Reliable Estimation of Influence Fields in Unreliable Sensor Networks

Sandip Bapat, Vinodkrishnan Kulathumani, Anish Arora
Department of Computer Science and Engineering
The Ohio State University
Columbus, OH 43210, USA

Abstract

The influence field is a commonly exploited feature in sensor network applications. By way of example, it is used in classification and tracking of different types of objects using a wireless sensor network. Being spatially distributed, this feature allows us to tradeoff nodal computation with network communication. By the same token, however, not only is its calculation subject to nodal failures but also to channel fading and channel contention. In this paper, we study how to accurately and efficiently estimate the influence fields of objects in such an unreliable setting so the objects can still be distinguished. We derive, for node and network fault models, the necessary nodal density for estimating the influence fields. We also identify conditions under which additional algorithmic techniques are needed to compensate for faults and present three such techniques: probabilistic reporting, temporal aggregation/segregation, and spatial reconstruction. We provide experimental corroboration of our analysis through 40-50 Mica2 sensor node experiments. Finally, we demonstrate how reliable and efficient classification and tracking were thus achieved in a fielded system of 90 Mica2 sensor nodes that we called "A Line In The Sand".

1 Introduction

The influence field of an object is, intuitively speaking, the region where the object is detectable (see Fig. 1). The influence field of an object thus depends on the type(s) of sensors being used for detection (magnetometers in Fig. 1), and it is essentially characterized by the area and the shape of the region. It is useful for a variety of sensor network applications, especially those of surveillance, where typical tasks include detection, classification, and tracking of various types of objects. This is because object types can often be distinguished based on their influence field areas and/or shapes and located based on some intrapolation of the influence field.

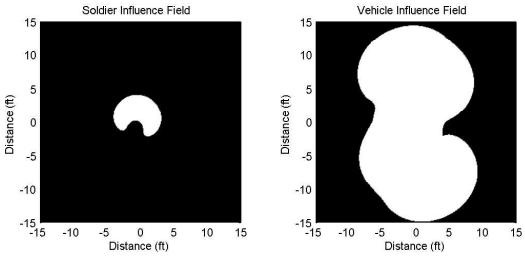

Figure 1. Magnetometer based influence fields for two object types.

In a system that we recently fielded in several outdoor settings, the influence field feature was successfully exploited as the primary basis for classification and tracking of people, people carrying a significant amount of metal on them (aka "soldiers") and cars, via a dense, wireless sensor network. The intended use of this system, which was called *A Line in the Sand*, is to use a sensor network to protect high-valued assets, to secure extended borders or perimeters, and to monitor activity of personnel or vehicles in remote or access-denied geographic areas. Our past deployments covered only a 20m by 10m area using 90 sensor nodes; currently, the influence field approach of *A Line in the Sand* is being tested for deployment in a 10km by 500m area using 10,000 sensor nodes.

The influence field approach should be contrasted to the traditional approach for classification and tracking using some form of Unattended Ground Sensors. The Remotely Monitored Battlefield Sensor System (REMBASS) is a representative example. REMBASS devices are equipped with high-fidelity acoustic, seismic and magnetic sensors and have sufficient processing power and memory to perform time-frequency analysis and

signature matching algorithms on the sensor outputs. These devices have a fairly large range for some object types, and provide high classification accuracy and low false positives. On the other hand, they are expensive, require careful and precise deployment as well as frequent remote monitoring. Their coverage area for relatively small objects, such as persons, is rather limited. And the failure of a single REMBASS device causes loss of coverage in its region, as they are too expensive to be redundantly deployed.

The drawbacks of the REMBASS approach motivate the use of a dense network of relatively inexpensive nodes in contexts such as *A Line in the Sand*. In turn, the resource constraints and limited sensing and communication range of such nodes —we used Mica 2 motes— motivate the suitability of using the influence field. Each sensor node merely has to detect a binary "presence" of an object; network-based aggregation of these bits yields the influence field without requiring substantial or complex node operation.

The key challenge in realizing the influence field, as our experience with *A Line in the Sand* showed, is the unreliability of wireless sensor networks. Event loss –both in nodes and in the network– is fundamental to wireless sensor node platforms and its impact on the application can be substantial. Thus both node and network unreliability have to be dealt with in estimating the influence field.

In the particular case of *A Line in the Sand*, network contention errors dominated the unreliability, intuitively because the number of messages being aggregated was large relative to the capacity of the network. This unreliability affected the design not just of the application, but also of every network service that we implemented. (As the reader will see, in some instances, the network service tuning has to be performed jointly with that of the application.) To the best of our knowledge, the impact of network unreliability in estimating the influence field has not been addressed before, nor has it been addressed in the context of distributed classification and tracking, which has led us to the present work.

Contributions of the paper. In this paper, we address the problem of reliably estimating the influence field of various objects types in the presence of various types of sensor network faults, in terms of how to preserve the differences in the area and/or shape across various object types.

For each fault type, we provide: (1) analytical results on the fault-affected estimates of the influence field, (2) procedures for calculating the ideal sensor node density for efficient and reliable estimation, (3) algorithms techniques to provide efficient and reliable estimation when the deployment density does not meet or exceeds the ideal density, and (4) where appropriate, experimental corroboration of our analysis or algorithmic techniques realized using 40–50 Mica2 sensor node experiments in our testbed.

Our results can be instantiated in a compositional way, and thus accommodate cases where multiple fault types occur. We show that shows that node and network faults raise competing (i.e., two-sided constraints) on the density of the network.

Finally, we show how we experimentally dealt with reliable estimation of influence fields to classify and track persons, soldiers, and cars in *A Line in The Sand*. The case study provides a data point for the importance of dealing with network unreliability in both the network and the application level. It also identifies a need for routing protocols in sensor networks to provide uniform reliability (at least for nodes that are equidistant from an aggregator node).

Organization of the paper. In Sec. 2, we describe previous related work. In Sec. 3, we present the system and fault models. We define and analyse the influence field feature, state its relation to the concept of sensor coverage, and formulate the problem of its reliable estimation in Sec. 4. Then, in Sec. 5, we analyze the impact of faults on estimation, derive necessary conditions for reliably estimating the area and shape of various objects' influence field in a manner that preserves their difference, and provide three algorithmic techniques that to deal with these faults whenever necessary. In Sec. 6, we provide some details of our implementation of a *A Line in the Sand* and how we dealt with different fault classes in our design. Finally, we make concluding remarks and discuss future work in Sec. 7.

2 Related Work

The notion of influence of an energy source is used in many science and engineering applications. In some formulations, the distribution of the intensity of the source at various points is considered

while modeling its influence. For example, Kellogg et al [1] model the temperature distribution of a heat source across a region as an *influence graph* and use the graph to design algorithms for distributed control. In other formulations, including Zhao et al [14] and ours, the distribution of the intensity is not modeled. Zhao et al [14] define an *influence area* as the number of sensors that "hear" an object. Our definition of the influence field also captures the shape of the influence of the object.

As mentioned in the introduction, we are not aware of prior work on reliably estimating influence field of an object in the presence of network faults. For the case of node faults, Krishnamachari et al [8] have presented probabilistic decoding mechanisms to detect regions of events in the presence of uncorrelated sensor faults with relatively low probability (around 10%). Our work accommodates the case of uniform nodal failures, we also present techniques to handle network faults whose impact is non-uniform across the network such as fading, and network faults whose failure probability grows with the event size such as contention.

With respect to applications of the concept of influence, Zhao et al [14] suggest that the influence area can be used to track multiple targets that are separated in space. To the best of our knowledge our work is the first to consider the robustness of the approach in classification and tracking in an unreliable network (the *A Line in the Sand* network has unreliability as high as 50%). We are not aware of prior work that applies the estimation of an influence field for robust classification.

Most approaches for classification and tracking of targets that are not based on influence use a centralized approach that typically involves pattern recognition or pattern matching using time-frequency signatures produced by the different types of objects [3]. Meesookho, et al [2] describe a collaborative classification scheme based on exchanging local feature vectors. The accuracy of this scheme improves only as the number of collaborating sensors increases, which imposes a high load on the network. By way of contrast, Duarte et al [5] describe a classifier in which each sensor extracts feature vectors based on its own readings and passes them through a local pattern classifier. Although the approach loads the network only slightly, it requires significant computational resources at each node. Most of the work on distributed tracking is based on collaborative signal and information processing, sequential Bayesian filtering, and extended Kalman filtering [4,9–11,14,15], that require significant node computation.

Our work relates the influence field to the notion of sensor coverage [6, 13]. Typically, sensor coverage is defined as independent of the type of objects at hand, whereas we consider a definition that is with respect to each target type.

3 System and Fault Models

In this section, towards defining and estimating the influence field we outline the system model. And towards analyzing the impact of faults on the reliability of the estimation of influence fields, we outline the fault models.

3.1 System model

The system consists of a number of wireless sensor nodes, N. Each node has a unique identifier and consists of a processing unit, memory, radio, power source, and one or more sensors of different types. A single node has limited processing power, memory, and energy which make the execution of complex, computation intensive algorithms on it infeasible. We assume a localization service that provides the relative or absolute position for each node and a time synchronization service that enables each node to timestamp its detection with a time value that is the same at all nodes in the network.

The sensor nodes are distributed uniformly over a geographic region that is to be monitored. We model this region as a large finite number, Ω, of perfectly spaced logical points that serve as the potential locations where nodes can be deployed. We denote the ratio N/Ω by ρ, which intuitively represents the sensor density in the network. When we refer to the *area* A of a subregion, we mean A is the number of these Ω points in the subregion.

We assume that the wireless network is connected, hence it is possible to aggregate messages from any subset of the nodes in the system, albeit that such aggregation may require multi-hop communications.

3.2 Fault models

Sensor networks are subject to a large class of faults, resulting from inexpensive hardware,

limited resources, unreliable communications and extreme environmental conditions. We consider both node and network fault types.

Node faults. Sensor nodes fail in a variety of ways, including hardware failure, software crash/deadlock/livelock, sensor or actuator debonding, battery exhaustion, or simply in the form of a transient event loss.

At any time, the net effect of a node fault can be modelled as missing the detection of a object, i.e. a false negative, or asserting the detection when there is no object, i.e. a false positive. Although the analysis and techniques described in the paper can be made to apply to the latter class, we henceforth consider only false negatives.

Node fault model: Formally, we assume that the probability that any node misses the detection of an event, whether it is due to transient, permanent or intermittent node fault, is $1 - p_n$. In other words, we assume that false negatives at nodes are independent of each other and independent of the objects.

Network faults. Wireless communications in the multi-hop network are subject to both fading and contention effects.

Due to the limited transmission power of each node, messages sent over the wireless channel are subject to propagation loss in the medium. This fading loss depends on link characteristics such as distance, relative orientation of sender and receiver, altitude differences, environmental conditions, etc.

Fading model: Formally, we assume that there is a uniform probability, $1 - p_f$, of message loss due to fading for any single hop communication in the network. It follows that in the absence of any other faults, the probability of message loss due to fading for any h hop communication is $1 - p_f^h$.

The broadcast nature of wireless communication implies that nodes share the same physical channel. This leads to loss due to channel contention when multiple senders try to transmit messages on the same frequency at the same time. The degree of message loss depends on several factors like the type of Medium Access Control (MAC) protocol used, the inherent synchronicity of message transmissions, and most importantly the number of nodes trying to send data simultaneously. We assume a standard CSMA/CA MAC protocol wherein nodes try to avoid collisions using random backoffs and channel sensing.

Contention model: Formally, we assume that there is a uniform probability, $1 - p_c$, of message contention loss that is a function of the number of nodes simultaneously trying to send a message each and of the number of slots available for a node to choose its random backoff from. (Characterization and experimental corroboration of this function is provided in Sec. 5).

4 Influence Field

The *influence field* of an object j with respect to a given sensing modality is the region surrounding j where j will be "detected" by a sensor of that type located at any point in that region. By detected, we mean that the signal-to-noise ratio of the disturbance created by the object at the sensor exceeds a fixed detection threshold.

Of course, the influence field of different objects may differ for a given sensor modality: e.g., a car may have a larger influence field area than that of a motorcycle with respect to magnetometer sensors. And the influence field may differ for different sensor modalities: e.g., a car may have a smaller influence field area than that of a motorcycle with respect to acoustic sensors.

Assumptions. For simplicity of presentation, we assume that the size of the influence field of j is invariant with respect to its location. Likewise, the shape of the influence field is invariant with respect to location, up to rotation. Also for simplicity, we limit our attention to one given sensing modality. The modality will remain implicit in our notation.

4.1 Estimating the influence field

Estimation consists of calculating the area and the shape of the influence field. Recall that the distribution of sensor nodes is uniform.

Regarding the former, if we assume sensor node density ρ exceeds some lower bound that depends on the targets at hand, the estimation of the area A of the influence field of j is effectively reduced to counting the number of nodes that detect j. With uniform distribution, the number of sensors in any region of area A follows a binomial distribution with parameters (A, ρ). For practical values of A and ρ, we can exploit the rule of thumb for

the normal approximation to this binomial distribution that the value of the random variable lies within 3 times the standard deviation of the expected value in 99% of the trials.

Proposition 1 Given uniform deployment of sensors, whp the number of sensors that lie in the influence field of j is in the interval

$$(A \times \rho) \; - \; 3 \times \sqrt{A \times \rho \times (1-\rho)} \; .. \\ (A \times \rho) \; + \; 3 \times \sqrt{A \times \rho \times (1-\rho)}$$

Relation to sensor coverage. Sensor coverage in a region is the *minimum* number of sensors that "cover" (i.e., will detect) each point in the region. While this typical definition of the concept is independent of the type of objects at hand, a more useful definition for our purposes would be one that is with respect to each object type. (Thus, the sensor coverage of object j in a region may be different from the sensor coverage of another object in that region.)

Consider the influence field of j. Since sensor nodes are distributed uniformly, the number of sensor nodes in the influence field of j has the same distribution as the number of sensor nodes that cover j when j is located at any point in a region of area A. We can now relate the area of the influence field of j to sensor coverage of j.

Proposition 2 Given a uniform deployment of sensor nodes and any region of area A. If K is the sensor coverage of j in that region, then whp

$$K \quad \geq \quad (A \times \rho) \; - \; 3 \times \sqrt{A \times \rho \times (1-\rho)}$$

Regarding the shape of the influence field, if we assume that the sensor density exceeds some lower bound that depends on the objects at hand, estimating the shape of the influence field of j is effectively reduced to calculating the shape of the region that smoothly bounds the sensor nodes that detect j.

4.2 The problem of reliable estimation

Node and network faults impact estimation of the area and the shape of the influence field of j. In particular, the impact may not be simply proportional to the area and the distribution of the sensors whose detections are successfully aggregated may no longer be uniform.

Classification is an example of an application that can exploit area estimation. Different object types may be classified via separation between the areas of their respective influence fields. Errors in area estimates can thus result in misclassifications.

Tracking is an example of an application that can exploit shape estimation. Object location may be tracked from the locations of the sensors that detect it, i.e., by computing the centroid of their locations. Shape distortion errors can thus result in inaccuracy of tracking. (The same argument would apply for classification based on the shape of the influence field.)

We are therefore led to the problem of how to reliably estimate the size and shape of the influence fields of objects so that they can still be compared. More specifically, we focus on two subproblems of reliable estimation:

1. How to ensure that for objects whose respective influence field areas are separable, the separation remains between the fault-affected estimates of their respective influence field areas?

2. How to preserve the shape of the influence field of an object in the fault-affected estimate of the object influence field, by preserving the uniformity of distribution of the sensors whose detections are not affected?

5 Reliable Estimation

In this section, we address problems 1. and 2. outlined above successively for each of the fault models discussed in Sec. 3. Specifically, for each fault model,

- we analytically identify necessary conditions for accurately preserving the separation between area estimates and the shape estimates,

- we derive an ideal density for node deployement with respect to the given sensing modalilty for the efficiency of estimation, and

- we characterize algorithmic techniques and parameter tuning options to deal with situations where the ideal density is not achievable, including cases where the deployed density is less than and more than ideal.

Also, where appropriate, we provide experimental corroboration of our analysis and algorithmic techniques via experiments realized on Mica2 sensor motes.

We note that, conceptually, the analyses for the node faults and the fading faults are similar, whereas that for the contention faults is different. Moreover, our approach is compositional and thus the analysis, which is presented separately for each fault type, can build upon the constraints/distributions identified for other fault types.

5.1 Nodal failures

Let A_1, A_2, ... A_k be the influence fields of k types of objects ranging from the smallest to the largest. Recall that p_n is the probability that a sensor node does not miss a detection. The number of non-faulty nodes in a region of area A_i thus has a binomial distribution with parameters (n_i, p_n), where n_i is the number of nodes in the area A_i. However, since the network is uniformly distributed with nodes, n_i itself is a random variable that a binomial distribution, with expected value $E(n_i)$ and variance $V(n_i)$ as follows:

$$E(n_i) = \rho \times A_i \tag{1}$$

$$V(n_i) = A_i \times \rho \times (1 - \rho) \tag{2}$$

The mean and variance of the number of nodes detecting object i, respectively $E(d_i)$ and $V(d_i)$, are then as follows:

$$E(d_i) = \rho \times A_i \times p_n \tag{3}$$

$$V(d_i) = A_i \times \rho \times p_n \times (1-\rho \times p_n) \tag{4}$$

For this distribution for variously chosen values of A_i, ρ and p_n, we heuristically observe that in 99% of the trials, the value of the random variable lies within three times the standard deviation of the expected value.

In order for separation between object types to be maintained whp, we require the following inequality to hold for each pair $(i, i+1)$.

$$E(d_{i+1})-(3\times\sqrt{V(d_{i+1})}) > E(d_i)+(3\times\sqrt{(V(d_i))}) \tag{5}$$

Using Eqs. 3, 4, and 5, the minimum density $\rho_{i(i+1)}$ required to maintain separation between objects i and $i+1$ whp, is given by the following condition.

$$\rho_{i(i+1)} > \frac{B}{(p_n + B \times p_n)} \tag{6}$$

where

$$B = \frac{9 \times (\sqrt{A_{(i+1)}} + \sqrt{A_i})^2}{(A_{(i+1)} - A_i)^2} \tag{7}$$

Let $\hat{\rho}$ be the maximum of all $\rho_{i(i+1)}$. ρ is then the minimum density required to maintain separation between object types whp in the presence of independent nodal failures with uniform probability p_n.

Dealing with inadequate density by temporal aggregation. If the density of the network is less than $\hat{\rho}$, temporal aggregation can be used to obtain the necessary separation of object types. In this technique, we aggregate the influence fields of object over a fixed interval of time t. The aggregated influence field is the area covered by the object in time t. This area depends on the size, shape and motion model of the object. Using Eq. 5, we can determine the interval over which we need to aggregate to be able to separate object types whp. The influence field is thus used in a spatio-temporal context. In this manner, we can also separate object types that have the same influence field but different motion models.

Dealing with excess density by probabilistic reporting. If the density of the network is greater than $\hat{\rho}$, we can improve system efficiency and lifetime by decreasing the number of messages sent over the network and still accurately estimate the influence field. Here, each node reports a message with uniform probability P_r. The number of reporting nodes in an area A_i is thus a random variable with mean and variance as given below:

$$E(r_i) = \rho \times A_i \times p_n \times P_r \tag{8}$$

$$V(r_i) = A_i \times \rho \times p_n \times P_r \times (1-\rho \times p_n \times P_r) \tag{9}$$

In order for separation between object types to be maintained whp, we require the following inequality to hold for each pair $(i, i+1)$and $(P_r)_{i(i+1)}$ is the probability of reporting.

$$E(r_{(i+1)})-(3\times\sqrt{(V(r_{(i+1)}))}) > E(r_i)+(3\times\sqrt{V(r_i)}) \tag{10}$$

Using Eqs. 8, 9 and 10, the minimum probability $(P_r)_{i(i+1)}$ required to maintain separation

between objects i and $i+1$ whp, is given by the following condition:

$$(P_r)_{i(i+1)} > \frac{B}{((p_n \times \rho) + (B \times p_n \times \rho))} \quad (11)$$

where

$$B = \frac{9 \times (\sqrt{A_{(i+1)}} + \sqrt{A_i})^2}{(A_{(i+1)} - A_i)^2} \quad (12)$$

Let P_r be the maximum of all $(P_r)_{i(i+1)}$. P_r is then the minimum probability required to maintain separation of object types whp in the presence of independent nodal failures.

5.2 Fading error over multiple hops

We now obtain a necessary condition to achieve separation of object types in the presence of fading over multiple hops en route to the aggregator. For simplicity, we assume that the object size is small as compared to the distance from the aggregator, hence all detections corresponding to the same object travel the same number of hops.

Recall that p_f is the probability of a message being received over a single hop in the presence of fading. Thus, over h hops, the probability of successful reception of a message equals $(p_f)^h$. Again, these failures being independent, the number of successful transmissions at any distance has a binomial distribution.

Since the messages originating from nodes closer to the aggregator have lower probability of failure, the influence field of a smaller object close to the aggregator can overlap with that of a larger object farther away. We apply the following techniques to maintain separation whp, between influence fields of two object types in the presence of fading errors.

Dealing with fading errors by distance dependent probabilistic reporting. In order to compensate for non-uniform reception probability, we once again use the probabilistic reporting technique. Let D be the maximum number of hops to the aggregator. The probability of reporting for a sensor at distance h from the aggregator is set to $p_f^{(D-h)}$. Thus, the probability of reporting for nodes D hops from the aggregator is 1, while for $h = 1$, the probability of reporting is $p_f^{(D-1)}$.

The number of messages successfully received for an object i which is h hops away from the aggregator, is a random variable $f(h)$ whose expected value and variance are obtained as follows:

$$E(f(h)_i) = A_i \times \rho \times p_f^{D-h} \times p_f^h = A_i \times \rho \times p_f^D \quad (13)$$

$$V(f(h)_i) = A_i \times \rho \times p_f^D \times (1 - (\rho \times p_f^D)) \quad (14)$$

In order for separation between object types to be maintained whp for each pair $(i,\ i+1)$, we require that the number of messages successfully received whp for the smaller object i located at $h = 1$ be less than the number of messages successfully received whp for the larger object $i+1$ located at $h = D$.

$$E(f(D)_j) - (3\sqrt{V(f(D)_j)}) > E(f(1)_i) + (3\sqrt{V(f(1)_i)}) \quad (15)$$

However, from Eq. 13 and Eq. 14, we observe that using distance based probabilistic reporting, the distribution of the number of successfully received messages is independent of the number of hops.

Thus, from Eqs. 13, 14 and 15, the minimum density $\rho_{i(i+1)}$ required to maintain separation between objects i and $i+1$ whp, is given by the following condition:

$$\rho_{i(i+1)} > \frac{B}{(p_f^D + (B \times p_f^D))} \quad (16)$$

where

$$B = \frac{9 \times (\sqrt{A_{(i+1)}} + \sqrt{A_i})^2}{(A_{(i+1)} - A_i)^2} \quad (17)$$

Let $\hat{\rho}$ be the maximum of all $\rho_{i(i+1)}$. $\hat{\rho}$ is then the minimum density required to maintain separation between object types whp in the presence of fading errors.

Note that we can deal with conditions where we have less than or more then this density ρ by using the techniques discussed in Sec.5.1.

Dealing with fading errors by spatial reconstruction. As an alternative to probabilistic reporting, we present the spatial reconstruction technique wherein the aggregator can compensate for multi-hop losses, given the source information for the messages it receives. Spatial reconstruction is the dual of probabilistic reporting. Since

the probability of receiving a message from a node h hops away is p_f^h, if k messages are received from distance h, the aggregator considers this as having received $\frac{k}{p^h}$ messages. In this case, all detecting nodes transmit with the same probability, which may be lower than 1 for reasons of efficiency as discussed earlier. The necessary conditions required to maintain separation between object types can be derived as before.

5.3 Channel contention over a single hop

Next, we analyze the effect of interference due to channel contention. Message losses due to collision occur in wireless sensor networks due to lack of a collision detection mechanism. In our event based traffic model, all nodes detect an object simultaneously and hence compete for the channel at the same instant. Thus, as the event size increases, the message losses increase too. We analyze the effect of channel contention on the aggregation, under the assumption of the following one hop model.

Suppose n nodes, all within one hop of each other and the aggregator, detect an object and want to convey this detection to a aggregator. Whenever a node detects an event, it randomly chooses one of c channel slots for transmitting the message. Let c be greater than n. If multiple nodes choose the same slot, their messages collide and all of them are assumed to be lost. In a sense, this models a random backoff MAC scheme.

The expected number of messages that will be successfully received by the aggregator equals the expected number of slots that are chosen by exactly one node. This is an instance of a classical occupancy problem in combinatorics. Let X_1, X_2, ... , X_n be random variables, so that X_i is 1 if a slot is chosen by exactly one node and 0 otherwise. The probability that a slot is chosen by exactly one node is equal to the probability that all other nodes choose different slots. This is the probability that a message does not get lost due to channel contention, which we denoted as p_c in Sec. 3.

$$p_c = (1 - 1/c)^{(n-1)} \tag{18}$$

The number of nodes with a slot for themselves, i.e., the number of messages that do not get lost due to channel contention is a random variable having a binomial distribution with parameters (n, p_c). The mean and variance of the distribution, denoted as $E(s)$ and $V(s)$, are as follows:

$$E(s) = n \times p_c = n \times (1 - \frac{1}{c})^{(n-1)} \tag{19}$$

$$V(s) = n \times p_c \times (1 - p_c) \tag{20}$$

From Eq. 19, it is seen that for a given c, as n increases, the expected number of successful messages reaches a maximum and then starts decreasing.

> *Definition:* The *inversion point* of a network with respect to a given observer is the number of senders for which the expected number of messages received is maximum.

The inversion point, denoted as n_{inv}, obtained by solving for the maxima of Eq. 19 is as follows:

$$n_{inv} = \frac{1}{ln(1 + \frac{1}{c})} \tag{21}$$

Due to inversion, the aggregator may receive fewer detection messages for a larger object than it receives for a smaller object, hence the separation between object types may not be preserved.

Experimental results.

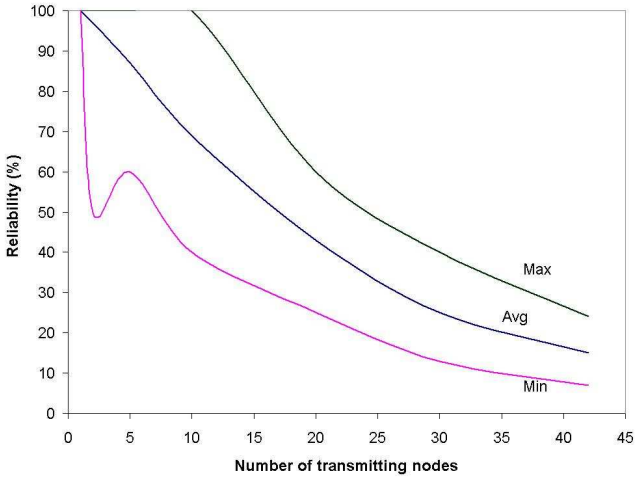

Figure 2. Network reliability under varying traffic load in a one hop network.

Fig. 2 shows the experimentally measured impact of increasing the number of transmitters on the network reliability of the single hop model. This experiment was performed using Mica2 motes running TinyOS [7] version 1.0, using globally synchronized time to generate concurrent messages. The nodes were placed in a circle around the aggregator within one hop of each other and of the aggregator. The transmission power of the nodes was set to be high enough to negate fading losses in the medium. The experimental results in Fig. 2 have been averaged over 50 trials for each of the traffic loads ranging over 1-40. For the traffic loads under consideration, not only does the reliability of the network decrease significantly as the number of nodes increases, but it leads to the inversion effect as can be seen in Fig. 3. Network unreliability also causes an overlap between the number of messages that can be received at the aggregator implying that previously separable influence fields will no longer be separable.

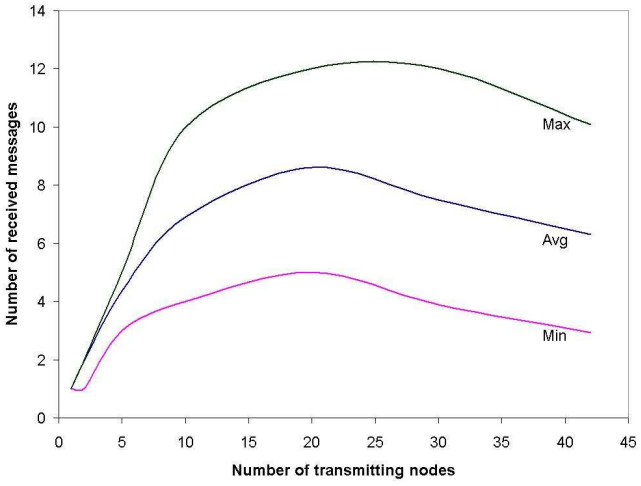

Figure 3. Inversion in a one hop network.

Dealing with inversion by probabilistic reporting. In order to eliminate inversion, we scale down the number of messages sent by reporting with a uniform probability P_r. For an object i, the number of nodes that will detect the object correctly is a random variable with expected value and variance as follows:

$$E(r_i) = A_i \times \rho \times P_r \tag{22}$$

$$V(r_i) = A_i \times \rho \times P_r \times (1 - \rho \times P_r) \tag{23}$$

Recall from Eq. 18 that the probability of a message being successfully received for a object i is dependent on the number of reporting nodes, which itself is a random variable. We make a simplifying assumption that while the number of reporting nodes for a object is a random variable, the probability of success is uniform and depends on the expected number of reporting nodes. This assumption results in a smaller traffic load being subjected to larger contention than it would really experience. Similarly, the larger traffic loads are subjected to lower contention than actual. Consequently, the interval over which the number of received messages is distributed, subsumes the interval that would be obtained in practice. Hence, the necessary conditions for maintaining separation between object types, resulting from our assumption are conservative.

We now have for object i

$$p_{c_i} = (1 - 1/c)^{(E(r_i)-1)} \tag{24}$$

Using Eqs. 22 and 24 the number of messages that are successfully received for this object is now a random variable whose mean and variance are given by:

$$E(s_i) = A_i \times \rho \times P_r \times p_{c_i} \tag{25}$$

$$V(s_i) = A_i \times \rho \times P_r \times p_{c_i} \times (1 - (\rho \times P_r \times p_{c_i})) \tag{26}$$

In order for separation between object types to be maintained whp, we require the following inequality to hold for each pair $(i, i+1)$:

$$E(s_{(i+1)}) - 3 \times \sqrt{V(s_{(i+1)})} > E(s_i) + 3 \times \sqrt{V(s_i)} \tag{27}$$

The following procedure can be used to mechanically select the probability of reporting such that the influence fields for all object types are separable.

select-p_r()
begin

- For each pair $(i, i+1)$, where $1 \leq i$ and $i < k$, using Eq. 25, 26 and 27, obtain a range of probabilities given by the closed interval $(min((P_r)_{ij}), max((P_r)_{ij}))$.
- Let $(P_r min, P_r max)$ denote the intersection of all such ranges.

- If the intersection is not empty, choose $P_r = P_r min$.

end
When all nodes report with uniform probability P_r obtained by the above procedure, separation between object types is maintained whp.

Dealing with empty intersection by temporal segregation. If the procedure described in the previous subsection returns an empty range of feasible reporting probabilities, there exist a, b, c such that $a < b < c$ and $max((P_r)_{ab}) < min((P_r)_{bc})$. Thus, there exist pairs for which eliminating inversion requires a small probability of reporting. In order to overcome this, the inversion point has to be increased. According to Eq. 21, this can be achieved by increasing the number of channel slots c. By doing this, we temporally segregate the messages. Temporal segregation can also be achieved by incorporating an additional application level backoff before reporting a message. Note that one can also eliminate the problem of channel contention by precisely scheduling the transmission of messages. One example of such a scheme is TDMA. The drawback of all these schemes is that they incur an additional delay overhead.

Experimental results.

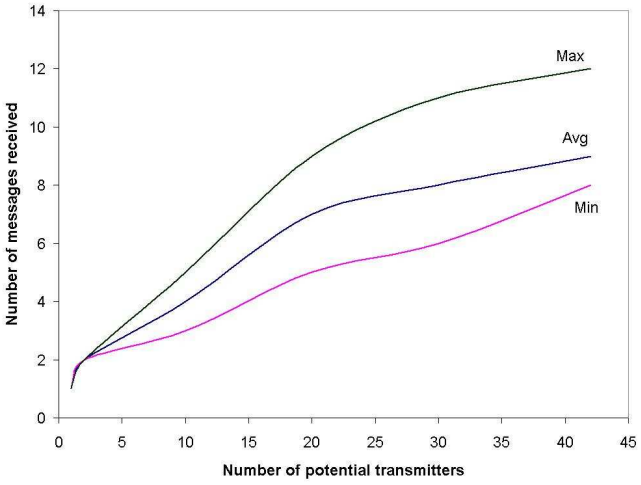

Figure 4. Probabilistic reporting using p=0.5.

Fig. 4 shows the results of using the probabilistic reporting scheme with probability for transmitting a message as 0.5 for the same experimental setup as in the previous subsection. At each concurrent sending event, all the potential transmitters independently decide whether or not to transmit their message. The graph shows that by using probabilistic reporting, inversion is avoided for the traffic loads under consideration. However, note that the intersection in this case is empty hence even with probabilistic transmission, we do not achieve disjoint ranges of message reception.

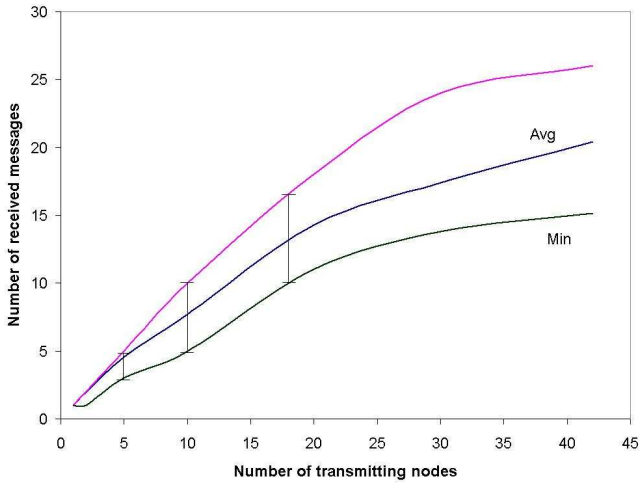

Figure 5. Temporal segregation by increasing number of channel slots.

Fig. 5 shows the results of increasing the number of slots available for backoff on the reliability of the single hop network. The experimental setup was similar to the first case wherein all nodes transmitted concurrently. However, in this experiment, the number of slots available for choosing when to transmit was quadrupled as compared to the previous case. Once again, the results have been averaged over 50 trials for each traffic load. As seen from the graph, by increasing the number of slots, we are able to avoid inversion for the traffic loads under consideration. The graphs also demonstrate that the overlap between messages received was eliminated for some traffic loads and significantly reduced for others. In fact, if the number of slots were increased even further, there would be no overlap.

Remark on tracking. As discussed in Sec. 4, the centroid of the estimated shape of the influence field for a object can be used to track it. In the absence of faults, the centroid approach still works because the nodes are deployed with uniform probability. Nodal faults occur independently at each node with a certain probability, hence the distribution of failed nodes over the influence field can be assumed uniform. Similarly, the one-hop contention model for faults is based on nodes selecting the same slot for transmission randomly, hence again the distribution of failures is uniform.

In the multi-hop case, however, the probability of failure is non-uniform because nodes farther away are subject to a higher loss rate than nodes nearby. However, the techniques of distance dependent probabilistic reporting and spatial reconstruction compensate for the non-uniformity of network failures. E.g., in the distance dependent probabilistic reporting scheme, at each hop farther from the aggregator, the probability of loss increases but the number of nodes reporting from that region is also scaled up accordingly so that messages are received uniformly over the entire influence field. Similarly, in the spatial reconstruction scheme, each message received from a region farther out is weighted accordingly to get a uniform spatial representation of the influence field. Due to the uniformity preserving nature of faults or the techniques described above, the shape of the influence field is preserved and we can still use the centroid approach to reliably track the movement of objects in the network.

6 Case study: A Line In The Sand

In this section, we describe the implementation of a distributed classification and tracking system using a wireless sensor network. The system, called *A Line In The Sand* consists of 90 MICA2 motes deployed in a grid-like manner with 1.5m inter-node separation to cover a 20m x 10m area. *A Line In The Sand* has been deployed in several outdoor settings to accurately distinguish between civilians, soldiers and cars by estimating their influence fields based on magnetometer and micropower impulse radar sensors. For simplicity of presentation, we only describe classification between a soldier and a vehicle using magnetometer based influence fields.

A magnetometer sensor measures the value of the magnetic field around it. As metallic objects approach these sensors, they create a disturbance in this magnetic field which can be detected by the sensors using techniques like variance-based thresholding. The amount of disturbance created by a vehicle containing a large amount of metal is significantly higher than that created by a soldier carrying a limited amount of metal in the form of weapons, etc. As explained in Sec. 4, the average inter-node separation of 1.5m and hence the density, of the network is dictated by the ability to sense and detect the smallest object in the system, i.e., a soldier. Consequently, in this high density network, a vehicle generates a large number of concurrent detections. Network loss due to contention, hence poses the key challenge in *A Line In The Sand*. To add to this, the limited communication range of each node necessitates the use of multi-hop routing to the classifier. Node communications are therefore subject to faults due to both contention and fading over multiple hops. Each of the services or components in our system had to be designed specifically to handle this network unreliability. In this section, we focus only on issues relating to the classification and tracking application.

Influence field measurements.

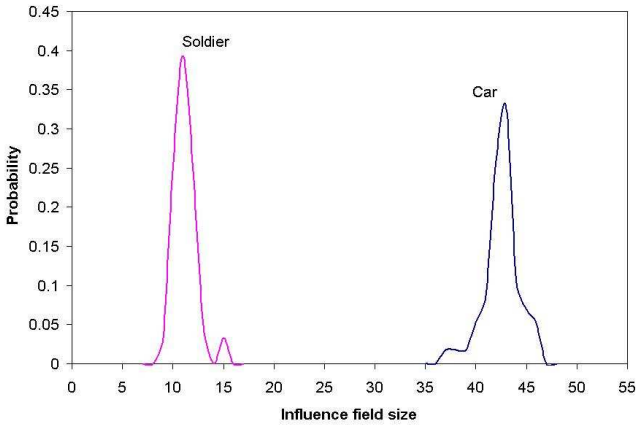

Figure 6. Probability Distribution Functions of the Influence field of a soldier versus a car measured experimentally at the motes.

Fig. 6 shows the probability distribution function for the influence fields when measured at the nodes. In this outdoor experiment, timestamped

detections were recorded in the non-volatile memory at each mote during 100 runs of a soldier and car each, moving through the network. These detections were then downloaded offline and time-correlated to recreate the measured influence fields and their probability distribution. While these nodal detections were subject to nodal and sensing faults, the measured influence fields were derived in the absence of network loss due to contention or fading.

Effect of network faults. As can be seen from the probability distributions in Fig. 6, there exists a clear separation between the measured influence fields for a soldier and a car, suggesting that classification between these object types should be feasible. However, the difference in message traffic generated by these object types affect network reliability differently. In the case of a soldier, the size of the influence field is relatively smaller and so the dominant losses are due to fading over multiple hops en route to the classifier. By tuning certain network parameters like transmission power of nodes relative to the average internode spacing of 1.5m, the effects of fading can be substantially mitigated. For a car though, the influence field is significantly larger, hence contention losses dominate multi-hop fading effects. Indeed, as predicted by the analysis in Sec. 5, we observe that reliability drops sharply for a car when compared to the soldier case.

Dealing with network faults. *Routing reliability.* The implications of network loss for protocol design are significant. From the analysis in Sec. 5 and the experimentally measured influence fields, we can derive the minimum reliability needed in the network to be able to separate the influence fields for a car and a soldier. In this particular case, we require that the number of messages received in the case where the least number of nodes, i.e. 37, detect the car, when hit by the minimum reliability should still be separable from the number of messages received in the case where the maximum number of nodes, i.e. 16, detect a soldier, and are subjected to the best network reliability which of course is bounded by 1. The minimum required reliability for our network is thus 17/37 or 45%.

Figs. 7 and 8 demonstrate the impact of network unreliability on the accuracy of estimating

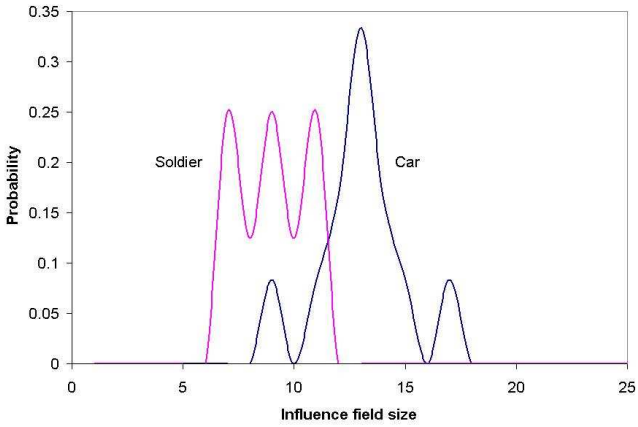

Figure 7. Influence field PDFs for soldier versus vehicle measured experimentally at the classifier using baseline routing protocol.

influence fields and hence on classification and tracking. Fig. 7 shows the probability distribution of the influence fields estimated by the classifier based on the messages received in each of 50 runs of a soldier and a car using a baseline routing protocol. Fig. 8 shows the distribution of the estimated influence fields for the same objects using the GridRouting protocol with appropriately tuned network parameters. As seen from the two graphs, the baseline version, which provides lower reliability, suffers from the inversion problem described in Sec. 5 whereas the GridRouting protocol meets the reliability requirements of the application and preserves separation between the object types.

Uniformity of routing reliability. Recall from Sec. 5 the requirement that for accurate classification and tracking, the reliability should not only be above a certain threshold, but also be uniform. Our experiments with different routing protocols have shown that the number of messages received under varying traffic loads and the distribution of these messages over the source region vary significantly for different routing protocols. Fig. 9 shows a comparison between the experimentally measured performance of two routing algorithms: GridRouting and MintRoute [12].

The experimental setup for this test consists of a 7x7 grid of motes spaced 1.5m apart in a manner similar to LITeS. Fig. 9 shows this layout with the classifier node at the bottom left corner.

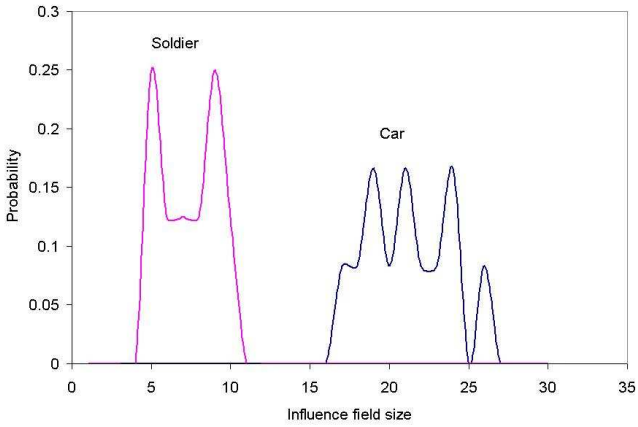

Figure 8. Influence field PDFs for soldier versus car measured experimentally at the classifier using GridRouting protocol.

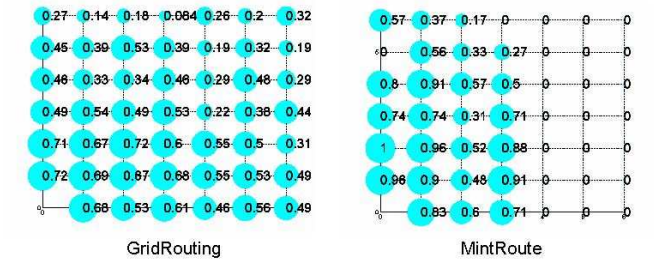

Figure 9. Comparison of uniformity of network reliability between GridRouting and MintRoute.

Two sets of experiments were conducted, one for each of the routing protocols running on TinyOS version 1.0. Each node in the figure has been labelled with the fraction of detection messages generated by that node which were received by the classifier. As can be seen from the figure, for the traffic model in *A Line In The Sand*, wherein a large burst of traffic is generated by detections for a car, GridRouting provides substantially more uniform reliability than MintRoute. Note that even in GridRouting, the overall reliability decreases as we consider nodes farther away from the classifier, as explained by the multi-hop analysis in Sec. 5. However, the nodes along the diagonal, which are equidistant from the classifier have fairly similar reliability in the case of GridRouting. The GridRouting protocol thus meets our requirement of uniformity of failure probabilities for nodes at the same distance from the classifier.

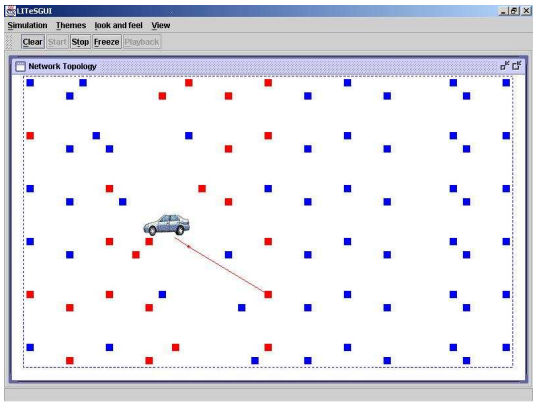

Figure 10. Classification and tracking of a car in *A Line In The Sand*.

System performance. Finally, we give some performance data for *A Line In The Sand*. By considering the influence field analysis and appropriately tuning the network protocols and parameters, we were able to achieve the desired classification accuracy of 99%. The accuracy of tracking is higher for the soldier (1-2m) as compared to a car (3-5m) providing further evidence of the claim that reliability and uniformity are dependent on the object type. The system is able to classify and track multiple objects moving concurrently through the network as long as they are separated by a minimum distance threshold. Fig. 10 shows a snapshot of the classification and tracking output produced by the system for a car moving through the network.

7 Conclusions and Future Work

In this paper, we considered the problem of reliably estimating the influence fields of different target types in a wireless sensor network subject to a variety of faults. We provided mechanical procedures for sensor node density selection as well as algorithmic techniques appropriate for dealing with each fault class. Corroboration of our results and techniques was provided through at-scale experiments.

We showed how reliable estimation was achieved to enable accurate classification and

tracking in *A Line In The Sand*. The case study also provided a data point for the significant impact of network unreliability on network and application design, as well as one for a need for routing protocols in sensor networks to provide uniform reliability.

Our work reveals a notable co-dependence between application design and network design. To achieve the desired estimation reliability, we needed in some cases to use both techniques that affected the network (such as tuning of MAC or routing protocol parameters) and that affected the application (such as tuning the probability of reporting and the rate of temporal aggregation). How to design stable and scalable systems when there are such cyclic dependencies involved is an issue of interest to us.

Although our compositional models allow us to reason about the effects of different types of node and network faults in isolation, there are some relevant and more complex fault models that we have not dealt with analytically. One such model, which we dealt with only experimentally in *A Line In The Sand* concerns multi-hop contention and fading errors. In future work, we seek to address this model analytically. We will also incorporate in our analysis consideration of fault positives and multiple concurrent targets that we dealt with experimentally, towards addressing the gap between existing theory and practice.

References

[1] C. Bailey-Kellogg and F. Zhao. Influence-based model decomposition for reasoning about spatially distributed physical systems. *Artificial Intelligence*, 130:125–166, 2001.

[2] C. S. Raghavendra C. Meesookho, S. Narayanan. Collaborative classification applications in sensor networks. *Second IEEE Sensor Array and Multichannel Signal Processing Workshop*, 2002.

[3] Michael J. Caruso and Lucky S. Withanawasam. Vehicle detection and compass applications using AMR magnetic sensors, AMR sensor documentation. `http://www.magneticsensors.com/datasheets/amr.pdf`.

[4] Ashwin D'Costa and Akbar Sayeed. Collaborative signal processing for distributed classification in sensor networks. *The 2nd International Workshop on Information Processing in Sensor Networks (IPSN '03)*, pages 193–208, 2003.

[5] Marco Duarte and Yu-Hen Hu. Vehicle classification in distributed sensor networks. *Journal of Parallel and Distributed Computing*, 2004, to appear.

[6] R. Hall. *Introduction to the Theory of Coverage Processes*. Wiley, 1988.

[7] Jason Hill, Robert Szewczyk, Alec Woo, Seth Hollar, David Culler, and Kristofer Pister. System architecture directions for network sensors. 2000.

[8] Bhaskar Krishnamachari and Sitharama Iyengar. Distributed bayesian algorithms for fault-tolerant event region detection in wireless sensor networks. *IEEE Transactions on Computers*, 53(3):241–250, 2004.

[9] Dan Li, Kerry Wong, Yu Hu, and Akbar Sayeed. Detection, classification and tracking of targets in distributed sensor networks. *IEEE Signal Processing Magazine*, 19(2):17–29, 2002.

[10] J. Liu, J. Reich, and F. Zhao. Collaborative in-network processing for target tracking. *Journal on Applied Signal Processing*, 2002.

[11] Donal McErlean and Shrikanth Narayanan. Distributed detection and tracking in sensor networks. *36th Asilomar Conf. Signals, Systems and Computers*, 2002.

[12] Alec Woo, Terence Tong, and David Culler. Taming the underlying issues for reliable multhop routing in sensor networks. *Proceedings of ACM Sensys 2003, Los Angeles, California*, 2003.

[13] Honghai Zhang and Jennifer C. Hou. On deriving the upper bound of alpha-lifetime for large sensor networks. *Proceedings of ACM Mobihoc 2004*, 2004.

[14] Feng Zhao, Jie Liu, Juan Liu, Leonidas Guibas, and James Reich. Collaborative signal and information processing: An information directed approach. *To appear in the proceedings of the IEEE*, 2003.

[15] Feng Zhao, Jaewon Shin, and James Reich. Information-driven dynamic sensor collaboration for tracking applications. *IEEE Signal Processing Magazine*, 2002.