.

3.1   Positioning using discrete angles

The problem of intersecting lines on a plane can be as simple as solving a linear system with one equation describing each line, were it not for the errors that can affect these lines, and for the fact that triangulation actually uses half lines. Since linear system solving typically optimizes for the sum of squares of vertical offsets, it may not be the best method to intersect lines affected by AOA errors. In our case a line is defined by the position of a known base station, and an angle affected by error. As will be shown later, this error has a normal unbiased distribution for measurements in a large departmental building.

Positioning using lines from several base stations can then be cast as an estimation problem. Let $\beta_i(\mathbf{x})$ be the true angles at which the base station i sees the mobile, α_i the measured angles, σ_a the variance of the AOA measurements, and $\mathbf{x}$ the position of the mobile. Based on the measurements

Figure 6: Signal peaks are ranked based on the signal strength of the peak.

Figure 7: Non optimal peak distribution

optimized for the sum of modules, as opposed to the sum of squares also because the distribution of the angle errors may be Laplacian, and not Gaussian. In the implementation, to optimize for the objective we used the Nelder-Meade simplex method that is provided by the command `nmsmax` in the free software package octave-forge [7].

3.1.1 Choosing the best angle

Most positioning schemes, including triangulation, assume that an estimate angle towards landmarks is available with a certain error. In our case however, we have two angles, and the best peak is among them in only 90% of the cases. Even if we work with two angles per base station, the question remains which of the two angles to use. In 10% of the cases, both angles will point away from the mobile, so we need a method to identify "bad" angles. For n of base stations, using all combinations yields a maximum of 2^n possible intersections to be optimized using objective function (2).

Fortunately, many of these intersections happen either outside the feasible space (outside the building), or do not corroborate among all base stations. In order to avoid the exponential number of intersections, we first compute in $O(n^2)$ time a $2n \times 2n$ boolean incidence matrix A describing whether any two directions intersect inside the feasible space. An additional $2n \times n$ matrix B is obtained by adding columns of the first matrix to describe incidence between a given direction and a base station (any direction from that base station). If $B(d, j) = 0$ it means that direction d from base station $\left\lceil \frac{d}{2} \right\rceil$ does not intersect any direction from base station j, therefore it can be eliminated completely as a candidate. But most combinations are eliminated during the direction assignment phase, when the matrices indicate that an assignment conflicts with a past or future assignment. When running with 5 base stations, the number of candidates obtained with this method was between 1 and 16, with a median of 5. For 7 base stations, there were a maximum of 24 candidates, with a median of 6.

In order to choose a candidate angle, we compare the ranking of the distances to the base stations with the ranking of the SS to the same base stations. The signal strength used for ranking is obtained by averaging over all the samples, by integration of the signal in Figure 4 (integration of $SS(\alpha)$).

3.1.2 Analysis of discrete angle positioning

An important question is how accurate a position can be obtained only with angles, and how it depends on the main parameters: density of base stations λ, and quality of the angles (variance σ_a^2). The answer helps in provisioning the deployment of the base stations in order to achieve a certain error.

available ($32\,points \times 7\,base\,stations$), the distribution of the angle error looks either Laplacian, or Gaussian, and we used the latter for the following analysis. For the implementation the Laplacian proved better for triangulation, because of the module being more robust to outliers.

Given that the measurements to different base stations are independent, the likelihood function is:

$$P\{\alpha_1, \alpha_2, ...|\mathbf{x}\} = \Pi P\{\alpha_i|\mathbf{x}\} = \frac{e^{-\frac{1}{2\sigma_a^2}\sum|\alpha_i - \beta_i|^2}}{\sigma_a^n\sqrt{(2\pi)^n}} \quad (1)$$

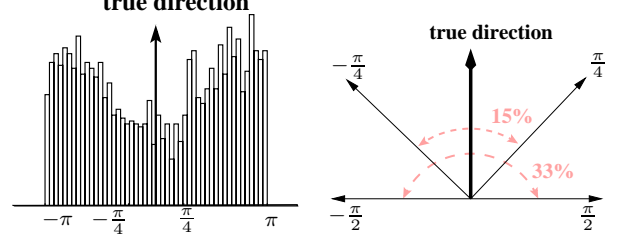The maximum likelihood estimate (MLE) is the solution to the equation:

$$\frac{\partial \ln P\{\alpha|\mathbf{x}\}}{\partial \mathbf{x}} = 0$$

which leads to least square estimate that minimizes the square error in the fit to the angle data α_i. The function to minimize is therefore

$$\sum(\alpha_i - \beta_i(\mathbf{x}))^2 =$$
$$= \sum(\alpha_i - atan2(y - y_i, x - x_i))^2 \quad (2)$$

where x_i, y_i are the coordinates of base station i, $\mathbf{x} = x, y$ is the candidate solution point, and `atan2` is a function that computes the polar angle of $(x - x_i, y - y_i)$. If large outliers are present, then m-estimators [6] can be used to optimize for a different objective giving less weight to the outliers. We

Figure 8: Position error variance depends linearly on σ_a^2 (angular error variance), on $1/\lambda$, and on $\ln^{-1}\frac{R}{R_m}$ (λ=base station density, R=mobile range, R_m=minimum distance to a base station).
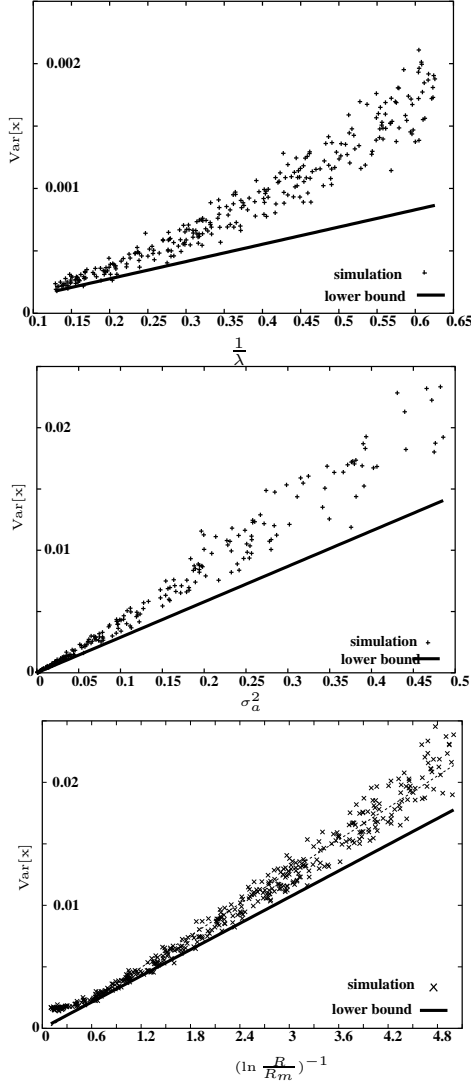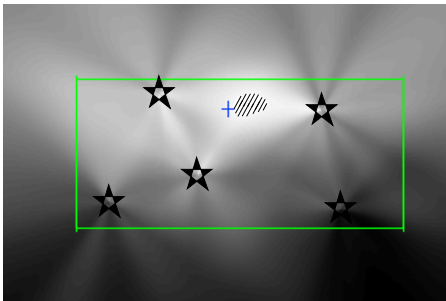

Figure 9: Map of position probabilities: lighter gray indicates higher likelihood.

To simplify the theoretical analysis, we assume the deployment area to be circular, with radius R, and base stations spread with a Poisson spatial distribution with rate λ. The first assumption is reasonable because the coverage will on average be circular in a large enough building (50-100m range for indoors 802.11). The second assumption however is more forced in the light of base station placement that might be preferential for reasons of coverage. It is nevertheless more general than the analysis of a particular placement such as in the corners of the building.

In the appendix it is shown that the error covariance of the position obtained by triangulation is bounded by:

$$Var[x] \quad > \quad \frac{\sigma_a^2}{\lambda \pi \ln \frac{R}{R_m}} \qquad (3)$$

where R_m is the minimum distance from the mobile to a base station. When $R_m \to 0$, the positioning error also becomes 0.
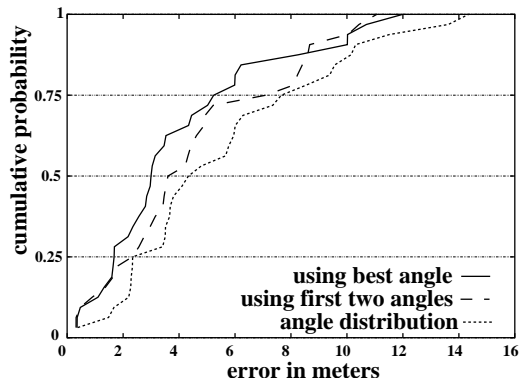
In order to verify the linearity of the positioning error with the angular error and with the inverse of the density, we ran a Monte Carlo simulation in a circle with $R = 1$. Varying λ so that the expected number of base stations is between 5 and 25, for $\sigma = 0.1$, $R_m = 0.1$, produces the points in Figure 8a, showing the standard deviation in the obtained position in a linear relation with $\frac{1}{\lambda}$ together with the line corresponding to equation (3). In Figure 8b, $\lambda = \frac{15}{\pi}$, $R_m = 0.1$ while σ_a - the angle measurement error, is varied in the interval $[0, 0.7]$ radians, or $[0, 45°]$ showing the same linear dependence for the variance of the position. Each point in this experiments is obtained through averaging over 500 different positions of the base stations. The dependence on R_m is also verified in Figure 8c, for $n = 45$, $\sigma = 0.4$, and $R_m \in [0, 0.8]$, which places $\frac{1}{\ln \frac{R}{R_m}}$ in $[0.1, 5]$. In all the cases, the simulation verifies the trend and the bound indicated by equation (3).

The importance of the result implied by equation (3) is for dimensioning and deploying a positioning infrastructure. For example, in order to cut the deviation of the position in half, we need to either half the deviation of the angle measurement, or quadruple the base station density. Increasing the density and reducing the range is in many cases the only way to scale up the data access for more users, cheaper than upgrading all the hardware. The third factor, the minimum distance to a base station R_m, is not as useful as a control knob, as the user does not know how close he is to a base station.

3.2 Positioning using angle distribution

The pattern in Figure 4 indicates directions in which SS is maximum, and can be used to approximate a continuous density of probability for the direction of the mobile. By translating the SS values in the $[0..1]$ interval and scaling so that they sum to 1 for all angles in $[0..2\pi]$, we obtain a probability for each angle around the base station. There is a zero probability for the minimum signal, and proportional values for the other values of the signal. For each base station, the mobile computes the corresponding probabilities for each point of a 25cm grid covering the entire area of interest. The probability of a given point depends on the value of $SS(\alpha)$ in that direction. For each base station we have a probability map for virtually all the points in the area.

Figure 10: Cumulative distribution of positioning error using discrete angles and angle distribution.

form factor, the angle of arrival obtained is quantized in multiples of $45°$. Another possible replacement for the mechanical part of VORBA is to just use eight different directional antennas at the base station. This would provide an angle of arrival quantized to $45°$, but will also increase the total power output of the base station.

Figure 11: Positioning with quantized angles.

(a) positioning with actual angle measurements

(b) positioning with measurements quantized to $45°$

(c) cumulative errors for quantizations at $22.5°$, $45°$, and $90°$.

Aggregating maps from different base stations produces a map like the one in Figure 9. In the map, we can distinguish the positions of the base stations, the same indicated by stars in Figure 5. The lighter areas have higher probability, and the true position in the upper corridor is indicated by a cross. By selecting the points with the highest probabilities, a maximum likelihood region is generated, shown as a hashed shape near the true position. The centroid of this region provides the estimated position.

The performance of positioning in the 32 points is shown as a cumulative distribution in Figure 10. The points in the main part of the building use the five main base stations, while the points separated by the elevators use three of the main base stations and two extra ones shown with hollow stars in Figure 5. The continuous line corresponds to an idealized performance that using the best measurement angle, the dashed line to using the two the angles and the heuristic in section 3.1.1, and the dotted line to the angle distribution method. The angle distribution method provides a slightly lower performance than the discrete method, but automatically deals with outliers. It has a lower complexity with respect to the number of base stations, but higher with the area. It effectively takes into consideration all the possible angles, but it only builds one grid of probabilities for each base station, which are then merged in a final probability map. For the discrete method we used five base stations for the left main part of the building, and three base stations for the left part, separated by elevators, because the heuristic of choosing between two angles is more sensitive to outliers. Only the five main base stations were used for the angle distribution results.
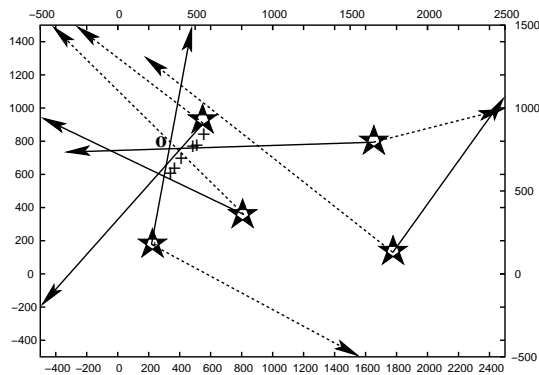
To summarize, the idealized method using the best angle achieves a 2.9m median error, the heuristic using best two angles 3.5m, and the angle distribution method 4.1m.
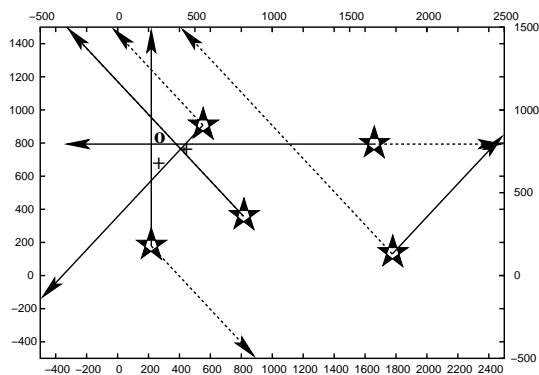
3.3 Positioning using quantized angles

It is convenient to achieve the AOA functionality in a system without moving parts as the one we used. Steerable and switched beam 802.11 antennas are appearing on the market (www.vivato.net and www.paratek.com). They electronically steer the beam to provide preferential amplification for certain directions. In many cases when using a phased antenna array, and also in order to achieve a small

Without hardware that provides quantized angles, we are using our measurements from the VOR base stations discretized after the peak selection phase. In Figure 11a, there is an example of a point that is computed using readings from the five main base stations, using the strongest two peaks. The stars indicate the positions of the five main VORBAs, as represented in figure 5. From each base station, a continuous arrow indicates the first candidate angle, and the dashed arrow indicates the second candidate. The "o" sign indicates the true position of the mobile, while the "+" signs are the candidates considered by the optimization process. In Figure 11b, each angle is replaced with the closest multiple of $45°$ and the position is recomputed.

The quality of positions does not decrease dramatically when using quantization at $22.5°$ (16 directions), or at $45°$ (8 directions) - as shown in Figure 11c. The continuous line is the same in Figure 10 repeated here for reference - it represents the idealized case in which the best angle is provided by an oracle, not by the heuristic proposed in section 3.1.1. The results are somewhat predictable, since the $45°$ quantization introduces an error that is maximum $22.5°$, similar to the measurement errors currently produced by the VOR base stations. For $90°$ quantization (4 directions), two of the points did not get a position, because of directions quantized to parallel lines, and we used the rest of the position errors to get a cumulative distribution. In this latter case, most of the points were optimized at values drawn from the coordinates of the base station themselves, meaning that the x coordinate is the x coordinate of some base station, and the y coordinate is drawn from possibly another one.

Using quantized measurements in 16 directions produces a 3m median error, a small degradation over the 2.9m of the idealized discrete method that uses the best of the two angles.

3.4 Positioning using ranges

In indoor situations, due to unpredictable propagation and fading effects, it is difficult to relate signal strength (SS) to distance. What motivated us to use SS for ranging is the fact that when using VOR base stations, SS can be more reliable being obtained from an integration over all angles, as opposed to a single arbitrary direction measurement. For a random pose of the mobile and for a fixed orientation of a regular base station, a SS sample would be a random value of $SS(\alpha)$ in Figure 4. Using VORBA, we can get a high resolution version of $SS(\alpha)$, and in this example there is a 15dB difference between the maximum and the minimum values that can be obtained. The mean value of $SS(\alpha)$ is a more accurate characterization of the SS from the mobile to VORBA simply because one factor (base station orientation) is eliminated.

We verified this hypothesis by fitting SS values averaged over a revolution on a curve describing distance ρ as a function of SS in dB. Starting with the distance attenuation relation:
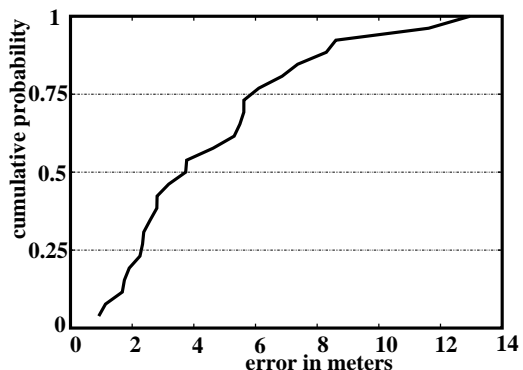
$$P[dBm] \quad = \quad P_0[dBm] - \log_{10}(\frac{\rho}{\rho_0})^n$$

we used the relation

$$\rho(P) \quad = \quad \rho_1 + \rho_0 \exp(\frac{P_0 - P}{n}) \tag{4}$$

where ρ_1 is a shifting factor introduced to allow for dif-

Figure 12: Cumulative distribution of positioning error using ranges from SS fitting

ferences in our hardware, whereas P_0 and ρ_0 are not needed both by the fitting procedure. Parameters ρ_1, ρ_0, and n were fairly similar for all base stations mainly because we used modulo residuals in the fitting procedure, in order to give less weight to large outliers. Especially the two base stations placed on the corridor had very large outliers - a much stronger signal with respect to distance.

The next step was five fold cross-validation: with parameters ρ_1, ρ_0, and n fitted from one base station we generate ranges from the SS readings for all the other base stations, and trilaterate using only these ranges. The procedure is repeated for the five main base stations (including the corner ones, which produce large outliers), and the positioning errors are cumulated into one distribution. Errors similar with those obtained from angles are shown in Figure 12.

The advantage of the method is that using some limited mapping of the signal strength, for example readings at various known points, a reliable curve (4) can be obtained that can be used for most other base stations. Distance measurements from signal strength can also be used to enhance angle measurements and assist candidate selection for discrete angle positioning (section 3.1.1).

The disadvantage is that some sampling of the signal map is needed for the fit, whereas the angle based methods provide positions just with placement of VOR base stations. This issue can be addressed using permanent stationary emitters as in [8, 4] which makes the resampling automatic.

The positioning errors for the range based methods are higher than those obtained with the angle based methods, with a median position error of 4m when fitting all the SS measurements, and 4.5m with cross validation.

3.5 Positioning using ranges and discrete angles

When the mobile knows both the angle α under which it is seen by a base station at (x_b, y_b), and the distance ρ to that base station, it can have an estimate of its position as

$$(x_b + \rho \cos(\alpha), y_b + \rho \sin(\alpha)) \tag{5}$$

What is the accuracy of this estimation? We know that the angle has an error with a standard deviation of $22°$, with the current performance of VORBA. This means that even with a perfect range at 50m, the mobile can easily be placed 20m away from its true location. Ranging from integration

Figure 13: Position covariance is represented as an ellipse. Its size depends on the range ρ, and its orientation on the angle α.

Figure 14: Cumulative distribution of positioning error using angles and ranges
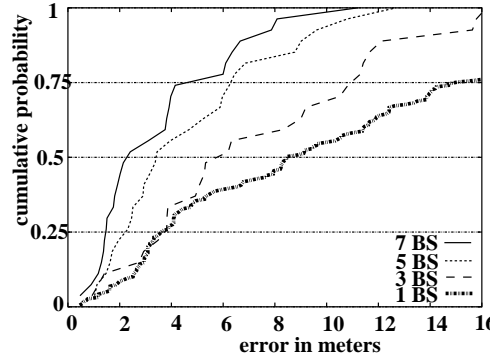

of $SS(\alpha)$ however is not perfect, in fact errors in ranging vary linearly with the actual range, as was found by other projects. After fitting the SS versus distance measurements as mention in section 3.4, we chose one of the non corridor base stations and used the fitting parameters obtained from it. Looking then at all measurements for all base stations, we found that standard deviation is 15-25% of the actual estimated range. We therefore assume $\sigma_r = 0.2$ as a relative deviation, modeling actual error as linear with the measured range $\sigma(range) = \sigma_r \rho$. The median positioning error that we get using relation (5) for a single base station under these conditions is 8.4m. This is an overall figure, as individual base stations provided median errors as low as 3.9m, and as high as 11m.

We can use estimations of positions from several base stations, which we can aggregate into a single estimate. In this case, it is necessary to weigh faraway base stations less than close ones, because for a fixed error in angle measurements, the uncertainty in position increases with distance. To see why that happens, it is convenient to characterize the uncertainty in the position estimation using a covariance matrix. When we consider the polar measurements (ρ, α), with standard deviations of $\sigma_r \rho = 0.2\rho$ for the range, and σ_a for the angle, the covariance matrix of the estimated position becomes:

$$U = \rho^2 \begin{bmatrix} \tan^2(\sigma_a^2) & 0 \\ 0 & \sigma_r^2 \end{bmatrix}$$

This matrix shows the particular case in which the base station is placed in $(0,0)$ and the mobile in $(0, \rho)$ so that angular error produces variation only on the x axis and range error variation only on the y axis. In reality, the covariance matrix describing the possible positions of the mobile is rotated to reflect the actual measurement α. In Figure 13, we see how the covariance matrix is shaped and placed depending on actual measurements. The tip of the arrow indicates the position estimate for the respective base station, and the size of the ellipse the size of the covariance. All the covariances have the same aspect ratio: the width is given by the angle error $\tan(22°) = 0.4$, and the

height by the range relative error $\sigma_r = 0.2$, both of them scaled with the actual measurement ρ. If each base station estimate is $[x_i \, y_i]$ and has covariance U_i, we can combine them in a weighted estimate using a simplified Kalman filter: $[x \, y] = (\sum_{i=1}^{n} [\, x_i \quad y_i] \, U_i^{-1})U$, where $U = (\sum_{i=1}^{n} U_i^{-1})^{-1}$ is the covariance of the estimated position.

In Figure 14, the cumulative distribution of positioning error is presented for 1,3,5, and 7 base stations. As expected, the error improves as we add base stations. For one base station, we used each individual station and cumulated the results, whereas for 3 and 5 we used configurations that favored base stations on the corner of the building. The five base stations in the main part of the building in Figure 5 provided a median error of 3.3m, better than the previous schemes using only angles and only ranges. The curve corresponding to 7 used all the measurements taken. The median error achieved in this latter case is of 2.1m, a clear improvement, provided the rather large error in angle and range estimation.

4. RELATED WORK

Indoor positioning schemes can be classified based on the infrastructure they use, and on the type of measurement medium they employ. In terms of infrastructure, some systems may require specialized instrumentation of the area, and possibly line of sight to the mobiles, whereas others rely on the base stations that are already used for data access. The measurement medium is usually a choice or a combination of RF, infrared, and ultrasound.

Active Badge [9] was one of the early indoor systems, that provided each user with an IR badge that can be read by an IR station that keeps updating user's position in a central database. The position granularity is limited by the density of stations, and the reliance on light can be detrimental in the presence of spurious infrared emissions such as sunlight.

Active Bat [10] is a more recent project of the same group that uses ultrasound instead of IR, and has centimeter accuracy. The bats are supported by a grid of sensors placed on the ceiling which communicate with a central location that performs sensor fusion tasks to provide position through trilateration and also orientation.

Cricket [11] is an MIT project that makes use of ultrasound and per room infrastructure to achieve indoor positioning. It uses the six order of magnitude difference in

the speeds of light and sound to achieve ranging to ceiling beacons. Mobiles then perform their own triangulation to position. Cricket compass [12], a followup project added the measurement of the orientation capability by using three ultrasound receivers on the mobile.

When compared to our VORBA system, these methods have the advantage of a higher accuracy. The disadvantage is that they require extensive support infrastructure, which translates in high cost of deployment. VORBA also requires extra infrastructure, but that is seen as an option for the data providing base stations rather than separate position providing infrastructure.

Another category of schemes relies on measurements of signal strength, and therefore can be used with existing wireless data infrastructure. RADAR [1] was the first system to propose the use of a signal map of the area. Average signal strength for each base station is stored as a fingerprint for each point in a dense grid covering the floor. When querying, a nearest neighbor match in the fingerprint space provides candidates for mobile's position. Ladd et all [2] improve on the the RADAR idea by using full distributions instead of just estimated expectations in each point. At query time, Bayesian inferences are used to search the grid for the most likely positions that fit the distributions of the query point. It improves on positioning performance of RADAR decreasing the median error.

The advantage of this class of systems is that they do not require any additional specialized infrastructure. The deployment cost however is dominated by the necessity of building the signal map of the floor. In addition, the measurements have to be re-taken when the propagation conditions change (people, furniture, etc). VORBA requires the specialized VOR base stations, but saves on the deployment and maintenance by not requiring major updates with changes in the environment. When only angles are used, no training is necessary, while for ranging, some limited training is required.

Landmarc [8] is an RFID based positioning scheme that is in a way similar to RADAR, except that the signal map is built on the fly by previously placed tags. A similar scheme of searching the nearest neighbor in signal strength space is used, but the system adapts more gracefully to changes. A similar idea is explored for 802.11 in [4]. The automatic way to sample the SS map using RFID tags in [8], and stationary emitters in [4] is applicable to VORBA for the range case, when some sampling is needed. However, because we use a parametric method (more robust because of rotation), less samples are needed, as explained in section 3.4.

The Lighthouse [13] location system is used for Smart Dust positioning, but is mentioned here for the similarity with our VOR approach in using a base station with a revolving antenna. Smart Dust are small sensors that only have optical communication capability, so they require line of sight to the base station. The lighthouse rotates its beam and based on the time a node senses the continuous light, it may infer its range to the base station.

Positioning in ad hoc and sensor networks [14, 15, 16, 17, 18, 19, 20, 21, 22] is the problem of positioning nodes based on ranges, angles, or simple connectivity. Most of these approaches face different issues such as the multihop nature of the ah hoc graph, and the requirement of having a distributed algorithm.

5. DISCUSSION AND FUTURE WORK

As presented in the theoretical analysis (section 3.1.2), the factors that affect accuracy of positioning in an idealized setup are density of base stations and quality of angles. Another factor is the actual triangulation procedure that in many cases has to deal with large outliers, therefore a question is how to eliminate an outlier base station reading from the process, without severely degrading the DOP [23]. DOP is dilution of precision due to the geometry of the base stations relative to the mobile, for example DOP is lower when the mobile is outside the convex hull of the base stations.

An aspect also related to the position resolution is the use of more than two discrete angles, but in current implementation we found it to be slower, and with reduced quality, because the extra angles introduce more noise (more candidates).

A number of facts that were found during experimentation may help identify more accurate mapping of the propagation patterns of a building using VOR base stations:

- there is no correlation between angle error and distance. We initially hoped to see such a correlation that could be used in the resolution of candidate positions.

- there is no correlation between angle error and mean signal strength. This is a direct consequence of the previous observation, since SS and distance are strongly correlated, as shown in section 3.4.

- corridors act as waveguides. It is better to place VOR base stations in rooms in order to achieve more accurate angles. This however would reduce "natural" amplification available in a building, which may be detrimental for data access, in order to enhance positioning.

Another way to enhance the accuracy of the $SS(\alpha)$ measurements would be to use a revolving signal at the mobile as well. This would not only simplify the self positioning procedure which currently requires measurements for four poses, but would enable a more accurate picture of the signal at both the base station and mobile.

Currently $SS(\alpha)$ is measured at the base stations because synchronization is easier. We found that the SS as a function of time has a similar shape when measured at the mobile, but with different phase. Provided that the mobile would have a solid frame of reference (compass), it would be possible to measure angles from the mobile towards the base stations, with respect to North. These are equivalent with the angles currently obtained, but the new method would offload some tasks from the base stations. On the other hand, if the mobile doesn't have a compass, it will report all angles to an arbitrary reference and a different type of triangulation must be applied. In this case the mobile only knows the angle under which it sees pairs of VORBAs, case which can be solved by a nonlinear optimization.

VORBA is supposed to be an extension to regular data access points, and not a separate infrastructure to support positioning, therefore an important issue that we are currently investigating is the data performance. The directional antenna used has an amplification of 14dB when facing the mobile, and a very small one in the opposite direction. While bit error rate is a function of signal strength, providing predictable effects on UDP traffic, it is not clear what the effects

of VORBA on TCP performance are. This may be less of an issue if the sweeping antenna pattern is achieved in a non mechanical manner, and is only activated on demand.

A related issue is the time to position, and the possibility to do mobile tracking. The limiting factors are currently the need for several poses for the measurements, and the large rotation period of the base station. Currently, a user has to take four poses, waiting three revolutions of two seconds each yielding a total of $3 \times 2 \times 4 = 24$ seconds, while the triangulation time is almost negligible on current laptops. While the revolution period factor can be greatly reduced using electronic beam forming, operating with fewer poses and with fewer revolutions is less promising because of the inherent variability of the SS with respect to pose and time. If for example, the user takes the measurement only in one pose, it may happen that his body or the screen of the laptop is blocking a direct path to VORBA.

The original VOR used for aircraft navigation is a 2D scheme. VORBA has also been designed to take advantage of the rotation of antennas in a single plane to provide positioning only in that plane. However, the information provided by a single VORBA amounts to placing the mobile on a plane passing through the vertical axis of rotation of the antenna (in fact it is the half plane created by the rotation axis and the direction of the mobile). Two VORBAs are sufficient for positioning in 2D because in addition to the two planes provided by the base stations, there is a third implicit plane of the 2D setup. VORBAs rotating around other axes (other than vertical) could theoretically provide additional planes for intersection. If the angle error is independent of the altitude with respect to the base station, the problem is a simple extension of the 2D case. In the more likely situation in which the error does depend on altitude, a problem for a 3D setup is to provision the directions of the rotation axes in order to minimize positioning errors. An additional aspect in 3D is the possible variation in polarization. In the current realization, the mobile has the card always in the same horizontal plane VORBA uses for rotation, but if base stations rotate in arbitrary planes, the polarization between the base station and the terminal may change in the course of every revolution.

6. CONCLUSIONS

VORBA (VOR base station) is a prototype base station that provides **angle and range measurements** using 802.11 signal strength. Its basic idea is to find the strongest maximums in the signal strength, and use them as the most likely directions in which the mobile can be. The current realization uses a revolving antenna, but a non mechanical implementation, even using quantized angles, would yield similar performance. A positioning architecture using VOR base stations and triangulation has the advantage of not requiring extensive measurements of the signal strength map, while providing performance similar (2.1m median error) with systems that require such sampling. When limited sampling is acceptable, the VOR base station can provide robust range estimations that can be used for trilateration.

7. ACKNOWLEDGMENTS

8. REFERENCES

[1] Paramvir Bahl and Venkata N. Padmanabhan. RADAR: An in-building RF-based user location and tracking system. In *INFOCOM*, Tel Aviv, Israel, March 2000.

[2] Andrew M. Ladd, Kostas E. Bekris, Algis Rudys, Guillaume Marceau, Lydia E. Kavraki, and Dan S. Wallach. Robotics-based location sensing using wireless ethernet. In *ACM MOBICOM*, Atlanta, GA, September 23–26 2002.

[3] Moustafa Youssef, Ashok Agrawala, and Udaya Shankar. WLAN location determination via clustering and probability distributions. Technical report, University of Maryland, College Park, March 2003.

[4] P. Krishnan, A. S. Krishnakumar, Wen-Hua Ju, Colin Mallows, and Sachin Ganu. A system for LEASE: System for location estimation assisted by stationary emitters for indoor rf wireless networks. In *IEEE Infocom*, Hong Kong, March 7–11 2004.

[5] Kamin Whitehouse. The design of calamari: an ad-hoc localization system for sensor networks. Technical report, University of California at Berkeley, 2002. Master's Thesis.

[6] Peter J. Huber. Robust estimation of a location parameter. *The Annals of Mathematical Statistics*, 35(1):73–101, March 1964.

[7] http://www.octave.org.

[8] Lionel M Ni, Yunhao Liu, Yiu Cho Lau, and Abhishek Patil. Landmarc: Indoor location sensing using active RFID. In *IEEE International Conference in Pervasive Computing and Communications 2003 (Percom 2003)*, Dallas, TX, March 2003.

[9] Roy Want, Andy Hopper, Veronica Falcao, and Jonathan Gibbons. The active badge location system. In *ACM Transactions on Information Systems*, volume 10, pages 91–102, January 1992.

[10] Andy Harter, Andy Hopper, Pete Steggles, Andy Ward, and Paul Webster. The anatomy of a context-aware application. In *ACM MOBICOM*, pages 59–68. ACM Press, 1999.

[11] N.B. Priyantha, A. Chakraborty, and H. Balakrishnan. The cricket location-support system. In *ACM MOBICOM*, Boston, MA, August 2000.

[12] N.B Priyantha, A. Miu, H. Balakrishnan, and S. Teller. The cricket compass for context-aware mobile applications. In *ACM MOBICOM*, Rome, Italy, July 2001.

[13] Kay Römer. The lighthouse location system for smart dust. In *ACM/USENIX Conference on Mobile Systems, Applications, and Services (MobiSys 2003)*, pages 15–30, San Francisco, CA, USA, May 2003.

[14] J. Hightower, G. Boriello, and R. Want. SpotON: An indoor 3D location sensing technology based on RF signal strength. Technical Report CSE Technical Report 2000-02-02, University of Washington, February 2000.

[15] Nirupama Bulusu, John Heidemann, and Deborah Estrin. GPS-less low cost outdoor localization for very small devices. In *IEEE Personal Communications Magazine*, Special Issue on Smart Spaces and Environments. October 2000.

[16] A. Savvides, C.-C. Han, and M. Srivastava. Dynamic

fine-grained localization in ad-hoc networks of sensors. In *ACM MOBICOM*, Rome, Italy, 2001.

[17] S. Čapkun, M. Hamdi, and J.-P. Hubaux. GPS-free positioning in mobile ad-hoc networks. *Cluster Computing*, 5(2), April 2002.

[18] Dragoş Niculescu and Badri Nath. Ad hoc positioning system (APS). In *GLOBECOM*, San Antonio, November 2001.

[19] L. Girod and D. Estrin. Robust range estimation using acoustic and multimodal sensing. In *International Conference on Intelligent Robots and Systems*, Maui, Hawaii, October 2001.

[20] L. Doherty, L. E. Ghaoui, and K. S. J. Pister. Convex position estimation in wireless sensor networks. In *IEEE INFOCOM*, Anchorage, AK, April 2001.

[21] Chris Savarese, Jan Rabaey, and Koen Langendoen. Robust positioning algorithms for distributed ad-hoc wireless sensor networks. Technical report, Delft University of Technology, 2001.

[22] Yi Shang, Wheeler Ruml, Ying Zhang, and Markus Fromherz. Localization from mere connectivity. In *ACM MOBIHOC*, Annapolis, MD, June 1-3 2003.

[23] R. Yarlagadda, I. Ali, N. Al-Dhahir, and J. Hershey. Geometric dilution of precision (GDOP) bounds and properties. Technical Report 97CRD119, GE Research & Development Center, 1997.

APPENDIX

A. LOWER BOUND FOR ANGLE ONLY POSITIONING

The Cramér-Rao lower bound is method that sets a lower bound on the variance of *any* unbiased estimator. In our case the triangulation problem is cast as an estimation problem by considering the true position $\mathbf{x}$ as the parameter to be estimated.

Given a circle with radius R, n VOR base stations are Poisson distributed with density λ. Coordinates of the base stations are (x_i, y_i) Cartesian, and (ρ_i, β_i) polar. Assume the mobile is in the center of the circle, $E[\mathbf{x}] = (0,0)$, but rhe angle readings $\alpha = [\alpha_1, ..., \alpha_n]$ are described by equation (1). The likelihood of being at position $\mathbf{x} = [x\ y]$, after having seen bearings from n base stations is:

$$
\begin{aligned}
L(\mathbf{x}|\alpha) &= \ln(p(\alpha|\mathbf{x})) \\
&= -\ln(\sigma^n \sqrt{(2\pi)^n}) - \frac{1}{2\sigma^2}\sum(\alpha_i - \beta_i)^2
\end{aligned}
$$

Using the likelihood, define

$$
I = -\int_{-\infty}^{\infty} \begin{bmatrix} \frac{\partial^2 L}{\partial x^2} & \frac{\partial^2 L}{\partial x \partial y} \\ \frac{\partial^2 L}{\partial x \partial y} & \frac{\partial^2 L}{\partial y^2} \end{bmatrix} p(\hat{\rho}; \mathbf{x}) d\hat{\rho}
$$

and the bound on the covariance of the position is given by I_{11}^{-1}. Second derivative of the likelihood $\frac{\partial^2 L(x|\alpha)}{\partial x^2}$ reduces to:

$$
\sum_{i=1}^{n} \frac{2(x - x_i)(y - y_i)(\alpha_i - \beta_i) - (y - y_i)^2}{\sigma^2 \rho_i^4}
$$

where $\rho_i = \sqrt{(x - x_i)^2 + (y - y_i)^2}$ is the distance form the point to the base station i.

$$
-\int_{-\infty}^{\infty} \frac{\partial^2 L}{\partial x^2} p(\alpha|\mathbf{x}) d\alpha = \sum_{i=1}^{n} \frac{(y - y_i)^2}{\sigma^2 \rho_i^4} \tag{6}
$$

$$
-\int_{-\infty}^{\infty} \frac{\partial^2 L}{\partial x \partial y} p(\alpha|\mathbf{x}) d\alpha = \sum_{i=1}^{n} \frac{(x - x_i)(y - y_i)}{\sigma^2 \rho_i^4} \tag{7}
$$

because $E[\alpha_i] = \int_{-\infty}^{\infty} p(\alpha_i|\mathbf{x}) d\alpha_i = \beta_i$.

To compute (6) and (7), we use polar coordinates (ρ_i, β_i) for all cartesian points (x_i, y_i): $\frac{y - y_i}{\rho_i} = \sin(\beta_i)$, $\frac{x - x_i}{\rho_i} = \cos(\beta_i)$, $(6) = \frac{n}{\sigma^2} E[\frac{\sin^2(\beta)}{\rho^2}]$ and $(7) = \frac{n}{\sigma^2} E[\frac{\sin(2\beta)}{2\rho^2}]$. It is known that the polar coordinates are independent, so we only need to compute $E[\frac{1}{\rho^2}]$, $E[\sin^2(\beta)]$, and $E[\sin(2\beta_i)]$. The distribution of β is uniform, with pdf $f_\beta = \frac{1}{2\pi}$, and it can be shown that $f_{\cos^2(\beta)}(s) = \frac{1}{\pi\sqrt{s(1-s)}}$, which yields $E[\sin^2(\beta_i)] = E[\cos^2(\beta_i)] = \frac{1}{2}$. Similarly, $f_{\sin(2\beta)}(s) = \frac{1}{\pi\sqrt{1-s^2}}$, with an expectation $E[\sin(2\beta_i)] = 0$.

The c.d.f. of distances to base stations is $F_\rho(s) = \frac{s^2}{R^2}$. Let $m = \frac{1}{\rho^2}$ a random variable with $m \in [\frac{1}{R^2}, \infty)$.

$$
F_m(s) = P\{m < s\} = P\{\frac{1}{\rho^2} < s\} = 1 - \frac{1}{sR^2}
$$

$$
E[m] = \int_{\frac{1}{R^2}}^{\infty} s f_m(s) ds = \frac{1}{R^2} ln(s)\Big|_{\frac{1}{R^2}}^{\infty} = \infty
$$

The interpretation of this is that error can become arbitrarily small when the mobile is getting infinitely close to the base station. For this reason, we use R_m - the minimum distance to the base station a mobile can have.

$$
E[m] = \frac{1}{R^2} ln(s)\Big|_{\frac{1}{R^2}}^{\frac{1}{R_m^2}} = \frac{2}{R^2} \ln \frac{R}{R_m}
$$

$$
A = nE[\frac{\sin^2(\beta)}{\rho^2}] = nE[\sin^2(\beta)]E[\frac{1}{\rho^2}]
$$

$$
A = \frac{n}{R^2} \ln \frac{R}{R_m} = \pi\lambda \ln \frac{R}{R_m}
$$

$B = 0$, which produces
$I = \frac{n}{\sigma^2 R^2} \ln \frac{R}{R_m} I_2$, where I_2 is the 2×2 identity matrix.

$$
Var(x) > I_{11}^{-1} = \frac{\sigma^2}{\pi\lambda \ln \frac{R}{R_m}}
$$