A Passive Approach to Sensor Network Localization

Rahul Biswas and Sebastian Thrun
Computer Science Department
Stanford University
Stanford, CA 94305; USA
Email: {rahul,thrun}@cs.stanford.edu

***Abstract*— Sensor networks present the opportunity to accurately localize phenomena of interest. To be able to do so however, sensor nodes need to themselves be accurately localized. We present herein an algorithm to do this based on uncontrolled sounds observed by each of the sensor nodes. A probabilistic generative model is presented and it is shown that the sensor nodes localization problem is equivalent to maximum likelihood estimation in the model. Experimental results are presented for both simulated sensor nodes and Crossbow MICA2 sensor nodes.**

I. Introduction

Recent developments in manufacturing technology have created a new generation of very small computers equipped with one of more sensors, limited computational capabilities, and wireless radio links. As the cost of manufacturing these so called sensor nodes decreases, the possibility of using large numbers of nodes in aggregate to form a sensor network has become possible. Such networks are essential for monitoring large environments unobtrusively and at a finer scale than was previously possible. Sensor networks can be used in a variety of civilian and military contexts, including factory automation, inventory management, environmental and habitat monitoring, health-care delivery, and battlefield awareness.

Most sensor network algorithms need to know the physical locations of individual sensors. While it is sometimes possible to measure these locations by hand, a rich body of sensor network localization algorithms has emerged to allow the sensor network to do this task autonomously. These techniques typically however take an active approach to localization. Some approaches place special beacon nodes with known location into the environment while others rely on expensive hardware or require the sensor nodes to beep to one another.

We offer a passive solution to the sensor node localization problem. Our algorithm allows sensor nodes to localize themselves from sounds naturally occurring in their environment. These sounds can be anything, from birds chirping to the sound of machinery. Our algorithm works despite knowing neither where nor when the external events generating these sounds take place. This passive approach expends less sensor node energy, does not require infrastructure outlay, and makes the sensor network less detectable than when employing active methods. Moreover, it is competitive with active localization methods and both more accurate and robust than other passive localization techniques.

Our approach structures the relationship between the sensor node locations, the locations and times of the external events generating the sounds heard by the sensor nodes, and the recorded times as a Bayesian network. We show how the localization problem reduces to maximum likelihood estimation in the Bayesian network and explain how we do this via gradient descent in an expression monotonic to the negative likelihood.

We synchronize the clocks of the sensor nodes with respect to one another and assume a time recorded at one sensor node is directly comparable to a time recorded at another sensor node [4]. We also assume that sounds are sufficiently well dispersed to prevent data association ambiguities between events.

II. Problem Definition

A. Overview

Let the sensor network consist of N nodes at locations $S = \{S_1 \dots S_N\}$. Let S_i^x refer to the x-coordinate of the location of the sensor node i and let S_i^y refer to the y-coordinate. We assume 2-D locations but the same technique will work in 3-D as well. These locations are not known to us.

Second, let there be M sound-generating events at locations $E = \{E_1 \dots E_M\}$. Moreover, let E_i^x and E_i^y represent the respective coordinates. These locations are not known to us.

Third, let the times of the sound-generating events be $T = \{T_1 \dots T_M\}$. These times are not known to us.

Fourth, let there exist $M \cdot N$ variables R_n^m that specify the time recorded at sensor node S_m for sound E_n. We will revisit these variables in Section IV when we structure the relations between these variables as a Bayesian network.

Given the sensor measurements R, we seek the sensor node locations S.

B. Sensor Model

We assume that R_n^m is distributed as follows:

$$\begin{aligned} R_n^m &= N(T_n + \frac{d(S_m, E_n)}{s}, \sigma) \qquad (1) \\ d(S_m, E_n) &= \sqrt{(S_m^x - E_n^x)^2 + (S_m^y - E_n^y)^2} \end{aligned}$$

where s is the speed of sound in free space, σ is the standard deviation of the error in recording times, and $N(\mu, \sigma)$ is a Normal distribution with mean μ and standard deviation σ.

III. Prior Work

A very popular technique often used outdoors is the Global Positioning System (GPS) [5]. It offers an inexpensive solution for many applications but lacks the relative accuracy generally required for sensor network applications.

Methods based on received signal strength (RSS) of radio signals as well as radio connectivity (e.g. [1], [2]) offer inexpensive solutions for both indoor and outdoor environments but are often not robust enough for many sensor network applications.

Another class of localization schemes assume the presence of so called beacon nodes that know their location with respect to a global reference frame [3], [9], [10]. While these methods are both robust and accurate, they require substantial cost and effort to set up beacon nodes and measure their locations by hand.

It is also possible to use ultrasonic ranging [11] or laser-range finders to perform ranging as an input to a localization system. These schemes require precise instrumentation or multiple microphones per sensor node because the sender and the receiver need to be precisely oriented to one another. The cost of this equipment substantially raises the cost of sensor network deployment.

Ranging in the audible domain [11] offers a solution that is both relatively inexpensive and reliable. It requires the sensors to beep to one another though. This expends both extra energy and may cause unwanted disturbance by or detection of the sensor nodes.

In terms of algorithms, our problem has strong analogues to the statistical problem of multi-dimensional scaling. There, dissimilarities between points are used to compute a planar embedding. Indeed, multi-dimensional scaling offers a direct solution to localization from ranging data but not knowing when the sounds occurred preempts its use in this domain.

The problem also resembles the affine structure from motion problem in computer vision when the sounds and the sensor nodes are sufficiently distant. There, corresponding points between multiple images are used to recover both camera position as well as the 3-D structure of a scene. In both cases, we seek to recover the rotations and translations of sensors observing common phenomena and with this information, complete the locations of points in a higher dimensional space.

IV. Our Approach

In this section, we encapsulate the relations between variables as a Bayesian network and show how the sensor node localization problem reduces to that of maximum likelihood estimation in this Bayesian network. We start with the formal definition of our Bayesian network, give conditions for when an unique solution exists,

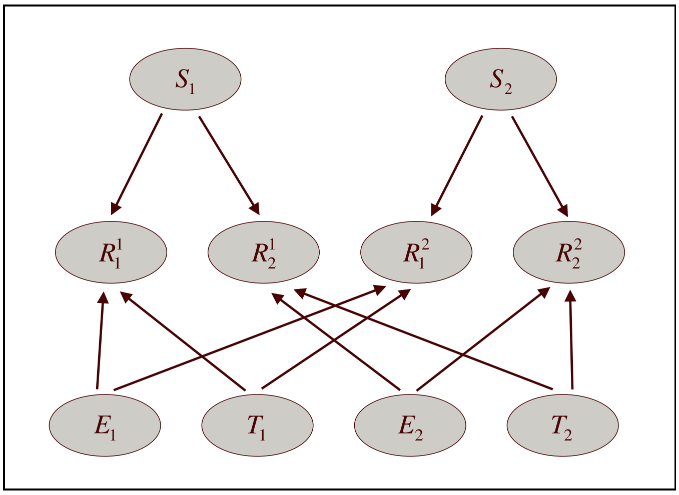

Fig. 1. Bayesian Network with 2 sounds and 3 sensor nodes

compute the likelihood, and show how gradient descent is used to perform likelihood maximization.

A. Bayesian Network

A Bayesian network [7] is characterized by its variables and the prior and conditional probabilities relating the variables to one another. We start with the variables and what they represent, then move on to the prior and conditional probabilities. Figure 1 shows a Bayesian network for the case of two sounds and two sensor nodes. Each of these variables are identical to the variables discussed in Section II-A.

1) Variable 1 : Sensor Node Locations: Let there be N sensor node location variables $S_1 \ldots S_N$ in our Bayesian network. These variables are not observed in the Bayesian network, i.e. their values are not known to the gradient descent algorithm. After doing gradient descent however, it is possible to form estimates regarding them.

2) Variable 2 : Sound Locations: Let there be M sound location variables $E_1 \ldots E_M$ in our Bayesian network. These variables are also unobserved.

3) Variable 3 : Sound Times: Let there be M sound time variables $T_1 \ldots T_M$ in our Bayesian network. These variables are unobserved as well.

4) Variable 4 : Recorded Times: Let there be P recorded time variables of the form

$$R(p) = R_{\phi(p)}^{\pi(p)}$$

where $\pi(p)$ maps the index of the acoustic event p to the sensor node that recorded this time and $\phi(p)$ maps the index of the sound p to the sound whose time was recorded. Not all sensor nodes will hear all sounds so the structure of the Bayesian network will depend on which observations actually occur. These variables are observed, i.e. their values are known.

5) Prior Probabilities: We assume that both sensor nodes and sounds are uniformly distributed in a rectangular region. Thus the prior probabilities are $P(S_i) = \frac{1}{A}$ and $P(E_j) = \frac{1}{A}$.

We also assume that sound times are uniformly distributed from the activation of the sensor network to when the last acoustic event is recorded. Thus the prior probability $P(T_j) = \frac{1}{T}$.

6) Conditional Probability : How Recorded Times are Generated: Recorded times are stochastically determined by sensor node locations, sound locations, and sound times as presented in Section II-B. Thus the conditional probability is:

$$P(R_j^i|S_i, E_j, T_j) \quad = \qquad (2)$$
$$\frac{1}{\sqrt{2\pi\sigma}} \cdot exp(\frac{(T_j + (d(S_i, E_j))/s - R_j^i)^2}{2 \cdot \sigma^2})$$

where $exp(x)$ stands for e^x.

B. Existence of an Unique Solution

Since there is no global reference frame, we impose the following arbitrary conditions to regularize our solution. S_1 is set to be (0,0). S_2 is set to lie on the positive x-axis. S_3 is set to lie above the x-axis, i.e. have a positive y-coordinate. This canonical form of the solution space is easily adapted if one or more sensor nodes have known location with respect to a global reference frame.

It is possible that R does not provide enough information to determine S. Sometimes, this lack of information is predictable given M, N, and P. Each variable S_i has components S_i^x and S_i^y so introduces two unknown variables. Each variable E_j does the same and each T_j variable introduces one unknown variable. Each variable R_j^i introduces one equation. Moreover, when $M \geq 2$, the canonicalization of the solution sets S_1 to (0,0) and constrains LS_2 to (x,0), thereby eliminating three unknown variables.

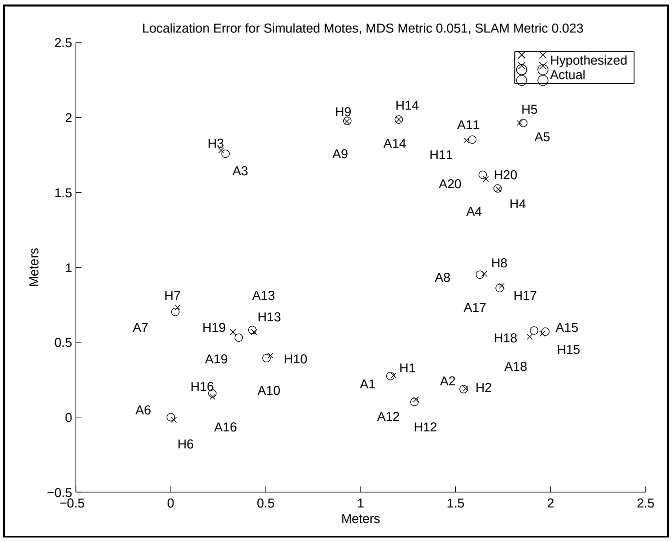

Fig. 2. Twenty Sensor Nodes in Simulation

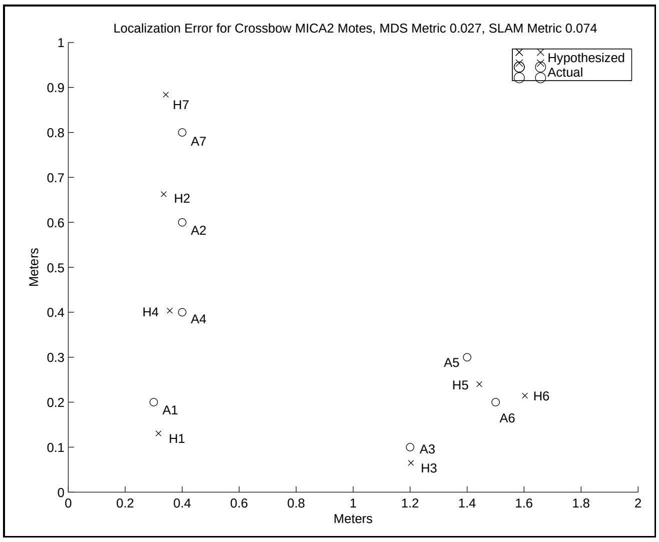

Fig. 3. Seven Crossbow MICA2 Sensor Nodes

Thus, only if:

$$P \geq 2 \cdot M + 3 \cdot N - 3 \tag{3}$$

can there exist an unique solution. In situations where the inequality is strict, we have an over-constrained set of equations. This information is necessarily not contradictory given the nature of the scenario. Even when the inequality is satisfied, there exist degenerate situations where the information is not sufficient for unique localization. For example, the Bayesian network may not be connected and this will lead to two sets of nodes whose relative positions may remain unknown.

C. Maximum Likelihood Estimation

Let $L(S,E,T,R)$ be the likelihood of the data in the Bayesian network. The likelihood of a Bayesian network is simply the product of each of the prior and conditional probabilities.

Doing this multiplication, we obtain:

$$\begin{aligned} & L(S,E,T,R) & (4)\\ &= \prod_{m=1}^{M} P(S_m) \cdot \prod_{n=1}^{N} P(E_n) \cdot \prod_{n=1}^{N} P(T_n) \cdot \\ & \prod_{p=1}^{P} P(R(p)|S_{\pi(p)}, E_{\phi(p)}, T_{\phi(p)}) \\ &= \prod_{m=1}^{M} \frac{1}{A} \cdot \prod_{n=1}^{N} \frac{1}{A} \cdot \prod_{n=1}^{N} \frac{1}{\hat{T}} \cdot & (5) \end{aligned}$$

$$\begin{aligned} & \prod_{p=1}^{P} \frac{1}{\sqrt{2\pi}\sigma} \cdot exp(\frac{(T_j + (d(S_i, E_j))/s - R_j^i)^2}{2 \cdot \sigma^2}) \\ &= \frac{1}{A^{M+N}} \cdot \frac{1}{\hat{T}^N} \cdot \frac{1}{(\sqrt{2\pi}\sigma)^P} \cdot & (6)\\ & exp(\frac{\sum_{p=1}^{P}(T_j + d(S_{\pi(p)}, E_{\phi(p)}))/s - R_{\phi(j)}^{\pi(i)})^2}{2 \cdot \sigma^2}) \end{aligned}$$

Since the log-likelihood is monotonic in the likelihood, we can take the log of Equation 6 to obtain:

$$\begin{aligned} l &= -(M+N) \cdot logA & (7)\\ & -N \cdot log\hat{T} - P \cdot log(\sqrt{2\pi}\sigma) + \\ & \frac{1}{2\sigma^2} \cdot \sum_{p=1}^{P}(T_j + \frac{d(S_{\pi(p)}, E_{\phi(p)})}{s} - R_{\phi(j)}^{\pi(i)})^2 \end{aligned}$$

Note that the first three terms are independent of the Bayesian network variables. We can remove these terms and multiply by $2\sigma^2$ to obtain another expression monotonic in the likelihood:

$$l' = \sum_{p=1}^{P}(T_j + \frac{d(S_{\pi(p)}, E_{\phi(p)})}{s} - R_{\phi(j)}^{\pi(i)})^2 \tag{8}$$

We use gradient descent with both line search and momentum to minimize $-l'$. Due to the monotonicity of l' with the likelihood, this process effectively finds the maximum likelihood estimate of the sensor node locations.

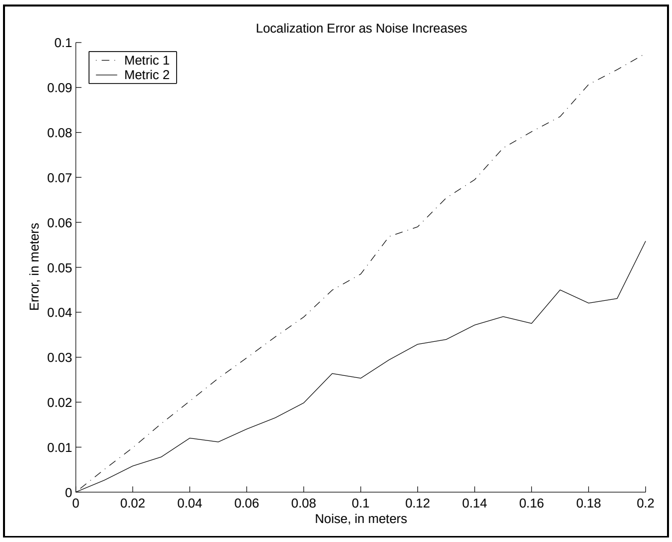

Fig. 4. Effects of Changing Sensor Noise

V. Experimental Results

We show in this section some examples of the sort of localization results the algorithm produces. First, we define the metrics we will use to evaluate the quality of our results. Second, we present the results of two specific networks. Third, we show how the error varies as we modify different parameters of the localization scenario.

A. Error Metrics

1) Metric 1: Scaled Likelihood: We use $\sqrt{-l'}\cdot s$ as our first metric. Intuitively, this metric measures the standard deviation of the expected distances between sounds and sensor nodes suggested by the data versus the distances stemming from the hypothesized solution.

2) Metric 2: Actual Error: For our second metric, we use the actual standard deviation of the predicted sensor node location from the actual sensor node locations, given optimal alignment between the two coordinate frames.

B. Specific Networks

1) Twenty Sensor Nodes in Simulation: Figure 2 shows localization results for a scenario with twenty sensor nodes, twenty sounds, and a standard deviation in recording error that corresponds to 10 cm. The predicted error, given by Metric 1, is 5.1 cm and the actual error is 2.3 cm. The circles in the figure represent the true sensor node location and the crosses represent the predicted location.

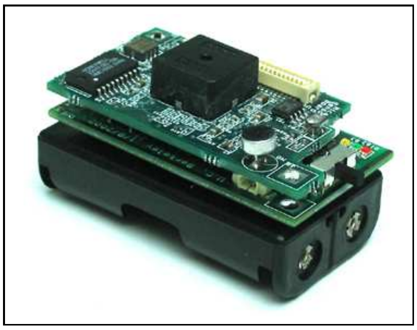

Fig. 5. Crossbow MICA2 Sensor Node

2) Seven Crossbow MICA2 Sensor Nodes: Figure 3 shows localization results using seven Crossbow MICA2 sensor nodes (see Figure 5) and seven sounds. The predicted error is 2.7 cm and the actual error is 7.4 cm.

C. Effects of Changing Network Parameters

In this section, we examine the effect of different parameters on the localization algorithm.

1) Effects of Changing Sensor Noise: Figure 4 shows the effect of varying the standard deviation of the recording time to the equivalent of 0 cm to 20 cm. Twenty sensor nodes and twenty sounds are used throughout. We find that both errors rise approximately linearly as the noise increases.

2) Number of Sensor Nodes: Figure 6 shows the effect of varying the number of sensor nodes being localized from four to twenty. Twenty acoustic events are used and the standard deviation in the recording error is set to correspond to 10 cm. We find that the algorithm is initially overconfident when dealing with less than seven sensor nodes. With more than eight sensor nodes, the error levels off to approximately 3 cm.

3) Number of Sounds: Figure 7 shows the effect of varying the number of sounds used for localization from four to twenty. Twenty sensor nodes are localized and the standard deviation in the recording error is set to correspond to 10 cm. We find that the error initially drops sharply and continues to decrease slowly as the number of sounds increases.

VI. Future Work

While we feel that we have made significant progress with the methods presented herein, the

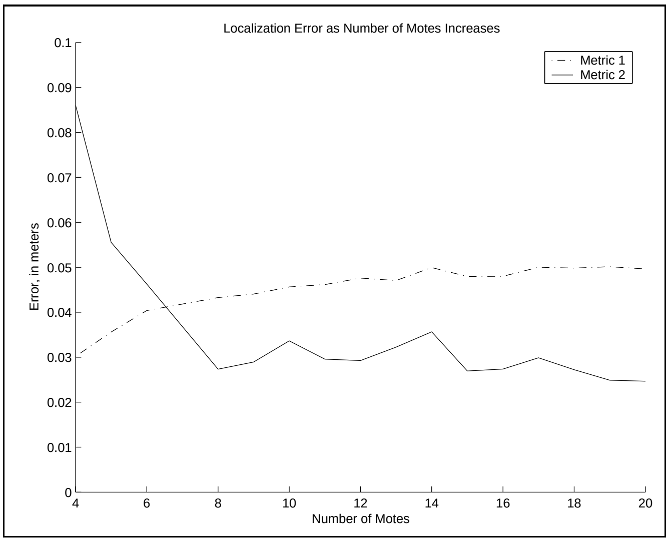

Fig. 6. Effects of Changing Number of Sensor Nodes

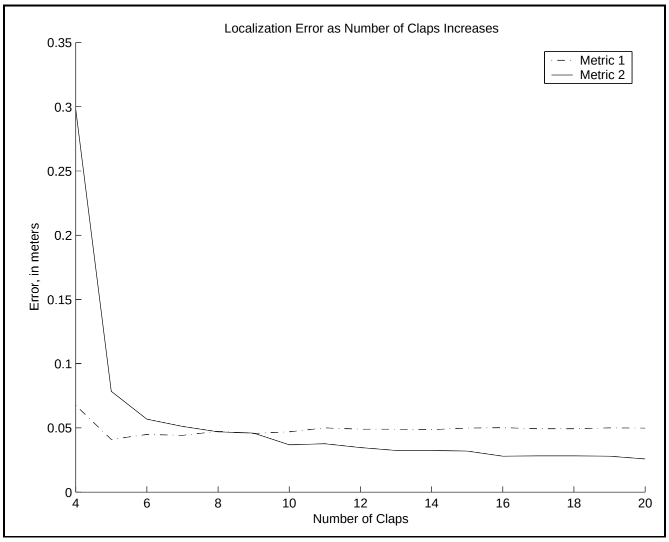

Fig. 7. Effects of Changing Number of Sounds

algorithm unfortunately does not fare as well when only a subset of the sensor nodes hears each sound. There are two typical failure modes.

First, the predicted error may be low but the actual error is high. This suggests that the problem is ill-conditioned in a way even when the condition of Equation 3 is met. We plan to explore the conditions under which the predicted error is or is not a reliable estimate of the actual error.

Second, gradient descent may get stuck in local minima, even with random restarts. We plan to explore more robust probabilistic reasoning algorithms to help deal with this case.

VII. Conclusion

We have presented herein a localization algorithm for sensor nodes that uses cheap, commodity hardware and relies only on uncontrolled sounds for localization. A probabilistic generative model is presented and it is shown how the localization problem reduces to maximum likelihood estimation on this model. Experimental results are presented for both simulated sensor nodes and Crossbow MICA2 sensor nodes.

VIII. Acknowledgments

This research is supported in part by DARPA's MARS Program (Contract number N66001-01-C-6018) and Intel Corporation, which is gratefully acknowledged. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing official policies or endorsements, either expressed or implied, of the United States Government or any of the sponsoring institutions.

References

[1] P. Bahl and V. Padmanabhan. “Enhancements to the RADAR user location and tracking system.” Technical Report MSR-TR-2000-12, Microsoft Research, Redmond, WA, USA, 2000.

[2] P. Bahl and V.Padmanabhan. “RADAR: An in-building user location and tracking system.” Proc. IEEE Infocom, 2000

[3] N. Bulusu. *Self-Configuring Localization Systems* Ph.D. Thesis, University of California at Los Angeles, 2002.

[4] J. Elson, L. Girod, and D. Estrin. “Fine-Grained Network Time Synchronization using Reference Broadcasts.” Proc. OSDI, 2000

[5] B. Hoffman-Wellenhof, H. Lichtenegger, and J. Collons. *Global positioning system: Theory and practice, 2nd ed.*” Springer-Verlag, 1992.

[6] J. Kahn, R. Katz, and K. Pister. “Mobile networking for smart dust” Proc. ACM/IEEE Intl. Conf. on Mobile Computing and Networking, 1999.

[7] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference.* Morgan Kaufmann Publishers, Inc., 1988.

[8] G. Pottie and W. Kaiser. “Wireless integrated network sensors.” Proc. ACM Communications, 2000.

[9] N. Priyantha, A. Chakraborty, and H. Balakrishnan. “The Cricket location support system.” Proc. ACM MobiCom, 2000.

[10] A. Ward, A. Jones, and A. Hopper. “A new location technique for the Active Office.” Proc. IEEE Personal Communications Magazine, 1997.

[11] K. Whitehouse *The Design of Calamari: an Ad-hoc Localization System for Sensor Networks* M.S. Thesis, University of California at Berkeley, 2002.