Scalable, Ad Hoc Deployable RF-based Localization

Nirupama Bulusu, Vladimir Bychkovskiy, Deborah Estrin John Heidemann
bulusu,vladimir,destrin@cs.ucla.edu johnh@isi.edu
Computer Science Department, UCLA USC/ISI

Abstract

Spatial localization or the ability to locate nodes is an important building block for next generation pervasive computing systems, but a formidable challenge, particularly, for very small hardware and energy constrained devices, for noisy, unpredictable environments and for very large ad hoc deployed and networked systems. In this paper, we describe, validate and evaluate in real environments a very simple *self* localization methodology for RF-based devices based only on RF-connectivity constraints to a set of beacons (known nodes), applicable outdoors. Beacon placement has a significant impact on the localization quality in these systems. To self-configure and adapt the localization in noisy environments with unpredictable radio propagation vagaries, we introduce the novel concept of *adaptive beacon placement*. We propose several novel and density adaptive algorithms for beacon placement and demonstrate their effectiveness through evaluations. We also outline an approach in which beacons leverage a software controllable variable transmit power capability to further improve localization granularity. These combined features allow a localization system that is scalable and ad hoc deployable, long-lived and robust to noisy environments. The unique aspect of our localization approach is our emphasis on *adaptive self-configuration*.

1 Introduction

Pervasive computing and sensor networks are emerging as key application drivers for wireless networks. *Pervasive computing* promises to bring an abundance of computation and communication in our lives, by simplifying collaboration, knowledge access and by automating repetitive tasks. Recent advances in miniaturization and low-cost, low-power design have led to active research in large-scale, highly distributed systems of small, wireless, low-power unattended sensors and actuators [9, 12] (referred to as *wireless sensor networks*). By allowing observation and control at unprecedented levels of detail, wireless sensor networks offer to fundamentally revolutionize the ways in which we understand and construct complex physical systems [9], from airplane wings to complex ecosystems. The most challenging of these applications require ad hoc deployable wireless networks that are scalable, long-lived and robust systems, despite being largely unattended, overcoming energy limitations and a lack of pre-installed infrastructure. Some applications with stringent device constraints in form factor and size are: environmental monitoring in the water and soil, tagging small animals unobtrusively, or tagging small and light objects in a factory or hospital setting.

The problem of estimating spatial coordinates is known as *localization*, and is of fundamental importance to pervasive computing and sensor networks. Many of these envisioned systems are embedded to monitor or control the behavior of physical systems (as compared with strictly virtual information systems), and therefore nodes often need to determine their action based on their physical location (am I the right sensor to monitor a particular object?). Localization is indispensable for context-aware applications that select services based on location [11], and for sensor networks that achieve power conservation by combining data from multiple sensors. Moreover, location information on a scale with transmission range can enable geographic routing algorithms that can propagate information efficiently through a multi-hop network [14].

Traditional information systems have not had to have such a location focus and as such our support for localization systems is relatively weak. Existing geolocation systems such as GPS do not always meet the operational (for example, low power), environmental (for example, indoors) or cost constraints. Although a number of localization systems have been proposed in the past few years [8, 11, 18, 20, 7, 17], none currently satisfies the requirements for ad hoc deployment of large scale sensor networks because no single existing localization system simultaneously provides (1) a self localization methodology for devices that is scalable to very large sensor networks (2) a low cost, hardware-independent localization approach for very small devices and (3) a self-configuring mechanism for the localization system to adapt to noisy environments. Our challenge then is to develop a localization system that meets the above requirements.

We assume all nodes we consider are equipped with a short range radio (RF) transceiver for wireless communications. Our work makes the following contributions toward *scalable*, *ad hoc deployable*, and *entirely RF-based* localization:

Localization Methodology: We have developed a simple RF-connectivity based self-localization methodology for devices in outdoor environments. In this paper, we present our localization approach, prototype implementation for two different devices, and experimental results that validate it in outdoor environments (Section 3).

Adaptive Beacon Placement: Ensuring a uniform threshold granularity across the terrain in a localization system requires not merely uniformly dense placement of known beacon nodes, but mechanisms to detect and adapt to radio propagation and terrain unpredictabilities[3]. To address this problem, we have formalized the problem of *adaptive beacon placement.* We discuss the design space and present our algorithms, simulation results and future work (Section 4).

2 Related Work

Research efforts related to our work fall into the following categories, (i) localization systems and (ii) techniques to improve node placement.

Localization Systems

Research on localization systems has garnered a lot of attention in recent years. We briefly cover a few important systems here, a more detailed discussion is given in [2].

Fine-grained localization systems that provide high precision location information, typically estimate ranges or angles relative to beacons (known nodes) and compute the location of the unknown node using trilateration (position estimation from distance to three points) or triangulation (position estimation from angles to three points).

GPS and Pinpoint[21] estimate distance from the RF signal time of flight using time difference of arrival (TDOA) techniques (or the amount of timing that the measured signal has to be advanced in order for the received signal to fit into the correct time slot to be in phase with an internally generated signal). Active Bat[11] and Cricket [18] make explicit time-of-arrival measurements based on two distinct modalities of communication, ultrasound and radio, which travel at vastly different speeds ($350m/s$ and $3 \times 10^8 m/s$ 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. In the RADAR indoor location system [1], distance is estimated from received signal strength by applying a Wall Attenuation Factor (WAF) based signal propagation model. The distance information is then used to locate a user by trilateration. The U.S Wireless Corporation's RadioCamera system[7] uses signal pattern matching techniques to compute location. Directionality based systems include VOR stations[17] and small aperture direction finding, used in cellular networks. As direction estimation is quite expensive, it can only be placed at the beacons. Responsibility for localization now lies with beacons. This approach does not scale well for larger numbers of such nodes.

Coarse-grained localization systems estimate unknown node location from proximity to beacons or landmarks. One of the earliest such systems was Active Badge [20], where, each person or object is tagged with an Active Badge. The badge transmits a unique IR signal every 10 seconds, which is received by sensors placed at fixed positions within the building. IR tends to perform poorly in the presence of direct sunlight and hence cannot be used outdoors. Another drawback is that it is a tracking rather than a self-localization system. Very recently, Doherty *et al*[8] also propose techniques for localization from RF-connectivity. However, their approach is centralized, unlike ours which is decentralized, low-cost and scalable.

Beacon placement

A significant component of localization error arises due to beacon placement[1]. Addressing beacon placement is orthogonal and complementary to many of the other proposed techniques to reduce localization error. Researchers have proposed guidelines based on knowledge of environment conditions and application requirements or proposed optimal approaches to placement.

[1] Measurement Errors also occur due to poor calibration of radios and sensors used.

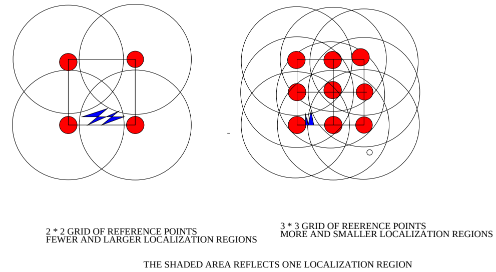

Figure 1: Localization Region Granularity vs. Range Overlap

Active Bat [11] proposes using ceiling mounted beacons to maximize likelihood of line of sight to beacons. Cricket [18], which is also proximity based, proposes deployment guidelines for beacons in indoor environments based on similarly practical considerations. [8] proposes placing beacons on the corners of the network, for best results with their centralized convex position estimation based. While guidelines work well for a specific localization method [8] or environmental setting [18, 11], they are not generalizable to a variety of environments and systems with unpredictable conditions, or to large scale ad hoc sensor network deployments.

Optimal placement problems have been studied in various contexts by researchers including facility location problems in theory [5] and pursuit evasion problems in robotics [10]. [16] proposes solutions to coverage problems in wireless ad hoc sensor networks given global knowledge of node positions using Voronoi diagrams to compute maximal breach paths and find gaps.

The fundamental limitation of these fixed approaches is that they do not take into account the unpredictable environmental conditions. The premise of our work is that an adaptive beacon placement is required to cope with noisy and unpredictable environmental conditions.

3 Localization Methodology

In this section, we describe our localization methodology using RF connectivity, its prototype implementation and evaluation.

Approach

Although approaches based on received signal strength (RSSI) of radios seem more attractive, we discarded this approach for several reasons relating to our short-range (10m) radios, detailed in [2]. Mainly, we found that received signal strength did not correlate well with distance for our radios due to multi-path, fading and other RF vagaries. These reasons caused us to focus on localization using RF connectivity. We use a mathematically simple, idealized radio model for predicting bounds on the quality of localization based on RF-connectivity. It makes two rather unrealistic assumptions. (i) Perfect spherical radio propagation and (ii) Identical transmission range (power) for all radios. To our surprise, this model compared quite well to outdoor radio propagation in uncluttered environments [2].

Beacons situated at known positions, (X_i, Y_i), transmit periodically with a time period T. Clients listen for a period $t >> T$ to evaluate connectivity. If the percentage of messages received from a beacon in a time interval t exceeds a threshold $CMthresh$, that beacon is considered connected. When the beacon placement is uni-

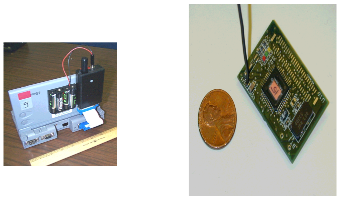

(1) Libretto with Radiometrix RPC radios (2) UCB Rene Mote

Figure 2: Experimental Testbeds

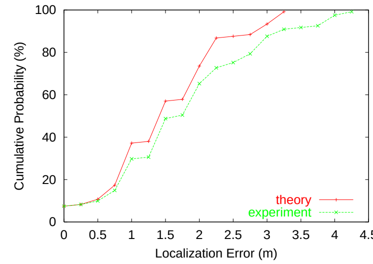

Figure 3: Cumulative Localization Error Distribution

form, the centroid of the positions of all connected beacons is a feasible solution in the region of connectivity overlap.For non-uniform placement, a feasible solution can be found using more general convex optimization techniques [8].
A client estimates its position (X_{est}, Y_{est}) to be the centroid of the positions of all connected beacons.

$$(X_{est}, Y_{est}) = \left(\frac{\Sigma_{i=1}^{k} X_i}{k}, \frac{\Sigma_{i=1}^{k} Y_i}{k}\right) \quad (1)$$

Given the actual position of the client (X_a, Y_a), we can compute the accuracy of the localization estimate or the *localization error* $LE_{\mathcal{B}}(X_a, Y_a)$, which is the distance between the client's estimated and actual positions.

$$LE_{\mathcal{B}}(X_a, Y_a) = [(X_{est} - X_a)^2 + (Y_{est} - Y_a)^2]^{\frac{1}{2}} \quad (2)$$

By increasing the density of the beacons that populate the grid, the overlap regions become smaller, and hence localization granularity improves (see Figure 1).

Prototype Implementation

We have implemented a prototype of our localization methodology on two experimental testbeds (1) Radiometrix RPC radios connected to laptops via. a serial interface (2) UC Berkeley Rene motes [12], completely integrated with RFM [19] radios completely, shown in Figure 2. Due to space restrictions, we only describe some experimental results on the RPC testbed here. Details on the second set of experiments with motes can be found in [4].

Our first experimental testbed consists of Radiometrix RPC 418 (radio packet controller) modules connected to a Toshiba Libretto running RedHat Linux 6.0. In our experiments, one of these modules is used as a receiver and four are used as beacons. A 3 inch antenna is used for the experimental purposes. The software for the Radiometrix RPC-418 modules consists of two components. (i) *Beacon*: The beacon periodically transmits a packet containing its unique ID and position. (ii) *Receiver*: The receiver obtains its actual position from user input and estimated position by listening to beacon packets.

Results

We evaluated our system indoors and outdoors, and found it to be feasible outdoors. In this section, we discuss the results of an outdoor experiment. We placed 4 beacons at the four corners of a $10m \times 10m$ square in an outdoor parking lot. This square was further subdivided into 100 smaller $1m \times 1m$ squares and we collected data at each of the 121 small square corners. We use the *localization error* metric defined previously to characterize the performance. Figure 3 shows the cumulative localization error distribution across all the square corners, from both a theoretical model(assuming Range = 8.9m (the measured median range)) 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% of the data points, the localization error falls within 3.0 meters. This result is based on 4 beacons. Since we observed a high correlation between our model and experiment, improved granularity can be expected with a higher overlap of beacons as Figure 4 shows. More details of our experimental results may be found in [2].

Our simple localization methodology is very effective in restricted domains (with ideal radio propagation). To generalize our scheme to noisy environments, we have developed techniques for measurement based adaptation of beacon placement, described next.

4 Adaptive Beacon Placement

Our approach to improving localization through beacon placement is based on *measurement based adaptation*. By *adaptation*, we mean we are improving the quality of localization by adjusting beacon placement or adding a few beacons rather than by completely redeploying all beacons. By *measurement based*, we mean the deployment of additional beacons is influenced by measurements of the operating localization system rather than by careful off-line analysis of a complete system model. We have explored three complementary approaches to adaptive beacon placement, For low beacon densities, we investigated algorithms to augment the existing beacon infrastructure by adding new beacons at *empirically determined points*, based on, (i) terrain exploration and measurements made by a mobile robot, which is centralized, and described in [3], (ii) simulated local exploration by beacons (HEAP), which is distributed, and described in [4]. For high beacon densities, the suitable approach is adaptive operational density (STROBE), which we highlight here.

4.1 STROBE Adaptive Density

Assume a localization system with static beacons, that will need to advertise their positions periodically during the system lifetime to support other mobile nodes. In proximity based localization systems using only local information, regardless of actual beacon placement, the localization granularity saturates at a certain threshold beacon density ν. This is verifiable through simulations (shown in Figure 4) and can also be proven mathematically, based on results of [15]. For approaches based on ranging, three line of sight (LOS) and

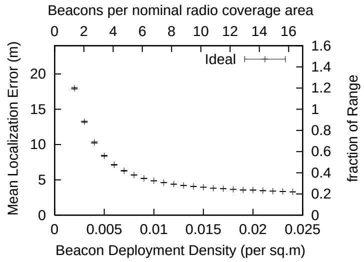

Figure 4: Mean localization error vs. Beacons per neighborhood. Localization granularity saturates at a certain number of beacons per neighborhood, around 6 in our case.

non-collinear beacons provide the critical threshold of beacons in 2D space.

In unattended sensor networks, where new beacons cannot be physically deployed as needed, we would begin with a very dense initial beacon deployment for redundancy. Tuning the operational beacon density by rotating functionality amongst beacons (by turning their radios on and off) (i) increases system lifetime without diminishing the localization granularity (ii) reduces the probability of self-interference amongst beacons by reducing overall transmissions ([4],[18]) (iii) allows adaptation to noisy environments when required (higher percentage of beacons could remain active in noisier obstructed parts of the terrain, whereas a smaller percentage of beacons need to be active in more benign parts of the terrain, achieving similar localization granularity overall). STROBE (Selectively TuRning Off BEacons) is our algorithm is to achieve such an adaptive operational beacon density to realize these benefits. Our design goal is to build STROBE using *localized algorithms, i.e.,* each beacon determines its role during a given time interval based on coordination with its neighbors rather than from an assignment by a central server.

4.1.1 STROBE duty cycle

Typically, each beacon transmits one position advertisement in a beaconing interval T_B and sleeps for the remaining part of the interval. Each position advertisement has four fields: beacon identifier, beacon position, sequence number, beacon status. Beacon status is usually set to be UP.

In STROBE, a beacon can be in one of three states: Voting (V), Designated (D) and Sleep (SL). The state transition diagram is depicted in Figure 5. All beacons start out in the Voting state, wherein, a beacon turns on its radio and broadcasts position advertisements every T_B seconds and also listens for advertisements from its neighboring beacons. When a beacon node enters Voting state, it sets a timer for T_V seconds. When the timer fires, it evaluates where it should go to sleep based on a decision making process explained in 4.1.2. If so, it broadcasts an advertisement with State set to be DOWN and transitions to the Sleep (SL) state. Otherwise, it transitions to the Designated (D) state. A beacon node in sleep state wakes up after a sleep time T_{SL} and transitions back to Voting (V) state. A beacon node in Designated state periodically advertises at intervals T_B for a time T_D and then transitions back to Voting (V) state. A beacon node in sleep state wakes up after a sleep time T_{SL} and transitions back to Voting (V) state.

Distinct Voting and Designated states are necessary in order to avoid the overhead incurred due to receiving advertisement messages from other neighbor beacons when in the Voting state. Three important parameters of STROBE that influence its energy usage and system lifetime are T_V, T_D, and T_{SL}.

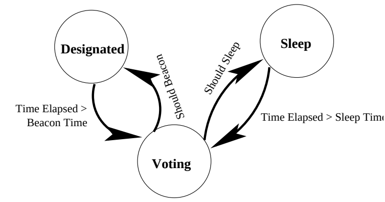

Figure 5: State Transition Diagram for STROBE

4.1.2 Beacon Decision Making

During the Voting (V) state, a beacon evaluates ζ, the number of currently active beacons that are its neighbors.

$$\zeta = |B_1 - B_2| \tag{3}$$

where B_1 is the set of all beacons it heard from whose most recent advertised state is UP and B_2 those whose most recent advertised state is DOWN. The number of active beacons in its neighborhood, including itself is $\zeta + 1$. Let ν be the threshold number of beacons in any given neighborhood at which the localization granularity saturates. If $(\zeta + 1) \leq \nu$, then it has to remain active. If $(\zeta + 1) > \nu$, then its transition probability p to the sleep state is given by:

$$p = \frac{\zeta - (\nu - 1)}{\zeta} \tag{4}$$

With probability $(1 - p)$ it transitions to the Designated (D) state.

More sophisticated decision making approaches would incorporate information such as energy reserve of a beacon and its neighbors, while needing to maintain additional state.

4.1.3 STROBE Evaluation

We conducted extensive performance evaluations of STROBE using simulations [4]. We report here on a representative experiment that demonstrates its benefits in improving system lifetime without degrading localization granularity.

The following metrics are used for evaluation. *Percentage of active beacons*, the percentage of total beacons that are in either *Voting* or *Designated* states at any given instant of time.
Percentage of alive beacons, the percentage of total beacons that possess energy reserves greater than zero at any given instant of time.
Median localization error in the terrain, as a function of time. *First node death*, is the time elapsed since the start before any single node in the terrain runs out of energy (dies).
System lifetime, is the time elapsed since the start before median localization error exceeds an operational threshold.

For our simulations, we choose an energy consumption model to mimic realistic sensor radios [13], summarized in Table 1.

The following simulation experiment demonstrates STROBE benefits for an applicable context (small beaconing interval, high beacon density). We simulate a terrain with 100 beacons distributed

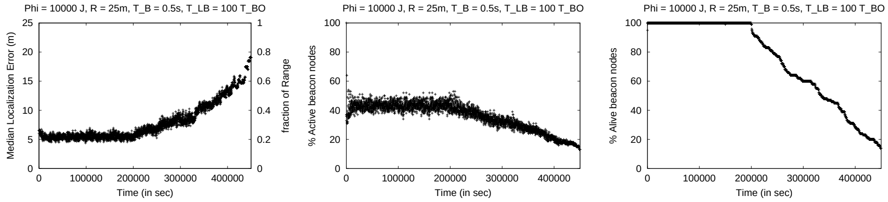

Figure 6: STROBE performance for N=100, R=25m, $T_B = 0.5s$, $T_V = 2T_B$, $T_D = 100T_V$, Φ=10000J.Lifetime of the algorithm using simple beaconing $L_B = 300000s$. Snapshot time = 100s.

Table 1: Energy Consumption parameters

Power Dissipation	Radio Operation Mode	Value
P_X	Transmit	660 mW
P_R	Receive	395 mW
P_I	Idle	35 mW
P_S	Sleep	0 mW

uniformly at random in a $100m \times 100m$ terrain. The nominal radio range of these beacons is 25m. The corresponding beacons per neighborhood μ = 19 = 3.1ν. We choose a reasonably small beaconing interval, $T_B = 0.5$ seconds. We set the various STROBE parameters as follows: $T_B = 0.5s$, $T_V = 2T_B$, $T_D = 100T_V$, Φ=10000J.

Figure 6 plots the median localization error, percentage of active beacons and percentage of beacons alive as a function of time. Snapshots are taken every 100 seconds. The degradation in median localization error as well as percentage of beacons alive is smooth. STROBE maintains a median localization error within $0.2 \times Range$ for up to 200000 seconds, $0.3 \times Range$ for up to 300000 seconds, and $0.5 \times Range$ for up to 400000 seconds. Actual system lifetime (L_{STROBE}) is increased to around 450000 seconds or $1.5L_B$. This is low compared to the best case lifetime of 850000 seconds or $2.8L_B$ possible[4], possible when energy usage can be uniformly load balanced. In practice, because STROBE transitions probabilistically from Voting (V) to Sleep (SL) states, a higher percentage of beacons than the threshold level remain active, leading to lower lifetimes. Leveraging auxiliary information can further improve these lifetimes. Lifetime gains can be improved further for higher beacon densities and energy dissipation rates in active state, and by augmenting STROBE with boundary detection mechanisms.

5 Conclusions and Future Work

Next generation pervasive computing applications that must scale to large numbers of hardware and energy-constrained devices, motivate scalable, ad hoc deployable approaches to localization that leverage existing device capabilities. In our work highlighted in this paper, we have (i) addressed *scalability* by developing an RF-connectivity based self-localization methodology for very small devices, implementing it on two different radio platforms; and validating it in outdoor environments (ii) addressed *ad hoc deployability* by developing and evaluating via simulation algorithms for automatically adapting beacon placement in more cluttered environments, and shown that the approach to beacon placement must take into account density of deployment. These contributions enable a low cost localization system that is both scalable and ad hoc deployable. The algorithms and methodology for beacon placement proposed here can be applied to other localization approaches (for instance, those that rely on multilateration/ranging), and also to other ad hoc wireless network problems influenced by node density and placement - for instance, topology maintenance and energy-efficient geographical routing. Now that we have validated our basic design of our adaptive beacon placement approach through simulation, we are currently testing our algorithms in real prototype systems with UC Berkeley motes.

Although our application space is rather forgiving in terms of localization granularity required, and does not include all applications, it covers a good range of important and useful applications. Finally, our localization approach is unique in its emphasis on *adaptive self-configuration i.e.*, adaptation to noisy environments through *decentralized, autonomous and measurement-based* techniques rather than careful instrumentation.

References

[1] P. Bahl and V. N. Padmanabhan. Radar: An in-building user location and tracking system. In *Proc. of the IEEE Infocom 2000*, volume 2, pages 775–84, March 2000.

[2] N. Bulusu, J. Heidemann, and D. Estrin. Gps-less low cost outdoor localization for very small devices. *IEEE Personal Communications Magazine*, 7(5):28–34, October 2000.

[3] N. Bulusu, J. Heidemann, and D. Estrin. Adaptive beacon placement. In *Proc. of ICDCS-21*, Phoenix, Arizona, USA, April 2001.

[4] N. Bulusu, J. Heidemann, and D. Estrin. Density-adaptive algorithms for beacon placement in wireless sensor networks. Technical Report UCLA CS TR 010013, UCLA Computer Science Department., May 2001. Available at http://www.isi.edu/ bulusu/papers/.

[5] M. Charikar, S. Guha, D. Shmoys, and E. Tardos. A constant-factor approximation algorithm for the k median problem. In *Proc. of ACM STOC 1999*, pages 1–10, May 1999.

[6] Radiometrix Corporation. http://www.radiometrix.com.

[7] U.S. Wireless Corporation. http://www.uswcorp.com/USWCMainPages/our.htm.

[8] L. Doherty, K. S.J. Pister, and L. El-Ghaoui. Convex position estimation in wireless sensor networks. In *Proc. of IEEE Infocom 2001*, Anchorage, Alaska, April " 22-26" 2001.

[9] D. Estrin, R. Govindan, J. Heidemann, and Satish Kumar. Next century challenges: Scalable coordination in sensor networks. In *Proc. of ACM/IEEE MobiCom 99*, pages 263–270, Seattle, WA, USA, August " 15–20" 1999. ACM Press.

[10] L. Guibas, D. Lin, J. C. Latombe, S. LaValle, and R. Motwani. Visibility-based pursuit evasion in a polygonal environment. *International Journal of Computational Geometry Applications*, 2000.

[11] A. Harter, A. Hopper, P. Steggles, A. Ward, and P. Webster. The anatomy of a context-aware application. In *Proc. of ACM/IEEE MobiCom 99*, Seattle, WA, USA, August " 15–20" 1999. ACM Press.

[12] J. Hill, R. Szewczyk, A. Woo, S. Hollar, D. Culler, and K. Pister. System architecture directions for networked sensors. In *Proc. of ASPLOS-IX*, pages 93–104, Cambridge, MA, USA, November 2000. ACM.

[13] W. J. Kaiser. Wins ng 1.0 transceiver power dissipation.

[14] B. Karp and H.T. Kung. Gpsr: Greedy perimeter stateless routing for wireless networks. In *Proc. of ACM/IEEE MobiCom 2000*, N.Y., August 2000. ACM Press.

[15] L. Kleinrock and J. Silvester. Optimum transmission radii for packet radio networks or why six is a magic number. In *Proc. of National Telecommunications Conference*, pages 4.3.1–4.3.5, 1978.

[16] S. Meguerdichian, F. Koushanfar, M. Potkonjak, and M. Srivastava. Coverage problems in wireless ad hoc sensor networks. In *Proc. of IEEE Infocom 2001*, Anchorage, Alaska, April " 22-26" 2001.

[17] VOR(Very High Frequency Omnirange). http://www.allstar.fiu.edu/aero/VOR.htm.

[18] N. B. Priyantha, A. Chakraborty, and H. Balakrishnan. The cricket location support system. In *Proc. of ACM/IEEE MobiCom 2000*, Boston, MA, August 2000.

[19] Inc. RF Monolitithics. http://www.rfm.com.

[20] R. Want, A. Hopper, V. Falcao, and J. Gibbons. The active badge location system. *ACM Transactions on Information Systems*, 10(1):91–102, January 1992.

[21] J. Werb and C. Lanzl. Designing a positioning system for finding things and people indoors. *IEEE Spectrum*, 35(9):71–78, sep 1998.