

UCLA

Papers

Title

Ad-Hoc Localization Using Ranging and Sectoring

Permalink

<https://escholarship.org/uc/item/51n0073n>

Authors

K. Chintalapudi
A. Dhariwal
R. Govindan
et al.

Publication Date

2004

Peer reviewed

Ad-Hoc Localization Using Ranging and Sectoring

Krishna Kant Chintalapudi, Amit Dhariwal, Ramesh Govindan, Gaurav Sukhatme
 Computer Science Department,
 University of Southern California, Los Angeles,
 California, USA, 90007.

Abstract—Ad-hoc localization systems enable nodes in a sensor network to fix their positions in a global coordinate system using a relatively small number of *anchor nodes* that know their position through external means (e.g., GPS). Because location information provides context to sensed data, such systems are a critical component of many sensor networks and have therefore received a fair amount of recent attention in the sensor networks literature. The efficacy of these systems is a function of the density of deployment and of anchor nodes, as well as the error in distance estimation (*ranging*) between nodes.

In this paper, we examine how these factors impact the performance of the system. This examination lays the groundwork for the main question we consider in this paper: Can the ability to estimate *bearing* to neighboring nodes greatly increase the performance of ad-hoc localization systems? We discuss the design of ad-hoc localization systems that use range together with either bearing or imprecise bearing (such as sectoring) information, and evaluate these systems using analysis and simulation.

I. INTRODUCTION

Sensor network localization has been an active area of research for the last few years. For sensor networks, and more generally for networks of embedded systems, the ability for nodes to determine their position through automatic means is recognized as an essential capability. The community has made great strides in ranging technologies, systems for infrastructure-based localization, and algorithms and techniques for ad-hoc localization (Section II). This last class is the subject of this paper.

In an ad-hoc localization system, nodes determine their position in a common coordinate system using a number of *anchor nodes* that already know their location (through some external means, such as GPS [1]) in that coordinate system. These systems assume all nodes possess a *ranging* capability (the ability to estimate distances to other nodes). Using their range estimates, nodes use one of several *distributed position fixing* techniques to determine their positions in the coordinate system.

There are two characteristics that are highly desirable in a distributed ad-hoc localization system¹; in fact, we assume that these are design *requirements* for such systems. The first requirement is that the performance of such a system be relatively insensitive to anchor placement, as long as the anchors are not placed in a degenerate configuration. From a sensor network perspective, this is desirable since it may often be difficult to engineer anchor placements in the environments

that these networks are deployed. Another way of saying this is that an ad-hoc localization system permits unplanned anchor placement. A second requirement is that relatively few anchors be necessary for obtaining good localization performance. This requirement is motivated by the fact that in some environments it may be difficult to obtain position estimates through external means (e.g., because GPS signals can be significantly attenuated/obstructed through foliage). We argue that ad-hoc localization systems should work well with *an order of magnitude* fewer anchors than nodes. This rule-of-thumb is motivated by a systems argument; if one in two or three nodes are required to be anchors, it will significantly constrain the deployment of such a system.

We begin this paper (Section III) by evaluating the performance of a range of ad-hoc localization techniques proposed in the literature. The performance of ad-hoc localization depends upon several factors: the accuracy of ranging, the density of node placement, the relative density (fraction) of anchors, as well as the particular position fixing schemes in use. Using both analysis and extensive simulations, we find that ad-hoc localization systems begin to perform acceptably only at node densities *well beyond the density required for network connectivity*. We find this to be rather pessimistic—being required to deploy more resources to get a component of a system working seems undesirable from an architectural standpoint. Moreover, we argue that this is a fundamental limitation of ad-hoc localization systems that use ranging devices only, rather than a shortcoming that can be remedied by designing better localization schemes.

We then consider whether adding the ability to estimate *bearing* to neighboring nodes can *qualitatively* improve the performance of ad-hoc localization schemes (Section IV). We show that there exists a highly accurate position fixing scheme that uses *both*² range and bearing information in order to localize nodes, at node densities comparable to that required for network connectivity. This is obviously an idealization, since it is unclear if accurate bearing estimation devices can be built at the form factors and energy-levels that sensor network nodes require.

What is more feasible, perhaps, is the ability to *approximately* detect bearing. Guided by this observation, we examine whether devices that enable nodes to place neighbors within *sectors* can enable acceptable ad-hoc localization performance at node densities that are sufficient for connectivity (Section V). We show that there exists a simple iterative scheme

¹Our focus in this paper is on distributed ad-hoc localization systems. Henceforth, when we use the term “ad-hoc localization system”, or simply “localization system”, we mean this class of systems.

²This is an important and unique contribution of the paper; the only other piece of work that uses bearing information [2] does *not* consider the estimation of node positions jointly using range and bearing.

that can provide attractive ad-hoc localization performance using both range and sector estimates. We conclude the paper by arguing that, given the importance of node localization in a sensor network, the community should invest some effort in sector estimation devices that will enable practical ad-hoc localization systems.

II. RELATED WORK

In recent years, there has been significant work in localization for sensor networks and networks of embedded devices. Localization has also been an active area of research within the robotics community for several years now. We briefly review the current state of knowledge in localization systems. Our intent is not to be exhaustive in our coverage of the literature; rather, we list representative pieces of work that help put this paper in context.

Ranging: An important focus of the localization literature has been robust techniques for estimating distances between nodes (*ranging*). The networking community has focused on two classes of ranging techniques: RF-based ranging and acoustic ranging. For the purposes of this paper, the exact choice of ranging technique is not important, but we include a discussion of these techniques because, as we shall see, they do have some bearing on our discussion of ad-hoc localization.

RF-based ranging, as exemplified by the SpotON [3] and Calamari [4] systems, is based on the premise that by measuring received signal strength a receiver can determine its distance to a transmitter. This presumes that RF propagation in an environment can be accurately characterized by a simple path loss model, whose parameters are known. Using this technique, nodes can estimate distances to all neighbors within radio range. Range errors upwards of 10% have been reported in the literature [4], usually after a fairly involved calibration step that estimates the path loss parameters and adjusts for variations in transceiver characteristics.

A second class of ranging schemes is based on measuring the time-of-flight of an acoustic or ultrasound signal [5], [6]. More precisely, these techniques measure the difference in arrival times of simultaneously transmitted radio and ultrasound signals, then estimate distances knowing the speed of sound. Some approaches in acoustic ranging use spread spectrum approaches for resilience to multipath, and employ techniques to correct for latencies induced by other system components [5]. Such techniques provide an order or magnitude better accuracy (1-2% error) over distances of 3-6 meters (*i.e.*, significantly lower than the nominal radio range of sensor platforms such as the Mica motes).

Finally, we note that the robotics community has, for some time now, used highly accurate laser range-finders (which exhibit less than 0.1% ranging error). It is unclear, however, whether such devices can be manufactured in the form factors and energy-consumption levels required of sensor nodes.

Position Fixing: A large body of work has examined algorithms for ad-hoc localization schemes. We taxonomize this literature later (Section III-A), and only briefly introduce the related work here. Perhaps the earliest pieces of work in the area of sensor network localization can be attributed to

Niculescu *et al.* [7] and Savvides *et al.* [8], [6]. In the former, nodes first estimate their distances to anchors using one of several techniques (DV-hop, DV-distance, and a Euclidean scheme), then fix their own position using these distances. The latter proposes an elegant N -hop multilateration scheme which we discuss later. That work also discusses a Kalman filtering based position refinement phase to improve position estimates. In more recent work, Savvides *et al.* discuss the error characteristics and the dependence on network size and anchor density of their schemes [9]. Savarese *et al.* [10] propose a three-phase scheme, where they obtain crude position estimates and use an iterative multi-lateration based scheme to refine the estimates. Each node uses the locations of its neighbors to multi-laterate and re-estimate its position. Finally, Langendoen *et al.* [11] discuss a fairly detailed comparison of all the above schemes in the face of ranging errors, different node densities and anchor fractions. Our work draws heavily upon theirs.

Somewhat tangentially related to our work is other literature on localization. Examples of such work include: position inference using specialized beacons [12], optimization methods for centrally estimating locations [13], in-building localization systems that enable position for context-aware applications [14], [15], [16], systems for estimating orientation of handheld devices [17], and systems for placing nodes within a relative coordinate system [18]. We do not survey this literature in detail.

Finally, predating the work in sensor network localization, there exists significant work in robotic node localization. This has discovered many tools that are applicable to ad-hoc localization systems: Kalman filter approaches to cooperative localization [19], refinement of probabilistic estimates of robot positions [20], and maximum-likelihood estimation of relative robot positions [21]. While not directly relevant to our work, they have influenced the design of some of our localization schemes.

III. AD-HOC LOCALIZATION USING RANGING

As we have seen in Section II, most of the *distributed* ad-hoc localization schemes studied to date leverage the ability of nodes to estimate distances to other nodes using ranging. Our goal in this section is to fairly extensively evaluate, through analysis and simulation, a representative sampling of such ad-hoc localization schemes, with the intent of understanding their error characteristics as a function of node and anchor density. Based on the simulations we gain some insight into the fundamental limitations (qualitatively and quantitatively) of range-only methods. In the rest of this section, for conciseness and to distinguish this class of localization approaches from those we consider later in the paper, we use the term *r*-only localization to describe these methods.

In some recent work, Langendoen *et al.* [11] have performed similar analyses. However, our evaluations (presented in this section) differ from theirs in several respects. In our simulations, we use larger topologies, examine a wider range of density and anchor ratios, and use a slightly more sophisticated ranging error model. We also introduce some

improved localization schemes of our own into the mix of schemes, and examine more metrics to explore the efficacy of ad-hoc localization.

A. A Taxonomy of Ad-Hoc Localization Schemes

Before describing the schemes we evaluate, we taxonomize r -only localization methods. Following Langendoen *et al.* [11], such localization schemes can be divided into three distinct sub-problems or stages: estimating distance to anchors; getting an initial position estimate; and iteratively refining the position estimate. When only ranging techniques are available, a sensor needs to estimate its distance to at least three anchors in order to localize itself into a global coordinate system common to those anchors. All r -only localization schemes begin with each sensor trying to estimate its distance (or some approximation thereof) to anchors — the first sub-problem. Having obtained distances to three anchors, obtaining an initial position estimate reduces to a simple geometric problem commonly referred to as *lateration*. With these estimates, some r -only localization approaches incorporate a refinement scheme that improves position estimates in the face of ranging error.

In the following subsections, we describe each of these three stages, and discuss the r -only localization schemes we evaluate.

Estimation of distance to an anchor: Broadly speaking, there are two classes of approaches to estimating distance to anchors: *topological* and *geometric*. Topological techniques are content to roughly estimate distance to anchors using information obtained by message flooding. Geometric techniques, on the other hand, more carefully compute distances to anchors and generally exhibit higher accuracy.

One example of a topological approach is the *DV-dist* approach by Niculescu *et al.* [7]. In this approach, each anchor initiates a (possibly constrained) flood. As the flooding message propagates, it accumulates the range measurement corresponding to each hop. Each node measures its distance to an anchor as the minimum accumulated distance. This approach consistently over-estimates the distance to anchors and is subject to accumulation of error over multiple hops. This error effects the second phase of ad-hoc localization significantly, since lateration is very sensitive to range errors. A slight variant of this approach estimates distance using hop-counts to anchors, a technique called DV-hop [7]. In Section III-B, we evaluate an ad-hoc localization scheme using the DV-dist approach.

We choose two representative qualitatively different geometric approaches in our evaluations: Savvides *et al.* [6] and Niculescu *et al.* [7]. In the former scheme each anchor defines a coordinate system locally, with itself as the origin. For the scheme to begin, the anchor should have at least three neighbors within its ranging neighborhood. One of the neighbors can be randomly chosen to lie on the x -axis, its range from the anchor being its coordinate. Knowing their range to the anchor and to each other, other neighbors of the anchor (one hop away) can calculate their locations in this coordinate system. Following this, nodes two hops away from the anchor now calculate their locations through lateration

using the locations (recently calculated) of the nodes one hop away and the ranges to them. This process can be repeated over multiple hops. While this local coordinate system can be a rotated and/or translated and/or mirror image version of the global coordinate system, the distance of nodes from the anchor defining the coordinate system is an invariant. Hence, nodes can determine their distance from anchors over multiple hops using their location in the global coordinate system. We use the term *relative localization* to describe the scheme. The second scheme differs from the first in that, instead of using lateration, each neighbor votes on a set of possible ranges that the node can have. The position with the maximum number of votes is chosen. Borrowing from [7] we use the term *euclidean scheme* to describe this scheme.

Both the geometric schemes fundamentally require three nodes in their neighborhood which know their ranges/relative positions to at least three different anchors. To formally define this requirement we introduce the notion of *3-connectedness* to an anchor. Let $G = (N, E)$ be a graph representing the connectivity of the sensor network, N be the set of all nodes, and $E = e_{ij}$ the set of all edges. An edge exists between two nodes which are within ranging neighborhood of each other. Further, let A_i represent the set of nodes that node n_i can range to. Let $x_i \in \mathfrak{R}^2$ be the location of node n_i . Also let $L \subset N$, be the set of all anchors, hence, the x_i s for $n_i \in L$ are known, and the remaining x_i s are unknown. Let the set of all nodes 3-connected to an anchor $a \in N$ be denoted by $N_a^3 \subset N$. Formally, $n_i \in N_a^3$ if and only if there exists $n_j, n_k, n_l \in N_a^3 \cap A_i$ such that the sub-graph comprising nodes n_j, n_k, n_l is connected. Furthermore, all the neighbors of anchor a are considered 3-connected to a . A node needs to be 3-connected to three different anchors to determine its position in the global coordinate system.

Such a requirement demands densities with very high average node degree. A significant fraction of the nodes may not be able to localize themselves at all at lower node deployment densities. We show this through extensive simulations in Section III-B.

Some researchers have pointed out that three nodes is not a fundamental requirement. Having two neighbors knowing their location gives rise to two possibilities. In some situations it may be possible to eliminate one of them using information beyond one hop. While this is true, in practice we found that the gain from this tends to be marginal in the face of ranging errors.

Initial position estimate and refinement: Once a node has estimated its distance to three anchors, multi-lateration seems to be the most commonly used approach for obtaining the node's initial position estimate. The performance of multi-lateration depends on the distances and relative locations of the anchors and nodes. It works well if the anchors are far apart and enclose the nodes within their convex hull [9].

Finally, four different refinement schemes have been proposed in the literature. The first is iterative multi-lateration, where nodes continuously estimate their positions using the estimated locations and ranges of their neighbors, and hoping to improve their locations [8]. The second is a Kalman filter based approach, which estimates the Jacobian of the

multi-iteration equations in an attempt to obtain a Gaussian representation for the error in the solution [6]. The third approach is a heuristic proposed in [10] where bounding boxes for the solutions are iteratively shrunk. Finally, the Weighted Least Square (WLS) method by Leon *et al.* [22], is designed to minimize the mean square error in the estimated positions of the nodes. This scheme assumes that the error in initial estimates of the nodes and the measured ranges be Gaussian. This is in general not true since errors due to lateration are non-gaussian and even more so over multiple hops. In addition, WLS requires calculating the pseudo-inverse of a matrix of the size of the number of nodes (which may be very large) and it is not clear that this could be implemented in a distributed manner.

In the following section, we describe a distributed refinement scheme designed on similar lines as WLS but which, instead of minimizing error in absolute node positions, attempts to minimize the relative position error between nodes. As a consequence, the scheme only corrects the node positions relative to each other and does not demand any assumptions regarding the error in initial estimates of the node positions. The scheme also allows for a completely distributed solution. This scheme has independent interest, but also lays the groundwork for some localization schemes we introduce later in the paper.

Iterative Least-Mean-Squared Refinement (LMSR): LMSR attempts to find locations for the nodes which “best” fits (in a least-mean-squares sense) the set of all range-measurements made in the sensor network. A range measurement r_{ij} taken at a node n_i to measure range to another node n_j , gives an equation,

$$\sqrt{(x_i - x_j)^T (x_i - x_j)} - r_{ij} = 0. \quad (1)$$

For every measurement taken in the network we would then have one such equation. Since the number of measurements (hence the number of equations) can far exceed the number of unknown locations, we have an over-determined system. A common way to solve such an over-determined system is to find locations which “best” fit the set of equations. Since each measurement could potentially have different error characteristics, some measurements could be deemed more reliable than others. In such a scenario, preferentially weighting the reliable equations in the equation set could give improved solutions. For this we adopt the most commonly used method for finding a best fit for locations—we find locations that *minimize the weighted mean square error taken over all the equations*.

We label such a scheme LMSR. LMSR basically weights measurements with their observed σ_{ij} , since these are a direct measure of the certainty of the measurements. As we describe below, LMSR is amenable to a distributed implementation. The theory behind LMSR assumes that ranging errors are Gaussian with known standard deviations σ_{ij} corresponding to a range measurement r_{ij} . This assumption is not critical to the scheme but we do find from measurements (see below) that the Gaussian assumption holds for range errors. Furthermore, in an implemented system, we would expect to use the nominal standard deviation obtained from a data sheet, standard deviations obtained by calibrating the devices prior to measurement,

or the standard deviations of multiple range measurements. Even on relatively impoverished hardware (*e.g.* the Mica motes [23]), it will be feasible to compute these quantities, particularly since the computation happens infrequently.

We now describe the theory behind LMSR.

Formally, LMSR attempts to minimize the objective function,

$$J = \sum_{e_{ij} \in E} \sigma_{ij}^{-2} \left(\sqrt{(x_i - x_j)^T (x_i - x_j)} - r_{ij} \right)^2. \quad (2)$$

Higher variance in the measure implies higher uncertainty in the corresponding equation. The objective function in Equation 2 weights each equation inversely proportional to its variance, so that more certain equations contribute to a greater extent in the solution. Note that instead of using σ_{ij}^2 for weighing, in J , other measures of uncertainty for example, measures based on entropy of the error distribution, could be used in case of non-gaussian range error distributions, without any significant changes in the algorithm.

The conditions for local minima can be written as, the set of $|N| - |L|$ simultaneous equations,

$$\sum_{n_j \in A_i} \left(\sigma_{ij}^{-2} + \sigma_{ji}^{-2} - \frac{\sigma_{ij}^{-2} r_{ij} + \sigma_{ji}^{-2} r_{ji}}{\sqrt{(x_i - x_j)^T (x_i - x_j)}} \right) (x_i - x_j) = 0, \quad (3)$$

$$\forall n_i \in N - L, n_j \in N.$$

The above equations can be re-written as,

$$x_i = \frac{\sum_{n_j \in A_i} \left((\sigma_{ij}^{-2} + \sigma_{ji}^{-2}) x_j + \frac{(\sigma_{ij}^{-2} r_{ij} + \sigma_{ji}^{-2} r_{ji})(x_i - x_j)}{\sqrt{(x_i - x_j)^T (x_i - x_j)}} \right)}{\sum_{n_j \in A_i} (\sigma_{ij}^{-2} + \sigma_{ji}^{-2})}, \quad (4)$$

$$\forall n_i \in N - L, n_j \in N.$$

Note that Equation 4 depends *only* on the locations estimates of nodes within n_i 's ranging neighborhood and the ranges they measure. This allows for a distributed implementation of LMSR, where nodes continuously exchange their current location estimates, then use Equation 4 to update their estimates until they converge. LMSR will converge to the nearest local minima in the objective function, hence an initial guess “sufficiently” close to the global minimum will force the nodes to a least-mean-squared fit of the measurements. This initial guess is, of course, obtained from the second stage of r -only localization.

Effect of refinement on the schemes: The geometric approaches, provide fairly accurate estimates for the range from anchors. This, coupled with such range estimates from several anchors provides a fairly accurate initial guess to work with. The unlocalized nodes do not provide an initial guess to work with and hence cannot be refined. As pointed out by Langendoen *et al.* [11], refinement does not improve performance significantly for geometric approaches.

On the other hand, the topological schemes for estimating distances to anchors leave enough room for improvement by refinement. For example, in the *DV-dist* approach at least, the estimates of ranges to anchors is incorrect often resulting in highly inaccurate initial guesses.

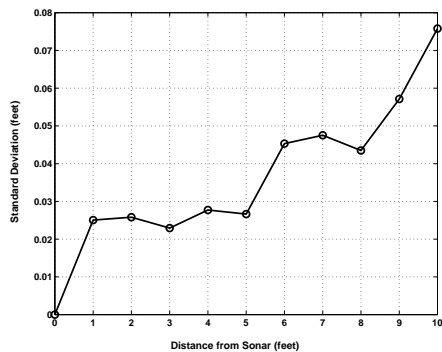


Fig. 1. Sonar ranging error standard deviation as a function of range

B. Evaluating Ad-Hoc Localization Systems

In the previous sections, we taxonomized the various r -only ad-hoc localization schemes known in the literature, and briefly discussed their expected performance. In this section, we use simulation to discuss the performance of r -only localization schemes.

The Schemes: According to our taxonomy (Section III-A), r -only localization can be classified into three phases. For the second of these phases, obtaining an initial position fix, almost all published schemes use multi-lateration. So do we.

For the last of these phases, refinement, we use the LMSR scheme described in Section III-A. This choice is somewhat different from that published in the literature, but we are fairly confident that LMSR is representative of the benefits that refinement can provide. For geometric schemes, refinement provides very little benefit anyway. We found that, for topological schemes (*DV-dist*), the performance of our refinement scheme “closely” matched the results reported by [11]. For these reasons, we do not believe that evaluating other refinements schemes will change the conclusions we make below.

The first, and most crucial, phase in range-only ad-hoc localization is estimating a node’s distance to anchors. This phase presents the most choices. In our evaluations, we choose three schemes, representative of three different classes of our taxonomy: the *Euclidean* [7] geometric scheme that does not estimate relative bearing to anchors; the *relative localization* technique that implicitly computes relative bearing to anchors; and the *DV-dist* topological scheme. We believe that this set adequately represents the state of the art in range-only ad-hoc localization mechanisms.

Simulation Methodology: All our simulations were carried out in a home-grown simulator. We carefully implemented many of the schemes described in the literature, and our results are in agreement with published results. Since our approach used an independent implementation, it represents a validation of most of the prior work in the field.

All our simulations use a field of 1000 randomly deployed nodes (node locations being drawn from a uniform distribution over the region). We chose the ranging distance to be 10m, and we also assume that this is the communication radius.

In our paper we define density as the expected number of nodes in a ranging neighborhood, hence *density is equivalent to node degree in our paper*. The reason for this choice is

that it provides a clearer framework for understanding the fundamental limitation of the schemes. The density of the nodes was controlled by increasing the area of deployment while keeping the number of nodes 1000. For example, an $252\text{m} \times 252\text{m}$ region results in a density such that density is 5, while a $200\text{m} \times 200\text{m}$ region results in a density is 8 nodes. A fraction of nodes (the anchor fraction) were chosen to be anchor nodes.

Most of these assumptions are consistent with the literature. One place where we deviate from the literature is in the use of a different model for ranging error. Most published work, at least in sensor network localization, has assumed that ranging error is independent of the ranging distance. The error model we use is a Gaussian with its standard deviation varying linearly with range, and given by $\sigma(r) = \alpha r$. Consider a uniform error model which uses a fixed standard deviation σ . In comparing this with our distance dependent error model where the standard deviation at the maximum distance was also σ , we see that our error model is more generous (or less pessimal) towards ad-hoc localization schemes; for smaller ranges, it assigns less ranging error. Our conclusions are not affected by this choice of error model; we have verified that using a uniform error model gives qualitatively the same result.

To validate our error model, we conducted measurements using a sonar ranging device on an off-the-shelf robot. Although the robot is significantly more capable than a sensor node, we anticipate that the error characteristics of the sonar will be similar to the ultra-sound ranging devices in consideration for sensor networks [5]. In fact, our error measurements in this device (about 1% over 3-4m) agrees qualitatively with ultrasound ranging error estimates measured elsewhere [5], [8]. Figure 1 shows the variation of standard deviation as a function of distance. Error variance increases with distance, and for the range we consider, a linear fit works reasonable well with an α of approximately 0.7.

Furthermore, the Gaussian assumption is, to a large extent, borne out by measurements. For example, we took measurements on a sonar ranging device at a 5 foot distance. The distribution of samples at this distance is shown in Figure 2 and is visually Gaussian with a slightly longer right tail. With increasing distance, this tail becomes slightly more pronounced.

Metrics: Most of the literature focuses on either *RMS error* or *mean error* as measures of the *accuracy* of ad-hoc localization. Both measures define measure the error of localization as the distance between the computed location and the actual (ground truth) location. RMS error can be skewed by nodes with large localization error, and often gives pessimistic estimates of localization error. Mean error weighs all values of localization errors equally and hence provides a more realistic estimate of the expected value of error. Mean error is defined as,

$$e_{\mu} = \frac{1}{n} \sum_{i=1}^n \sqrt{(x_i - \hat{x}_i)^T (x_i - \hat{x}_i)}. \quad (5)$$

where \hat{x}_i are the estimates of x_i found by the algorithm. In this paper, we use mean error as a measure of localization accuracy, both because it is less skewed than RMS error by

large error values, but also because it is more easily understood as an error measure. We compute mean error only on nodes that were able to obtain a position estimate at all.

Also motivated by our discussion in Section III-A, we introduce a new metric that measures the *extent* of localization. This is essentially a quantile metric and is defined as the fraction of non-anchor nodes that are localized to within 2 meters of their ground-truth position. In our simulations, 2m represents a fairly conservative estimate (about 20% of range). The extent metric serves to differentiate geometric schemes (which rely on 3 nodes to range themselves from an anchor) and topological schemes (which do not). Much of the literature has ignored this metric, choosing to evaluate localization error only in regimes with high localization extent. We believe that this metric brings out a more complete picture of the dynamic performance range of ad-hoc localization schemes.

Simulation Results: We have evaluated our three representative schemes using simulation for three different ranging errors (0.1%, 1% and 5%) and three different anchor fractions (5%, 10% and 20%). Given space constraints, we present results from one combination (20% anchors, 1% error) that represents the state of the art ranging error using ultrasound ranging, and a 20% anchor fraction that represents a fairly generous sprinkling of anchors. Other combinations of errors and anchor densities do not affect our conclusions of Section III-C.

Figure 3 shows the localization extent for our three schemes. Qualitatively, there seem to be two classes of performance: the DV-dist approach (a topological scheme) has higher localization extent at lower node densities than the Euclidean and the relative localization schemes (the geometric schemes), but those two catch up at about a density of 12. Clearly, the 3-connectedness requirement greatly affects the performance of the second class; only 20 to 30% of the nodes are able to get a position fix at all at low densities. In the DV-dist approach, even at low densities, more than 95% of the nodes are able to get a position fix (as predicted by our analysis), but their estimates are very poor in quality (a greater than 20% error in position).

The same two classes are visible when we examine the localization error (Figure 4). The geometric schemes exhibit low mean error (and also low RMS error – we have computed this metric but do not include it in our evaluations) across the entire range of densities because if a node gets a position estimate, it gets a good one. On the other hand, the topological schemes start with high error at low densities and achieve good performance only at densities of 9 or 10. There is some performance difference between the two geometric schemes (Euclidean outperforms relative localization), but we suspect that their performance could be made comparable by applying well-known optimizations. We did not do this since our intent was not to compare similar schemes, but understand *classes* of performance).

In our simulations, we have assumed that anchor nodes have perfect location information since our intent was to quantify the intrinsic error of our localization schemes. In practice, of course, anchors may have imperfect location information. Another potential source of error that our paper does not model is topological anisotropies (e.g. caused by obstructions).

The *qualitative* conclusion we draw from this is that *r-only localization schemes require an average degree of 11-12 nodes within the ranging neighborhood in order to achieve 90% localization and 5% accuracy.*

C. The Feasibility of Ad-Hoc Localization Systems

Our conclusion above represents a somewhat subjective choice (90% extent, 5% mean error). However, we believe these thresholds to be fairly generous from the perspective of a deployed system. Also our evaluations are fairly extensive and borne out by analyses. We argue that *this result is pessimistic*, because in order for *r-only* localization schemes to work, almost *twice as many nodes are required as for mere connectivity* (it is well known that a radio network with uniformly distributed nodes is connected, with high likelihood, if it has an expected number of 5-6 neighbors per node). This is an architectural oddity, since it isn't clear one should *have* to deploy twice as many resources for only one component (albeit a very important one) of a system.

We should note that the literature has long hinted at this pessimism, preferring to evaluate their schemes at relatively high densities. However, we could not find an explicit statement of this pessimism, rather a general acceptance of the state of affairs. We do not mean to disparage existing pieces of work; many of the localization schemes we have referred to have very elegant and intricate algorithms, and we could not (despite trying) improve any of them in any significant way. Rather, we think our observation is *fundamental* of *r-only* localization and optimizations and improved refinement schemes cannot change this. Euclidean schemes require 3-connectedness, which occurs only at high densities. Topological schemes on the other hand rely on the congruency between topological and physical distance for low error and this sets in only at high densities.

There are four counter-arguments to our position, which we now discuss.

Argument 1: Why not just deploy more anchors? This is a reasonable position, and when it is feasible to make 30% or more of the nodes as anchors, then with high likelihood *r-only* localization will begin to work well. We have seen this in our simulations. However, this goes against the basic rationale for ad-hoc localization systems (Section I).

Argument 2: What's wrong with deploying more resources: we need them anyway for increased network lifetimes? On the surface, this is an attractive argument. From an architectural perspective, however, we believe that *requiring* deployments to be redundant so that ad-hoc localization schemes work is undesirable.

But, the question of higher density deployments is a bit more subtle than just that. Note that in our simulations when we say "density" we mean the number of nodes within ranging radius. In drawing our conclusions of the previous section, we *implicitly equated* ranging radius and communication radius. While this assumption is true of RF-based ranging, the ranging errors for this technique are quite high (10% even with significant calibration [4]). The more accurate ultrasound ranging devices have a reported range of 3-6m [8], [5], while

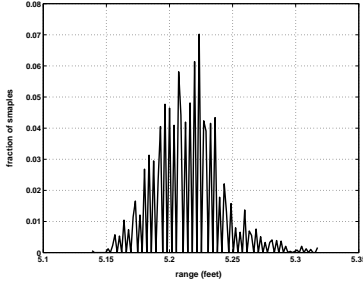


Fig. 2. Sonar Ranging Measurement Distribution

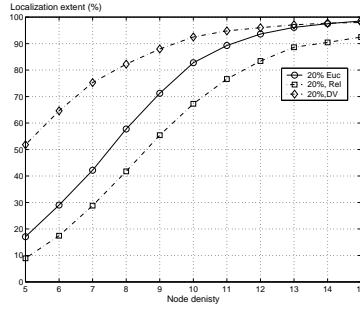


Fig. 3. Localization extent

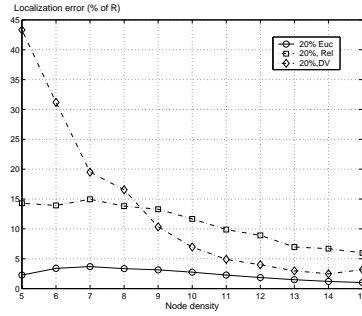


Fig. 4. Localization Error

the current radios on the Mica motes already have a range of tens of meters and the new Chipcon radios are reported to have a greater range. This means that r -only schemes might, in theory, require impossibly high density of nodes within the communication radius. For example, if the radio range is 20 m, while the ranging radius is only 10 m. About $11 \times 2^2 = 44$ nodes within radio will be needed to successful localization, compared to about 5-6 needed for mere radio connectivity for an r -only scheme. This is not a fundamental limitation, however, and there are many possible approaches to circumvent this: controlling the transmit power on the radios (which might be necessary if node deployments need to be fairly dense anyway for application reasons), using UWB localization (the current generation UWB localization methods are reputed to exhibit high error), or finding low-power ranging devices which can range at higher distances (this may be difficult: a fairly good off-the-shelf sonar, the Polaroid 6500, which can range up to 35 ft. requires a 2A current draw when taking measurements).

Argument 3: Random anchor placement makes no sense, so we should perhaps consider more constrained placements? While more constrained placements, like placement of anchors along the periphery of the network can improve localization error, we do not believe they can increase localization extent at low densities. Topological schemes rely on the congruency of topology and physical distance to get reasonable error, and again this cannot be significantly improved by preferentially placing more anchors. We have also verified this using our simulations, and the graphs are omitted for brevity.

Argument 4: There may be accurate schemes that do not require 3-connectedness which are yet to be discovered. While we cannot categorically deny such a possibility, based on our experience we do not believe that an r -only scheme will be able to out-perform the existing schemes by orders of magnitude.

Having painted this rather pessimistic picture, we now ask: Is there an alternative solution?

IV. USING BEARING INFORMATION FOR HIGHER LOCALIZATION EXTENT

We have argued that r -only localization schemes exhibit fairly poor performance. In this section, we discuss whether ad-hoc localization schemes could exhibit improved localization performance (low mean error and high localization extent

even at low node densities) if each node had the ability to estimate bearing to its neighboring nodes, in addition to range. To our knowledge, this question has not been considered in the literature before.

In this section, we first describe a distributed algorithm for range-and-bearing based ad-hoc localization—we use the term $r-\theta$ localization to describe this class of ad-hoc localization systems. We end this section with a discussion of the performance and practical feasibility of $r-\theta$ localization.

Iterative Least-Mean-Squares Refinement using Range and Bearing: The key intuition for the efficacy of $r-\theta$ localization can be seen from a simple geometric argument. Suppose that a node n_i has computed its position x_i (note that, per our notation in Section III-A x_i is a *vector* that represents the node's position). Now, if n_i is able to measure its range (r_{ij}) and its bearing (θ_{ij}) to another nearby node n_j , then the coordinates of n_j would given by

$$x_j = x_i + \begin{bmatrix} r_{ij} \cos(\theta_{ij}) \\ r_{ij} \sin(\theta_{ij}) \end{bmatrix}. \quad (6)$$

Thus, a node needs only *one* neighbor that has localized itself to accurately estimate its own position. Contrast this with the 3-connected requirement of r -only schemes that use geometric methods for estimating distances to anchors.

But, in the presence of range and bearing measurement errors, this simple approach could propagate these errors throughout the network. How, then, can we accurately estimate position in a distributed fashion?

The approach we use (Least-Mean-Squares refinement using Range and Bearing, or LMSRB) is very similar to the iterative refinement scheme described for r -only localization in Section III-A. Each $r-\theta$ measurement leads to the equation,

$$x_i - x_j - \Delta x_{ij} = 0, \quad (7)$$

here Δx_{ij} is given by the second term on the right hand side of Equation 6.

As with LMSR, LMSRB finds the locations which ‘best’ fit the set of all readings. The basic idea is to cast the position estimation problem within an optimization framework, and then observe that minimizing the objective function is amenable to re-formulation in a manner that would permit distributed, iterative, computation. LMSRB also weights $r-\theta$ measurements based on their level of certainty. For this we assume the $r-\theta$ measurements to have zero mean Gaussian

error. Assuming relatively small errors (up to about 10° cones in bearing errors) and ignoring third and higher order moments, Δx_{ij} can be approximated by a Gaussian random vector with covariance matrix:

$$\Delta M_{ij} = \begin{bmatrix} \sigma_{r_{ij}}^2 \cos^2 \theta_{ij} + & \frac{(\sigma_{r_{ij}}^2 + r_{ij}^2 \sigma_{\theta_{ij}}^2) \sin 2\theta_{ij}}{2} \\ \sigma_{\theta_{ij}}^2 r_{ij}^2 \sin^2 \theta_{ij} & \\ \frac{(\sigma_r^2 + r_{ij}^2 \sigma_{\theta_{ij}}^2) \sin 2\theta_{ij}}{2} & \sigma_{r_{ij}}^2 \sin^2 \theta_{ij} + \\ & \sigma_{\theta_{ij}}^2 r_{ij}^2 \cos^2 \theta_{ij} \end{bmatrix}. \quad (8)$$

One can use the covariance matrix as a weighting function to formulate the localization problem as a weighted least-squares optimization.

Our approach seeks to minimize the objective function,

$$J = \sum_{e_{ij} \in E} (x_i - x_j - \Delta x_{ij})^T \Delta M_{ij}^{-1} (x_i - x_j - \Delta x_{ij}). \quad (9)$$

The objective function in Equation. 9 is a commonly used weighted least mean square formulation, based on the Mahalanobis norm. We expect that LMSRB will not be significantly perturbed by any slight deviation from the Gaussian, since the Gaussian assumption merely provides a weighting mechanism. Let the standard deviations in measurements, r_{ij} and θ_{ij} be $\sigma_{r_{ij}}$ and $\sigma_{\theta_{ij}}$ respectively. Note, that each measurement could potentially have a different standard deviation. This allows for the use of sensors with different characteristics and variation of standard deviation with range.

It turns out that the conditions for local minima in this formulation can be written as:

$$x_i = \frac{\left(\sum_{n_j \in A_i} \left(\frac{\Delta M_{ij}^{-1} (x_j + \Delta x_{ij}) + \Delta M_{ji}^{-1} (x_j - \Delta x_{ji})}{\left(\sum_{n_j \in A_i} \Delta M_{ij}^{-1} + \Delta M_{ji}^{-1} \right)} \right) \right)}{\forall i, j | n_i \in N - L, n_j \in N}. \quad (10)$$

In this equation, x_i depends only on the measured ranges and angles from its neighbors. Solving these equations using the Gauss Seidel iterative technique corresponds to solving the equations from an initial guess successively. It can be shown that the scheme is guaranteed to converge to a unique global minima from ‘‘any’’ initial guess. The proof of convergence is provided in [24].

Thus, LMSRB would simply involve each node broadcasting its current position estimate locally to its neighbors (together with its own estimates of its distances and bearings to them). Each node would then use Equation 10 to update its current estimate and repeat the process until the estimate converges.

Results: We evaluated the $r - \theta$ scheme for different node densities, three different anchor densities (5, 10 and 20%) and three different error values (0.1%, 1% and 5%). We used the same error model for ranging errors as used in evaluating r -only schemes. We evaluated the $r - \theta$ scheme for three different kinds of bearing error: low (a 2° cone), moderate (a 4° cone), and high (an 8° cone). The low values were motivated by laser rangefinders (which are known to provide millimeter level accuracies at 10m, within a 2° cone). The

higher values were motivated by a sonar ranging device or a less expensive laser device for bearing.

We use the same simulation framework as for the r -only localization schemes (Section III-B). The only new addition is the error model for bearing which we assume to be a zero mean Gaussian. A bearing measurement having a Gaussian error with standard deviation σ_θ , results in a cone of $4\sigma_\theta$. LMSRB scheme usually stabilizes within 20-30 iterations (message exchanges), depending on anchor ratios; however, we ran the experiments to 100 iterations to ensure asymptotic behavior. All the results are averaged over 50 simulations and the 99% confidence interval for localization extent was 0.5% and 5cm for localization error.

As seen from Figure 5, even with an anchor fraction of 5% and node density of 5 the localization extent is more than 90% and reaches 98% by node density 6 for moderate error conditions. Furthermore, localization extent only improves with higher anchor fractions and node densities. Figure 6 depicts the localization error as a function of density at moderate error levels (4° cone). Even at an anchor fraction of 5% and node density (mean degree) of 5, mean localization error is less than 1% of ranging neighborhood (10cm in this case). As expected, performance only improves with increase in node densities and higher anchor fractions.

Figure 7 captures how errors in range and bearing effect the algorithm at 10% anchor fraction. An error of 2.5° translates to an average position error of $\frac{2R}{3} \sin 2.5 = 30\text{cm}(3\%)$, since the expected range of a node is $\frac{2R}{3}$ and an angular deviation of θ at this distance implies a physical deviation of $\frac{2R}{3} \sin \theta$. Moreover, a σ_θ of 2.5° implies an 8° cone. Also, considering that error might accumulate over multiple hops, it actually extremely encouraging that at high error, and a node density (degree) of 5, the mean error stays below 2.5% (25cm). The explanation for this is that the LMSRB is a global minimization and takes into account all the measurements in the sensor network. We did not, however, find any significant difference in localization extent with much higher anchor densities.

In conclusion, LMSRB provides a high localization extent and low localization errors even at low node densities and low anchor fractions.

Discussion: Clearly, then, LMSRB shows that it is possible to achieve very good localization performance even at low densities (where, as before, by density we mean the number of nodes within ranging radius). Our LMSRB scheme is a contribution to the literature; to our knowledge no other work has attempted to use range *and* bearing information to obtain location (although a recent piece of work [2] attempts to use only bearing information to estimate position, it does not use any explicit bearing measurements).

The key question is: Is $r - \theta$ localization practical? That is, can accurate bearing estimation devices be built at the form and energy factors that sensor nodes require? We do not know the answer to this question. Certainly, at larger form-factors, laser range finders can also estimate bearing, but they are also power-hungry and expensive.

One possible way to estimate bearing cheaply would be to use a ring of charge-coupled devices (CCDs) that detects light at certain frequencies. A node can emit a light pulse, and

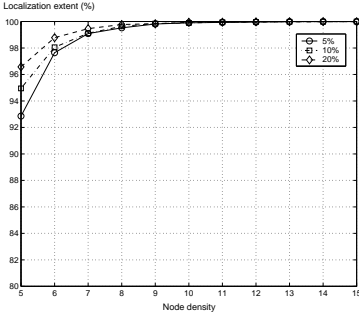


Fig. 5. Localization Extent

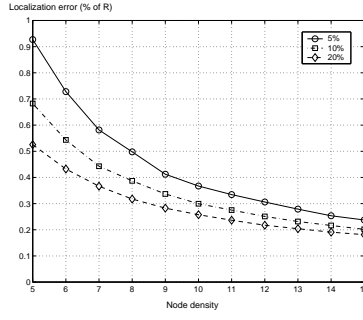


Fig. 6. Localization Error

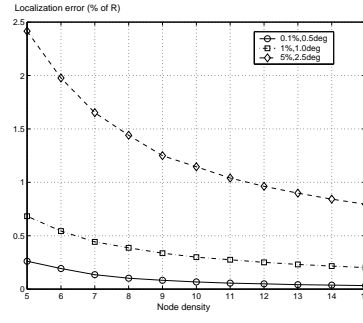


Fig. 7. Localization Error for Higher Bearing Error

the CCD array on a receiving node can estimate bearing to the transmitter using little energy. However, while this seems plausible to us, we have not investigated how such a device might actually be engineered.

V. RANGE-SECTOR BASED LOCALIZATION SYSTEMS

Given that the feasibility of building accurate bearing estimation devices is an open question, we ask: if nodes could estimate imprecise bearing to each other, how would the resulting ad-hoc localization scheme perform? More precisely, if we allowed nodes to have *sector* estimation devices (those that could place a neighbor within, say a 45° sector), would the performance of *r*-sector localization schemes be closer to *r*- θ scheme or *r*-only schemes?

We now describe an *r*-sector distributed ad-hoc localization scheme, and evaluate its performance.

Iterative Least-Mean-Square Refinement with Range and Sector: LMSRB cannot be directly extended to perform *r*-sector localization, for two reasons. First, sectoring devices are more likely to exhibit uniform, rather than Gaussian (or “close” to Gaussian) errors. Second, the first order approximation that we used to derive the covariance matrix in Equation 8 does not hold since the higher order terms will have a significant contribution beyond cones of 10° .

Instead, since the only relatively precise information available is the range information, our LMSRS (Iterative Least-Mean-Squares refinement with Range and Sector) algorithm essentially uses the LMSR algorithm described in Section III-A. LMSR requires a good initial guess for the locations of the nodes, and for this we use the sector information. A node gets an initial guess using Equation 7 and the center of the sector as an estimate for bearing to the neighbor.

It is possible, with a little communication, to obtain a better initial guess; this is not crucial for LMSRS but can speed up its convergence. Consider Figure 10 where A has an initial position estimate, and node B is trying to get a position for itself using Equation 7. A and B announce to each other the sector that each thinks the other lies in (e.g., B might say “A lies in my 45° to 90° sector”). Let CAD represents the sector which A thinks B is in. Let EBF be the sector which B thinks A is in. Suppose BF were the true bearing of A measured from B, then B would lie on the line AF’. Similarly if BE were the true bearing of A measured from B, then B would lie on somewhere on the line AE’. This means that

using the sector information of B, B must lie in the sector F’AE’. The intersection of sectors CAD and F’AE’ represents, a tighter bound on the region where B lies relative to A. Hence, now, B can initialize itself along the angle bisector of F’AD (AB’) at a range $\frac{r_{AB}\sigma_{AB}^{-1} + r_{BA}\sigma_{BA}^{-1}}{\sigma_{AB}^{-1} + \sigma_{BA}^{-1}}$, as depicted in Figure 10. Here, r_{AB} and r_{BA} (and respectively the σ s) are ranges and standard deviations measured at A to B and vice versa. This approach provides a better initial guess only if the orientations of A and B are *not* identical, which is highly likely even in loosely planned deployments.

Our simulation setup for LMSRS is identical to that for LMSRB. In general, LMSRS takes about 40-50 iterations (recall that both LMSRB and LMSRS schemes are *distributed* schemes; when we say iterations we mean the number of message exchanges required to converge). All our simulations actually run over 100 iterations and reflect asymptotic results. Our results are averaged over 50 simulations to ensure that the 99% confidence intervals for localization extent are less than 1% and for localization error are less than 5cm.

Results: In our simulations, we found that that sensitivity of LMSRS to ranging error is relatively small at all anchor fractions and node densities. Accordingly (given the limited space), we discuss the how LMSRS performs for three different sector angles (30, 45 and 60) at three different anchor ratios (5%, 10% and 20%). In Figure 8 and Figure 9 we plot the variation of localization extent and error for these combinations. The most notable results are that at even as large 30° sectors, 10% anchor ratio and a density of 5 nodes, more than 90% nodes are localized within 2m and the mean error is around 5.5% (55cm). The corresponding values at 20% anchor ratio are 95% and 3% (30cm). The performance improves at higher node densities and is expected to increase at higher anchor ratios. While these are extremely encouraging results, we examine the effect increased sector angle has on the performance. At 60° sectors and 20% anchor ratio and node density (mean degree) of 5, the localization extent is around 90% and the localization error is about 6% (60cm). This implies that even with sectors as large as 60° one could achieve “satisfactory” localization if 20% anchors were deployed.

Discussion: These results imply that an ad-hoc localization system that uses range and sector estimates can approach the performance offered by our *r*- θ scheme. More importantly, the localization extent and error are acceptable even at

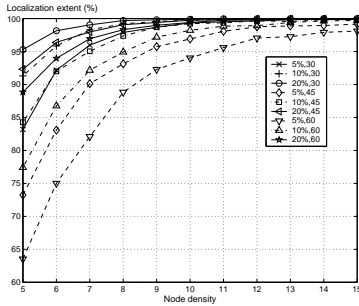


Fig. 8. Localization extent for LMSRS

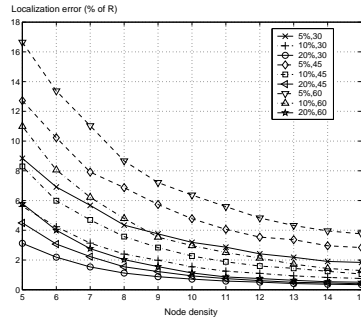


Fig. 9. Localization error for LMSRS

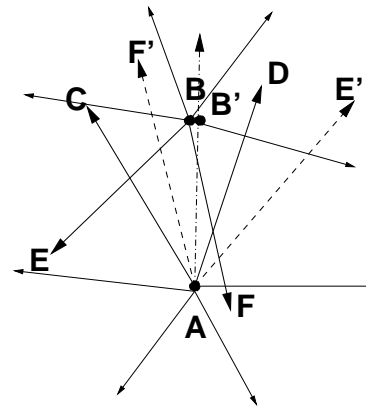


Fig. 10. Obtaining an initial guess

low densities.

As with $r - \theta$ localization, the question is: Are sector estimation devices really practical, particularly for sensor networks? We don't know, given that we know of no devices explicitly built to estimate sectors (directional antennas, sonars with a narrow spread etc. might, for example, be able to estimate sector bearing, but are not explicitly designed for that purpose). As a preliminary step we used the Pioneer P3-DX8 robots and implemented LSMRS. These robots are equipped with a belt of 16 sonars, each able to range obstacles in a roughly 22.5 degree cone. Of course, there hasn't been, till date, a pressing need to develop such a device. We believe our finding indicates that there is now a need; r -sector localization will enable ad-hoc localization systems at reasonable deployment densities, and ad-hoc localization seems to be an essential component of wireless sensor networks.

Certainly, it is more reasonable to expect that sector estimation devices will be easier to engineer, and perhaps cheaper, than bearing estimation devices. As before, a simple CCD array might actually be quite feasible, and we intend to examine the likelihood of manufacturing such a device.

VI. CONCLUSIONS AND FUTURE WORK

This paper makes the argument that the class of ad-hoc localization schemes considered in the literature so far fundamentally require high node densities in order to get acceptable performance. By contrast, new schemes that can use range and bearing estimates, or even range and sector estimates, can give good performance even at densities that ensure node connectivity.

Clearly much work is needed before ad-hoc localization systems that use ranging and sectoring devices become a reality. Sector estimation devices have to be built, our algorithms have to be validated by implementation and actual deployment in realistic environments. We are continuing to pursue these directions.

REFERENCES

- [1] P. Enge and P. Misra, "Scanning the Issue/Technology," *Proceedings of IEEE Special Issue on Global Positioning System*, vol. 87, no. 1, 1999.
- [2] D. Niculescu and B. Nath, "Ad-hoc Positioning System (APS) Using AOA," in *Proceedings of IEEE Infocom*, San Francisco, CA, 2003.
- [3] J. Hightower, C. Vakili, G. Borriello, and R. Want, "Design and Calibration of the SpotON Ad-Hoc Location Sensing System," 2001.
- [4] K. Whitehouse and D. Culler, "Calibration as a Parameter Estimation Problem in Sensor Network," in *Proc. ACM Workshop on Sensor Networks and Applications*, Atlanta, GA, 2002.
- [5] L. Girod and D. Estrin, "Robust Range Estimation Using Acoustic and Multimodal Sensing," in *Proc. IEEE International Conference on Intelligent Robots and Systems*, 2001.
- [6] A. Savvides, H. Park, and M. Srivastava, "The Bits and Flops of the N-hop Multilateration Primitive for Node Localization Problems," in *Proceedings of the First International Workshop for Wireless Sensor Networks and Applications (WSNA)*, 2002.
- [7] D. Niculescu and B. Nath, "Ad-hoc Positioning System," in *Proceedings of IEEE Globecom*, 2001.
- [8] Andreas Savvides, Chih-Chien Han, and Mani Srivastava, "Dynamic Fine-Grained Localization in Ad-Hoc Networks of Sensors," in *Proc. of Seventh ACM International Conference on Mobile Computing and Networking (MOBICOM)*, Rome, Italy, July 2001, ACM.
- [9] A. Savvides, W. Garber, S. Adlakha, R. Moses, and M. Srivastava, "On the Error Characteristics of Multihop Node Localization in Wireless Sensor Networks," in *Proceedings of First International Workshop on Information Processing in Sensor Networks*, 2003.
- [10] C. Savarese, K. Langendoen, and J. Rabaey, "Robust Positioning Algorithms for Distributed Ad-Hoc Wireless Sensor Networks," in *Proc. Usenix Annual Technical Conference*, Monterey, CA, June 2002.
- [11] K. Langendoen and N. Reijers, "Distributed Localization in Wireless Sensor Networks: A Quantitative Comparison," Tech. Rep. PDS-2002-003, Technical University, Delft, November 2002.
- [12] Nirupama Bulusu, J. Heidemann, and D. Estrin, "GPS-less Low Cost Outdoor Localization for Very Small Devices," *IEEE Personal communications magazine, Special Issue on Smart Spaces and Environments*, vol. 7, no. 5, 2000.
- [13] L. Doherty, K. Pister, and L. El-Ghaoui, "Convex Position Estimation in Wireless Sensor Networks," in *Proceedings of IEEE Infocom*, 2001.
- [14] Nissanka B Priyantha, Anit Chakraborty, and Hari Balakrishnan, "The Cricket Locations Support System," in *Proc. of Sixth ACM International Conference on Mobile Computing and Networking (MOBICOM)*, Boston, Massachusetts, USA, 2000.
- [15] A. Harter, A. Hopper, P. Steggle, A. Ward, and P. Webster, "The Anatomy of a Context Aware Application," in *Proc. of ACM International Conference on Mobile Computing and Networking (MOBICOM)*, Seattle, WA, USA, 1999.
- [16] P. Bahl and V.N. Padmanabhan, "RADAR: An In-Building RF-Based User Location and Tracking System," in *Proc. of IEEE Infocom*, Tel Aviv, Israel, 2000.
- [17] Nissanka B Priyantha, Alen K. L. Miu, Hari Balakrishnan, and Seth Teller, "The Cricket Compass for Context-Aware Mobile Applications," in *Proc. of Sixth ACM International Conference on Mobile Computing and Networking (MOBICOM)*, Rome, Italy, July 2001.
- [18] S. Capkun, Maher Hamdi, and J. P. Hubaux, "GPS-Free Positioning in Mobile Ad-Hoc Networks," *Cluster Computing*, vol. 5, no. 2, April 2002.
- [19] S. Roulmoutis and G. Bekey, "Collective Localization: a Distributed Kalman Filter Approach," in *Proc. IEEE ICRA*, San Francisco, May 1999, vol. 3, pp. 2958-2965.

- [20] D. Fox, W. Burgard, H. Kruppa, and S. Thrun, "A Probabilistic Approach to Collaborative Multi-robot Localization," *Autonomous Robots*, vol. 8, no. 3, pp. 325–344, 2000.
- [21] A. Howard, M. Mataric, and G. S. Sukhatme, "Localization for Mobile Robot Teams Using Maximum Likelihood Estimation," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Lausanne, Switzerland, October 2002, pp. 434–459.
- [22] L. Evers, W. Bach, D. Dam, M. Jonker, H. Scholten, and P. Havinga, "An Iterative Quality-based Localization Algorithm for Ad Hoc Networks," in *Proc. 1st International Conference on Pervasive Computing*, Zurich, Switzerland, August 2002, pp. 55–61.
- [23] J. Hill and D. Culler, "A wireless embedded sensor architecture for system-level optimization," Tech. Rep., Computer Science Department, University of California Berkeley, 2002.
- [24] K. Chintalapudi, A. Dhariwal, R. Govindan, and G. Sukhatme, "On the feasibility of ad-hoc localization systems," Tech. Rep., Computer Science Department, University of Southern California, Los Angeles, 2003.