

Localization of Wireless Sensor Networks with a Mobile Beacon

Mihail L. Sichitiu and Vaidyanathan Ramadurai
Dept. of Electrical and Computer Engineering,
North Carolina State University,
Raleigh, NC 27695

Abstract

Wireless sensor networks have the potential to become the pervasive sensing (and actuating) technology of the future. For many applications, a large number of inexpensive sensors is preferable to a few expensive ones. The large number of sensors in a sensor network and most application scenarios preclude hand placement of the sensors. Determining the physical location of the sensors after they have been deployed is known as the problem of localization. In this paper, we present a localization technique based on a single mobile beacon aware of its position (e.g. by being equipped with a GPS receiver). Sensor nodes receiving beacon packets infer proximity constraints to the mobile beacon and use them to construct and maintain position estimates. The proposed scheme is radio-frequency based, and thus no extra hardware is necessary. The accuracy (on the order of a few meters in most cases) is sufficient for most applications. An implementation is used to evaluate the performance of the proposed approach.

I. Introduction

Large numbers of untethered sensing devices are bound to revolutionize the way we interact with the physical world [1]. Recent advances in sensing, processing and communication made possible tight integration of a complete sensor node on a single chip [2]. On-chip integration enables inexpensive production of large numbers of such sensors. Being deployed in large numbers results

in better coverage of a geographical area, but it also poses numerous challenges to the communication protocols.

From tactical surveillance and target tracking to environmental monitoring and space exploration, the applications of sensor networks are limited only by our imagination [1]. For most applications, sensed data without spatial and temporal coordinates is of very limited use. Sensor nodes have to be aware of their location to be able to specify “where” a certain event takes place. Therefore, the problem of localizing the sensors is of paramount importance for many classes of sensor network applications.

Sensors aware of their position can also improve routing efficiency [3]–[6] by selective flooding or selective forwarding data only in the direction of the destination. Sensor nodes may not have an individual identifier (i.e. address); the location of the sensor may be (part of) the address of the sensors. Various algorithms that use the location as part of the address have been proposed [7]–[10].

The position of each sensor can be manually introduced if the sensors are hand-placed; however, when the number of sensors is large, this becomes a tedious and error-prone method of localization. In many applications, hand-placing the sensors is not an option. If the sensors are scattered from a plane or from a mortar shell, a different localization method has to be employed. If each sensor node has a global positioning system (GPS) [11] receiver, the problem becomes trivial. However, having a GPS receiver on every node is currently a costly proposition in terms of power, volume and money.

We propose a localization method using

Bayesian inference for processing information from one mobile *beacon*. A beacon is a node aware of its location (e.g. equipped with GPS). The nodes of initially unknown positions will be called *unknown* nodes. After the sensor node has been deployed, the mobile beacon assists the unknown nodes in localizing themselves. The mobile beacon can be a human operator, an unmanned vehicle deployed with the sensor network, or in the case of a deployment from a plane, the plane itself. Our approach is described in detail in Section III.

The method presented in this paper is radio-frequency (RF) based: the received signal strength indicator (RSSI) is used for ranging, although other methods of ranging can also be used [12], [13]. The advantage of the RSSI ranging is its ubiquitous availability in practically all available receivers on the market; the disadvantage is its inaccuracy. If the sensor network is deployed indoors, walls would severely reduce the precision of the method due to nonlinearities, noise, interference and absorption [14]–[18]. However, most sensor networks will likely be deployed outdoors and will be able to take advantage of the proposed approach.

II. Related Work

The problem of localization is very important for many engineering fields and has been researched for many years. In robotics, the exact location, as well as the orientation, of a robot was extensively considered [19]–[23].

Many of the outdoor localization systems rely heavily on infrastructure. In 1996, the US Federal Communications Commission (FCC) required that all wireless service providers be able to provide location information to Emergency 911 services. Cellular base stations are now used to determine the position of the users. In 1993, the global positioning system, based on 24 NAVSTAR satellites, was deployed. Since that time, it has become the standard way to do localization whenever a GPS receiver can be used. LORAN operates in a similar way to GPS but uses ground base stations instead of satellites.

Recently, interest has been increasing for indoor localization systems. The RADAR system [15] can

track the location of users within a building and is based on RF signal strength measurements. The Cricket [24] location support system is also designed for indoor localization but uses ultrasound instead. Indoors, the walls severely reduce the precision of RSSI measurements, and practically accurate results can only be obtained through the use of some form of calibration [14]–[18].

In the ad hoc domain, fewer localization systems have been proposed and implemented. In [25] a rectangular grid of beacons and RF connectivity constraints are considered. The performance of the system is carefully analyzed in [26]. An iterative multilateration is investigated in [27]. The presented algorithm performs well when a large percentage of beacons are available, the graph connectivity is high and precise range measurements can be determined.

A nice paper [13] tackles the thorny problem of cooperative multilateration. The presented approach relies on precise range determination (using an ultrasonic ranging technique [12]) and then on solving a least square problem on a large order system. The information necessary for the complete formulation of the least square problem needs to be transmitted from a potentially distant location through multiple hops. The follow-up [28] improves the scalability of the approach.

An interesting idea is explored in [29]. The authors consider the problem of localization in the absence of any beacons. The nodes build local coordinate systems and further aggregate them into a unique network coordinate system. Directional approaches considering steerable directional antennas have been explored [30].

A method for estimating unknown node positions in a sensor network based exclusively on connectivity-induced constraints is described in [31]. Known peer-to-peer communication in the network is modeled as a set of geometric constraints on the node positions. The global solution of a feasibility problem for these constraints yields estimates for the unknown positions of the nodes in the network. One drawback of the method in [31] includes a central point of computation with the associated traffic overhead, scalability and reliability issues.

Another original idea is presented in [32] where

auto-calibration is used to improve the accuracy of a localization algorithm. The authors impose common sense constraints (e.g. the distance from A to B equals the distance from B to A, as well as the triangle inequality) on the position of the nodes, and thus “auto-correct” the range measurements. We believe that similar ideas can be used to improve the accuracy of the proposed method, but we will not pursue them in this paper.

As it is correctly pointed in [33], all localization problems can be seen as part of a larger problem that attempts to find the position of unknown nodes using as much of the available information as possible. An interesting approach (mesh relaxation) is used in [33] to solve several of those sub-problems. The approach in [33] is best suited for centralized post-processing of the gathered data.

While the proposed solution shares some of the ideas presented in related work (several Bayesian approaches and localization/mapping algorithms using mobile robots are published), to the best of our knowledge, no other localization system based on a single mobile beacon employing a Bayesian approach has been previously investigated. The performance evaluation based on an experimental testbed shows that the presented approach outperforms existing localization approaches.

III. Proposed Approach

The idea for this paper came about in a deliberate attempt to eliminate some of the drawbacks of existing localization systems. With one exception, all localization systems for sensor networks rely on several beacons scattered throughout the sensor network. The system in [29] does not use any beacons, and it is able to localize itself with respect to an arbitrary local coordinate system. While this relative localization is useful for some applications (e.g. location aware routing), most systems require localization with respect to a fixed coordinate system (e.g. latitude and longitude); therefore, at least one beacon node is necessary. It is shown [13], [25]–[27] that the precision of the localization increases with the number of beacons. The main problem with an increased number of beacons is that they are more expensive than the rest of the sensor nodes. Indeed, if a GPS receiver

is available for each beacon, a beacon node can be two orders of magnitude more expensive than an unknown node [34]. This means that, even if only 10% of the nodes are beacons, the price of the network will increase tenfold. Another observation is that after the (stationary) unknown nodes have been localized, the beacons become useless; they no longer use their (expensive) GPS receivers. The reasoning mentioned above leads us to believe that a single mobile beacon can be used to localize the entire network.

The only way to ensure scalability is to make the necessary computations at the unknown nodes using only local information. If the computation has to be centralized, and data gathered from the entire network, the approach is likely to consume significant resources in very large networks. Addressing this concern in the proposed system, each unknown node locally computes its position estimate without any information from other unknown nodes, and without a single transmission from any of the unknown nodes. Thus, the positioning system can scale to *any* number of unknown nodes.

We could have used the acoustic method of ranging [12], [13] to obtain precise range estimation; however, we believe that, for most applications, an accuracy of a few meters is sufficient. Therefore, we opted for an RSSI ranging method, which is readily available on most transceivers. Since sensor nodes were not available for experimentation, we used HP iPAQ Pocket PCs equipped with Lucent Orinoco Gold 802.11b cards. We will show later that, despite the significant bandwidth, storage and computing power difference between a Berkeley mote and a Compaq iPAQ, the algorithm that we proposed can be easily implemented on a system with capabilities similar to the Berkeley motes.

A. System Calibration

To calibrate the system, we took a series of measurements between a pair of iPAQs at different distances. The results of those measurements are presented in Fig. 1. We measured the RSSI values every 2.5m. The mean values and three times the standard deviation are depicted for every distance.

The signal strength measurement units are arbitrary. We found that the antennas of the Lucent cards mounted on the iPAQs have fairly directional radiation patterns, so we had to take measurements for different orientations of the sender and the receiver.




Fig. 1. Signal strength measurements as a function of the distance. The mean and plus/minus three times the standard deviation of the signal strengths for each distance.




Fig. 2. Probability distribution function that a node receiving a packet with a signal strength of 77 will be at a certain distance from the sender of the packet.

From the signal strength vs distance measurements, we derived the inverse function: distance vs signal strength. In other words, for a given signal strength, we wanted to determine the probability distribution function (pdf) as a function of the distance. Figure 2 depicts the pdf of the ranges for an RSSI of 77.

Trading accuracy for simplicity, we fitted Gaussian curves to the ranging results (see Fig. 2). To our surprise, the Gaussians were a good fit (they passed the Kolmogorov-Smirnov normality test). The standard deviation of the RSSI measurements does not vary significantly with the range; however, the standard deviation of the inverse function does. Indeed, an RSSI of 90 indicates the range much more precisely than an RSSI of 70. The standard deviation of the range vs the RSSI is depicted in Fig. 3.




Fig. 3. Standard deviation of the ranges as a function of the signal strength.

The measurement results depicted in Fig. 1 were obtained from an outdoor, open environment. To assess the reliability of the measurements in different environments, we repeated the experiments in a heavily wooded area. The results were surprisingly similar to the ones in the open environment (except for a slightly larger standard deviation). This increases our confidence that, while our proposed approach will likely perform poorly indoors, it will work well in various outdoor environments.

B. Localization Algorithm

Figure 4 depicts a sensor network deployed over a geographical area. After deployment, a mobile beacon traverses the sensor network while broadcasting beacon packets. A beacon packet contains the coordinates of the beacon. Any node receiving the beacon packet will be able to infer that it must be somewhere around the mobile beacon with a certain probability. This information




Fig. 4. One mobile beacon assisting in the localization of a sensor field.

constrains the possible locations of a node. The RSSI is measured for each beacon that is received. Corresponding to the RSSI measurement and the position of the beacon (x_B, y_B) (included in the beacon packet), each node receiving the beacon constructs a constraint on its position estimate:

$$C(x, y) = PDF_{RSSI}(d((x, y), (x_B, y_B))) \quad (1)$$

$$\forall (x, y) \in [(x_{\min}, x_{\max}) \times (y_{\min}, y_{\max})]$$

where PDF_{RSSI} is the probability distribution function of the distance corresponding to the RSSI of the beacon packet, $d(A, B)$ is the Euclidean distance between points A and B , and $x_{\min}, x_{\max}, y_{\min}$ and y_{\max} are the bounding coordinates of the deployment area.

Figure 5 depicts the constraint imposed on the position of the unknown node 2 when a beacon packet of coordinates (30,35) and RSSI 75 is received from mobile beacon 1.

Once the constraint is computed, each node applies Bayesian inference to compute its new position estimate NPE from its old position estimate OPE and the new constraint C :

$$NPE(x, y) = \frac{OPE(x, y) \times C(x, y)}{\int_{x_{\min}}^{x_{\max}} \int_{y_{\min}}^{y_{\max}} OPE(x, y) \times C(x, y)} \quad (2)$$

$$\forall (x, y) \in [(x_{\min}, x_{\max}) \times (y_{\min}, y_{\max})]$$

The initial position estimate is initialized to a constant value, as in the beginning, all positions in the deployment area are equally likely.

Figure 6 depicts the evolution of the position estimate of an unknown node as the mobile beacon broadcasts beacon packets. In Fig. 6(a) the position of the unknown node and the mobile beacon

trajectory are shown. The beacon sends beacon packets at each of the positions marked on the trajectory. Only nine beacon packets are received by the unknown node (in the order shown in Fig. 6(a)). The evolution of the position estimate of the unknown node as beacon packets 1,2,4,5 and 9 are received is shown in Figs. 6(b-f) respectively.

Once the final position estimate $PosEst$ pdf is computed, (i.e. after all beacon packets have been received), the best position coordinates (\hat{x}, \hat{y}) can be determined as a weighted average:

$$(\hat{x}, \hat{y}) = \left(\int_{x_{\min}}^{x_{\max}} \int_{y_{\min}}^{y_{\max}} x \times PosEst(x, y) dx dy, \int_{x_{\min}}^{x_{\max}} \int_{y_{\min}}^{y_{\max}} y \times PosEst(x, y) dx dy \right) \quad (3)$$

C. Beacon Trajectory

An interesting question is “What is the optimum beacon trajectory and when should the beacon packets be sent ?” Notice that the problem is quite difficult since the position of the unknown nodes is not known *a priori*. Once the nodes are (at least partially) localized, the beacon can be steered to assist nodes with large uncertainties. We will not try to answer the optimality question in this paper. Instead, we will make some remarks regarding some properties that the trajectory should have.

First, a node is best localized if the beacon trajectory is close to that node. This observation stems from the calibration data; as shown in Fig. 3, the standard deviation for close range measurements is significantly lower than for higher ranges. Therefore, the trajectory of the beacon node should pass closely to as many potential node positions as possible.

Second, even if the beacon node passes close to a node in a straight line, the proposed localization algorithm will not be able to determine on which *side* of the line the node lies. Indeed, positions symmetric to the line will be equally probable. Figure 6(d) depicts the position estimate after four approximately collinear beacon packets have been received by the unknown node. To eliminate one of the candidates, at least one non-collinear beacon packet must be received. Figure 6(e) shows the position estimate after just one more beacon has been received.

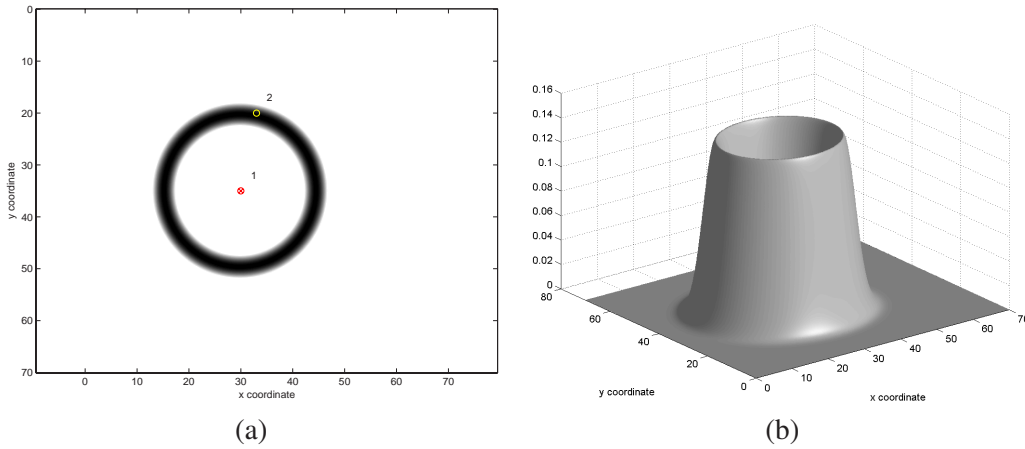


Fig. 5. (a)The constraint imposed on the position of the unknown node 2 when a beacon packet of coordinates (30,35) and RSSI 75 is received from the mobile beacon 1; (b) A 3D view of the constraint.

Therefore, the beacon trajectory should be designed in such a way that all possible positions are fully covered by at least three non-collinear beacons, and the “grid” formed by the beacons should be as tight as possible (to increase precision).


D. Implementation Issues

The computational complexity and storage requirements of the proposed approach depend on how the position estimates are represented. There are several methods to represent position estimates and constraints.


Perhaps the simplest method is to simply record the beacon coordinates and their associated RSSI, and send this information as the position of the node. For example, if seven beacons are received by an unknown node and both the coordinates and the RSSI are represented using a byte, then 21 bytes will be required to represent the position estimate of that node. In this case, the unknown nodes will do absolutely *no computations*. Upon receipt of the 21 bytes representing the node’s position estimate, a base station (or the monitoring station) will perform the required computations to determine the best estimate of the coordinates (\hat{x}, \hat{y}) . In this approach, the computational burden is transferred to the base station, which usually has more computational and storage resources than the sensor nodes.

However, computing (\hat{x}, \hat{y}) on the sensor node


itself is entirely possible even on sensor nodes similar to the Berkeley motes [35]. Assume that the sample space is represented by a grid of $n \times n$ squares. Each of the operations in (1-3) has $O(n^2)$ complexity. Processing associated with each received beacon using an 100×100 grid of bytes (type-casted to floats to avoid quantization effects and overflows) and an Atmel Atmega103 microcontroller (the same microcontroller used in the first generation of MICA Berkeley motes) at 16MHz takes approximately 15s. A microprocessor with multiplication support (e.g. Atmega128) will perform significantly faster. The main time consuming procedure is the square root used to compute the Euclidean distance in (1). Additional software optimizations can further reduce the computation time. For example, computing a constraint for a coordinate (x, y) where the current estimate is equal to zero is useless, as the product in (2) will be zero regardless of the value of the constraint. In short, considering that for static sensor networks, the localization is only done once (immediately after deployment), spending a couple of minutes on computing the position estimates is perhaps reasonable. The node storage requirements are also $O(n^2)$. For example, if a 100×100 grid is used, almost 10KB of RAM will be used to store a node’s position estimate. Once the best estimate coordinates (2 bytes - \hat{x}, \hat{y}) are computed, they can be used in the application




(a)




(b)




(c)



(d)



(e)



(f)

Fig. 6. (a) The trajectory of the mobile beacon and the position of the nine beacon transmissions received by the unknown node. (b-f) The evolution of the position estimates of the unknown node after 1,2,4,5 and respectively 9 beacon packets have been received.

running on the sensor network (e.g. tracking), and the 10KB of memory allocated for the position estimate can be freed for other purposes (e.g. buffering for the forwarding layer).

IV. Experimental Results

To evaluate the performance of the proposed approach, we implemented the system on HP iPAQ PocketPCs equipped with Lucent Orinoco Gold (IEEE 802.11b compatible) cards. As we mentioned in Section III-D, the implementation can be easily ported on resource constrained sensor nodes. To avoid human interference, we used a radio-controlled truck to carry the mobile beacon (also an HP iPAQ) and GPS receiver used to determine the coordinates of the beacon.

The beacon periodically reads its current coordinates from the GPS receiver (through a serial interface) and broadcasts beacon packets. Figure 7 depicts the the real as well as the estimated position of the unknown nodes and the trajectory of the mobile beacon. Notice that for clarity the two axes of Fig. 7 have different scales.




Fig. 7. The real (denoted by an 'o') and the estimated (denoted by an 'x') positions of the unknown nodes and the trajectory of the mobile beacon used in the experimental setup.

Figure 8 depicts the final position estimates of the 12 unknown nodes. Some of the estimates, especially for the nodes far from the beacon trajectory, have a larger standard deviation, which is expected considering the results in Fig. 3.

The localization error, (i.e. the distance between the coordinate of the position estimate and the




Fig. 8. The final position estimates of the unknown nodes.




Fig. 9. The localization error for each of the unknown nodes when using the proposed probabilistic approach.

actual position of each node) is shown in Fig. 9. All 12 nodes have a localization error less than 3 m. Considering that we used GPS to measure both the position of the beacon and the positions of the nodes, and that the GPS accuracy was around 3-4m (as reported by the GPS receiver), the accuracy of the algorithm is exceedingly good. Given the GPS accuracy we expected errors on the order of 5-10m. The unexpected accuracy may be due to the differential precision of the GPS receiver, which, in certain situations, can be significantly better than the absolute precision.

Figure 10 depicts the variation in the localization error for unknown node with the increase in the number of beacon messages received. There are several interesting observations that result from Fig. 10. First, some of the nodes received more beacons than others. Secondly, while the tendency of the error is to decrease with an increase in the number of beacons, the error graphs are not mono-




Fig. 10. Evolution of the localization error with the number of beacons received.

tonic. The explanation stems from the statistical nature of our approach: a “bad” beacon message with an eccentric signal strength (i.e. far from the mean) will actually hurt the position estimate. In a similar way, a “good” beacon message, can dramatically reduce the localization error. Finally, there is more than an order of magnitude difference between the localization precision after the first beacon has been received (also due to the “quality” of the beacon).




Fig. 11. The localization error for each of the unknown nodes when using multilateration.

Figure 11 depicts the results that result from a multilateration approach [13] using the same data as used for the proposed probabilistic approach (results depicted in Fig. 9). As it can be seen, the errors are significantly higher: the average error for the probabilistic approach is 1.4 m, the average error for the multilateration approach is 10.67 m, almost an order of magnitude higher. This relatively high error rates resulting from a multi-

lateration approach results from the inaccuracy of the RSS measurements: a single RSS measurement can result in a large multilateration error. It is thus clear, that for environments with large ranging errors (typical case for RSS ranging), the proposed probabilistic approach performs significantly better.

V. Conclusion

An outdoor, RF-based localization algorithm based on a mobile beacon is presented. The algorithm scales well to any number and density of unknown nodes and uses a single mobile beacon. The mobile beacon can be controlled by a human operator, or an automatic unmanned aerial or ground vehicle. The system requires an initial calibration phase before deployment. Calibration data in various environments was very consistent, increasing the confidence for a wide applicability of this approach. The precision of the localization is good and uniform as long as the trajectory of the beacon covers the entire deployment area in such a way that each point receives at least three non-collinear beacon messages. Data from the mobile beacon is aggregated using a Bayesian approach. The performance of the approach has been evaluated using a real implementation. The experimental results reveal an unexpectedly good accuracy, almost an order of magnitude better than existing approaches.

References

- [1] I. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci, “A survey on sensor networks,” *IEEE Communication Magazine*, vol. 40, pp. 102–116, Aug. 2002.
- [2] J. M. Kahn, R. H. Katz, and K. S. J. Pister, “Next century challenges: Mobile networking for Smart Dust,” in *Proc. of ACM/IEEE Mobicom*, (Seattle, WA), pp. 271–278, Aug. 1999.
- [3] Y. B. Ko and N. H. Vaidya, “Location aided routing (LAR) in mobile ad hoc networks,” *Wireless Networks*, vol. 6, pp. 307–321, Sept. 2000.
- [4] S. Basagni, I. Chlamtac, V. Syrotiuk, and B. Woodward, “A distance routing effect algorithm for mobility (DREAM),” in *Proc. of ACM Mobicom’98*, (Dallas, TX), pp. 76–84, Oct. 1998.
- [5] J. Li, J. Jannotti, D. DeCouto, D.R.Karger, and R.Morris, “A scalable location service for geographic ad-hoc routing,” in *Proc. of ACM Mobile Communications Conference*, (Boston, MA), Aug. 2000.

- [6] K. Amouris, S. Papavassiliou, and M. Li, "A position-based multi-zone routing protocol for wide area mobile ad-hoc networks," in *Proc. of IEEE Vehicular Technology Conference*, (Houston, TX), pp. 1365–1369, 1999.
- [7] D. Estrin, R. Govindan, J. S. Heidemann, and S. Kumar, "Next century challenges: Scalable coordination in sensor networks," in *Mobile Computing and Networking*, pp. 263–270, 1999.
- [8] C. Intanagonwiwat, R. Govindan, and D. Estrin, "Directed diffusion: a scalable and robust communication paradigm for sensor networks," in *Mobile Computing and Networking*, pp. 56–67, 2000.
- [9] J. S. Heidemann, F. Silva, C. Intanagonwiwat, R. Govindan, D. Estrin, and D. Ganesan, "Building efficient wireless sensor networks with low-level naming," in *Symposium on Operating Systems Principles*, pp. 146–159, 2001.
- [10] Y. Ko and N. Vaidya, "Geocasting in mobile ad hoc networks: Location-based multicast algorithms," in *Proc. of WMCSA*, (New Orleans), 1999.
- [11] B. Hofmann-Wellenhopf, H. Lichtenegger, and J. Collins, *Global Positioning System: Theory and Practice*. Springer-Verlag, 4th ed., 1997.
- [12] L. Girod and D. Estrin, "Robust range estimation using acoustic and multimodal sensing," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2001)*, (Maui, Hawaii), Oct. 2001.
- [13] A. Savvides, C. C. Han, and M. B. Srivastava, "Dynamic fine-grained localization in ad-hoc networks of sensors," in *Proc. of Mobicom'2001*, (Rome, Italy), pp. 166–179, July 2001.
- [14] P. Bahl and V. N. Padmanabhan, "Enhancements to the radar user location and tracking system," in *Microsoft Research Technical Report MSR-TR-2000-12*, 2000.
- [15] P. Bahl and V. Padmanabhan, "RADAR: An in-building RF-based user location and tracking system," in *Proc. of Infocom'2000*, vol. 2, (Tel Aviv, Israel), pp. 775–584, Mar. 2000.
- [16] A. M. Ladd, K. E. Bekris, G. Marceau, A. Rudys, L. E. Kavraki, and D. Wallach, "Robotics-based location sensing using wireless ethernet," in *Proc. of Eighth ACM International Conference on Mobile Computing and Networking (MOBICOM 2002)*, (Atlanta, Georgia), Sept. 2002.
- [17] M. Helén, J. Latvala, H. Ikonen, and J. Niittylahti, "Using calibration in RSSI-based location tracking system," in *Proc. of the 5th World Multiconference on Circuits, Systems, Communications & Computers (CSCC20001)*, 2001.
- [18] J. Hightower, R. Want, and G. Borriello, "SpotON: An indoor 3D location sensing technology based on RF signal strength," Tech. Rep. CSE 2000-02-02, University of Washington, Seattle, WA, Feb. 2000.
- [19] D. Fox, W. Burgard, F. Dellaert, and S. Thrun, "Monte Carlo localization: Efficient position estimation for mobile robots," in *In Proc. of the Sixteenth National Conference on Artificial Intelligence (AAAI-99)*, (Orlando, Florida), pp. 343–349, 1999.
- [20] D. Fox, W. Burgard, H. Kruppa, and S. Thrun, "A probabilistic approach to collaborative multi-robot localization," *Autonomous Robots*, vol. 8, pp. 325–344, June 2000.
- [21] D. Fox, W. Burgard, and S. Thrun, "Markov localization for mobile robots in dynamic environments," *Journal of Artificial Intelligence Research (JAIR)*, vol. 11, pp. 391–427, Nov. 1999.
- [22] S. Thrun, "Probabilistic algorithms in robotics," *AI Magazine*, vol. 21, no. 4, pp. 93–109, 2000.
- [23] S. Thrun, W. Burgard, and D. Fox, "A probabilistic approach to concurrent mapping and localization for mobile robots," *Machine Learning*, vol. 31, no. (1-3), pp. 29–53, 1998.
- [24] N. Priyantha, A. Chakraborty, and H. Balakrishnan, "The cricket location-support system," in *Proc. of International Conference on Mobile Computing and Networking*, (Boston, MA), pp. 32–43, Aug. 2000.
- [25] N. Bulusu, J. Heidemann, and D. Estrin, "GPS-less low cost outdoor localization for very small devices," *IEEE Personal Communications Magazine*, vol. 7, pp. 28–34, Oct. 2000.
- [26] N. Bulusu, J. Heidemann, D. Estrin, and T. Tran, "Self-configuring localization systems: Design and experimental evaluation," *ACM Transactions on Embedded Computing Systems (ACM TECS)*, *Special issue on networked embedded systems*, 2003.
- [27] C. Savarese, J. M. Rabaey, and J. Beutel, "Locating in distributed ad-hoc wireless sensor networks," in *Proc. of ICASSP'01*, vol. 4, pp. 2037–2040, 2001.
- [28] A. Savvides, H. Park, and M. Srivastava, "The bits and flops of the n-hop multilateration primitive for node localization problems," in *First ACM International Workshop on Wireless Sensor Networks and Applications*, (Atlanta, GA), Sept. 2002.
- [29] S. Capkun, M. Hamdi, and J. P. Hubaux, "GPS-free positioning in mobile ad-hoc networks," *Cluster Computing*, vol. 5, April 2002.
- [30] A. Nasipuri and K. Li, "A directionality based location discovery scheme for wireless sensor networks," in *First ACM International Workshop on Wireless Sensor Networks and Applications*, (Atlanta, GA), Sept. 2002.
- [31] L. Doherty, K. S. J. Pister, and L. E. Ghaoui, "Convex position estimation in wireless sensor networks," in *Proc. IEEE Infocom 2001*, vol. 3, (Anchorage AK), pp. 1655–1663, Apr. 2001.
- [32] K. Whitehouse and D. Culler, "Calibration as parameter estimation in sensor networks," in *First ACM International Workshop on Wireless Sensor Networks and Applications*, (Atlanta, GA), Sept. 2002.
- [33] A. Howard, M. Mataric, and G. Sukhatme, "Relaxation on a mesh: A formalism for generalized localization," in *Proc. IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, (Wailea, Hawaii), Oct. 2001.
- [34] "Spec: Smartdust chip with integrated RF communications." <http://www.cs.berkeley.edu/~jhll/spec/index.htm>.
- [35] J. Hill, R. Szwedczyk, A. Woo, S. Hollar, D. E. Culler, and K. S. J. Pister, "System architecture directions for networked sensors," in *Architectural Support for Programming Languages and Operating Systems*, pp. 93–104, 2000.