Abstract

Instrumenting the physical world through large networks of wireless sensor nodes, particularly for applications like environmental monitoring of water and soil, requires that these nodes be very small, lightweight, untethered, and unobtrusive. The problem of localization, that is, determining where a given node is physically located in a network, is a challenging one, and yet extremely crucial for many of these applications. Practical considerations such as the small size, form factor, cost and power constraints of nodes preclude the reliance on GPS of all nodes in these networks. In this article we review localization techniques and evaluate the effectiveness of a very simple connectivity metric method for localization in outdoor environments that makes use of the inherent RF communications capabilities of these devices. A fixed number of reference points in the network with overlapping regions of coverage transmit periodic beacon signals. Nodes use a simple connectivity metric, which is more robust to environmental vagaries, to infer proximity to a given subset of these reference points. Nodes localize themselves to the centroid of their proximate reference points. The accuracy of localization is then dependent on the separation distance between two adjacent reference points and the transmission range of these reference points. Initial experimental results show that the accuracy for 90 percent of our data points is within one-third of the separation distance. However, future work is needed to extend the technique to more cluttered environments.

GPS-less Low-Cost Outdoor Localization for Very Small Devices

NIRUPAMA BULUSU, JOHN HEIDEMANN, AND DEBORAH ESTRIN
UNIVERSITY OF SOUTHERN CALIFORNIA/INFORMATION SCIENCES INSTITUTE

Wireless networks of sensors greatly extend our ability to monitor and control the physical world. The availability of microsensors and low-power wireless communications enables the deployment of densely distributed sensor/actuator networks for a wide range of biological and environmental monitoring applications, from marine to soil and atmospheric contexts. Networked sensors can collaborate and aggregate the huge amount of sensed data to provide continuous and spatially dense observation of biological, environmental, and artificial systems. Applications include environmental monitoring of water and soil, tagging small animals unobtrusively, and tagging small and lightweight objects in a factory or hospital setting. Instrumenting the physical world, particularly for such applications, requires that the devices we use as sensor nodes be small, lightweight, unobtrusive, and untethered. This imposes substantial restrictions on the amount of hardware that can be placed on these devices.

In these large sensor network systems, we need nodes to be able to locate themselves in various environments and on different distance scales. This problem, to which we refer as *localization*,[1] is a challenging one, yet extremely crucial for many applications of very large networks of devices. For example, localization opens up new ways of reducing power consumed in multihop wireless networks. In context-aware applications, localization enables the intelligent selection of appropriate devices, and may support useful coordination among devices. The desired granularity of localization is itself application-dependent.

The Global Positioning System (GPS) [1] solves the problem of localization in outdoor environments for PC-class nodes. However, for large networks of very small, cheap, low-power devices, practical considerations such as size, form factor, cost, and power constraints of the nodes preclude the use of GPS on all nodes. In this article we address the problem of localization for such devices, with the following design goals:

- *RF-based*: We focus on small nodes that have some kind of short-range radio frequency (RF) transceiver. Our primary goal is to leverage this radio for localization, thereby eliminating the cost, power, and size requirements of a GPS receiver.
- *Receiver-based*: In order to scale well to large distributed networks, the responsibility for localization must lie with the receiver node that needs to be localized and not with the reference points.
- *Ad hoc*: In order to ease deployment, we desire a solution that does not require preplanning or extensive infrastructure.
- *Responsiveness*: We need to be able to localize within a fairly low response time.
- *Low energy*: Small untethered nodes have modest processing capabilities and limited energy resources. If a device uses all its energy localizing itself, it will have none left to perform its task. Therefore, we desire to minimize computation and message costs to reduce power consumption.
- *Adaptive fidelity*: In addition, we want the accuracy of our localization algorithms to be adaptive to the granularity of available reference points.

This article uses an idealized radio model and proposes a simple connectivity-based localization method for such devices in unconstrained outdoor environments. It leverages the inherent RF communications capabilities of these devices. A fixed number of nodes in the network with overlapping

This research is supported by the SCOWR project through NSF grant ANI-9979457.

[1] We borrow the term localization from robotics, where it refers to the problem of determining the position of a mobile robot in some coordinate system.

1070-9916/00/$10.00 © 2000 IEEE

regions of coverage serve as *reference points* and transmit periodic beacon signals. Nodes use a simple connectivity metric to infer proximity to a given subset of these reference points and then *localize* themselves to the centroid of the selected (proximate) reference points.

The article makes the following contributions:
- It presents a detailed exploration and classification of the design space and work done in the area of localization.
- It proposes a method for coarse-grained localization based on an idealized radio model, and demonstrates its validity and applicability in outdoor unconstrained environments.
- It describes a simple implementation of the model and presents initial results.

Related Work

Localization approaches typically rely on some form of communication between reference points with known positions and the receiver node that needs to be localized. We classify the various localization approaches into two broad categories based on the granularity of information inferred during this communication. Approaches that infer fine-grained information such as the distance to a reference point based on signal strength or timing measurements fall into the category of *fine-grained localization* methods; those that infer coarse-grained information such as proximity to a given reference point are categorized as *coarse-grained localization* methods.

Fine-Grained Localization

Fine-grained localization methods can be classified further into range-finding and directionality-based methods, depending on whether ranges or angles relative to reference points are being inferred. Additionally, signal pattern matching methods are also included in fine-grained localization methods.

In range-finding methods, the ranges of the receiver node to several reference points are determined by one of several timing- or signal-strength-based techniques. The position of the node can then be computed using multilateration (e.g., see [2]). We discuss timing- and signal-strength-based range-finding methods separately.

Timing — The distance between the receiver node and a reference point can be inferred from the time of flight of the communication signal.

The time of flight may be calculated using the timing advance technique which measures the amount the timing of the measuring unit has to be advanced in order for the received signal to fit into the correct time slot. This technique is used in GPS [1] and Pinpoint's Local Positioning System (LPS) [3]. GPS measures one-way flight time, whereas LPS measures round-trip time (thereby eliminating the need for time synchronization).

GPS [1] is a wide-area radio positioning system. In GPS each satellite transmits a unique code, a copy of which is created in real time in the user-set receiver by the internal electronics. The receiver then gradually time shifts its internal clock until it corresponds to the received code, an event called *lock-on*. Once locked on to a satellite, the receiver can determine the exact timing of the received signal in reference to its own internal clock. If that clock were perfectly synchronized with the satellite's atomic clocks, the distance to each satellite could be determined by subtracting a known transmission time from the calculated receive time. In real GPS receivers, the internal clock is not quite accurate enough. An inaccuracy of a mere microsecond corresponds to a 300-m error.

Pinpoint's 3D-iD system [3] is an LPS that covers an entire three-dimensional indoor space and is capable of determining the 3-D location of items within that space. The LPS subdivides the interior of the building into cell areas that vary in size with the desired level of coverage. The cells are each handled by a cell controller which is attached by a coaxial cable to up to 16 antennas. It provides an accuracy of 10 m for most indoor applications, although some may require accuracy of 2 m. The main drawback of this system is that it is centralized, and requires significant infrastructural setup.

Alternately, the time of flight can be calculated by making explicit time-of-arrival measurements based on two distinct modalities of communication, *ultrasound* and *radio*, as in the Active Bat [2] and more recently in [4]. These two modalities travel at vastly different speeds (350 ms^{-1} and 3×10^{-8} ms^{-1}, respectively), enabling the radio signal to be used for synchronization between the transmitter and the receiver, and the ultrasound signal to be used for ranging. The Active Bat system, however, relies on significant effort for deployment indoors. Ultrasound systems may not work very well outdoors because they all use a single transmission frequency (40 kHz), and hence there is a high probability of interference from other ultrasound sources.

Signal Strength — An important characteristic of radio propagation is the increased attenuation of the radio signal as the distance between the transmitter and receiver increases. Radio propagation models [5] in various environments have been well researched and have traditionally focused on predicting the average received signal strength at a given distance from the transmitter (large-scale propagation models), as well as the variability of the signal strength in close spatial proximity to a location (small-scale or fading models). In the RADAR system [6], Bahl *et al.* suggest estimating distance based on signal strength in indoor environments. They compute distance from measured signal strength by applying a wall attenuation factor (WAF) based signal propagation model. The distance information is then used to locate a user by triangulation. This approach, however, yielded lower accuracies than RF mapping of signal strengths corresponding to various locations for their system. Their RF-mapping-based approach is quite effective indoors, unlike ours, but requires extensive infrastructural effort, making it unsuitable for rapid or ad hoc deployment.

Signal Pattern Matching — Another fine-grained localization technique is the proprietary Location Pattern Matching technology, used in U.S. Wireless Corporation's RadioCamera system [7]. Instead of exploiting signal timing or signal strength, it relies on signal structure characteristics. It turns the multipath phenomenon to surprisingly good use: by combining the multipath pattern with other signal characteristics, it creates a signature unique to a given location. The RadioCamera system includes a signal signature database for a location grid of a specific service area. To generate this database, a vehicle drives through the coverage area transmitting signals to a monitoring site. The system analyzes the incoming signals, compiles a unique signature for each square in the location grid, and stores it in the database. Neighboring grid points are spaced about 30 m apart. To determine the position of a mobile transmitter, the RadioCamera system matches the transmitter's signal signature to an entry in the database. The system can use data from only a single point to determine location. Moving traffic and changes in foliage or weather do not affect the system's capabilities. The major drawback of this technique, as with RADAR [6], is the substantial effort needed for generation of the signal signature database. Consequently, it is not suited for the ad hoc deployment scenarios in which we are interested.

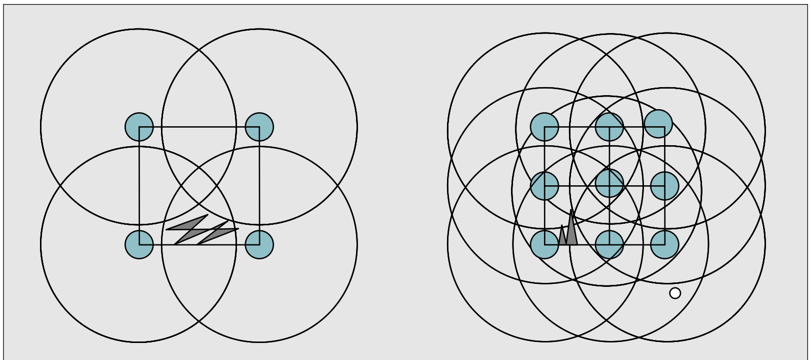

Figure 1. *Granularity of localization regions vs. range overlap.*

Directionality — Another way of estimating location is to compute the angle of each reference point with respect to the mobile node in some reference frame. The position of the mobile node can then be computed using triangulation methods.

An important example of directionality-based systems are the VOR/VORTAC stations [8], which were used for long distance aviation navigation prior to GPS. The VOR station transmits a unique omnidirectional signal that allows an aircraft aloft to determine its bearing relative to the VOR station. The VOR signal is electrically phased so that the received signal is different in various parts of the 360° circle. By determining which of the 360 different radials it is receiving, the aircraft can determine the direction of each VOR station relative to its current position.

Small aperture direction finding is yet another directionality-based technique used in cellular networks. It requires a complex antenna array at each cell site location. The antenna arrays can in principle work together to determine the angle (relative to the cell site) from which a cellular signal originated. When several cell sites can determine their respective angles of arrival, the cell phone location can be estimated by triangulation. There are two drawbacks of this approach which make it inapplicable to our application domain. The cost of the complex antenna array implies that it can be placed only at the cell sites. Second, the cell sites are responsible for determining the location of the mobile node, which will not scale well when we have a large number of such nodes and desire a receiver-based approach.

Directionality-based methods are not very effective in indoor environments because of multipath effects.

Coarse-Grained Localization

The work we describe in this article is perhaps most similar to earlier work done in coarse-grained localization for indoor environments using infrared (IR) technology.

The *Active Badge* [9] system was one of the earliest indoor localization systems. Each person or object is tagged with an Active Badge. The badge transmits a unique IR signal every 10 s, which is received by sensors placed at fixed positions within a building and relayed to the location manager software. The location manager software is able to provide information about the person's location to the requesting services and applications.

Another system based on IR technology is described in [10]. This system requires IR transmitters to be located at fixed positions inside the ceiling of the building. An optical sensor sitting on a head-mounted unit senses the IR beacons, and system software determines the position of the person.

Both these IR-based solutions perform quite well in indoor environments, because IR range is fairly small and can be limited to the logical boundaries of a region, such as a room (by walls). On the other hand, the same technique cannot be applied using RF in indoor environments, because RF propagation in indoor environments suffers from severe multipath effects that make it impossible to limit the RF range to exactly within a room. The short range of IR, which facilitates location, is also a major drawback of these systems because the building has to be wired with a significant number of sensors. In the few places where such systems have been deployed, sensors have been physically wired in every room of the building. Such a system scales poorly, and incurs significant installation, configuration, and maintenance costs. IR also tends to perform poorly in the presence of direct sunlight and hence cannot be used outdoors.

An Idealized Radio Model and Localization Algorithm

We considered two approaches to engineer an RF-based localization system, based on measurements of received signal strength and connectivity, respectively. The signal-strength-based approach did not work very well, while the connectivity-based approach proved quite effective outdoors.

Signal Strength Approaches

One approach to RF-based localization is to use measured signal strength of received beacon signals to estimate distance, as in the RADAR system [6], with an outdoor radio signal propagation model. We discarded this approach for several reasons relating to our short-range (10 m) radios. First, signal strength at short ranges is subject to unpredictable variation due to fading, multipath, and interferences; therefore, it does not correlate directly with distance. Moreover, short range does not allow much gain in density of reference points when considering signal strength. Finally, our commercial off-the-shelf (COTS) radios did not provide software-accessible signal strength readings. These reasons caused us to focus on connectivity-based localization, described next.

An Idealized Radio Model

We have found an idealized radio model useful for predicting bounds on the quality of connectivity-based localization. We chose this model because it was simple and easy to reason about mathematically. This section presents this idealized model. To our surprise, this model compares quite well to outdoor radio propagation in uncluttered environments, explored in the next section.

We make two assumptions in our idealized model:

- Perfect spherical radio propagation
- Identical transmission range (power) for all radios

A Localization Algorithm

Multiple nodes in the network with overlapping regions of coverage serve as reference points (labeled R_1 to R_n). They are situated at known positions, (X_1, Y_1)–(X_n, Y_n), that form a regular mesh and transmit periodic beacon signals (period = T) containing their respective positions. We assume that neighboring reference points can be synchronized so that their beacon signal transmissions do not overlap in time. Furthermore, in any time interval T, each reference point would have transmitted exactly one beacon signal.

First, we define a few terms:

d: Separation distance between adjacent reference points

R: Transmission range of the reference point
T: Time interval between two successive beacon signals transmitted by a reference point
t: Receiver sampling or data collection time
$N_{sent}(i, t)$: Number of beacons sent by R_i in time t
$N_{recv}(i, t)$: Number of beacons sent by R_i received in time t
CM_i: Connectivity metric for R_i
S: Sample size for connectivity metric for reference point R_i
$CMthresh$; Threshold for CM
(X_{est}, Y_{est}): Estimated location of the receiver
(X_a, Y_a): Actual location of the receiver

Each mobile node listens for a fixed time period t and collects all the beacon signals it receives from various reference points. We characterize the information per reference point R_i by a *connectivity* metric (CM_i), defined as

$$CM_i = \frac{N_{recv}(i,t)}{N_{sent}(i,t)} \times 100.$$

In order to improve the reliability of our connectivity metric in the presence of various radio propagation vagaries, we would like to base our metric on a sample of at least S packets, where S is the sample size, a tunable parameter of our method (i.e., $N_{sent}(i, t) = S$). Since we know T to be the time period between two successive beacon signal transmissions, we can set t, the receiver's sampling time, as

$$t = (S + 1 - \varepsilon)T \quad (0 < \varepsilon \ll 1).$$

From the beacon signals it receives, the receiver node infers proximity to a collection of reference points for which the respective connectivity metrics exceed a certain threshold, CM_{thresh} (say 90 percent). We denote the collection of reference points $R_{i1}, R_{i2}, \ldots, R_{ik}$. The receiver localizes itself to the region which coincides with the intersection of the connectivity regions of this set of reference points, which is defined by the *centroid* of these reference points:

$$(X_{est}, Y_{est}) = \left(\frac{X_{i1} + \ldots + X_{ik}}{k}, \frac{Y_{i1} + \ldots + Y_{ik}}{k} \right).$$

We characterize the accuracy of the estimate by the localization error LE, defined as

$$LE = \sqrt{(X_{est} - X_a)^2 + (Y_{est} - Y_a)^2}.$$

By increasing the *range overlap* of the reference points that populate the grid (i.e., increasing the ratio R/d), the granularity of the localization regions becomes finer, and hence the accuracy of the location estimate improves. This is illustrated in Fig. 1.

Validation

Since our localization model depends on the spherical radio propagation assumption described in the previous section, we checked the validity of our assumption in both outdoor and indoor environments.

In outdoor environments, we evaluated the effectiveness of our idealized radio model by comparing its accuracy to experimental measurements. We evaluated propagation between two Radiometrix radio packet controllers (model RPC-418) operating at 418 MHz. A node periodically sent 27-byte beacon signals; we define a 90 percent packet reception rate as *connected* and empirically measured an 8.94 m spherical range for our simple model.

To evaluate how well our simple model compares to a real-world scenario, we placed a radio in the corner of an empty parking lot,, at the origin (0, 0), and then measured connectivity at 1 m intervals over a 100 m^2 quadrant.

Figure 2 compares these measurements with connectivity as predicted by the model. Among the 78 points measured, the simple spherical model matches correctly at 68 points (87 percent correlation) and mismatches at 10, all at the edge of the range. Error was never more than 2 m. No dead spots were observed.

As expected, our simple idealized radio model approximation is not appropriate for indoor environments where reflection and occlusion are common. Our indoor measurements of propagation range varied widely from 4.6 to 22.3 m, depending on walls and exact node locations and orientations. Furthermore, these measurements were not time-invariant. We found that connectivity could vary from 0 to even 100 percent for the same transmitter receiver positions at different times of the day.

Hence, the idealized radio model may be considered valid for outdoor unconstrained environments only.

Experimental Results

The Experimental Testbed

Our experimental testbed [11] consisted of five Radiometrix RPC 418 (radio packet controller) modules connected to a Toshiba Libretto running RedHat Linux 6.0. One of these modules is used as a receiver, and the rest are used as reference points. A 3 in antenna is used for experimental purposes.

The software for the Radiometrix RPC-418 modules consists of two components:

- *Beacon*: The reference point periodically transmits a packet (every 2 s in our experiment) containing its unique ID and position.
- *Receiver*: The receiver obtains its current measured position based on an input from the user. For each measured position, it samples for a time period t determined by sample size S, and logs the set of reference points from which it hears and its current localization estimate.

For our experiment, we placed the four reference points at the four corners of a 10 m x 10 m square in an outdoor parking lot. This square was further subdivided into 100 smaller 1

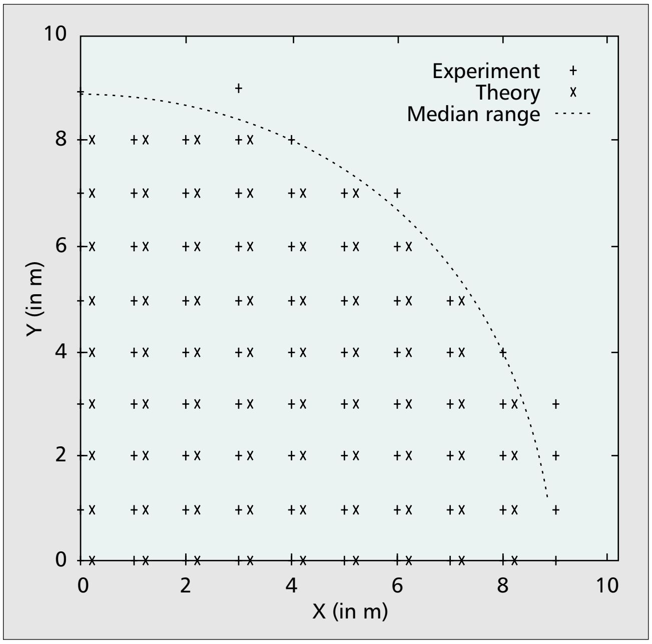

■ **Figure 2.** *90 percent connectivity ranges for reference point (0,0).*

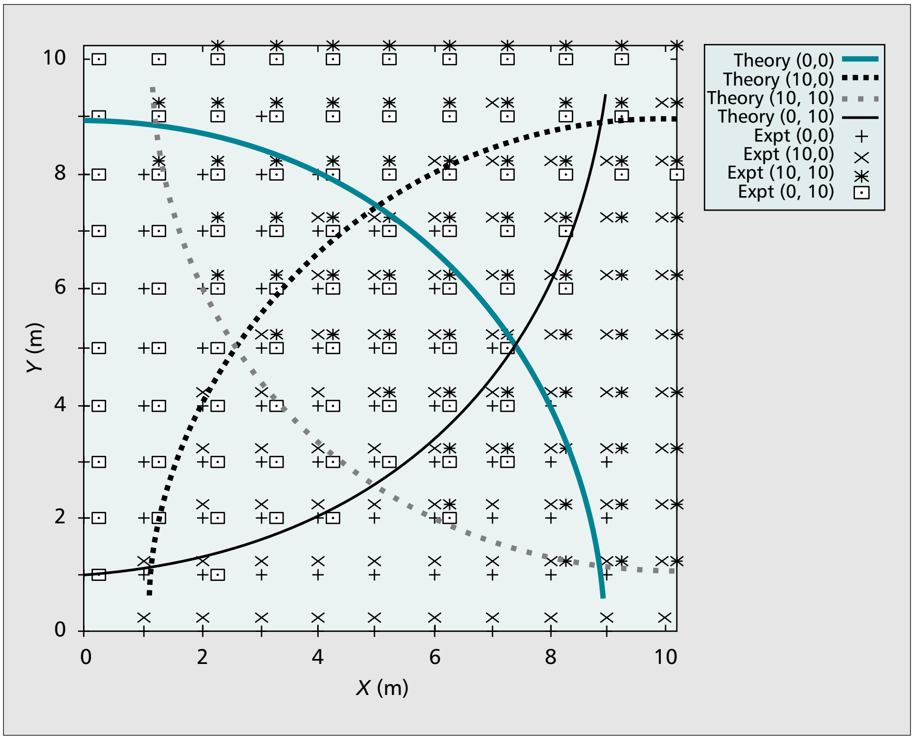

■ **Figure 3.** *Experimental vs. theoretical 90 percent connectivity ranges for the four reference points.*

m x 1 m grids, and we collected data at each of the 121 small grid intersection points.

Outdoor Results

In this section we discuss the results obtained from our outdoor experiments. Our experimental parameters were $T = 2$ s, $S = 20$, $t = 41.9$ s.

Figure 3 shows the areas of connectivity of the four reference points in the grid. We see several distinct regions in the grid, based on the areas of overlap. Each distinct region constitutes an equivalence class, defined by the centroid of the reference points in the region. These can be contrasted with the theoretically predicted overlap regions, also seen in Fig. 3.

The location estimate at each grid point is the centroid. We use the *localization error* metric defined previously to characterize the performance.

In Fig. 4, the *localization error* obtained from the experiment is plotted as a function of the position. The localization error is lowest at the position corresponding to the centroid of the region and increases toward the edges of the region. The average localization error was 1.83 m and the standard deviation 1.07 m. The minimum error was 0 m and the maximum error 4.12 m across 121 grid points.

Figure 5 shows the cumulative localization error distribution across all the grid points, from both the theoretical model and the experiment. They track each other closely, including plateaus in the error levels, although the spherical model is consistently more optimistic. In our experimental results, for over 90 percent of the data points the localization error falls within 3.0 m (i.e., within 30 percent of the separation-distance between two adjacent reference points). This result is based on four reference points only. Since we observed a high correlation between our model and experiment, improved granularity can be expected with a higher overlap of reference points.

Based on our validated outdoor model, we performed numerical simulations to predict how the granularity of localization could be expected to vary using our scheme when the overlap of reference points is increased.

In our simulation, we assume an infinite two-dimensional mesh of reference points, with any two adjacent reference points spaced a distance d apart and transmission range R. Our coordinate system is centered at one such reference point, which is assumed to be at (0, 0).

The localization estimate of any point (X, Y) in the mesh can be obtained in two steps:

- *Step 1*: Determine all the reference points that are within range R of (X, Y), by considering the reference points between $(X - R, Y - R)$ and $(X + R, Y + R)$.
- *Step 2*: Localize (X, Y) to the centroid of the selected reference points and compute the corresponding localization error.

For a given d, we increase the overlap R/d from 1 to 4. We consider the average and maximum localization errors of the localization estimates for 10,201 uniformly spaced points within one grid in the mesh for each R/d value. Figure 6 presents the *simulation-based* scaling result of the localization error behavior. Although the maximum and average error do not decrease monotonically, nontrivial increments to R/d (for instance, an increment of 1) lead to lower maximum and average localization errors on the whole. In particular, the maximum localization error experiences a substantial drop (from $0.5d$ to $0.25d$) when the overlap R/d is increased from 1 to 4.

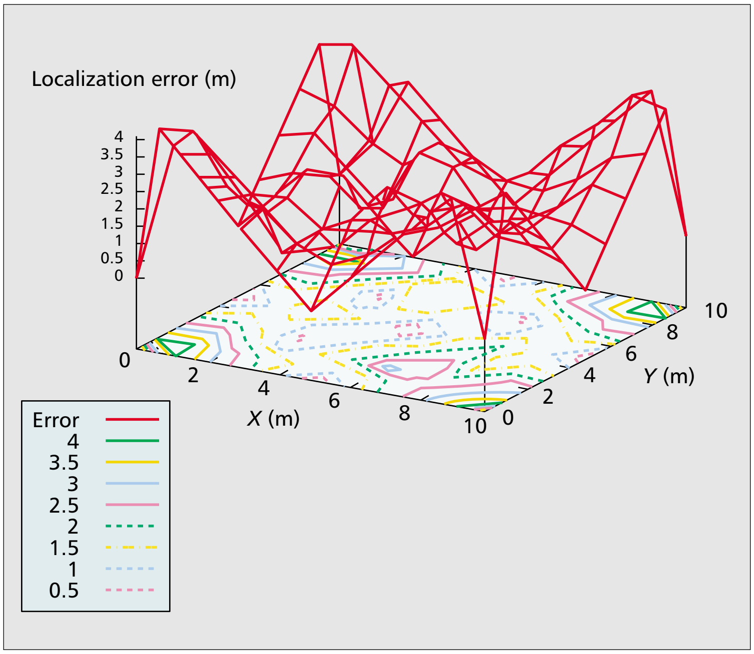

■ **Figure 4.** *Localization error vs. position.*

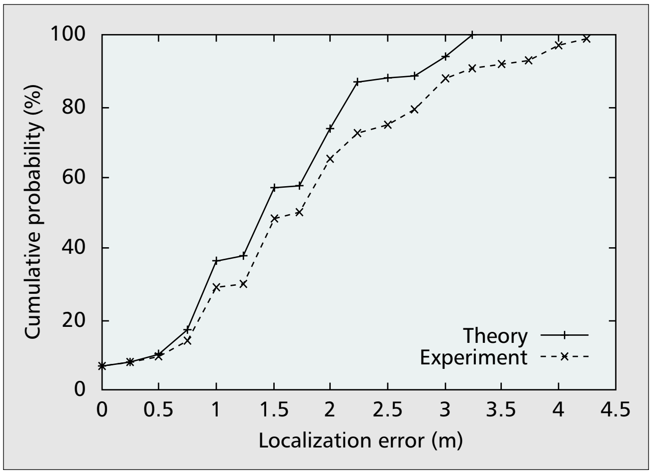

Figure 5. *Cumulative localization error distribution.*

Discussion and Future Work

In this section we discuss some general problems that arise in deploying our localization method and present some of our ideas on solving them.

Collision Avoidance — For our method to work well, neighboring reference points need to synchronize their beacon signal transmissions so as to avoid collisions. To achieve this, we propose the following randomized scheme. Each time interval T is subdivided into several smaller slots. Each reference point then chooses a slot randomly with a uniform distribution to transmit its beacon signal. We need to study this further, although randomized schemes have proven extremely effective in designing various network protocols to avoid contentions.

Tuning for Energy Conservation — The parameters T, the time period for beacon transmissions, and S, the sample size, must be tuned to avoid collisions and ensure the consistency of the connectivity metric while reducing power consumption. Since we use the connectivity metric as a coarse-grained measure, our experiences with our experimental testbed proved that a small value of S (e.g., 10) would suffice to establish connectivity. The value of T can be determined based on the reference point overlap R/d and the efficacy of the collision avoidance scheme.

Nonuniform Reference Point Placement — Our localization method assumes that the reference points are placed in a regular mesh structure. We controlled the placement to bound the quality of localization. In practice, it may not always be feasible to achieve a strictly uniform placement of reference points. To understand the effect of nonuniform placement, we simulated several scenarios with reference points placed randomly in a uniform distribution in a square grid. Uniform placement consistently yields superior quality of localization across the grid compared to random placement of an equal number of reference points. However, only a small fraction of grid points (less than 15 percent) experience significantly worse localization due to nonuniform placement. The trade-off here is to use a random, but slightly more dense, distribution of reference points to achieve the same quality of localization as uniform placement.

Reference Point Configuration — We have left open the issue of how the reference point coordinates are configured and deployed. This could be achieved through limited human intervention. The reference points themselves can determine their positions through the use of GPS or other fine-grained localization methods, since they do not have similar constraints as other nodes. Initially, the reference points may be deployed manually or scattered randomly across the terrain. We are working on automated algorithms to select good places to deploy additional nodes as reference points.

Robustness — Since the success of our localization method depends on the node reliably inferring connectivity, and hence proximity to its neighboring reference points, it must be tolerant to reference point failures (and also to nonuniform reference point placement). Reference points should monitor themselves and failstop when their battery power drops. Some amount of redundancy (additional nodes that can serve as reference points if needed) should be incorporated into the system to tolerate reference point failures.

Adaptation to Noisy Environments — Our simple localization method is very effective in restricted domains with idealized radio conditions. Idealized radio conditions do not hold in noisy environments characterized by severe multipath phenomenon, fading, obstructions, dead spots, and so on. In order to generalize our scheme to noisy environments, we are currently investigating techniques for empirical adaptation of reference point placement.

Conclusion

This article addressed localization in unconstrained outdoor environments for very small low-cost devices that do not have GPS. We characterized existing localization techniques and explored an RF-based localization method in which the receiv-

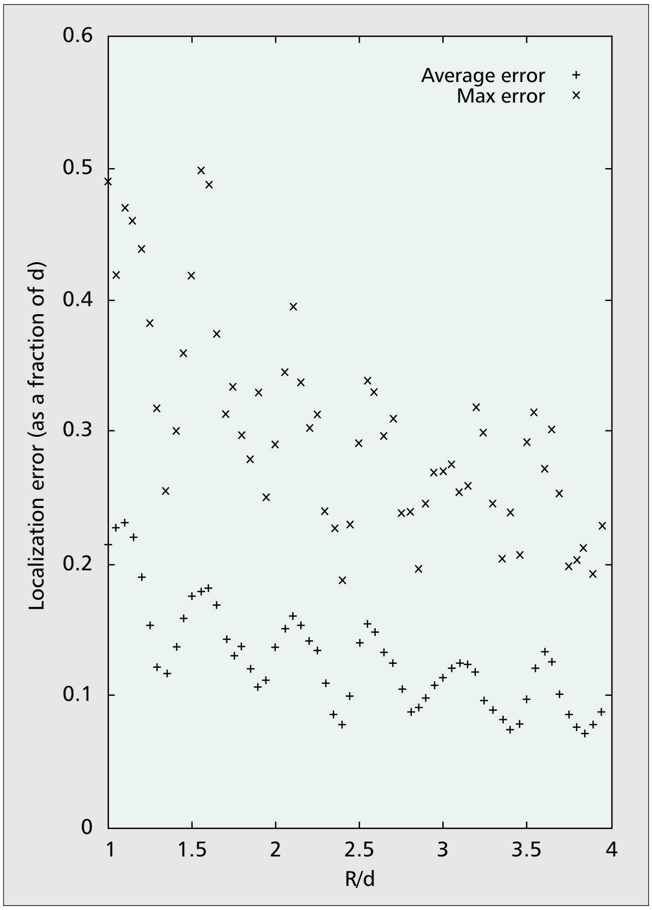

Figure 6. *Localization error vs. overlap, R/d. (simulations).*

er localizes itself with high confidence (under an idealized radio model) to the centroid of a set of proximate reference points using a connectivity metric. Although our approach uses a very simple radio model, in outdoor environments our model correlated very well with reality (87 percent).

Our approach is simple, entirely RF-based, receiver-based, and adaptive to the granularity of reference points available. Additionally, it requires no coordination among reference points or sensor nodes. It is therefore potentially scalable to very large distributed networks of devices.

Initial experiments have shown promising results, with our simple scheme, for a small number of reference points. Our simulation results suggest that the granularity of localization can be further improved by increasing the overlap of reference points. While our approach is essentially coarse-grained, it is nevertheless useful for several applications with less stringent accuracy requirements.

We also outlined some general problems which need to be tackled for large-scale deployment. In particular, our future work includes adapting our localization method to noisy environments.

Acknowledgments

We wish to thank Jeremy Elson, Lewis Girod, and Gaurav Sukhatme for their suggestions and feedback, and also for their contributions to our experimental testbed. This research is supported by the SCOWR project through NSF grant ANI-9979457.

References

[1] B. Hofmann-Wellenhof, H. Lichtenegger, and J. Collins, *Global Positioning System: Theory and Practice*, 4th ed., Springer Verlag, 1997.

[2] A. Ward, A. Jones, and A. Hopper, "A New Location Technique for the Active Office," *IEEE Pers. Commun.*, vol. 4, no. 5, Oct. 1997, pp. 42–47.

[3] J. Werb and C. Lanzl, "Designing a Positioning System for Finding Things and People Indoors," *IEEE Spectrum*, vol. 35, no. 9, Sept. 1998, pp. 71–78.

[4] L. Girod, "Development and Characterization of an Acoustic Rangefinder," Tech. rep. 00-728, Computer Science Department, USC, Apr. 2000.

[5] T. S. Rappaport, *Wireless Communications — Principles and Practice*, Prentice Hall, 1996.

[6] P. Bahl and V. N. Padmanabhan, "Radar: An In-Building RF-Based User-Location and Tracking System," *Proc. IEEE INFOCOM 2000*, Tel Aviv, Israel, vol. 2, Mar. 2000, pp. 775–84.

[7] http://www.uswcorp.com/USWCmainpages/our.htm

[8] http://www.arundel.net/xplane/html/vor.html

[9] R. Want *et al.*, "The Active Badge Location System," *ACM Trans. Info. Sys.*, vol. 10, no. 1, Jan. 1992, pp. 91–102.

[10] R. Azuma, "Tracking Requirements for Augmented Reality," *Commun. ACM*, vol. 36, no. 7, July 1993, pp. 50–55.

[11] http://netweb.usc.edu/scadds/testbeds.html

Biographies

NIRUPAMA BULUSU (bulusu@isi.edu) is currently a graduate student in the UCLA Computer Science Department and a graduate research assistant at the USC Information Sciences Institute (ISI). She received her M.S. (2000) from the University of Southern California and her B.Tech. (1997) from the Indian Institute of Technology, Madras, both in computer science.

JOHN HEIDEMANN [M] is a computer scientist at USC/ISI and a research assistant professor at USC. At ISI he investigates networking protocols and simulation as part of the VINT project, and embedded networking and data diffusion in sensor networks as part of SCADDS. He received his B.S. from the University of Nebraska-Lincoln, and his M.S. and Ph.D. from the University of California at Los Angeles (UCLA). He is a member of ACM and Usenix.

DEBORAH ESTRIN is currently a professor in the UCLA Computer Science Department and a project leader at USC/ISI. From 1986 to 2000 she was on the faculty of computer science at USC. She received her Ph.D.(1985) and M.S.(1982) from the Massachusetts Institute of Technology and her B.S. (1980) from the University of California at Berkeley.