Reliable Estimation of Influence Fields for Classification and Tracking in Unreliable Sensor Networks

Sandip Bapat, Vinodkrishnan Kulathumani, Anish Arora
Department of Computer Science and Engineering
The Ohio State University
Columbus, OH 43210, USA
{bapat,vinodkri,anish}@cse.ohio-state.edu

Abstract

The influence field of an object, a commonly exploited feature in science and engineering applications, is the region where the object is detectable by a given sensing modality. Being spatially distributed, this feature allows us to tradeoff nodal computation with network communication. By the same token, not only is its calculation subject to nodal failures and false detections, 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 and how this reliable estimation of influence fields can be used to classify and track different types of objects. We derive, for node and network fault models, the necessary nodal density for reliably estimating the influence fields so that objects can be classified and tracked. We present four algorithmic techniques: Temporal Aggregation, Probabilistic Reporting, Temporal Segregation and Spatial Reconstruction, to deal with cases where the effective network density differs from this minimum. We provide corroboration of our analysis through field experiments with Mica2 sensor nodes wherever appropriate. Finally, we demonstrate how these results and techniques were applied to achieve reliable and efficient classification and tracking in a fielded system of 90 Mica2 sensor nodes that we called "A Line In The Sand".

Keywords: *wireless sensor networks, reliability, fault-tolerance, intrusion detection*

1 Introduction

The *influence field* of an object j with respect to a given sensing modality is the region surrounding j where it can be "detected" by a sensor of that modality. This region depends on both the characteristics of the object, such as its size and shape, as well as the sensing modality being used. Differences in the area and/or shape of the influence fields of different objects can often be used to distinguish between them. Fig. 1 illustrates the differences between magnetic influence fields for two objects, a person carrying a metal rod and a vehicle. The size and shape of the influence fields shown in this figure depend on the amount and distribution of metallic content in each object type and the orientation of the object (e.g., the dumbell shape of the vehicle influence field is attributed to the positions of its axles). The influence field feature is thus useful in sensor network applications for surveillance, where typical tasks include detection, classification, and tracking of various types of objects.

To estimate the influence field, each node merely has to detect a binary "presence" of an object; network-based aggregation of these bits yields the influence field without substantial or complex node operation. The influence field feature is thus well suited for wireless sensor network applications where individual nodes are constrained due to limited processing, sensing and communication capabilities. The key challenge in realizing the influence field is the unreliability of wireless sensor networks. Event losses –both in nodes and in the network– are fundamental to wireless sensor node platforms and their impact on the application can be substantial. Thus, both node and network unreliability have to be dealt with while estimating the influence field.

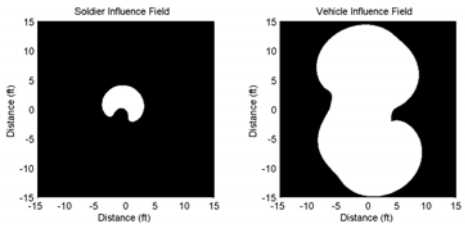

Figure 1. Magnetometer based influence fields for two object types.

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 significant amount of metal (aka "soldiers") and vehicles, via a dense, wireless sensor network. The intended use of this system, which we called *A Line in the Sand* [1], is to use a sensor network to protect high-valued assets, secure extended perimeters, and monitor activity of personnel or vehicles in remote or access-denied areas. *A Line in the Sand* was originally demonstrated in a 18m x 7m area using 90 Mica2 [2] sensor nodes. More recently, the same approach has also been demonstrated to achieve accurate classification and tracking of intruders in a large scale network, which we called *ExScal* [3], of over 1000 nodes deployed over a 1.3km x 300m area.

Contributions of the paper. In this paper, we address the problem of reliably estimating the influence fields of various object types in the presence of multiple sensor network faults, in terms of preserving the differences in the area and/or shape across various object types. For each fault type, we provide: (1) analytical results about the fault-affected influence field estimates, (2) procedures for calculating the ideal sensor node density for efficient and reliable estimation, (3) algorithmic 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 experiments with 40–50

Proceedings of the 2005 24th IEEE Symposium on Reliable Distributed Systems (SRDS'05)
0-7695-2463-X/05 $20.00 © 2005 IEEE

Mica2 sensor nodes in our testbed.

We show that node and network faults impose competing (i.e., two-sided) constraints on network density. However, since our results can be instantiated in a compositional way, we can deal with cases where multiple types of faults occur.

Finally, we show how we experimentally validated the reliable estimation of influence fields so as to accurately classify and track persons, soldiers, and vehicles in *A Line in The Sand*. This case study also provides a data point for the importance of dealing with network unreliability at both the network and the application level. It identifies a need for routing protocols in sensor networks to provide uniform reliability (at least for nodes that are equidistant from the receiver). We also describe possible extensions to estimation of influence fields for classification and tracking of objects.

Organization of the paper. In Sec. 2, we present the system and fault models and formally define the problem of reliable estimation of influence fields. Then, in Sec. 3, 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 algorithmic techniques to deal with these faults whenever necessary. Sec. 4 describes how these results and techniques can be composed and applied to identify and track objects. In Sec. 5, we provide details of our implementation of a *A Line in the Sand* and how we dealt with different fault classes in our design. Some extensions to the influence field approach are noted in Sec. 6. We discuss related work in Sec. 7 and make concluding remarks and discuss future work in Sec. 8.

2 System and Fault Models

In this section, we outline the system and fault models and define the problem of reliably estimating influence fields.

2.1 System model

The system consists of N wireless sensor nodes, each with a unique identifier. We assume a localization service that provides the relative or absolute position for each node and a global time synchronization service that enables each node to timestamp its detections.

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. When we refer to the *area* A of a subregion, we mean A is the number of these Ω points in the subregion. We denote the ratio N/Ω by ρ, which represents the sensor density in the network.

We assume that the wireless network is connected, hence it is possible to aggregate messages from any subset of nodes in the system, possibly over multiple hops. We assume a central aggregator node known to the rest of the nodes for simplicity, however our results are applicable even for dynamically chosen aggregators.

Recall that the *influence field* of an object j with respect to a sensing modality is the region surrounding j wherein j will be "detected" by a sensor of that type. 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. We limit our attention to one given sensing modality which will remain implicit in our notation.

2.2 Fault model

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 and software failures, or simply in the form of a transient event loss. They also occasionally generate false positives due to unreliable hardware, environmental perturbations and transient state corruption. Thus at any time, the net effect of a node fault can be modelled as missing the detection of an object, i.e. a false negative, or asserting a detection when there is no object, i.e. a false positive.

> Node fault model: The probability that any node misses the detection of an object, whether it is due to a transient, permanent or intermittent node fault, is $1 - p_n$, while the probability that any node generates a false positive is p_{fp}.

In other words, we assume that false negatives and false positives at nodes are independent of each other and of the objects. It follows trivially that for any meaningful detection to be possible, the probability of a node detecting an event, p_n, must be greater than the probability of a node generating a false positive, p_{fp}.

Network faults. Wireless communications in a multi-hop network are subject to both fading and contention effects. Channel fading loss depends on link characteristics such as distance, relative orientation of sender and receiver, environmental conditions, etc.

> Fading model: The probability of message loss due to fading for any single hop communication in the network is $1 - p_f$.

It follows that in the absence of any other faults, the probability of message loss due to fading for any h hop communication is $1 - p_f^h$.

Channel contention losses occur when multiple senders try to transmit messages at the same time. The degree of message loss depends on factors such as the Medium Access Control (MAC) protocol used, the synchronicity of message transmissions, and 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: The probability of message contention loss, which is a function of the number of nodes simultaneously trying to send a message and of the number of slots available for a node to choose its random backoff from, is $1 - p_c$.

The end-to-end network reliability of message reception for a given source depends on several factors including the size of the traffic load, the number of hops traversed by each message and the routing protocol. We use the term traffic source to indicate a set of nodes that detect a phenomenon and report it to the aggregator concurrently. We assume a convergecast routing model rooted at the aggregator.

> End-to-end reliability model: The probability of end-to-end message loss for a traffic source i, which is a function of the number of senders in i, and the distance or number of hops h between the traffic source and the aggregator, is $1 - p_{rcv_{ih}}$.

Our model postulates that network reliability is uniform across nodes equidistant from the aggregator. It is possible to realize this end-to-end model by careful design of the routing protocol. The GridRouting protocol [4], used in *A Line In The Sand* achieves this by uniformly balancing the traffic load over paths including stable, reliable links. An experimental validation of the end-to-end reliability model can found in a related technical report [5].

All the fault classes discussed above and their associated probabilities, are the result of several factors ranging from hardware failures to protocol design. In this paper, we analytically characterize these fault models wherever possible (e.g. an analytical characterization of the contention model is presented in Sec. 3). However, not all fault models can be characterized in this manner as they may depend on factors which are hard to model (e.g. environmental perturbations or choice of protocol). In such cases, we resort to empirical characterizations of these faults.

2.3 Problem Definition

Estimation consists of calculating the area and the shape of the influence field. If we assume that sensor node density ρ exceeds some lower bound that depends on the objects at hand, estimating the area A of the influence field of object 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 more than 99% of the trials. We use the 3 standard deviation rule in the rest of this paper as it satisfies the desired classification accuracy of 99%, and denote it as whp (for 'with high probability').

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

$$[(A\times\rho)-3\times\sqrt{A\times\rho\times(1-\rho)}\ ,\ (A\times\rho)+3\times\sqrt{A\times\rho\times(1-\rho)}]$$

Similarly, 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.

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

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 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 sensors that detect it. Shape distortion errors can thus result in inaccuracy of tracking. (The same argument applies for classification based on the shape of the influence field.) Node and network faults impact estimation of the area and shape of the influence field of an object. More importantly, these faults may produce a non-uniform distribution of sensors whose detections are aggregated, resulting in area and/or shape distortion errors.

We are thus led to the problem of how to reliably estimate the size and shape of the influence fields of objects so that they can be distinguished from each other and localized. More specifically, we focus on two subproblems of reliable estimation:

- **Problem 1:** How to ensure for objects whose respective influence field areas are separable, that the fault-affected estimates of their respective influence field areas are also separable?
- **Problem 2:** How to preserve the influence field shape of an object in its fault-affected estimate by preserving the uniformity of distribution of the sensors whose detections are not affected?

3 Reliable Estimation

In this section, we address problems 1 and 2 outlined above successively for each of the fault models discussed in Sec. 2. It should be noted that 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.

3.1 Node faults

Recall from Sec. 2 that the net effect of node faults is modelled as false negatives and false positives.

3.1.1 False Negatives

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 detects an object. 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, recall from Proposition 1 that n_i itself is a random variable that has a binomial distribution, characterized by:

$$E(n_i) = \rho \times A_i \quad (1)$$

$$V(n_i) = A_i \times \rho \times (1-\rho) \quad (2)$$

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

$$E(d_i) = \rho \times A_i \times p_n \quad (3)$$

$$V(d_i) = A_i \times \rho \times p_n \times (1-\rho \times p_n) \quad (4)$$

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 standard deviations of the mean.

In order for separation between estimated influence fields of object types to be maintained whp, we require the following inequality to hold for each pair $(i, i{+}1)$ of objects.

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

Let $\rho_{fn_{i(i+1)}}$ be the minimum density required to distinguish between objects i and $i+1$ whp. By solving Eqns. 3, 4, and 5, we obtain $\rho_{fn_{i(i+1)}}$ as:

$$\rho_{fn_{i(i+1)}} = \frac{B}{(p_n+B\times p_n)}$$

where

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

Theorem 1*: The minimum network density, $\hat{\rho_{fn}}$, required to maintain separation between the estimated influence fields of all objects in the presence of false negatives is the maximum of all pairwise densities $\rho_{fn_{i(i+1)}}$.*

3.1.2 False positives

We now study the impact of false positives on estimation. In order for separation to be maintained between the estimated influence fields of objects i and $i+1$, we require that the number of detections in area A_i, when combined with the false positives in the region of area $A_{i+1} - A_i$, should be lower than the number of detections in area A_{i+1}. The mean and variance of the number of false positives in the area $A_{i+1} - A_i$, respectively $E(fp_i)$ and $V(fp_i)$, are:

$$E(fp_i) = \rho_{i(i+1)} \times (A_{i+1} - A_i) \times p_{fp} \quad (6)$$

$$V(fp_i) = E(fp_i) \times (1-\rho_{i(i+1)} \times p_{fp}) \quad (7)$$

We thus require the following inequality to hold in order to maintain separation between the estimated influence fields of two objects A_i and A_{i+1}.

$$E(n_i)+(3\times\sqrt{(V(n_i))})+E(fp_i)+(3\times\sqrt{V(fp_i)}) < E(n_{i+1})-(3\times\sqrt{V(n_{i+1})}) \quad (8)$$

Note that the term on the right hand side does not contain any expression for false positives. This is because false positives outside the area A_{i+1} would affect the estimation of both objects equally. Let $\rho_{fp_{i(i+1)}}$ be the minimum density necessary to distinguish between objects i and $i+1$, obtained by solving Eqns. 6, 7 and 8.

Theorem 2*: The minimum network density, $\hat{\rho_{fp}}$, required to maintain separation between the estimated influence fields of all objects in the presence of false positives is the maximum of all pairwise densities $\rho_{fp_{i(i+1)}}$.*

3.1.3 Density compensation techniques

The conditions derived above state the minimum density required for preserving separation between the estimated influence fields of objects. However, in many cases, network density is a function of other factors such as cost and communication range, and thus may not be a parameter of choice. We therefore present techniques for dealing with both inadequate and excess network density.

Temporal Aggregation. This technique is used to obtain the necessary separation of estimated influence fields when the network density does not meet the minimum requirements. *Temporal Aggregation* involves the following steps:

1. Using Theorems 1 and 2, compute minimum density $\hat{\rho}$ required to distinguish all objects
2. Given network density ρ, choose the *aggregation interval* t such that $\rho \times t > \hat{\rho}$
3. Aggregate node detections over time t to estimate object influence fields.

The aggregated influence field, used in a spatio-temporal context, is the area covered by an object in time t and it depends on the size, shape and motion model of the object. The proof of why temporal aggregation helps achieve separation can be deduced by rewriting Eq. 5 as

$$E(d_{i+1})-E(d_i) > (3\times\sqrt{V(d_{i+1})})+(3\times\sqrt{(V(d_i))}) \quad (9)$$

We can see that when aggregated over time, expected values grow faster than standard deviations, hence the desired inequality can be satisfied. Thus, temporal aggregation can also be used to distinguish between objects that have the same influence field but different speeds or motion models.

Theorem 3*: Separation between estimated influence fields of objects is preserved in networks with insufficient density or for objects with different speeds, by aggregating detections over interval t.*

Probabilistic Reporting. This technique is used to improve system efficiency and lifetime in cases where network density exceeds the minimum specified by Theorems 1 and 2. *Probabilistic Reporting* involves the following steps:
For each node:

1. Compute probability p_r of reporting a detection.
2. For each detected object, send message to aggregator with probability p_r

Computing p_r: Consider the analysis presented earlier wherein each node detects an object with probability p_n. If each detecting node reports with probability p_r, the number of reporting nodes in an area A_i is a random variable with mean and variance as given below:

$$E(r_i) = \rho \times A_i \times p_n \times p_r \quad (10)$$

$$V(r_i) = A_i \times \rho \times p_n \times p_r \times (1-\rho \times p_n \times p_r) \quad (11)$$

In order for separation between estimated influence fields to be maintained whp, the following inequality must hold for each pair (i, i+1) where $(p_r)_{i(i+1)}$ is the probability of reporting.

$$E(r_{(i+1)})-(3\times\sqrt{(V(r_{(i+1)}))}) > E(r_i)+(3\times\sqrt{V(r_i)}) \quad (12)$$

Let $(p_r)_{i(i+1)}$ be the minimum probability required to distinguish between objects i and $i+1$. By solving Eqns. 10, 11 and 12, we get

$$(p_r)_{i(i+1)} = \frac{B}{((p_n\times\rho)+(B\times p_n\times\rho))}$$

where

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

Theorem 4*: The minimum probability, p_r, of reporting a detection required to distinguish between all object types is the maximum of all pairwise reporting probabilities $(p_r)_{i(i+1)}$.*

3.2 Network faults

In this subsection, we discuss the impact of fading and contention faults on reliable estimation.

3.2.1 Channel fading

Recall that p_f is the probability of message reception over a single hop in the presence of fading. Thus, over h hops, the probability of successful reception of a message equals $(p_f)^h$. Since messages from nodes closer to the aggregator have lower probability of failure, the estimated influence field of a small object close to the aggregator can overlap with that of a large object far away. We therefore need to compensate for the effect of distance of an object from the aggregator during its estimation.

3.2.2 Distance compensation techniques

We now present necessary conditions to maintain separation between the estimated influence fields 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.

Probabilistic Reporting. To compensate for non-uniform reception probability, we use the probabilistic reporting technique presented earlier. In this case however, the reporting probability p_r is not uniform for all nodes, rather it depends on the distance to the aggregator.

Selecting p_r: Let D be the maximum number of hops to the aggregator in the network . The probability of reporting for a sensor at distance of h hops from the aggregator is chosen to be $p_f^{(D-h)}$. Thus, the probability of reporting for nodes D hops from the aggregator is 1, while for $h = 1$, the probability of reporting is $p_f^{(D-1)}$.

We first show that distance-dependent probabilistic reporting compensates for the effect of distance on estimation. Since fading errors are independent, the number of successful transmissions at any distance has a binomial distribution. The number of messages successfully received for an object i which is h hops away from the aggregator, is thus a random variable $f(h)$ whose mean and variance are obtained as follows:

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

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

From Eq. 13 and Eq. 14, we observe that the distribution of the number of successfully received messages is now independent of the number of hops.

We now derive a necessary condition for distinguishing objects whp. To achieve this, it can be shown that for each pair of objects (i, $i+1$), the number of messages successfully received whp for the smaller object i located at $h = 1$ should be less than the number of messages successfully received whp for the larger object $i+1$ located at $h = D$, as this represents the worst case. We thus have:

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

Let $\rho_{i(i+1)}$ be the minimum density required to maintain separation between estimated influence fields of objects i and $i+1$. Solving Eqns. 13, 14 and 15, we get:

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

where

$$\text{B} = 9 \times \left(\frac{\sqrt{A_{(i+1)}} + \sqrt{A_i})^2}{(A_{(i+1)} - A_i)^2}\right.$$

Theorem 5*: Let each node h hops from the aggregator report its detections with probability $p_f^{(D-h)}$ where D is the maximum number of hops to the aggregator. The minimum density $\hat{\rho}$ required to maintain separation between the estimated influence fields for all object types in the presence of multi-hop fading faults, is the maximum of all pairwise densities $\rho_{i(i+1)}$.*

Note that since our techniques are compositional, we can deal with conditions where we have less than or more than this density $\hat{\rho}$ by using the techniques discussed earlier.

Spatial Reconstruction. As an alternative to probabilistic reporting, we present the spatial reconstruction technique which involves the following steps:

1. Upon detecting an object, each node sends a message to the aggregator.
2. For each distance h, aggregator scales number of received messages from that distance by $\frac{1}{p^h}$.

Thus, if the aggregator receives k messages from distance h, it considers this as having received $\frac{k}{p^h}$ messages. The scaling factor is a result of the uniform probability of receiving a message from a node h hops away being p_f^h. Spatial Reconstruction is the dual of distance-dependent Probabilistic Reporting. In this case, all detecting nodes transmit with the same probability, which may be lower than 1 for reasons of efficiency. The minimum network density required to distinguish between object types is the same as in Theorem 5 because the probability of fading loss is independent of the number of messages being transmitted.

3.2.3 Channel contention

In this subsection, we analyze the effect of interference due to channel contention on a single hop. In our event based traffic model, all nodes detect an object and hence compete for the channel at nearly the same instant. Thus, as the event size increases, the message losses increase too. We analyze the effect of channel contention on aggregation, under the assumption of the following one hop model.

Suppose n nodes, all within one hop of each other and the aggregator, want to send a detection message to the aggregator. Each node randomly chooses one of c time 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 lost. This models a random backoff MAC scheme, commonly used in wireless communications.

The expected number of messages successfully received by the aggregator is the expected number of time slots that are chosen by exactly one node. This is an instance of a classical occupancy problem in combinatorics. 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. 2.

$$p_c = (1 - 1/c)^{(n-1)} \quad (16)$$

The number of nodes with a time 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:

$$E(s) = n \times p_c = n \times (1 - \frac{1}{c})^{(n-1)} \quad (17)$$

$$V(s) = n \times p_c \times (1 - p_c) \quad (18)$$

From Eq. 17, 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. 17 is as follows:

$$n_{inv} = \frac{1}{ln(1 + \frac{1}{c})} \quad (19)$$

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 the estimated influence fields of the object types may not be preserved.

Experimental results.

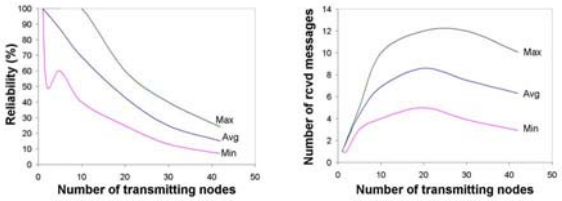

(a) (b)

Figure 2. Inversion 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 [6], using globally synchronized time to generate concurrent messages. The nodes were placed within one hop of each other and of the aggregator and their transmit power was set to be high enough to negate fading losses. The experimental results in Fig. 2 have been averaged over 50 trials for each of the traffic loads of size 2,5,10,20,30 and 40 sources. For these traffic loads, not only does the reliability of the network decrease significantly as the number of nodes increases as seen in Fig. 2(a), but it leads to the inversion effect seen in Fig. 2(b). This inversion may cause an overlap between the number of messages received at the aggregator implying that previously separable influence fields will no longer be separable.

3.2.4 Contention compensation techniques

In this subsection, we describe two techniques to compensate for inversion effects produced by network contention to preserve separation between estimated influence fields.

Probabilistic Reporting. We use the same technique described earlier to compensate for contention effects, however the constraints for choosing the reporting probability p_r differ as follows.

If each detecting node reports with probability p_r, the number of reporting nodes is a random variable with expected value and variance as follows:

$$E(r_i) = A_i \times \rho \times p_r \quad (20)$$

$$V(r_i) = A_i \times \rho \times p_r \times (1 - \rho \times p_r) \quad (21)$$

Recall from Eq. 16 that the probability of a message being successfully received for an object i is dependent on the number of reporting nodes, which itself is a random variable. We make a simplifying, yet conservative, assumption that while the number of reporting nodes for an object is a random variable, the probability of successful reception is uniform and depends on the expected number of reporting nodes. This assumption results in a smaller traffic load than the expected value being subjected to larger contention than it would really experience. Similarly, 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

$$p_{c_i} = (1 - 1/c)^{(E(r_i)-1)} \quad (22)$$

Using Eqns. 20 and 22, the number of messages that are successfully received for this object is now a random variable whose mean and variance are given by:

$$E(s_i) = A_i \times \rho \times p_r \times p_{c_i} \quad (23)$$

$$V(s_i) = A_i \times \rho \times p_r \times p_{c_i} \times (1 - (\rho \times p_r \times p_{c_i})) \quad (24)$$

For separation between estimated influence fields to be maintained whp, we require the following inequality to hold for each pair $(i, i+1)$:

$$E(s_{(i+1)}) - 3 \times \sqrt{V(s_{(i+1)})} > E(s_i) + 3 \times \sqrt{V(s_i)} \quad (25)$$

Selecting p_r: Solving the above inequality yields a quadratic whose solutions denote the minimum and maximum probabilities of reporting for which the two object types can be distinguished. The following procedure can be used to select the probability of reporting such that the estimated influence fields for all object types are separable.

1. For each pair $(i, i+1)$, where $1 \leq i$ and $i < k$, using Eq. 23, 24 and 25, obtain a range of probabilities given by the closed interval $(min((p_r)_{ij}), max((p_r)_{ij}))$.
2. Let $(p_r min, p_r max)$ denote the intersection of all such ranges.
3. If the intersection is not empty, choose $p_r = p_r min$.

Theorem 6*: Assume p_r is the probability determined by the selection procedure. Separation between estimated influence fields of all objects is achieved when each node reports its detections with probability p_r.*

Note that there may be cases where the selection procedure returns an empty intersection in Step 2. We now describe an additional algorithmic technique to deal with such cases.

Temporal Segregation. If the procedure described above returns an empty range of feasible reporting probabilities, it means that there exist objects a, b, c in order of increasing influence field sizes such that $max((p_r)_{ab}) < min((p_r)_{bc})$. Thus, there exist pairs for which eliminating inversion requires such a small probability of reporting that other objects are no longer distinguishable. To overcome this problem, the *inversion point* n_{inv}, has to be increased. According to Eq. 19, this can be achieved by increasing the number of time slots c. In other words, we temporally segregate the messages. Temporal segregation is also achieved by using additional application level backoffs before reporting a detection. 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 such schemes is that they incur an additional delay overhead.

Experimental results. Recall the inversion effect demonstrated experimentally in Fig. 2. We now demonstrate how the techniques described above compensate for the inversion effect, allowing us to distinguish between object types under consideration. The experimental setup is the same as in the previous experiments with the same internode distances, same transmit power and the same traffic source sizes. Fig. 3(a) demonstrates the results of using Proba-

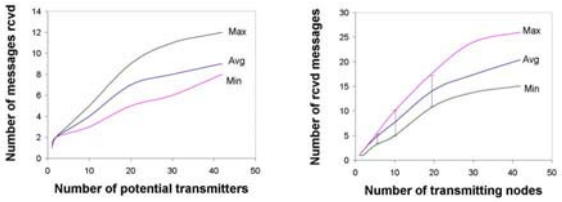

(a) Probabilistic reporting (b) Temporal segregation

Figure 3. Dealing with inversion using Contention compensation techniques.

bilistic Reporting with probability 0.5. 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 for sources of size 5, 10 and 20. Thus, even with probabilistic transmission, we do not achieve disjoint ranges of message reception.

Fig. 3(b) demonstrates how Temporal Segregation helps achieve this separation. In this experiment, the number of slots available for choosing when to transmit was increased 4 times as compared to the previous case. 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 is eliminated for sources of size 5, 10 and 20. In fact, if the number of slots were increased even further, there would be no overlap for any source size.

4 Putting it all together: Classification and Tracking

In this section, we show how the analysis and techniques for isolated fault classes presented earlier are composed, and discuss the effect of end-to-end unreliability on reliable estimation. We also present a procedure for distinguishing objects from false positives and discuss how object locations can be tracked by shape estimation of influence fields.

4.1 Composing fault classes

In this subsection, we describe how the necessary conditions for maintaining separation between estimated influence fields of objects can be derived when multiple faults can occur simultaneously. We illustrate this compositional approach through an example of node faults; other fault classes can be dealt with similarly.

Recall from Sec. 3.1, Eqns. 5 and 8, which specify conditions for maintaining separation between the estimated influence fields of two objects i and $i+1$ in the presence of false negatives and false positives respectively. However, if both these faults can occur, neither of these conditions is sufficient. The necessary condition for distinguishing between two objects i and $i+1$ in the presence of both type of node faults can thus be stated as:

$$E(d_i)+(3\times\sqrt{(V(d_i)})+E(fp_i)+(3\times\sqrt{V(fp_i)}) < E(d_{i+1})-(3\times\sqrt{V(d_{i+1})}) \quad (26)$$

Let $\rho_{node_{i(i+1)}}$ be the minimum density required to distinguish between i and $i+1$, obtained using Eq. 26.

Theorem 7*: The minimum density ρ_{node} required to distinguish between all objects in the presence of both false negatives and false positives is the maximum of all pairwise densities $\rho_{node_{i(i+1)}}$.*

End-to-end reliability. The method described above can be used to analyze the effects of multiple fault classes affecting reliable estimation. We now present a unified analysis for studying the impact of end-to-end reliability. This net reliability encapsulates all losses that may be encountered including false negatives, fading and contention losses. We do distinguish false positives in this analysis as they are additive faults. As discussed in Sec. 2 and as we will demonstrate in the next section, we characterize end-to-end reliability empirically. We denote the uniform probability of receiving a detection from a node in the influence field as $p_{rcv_{ih}}$. Using distance-dependent Probabilistic Reporting or Spatial Reconstruction, we can compensate for the effect of distance on reliability. The mean and variance of the number of detections received by the aggregator for object i, respectively $E(rcv_i)$ and $V(rcv_i)$, then are as follows:

$$E(rcv_i) = \rho \times A_i \times p_{rcv_i} \quad (27)$$

$$V(rcv_i) = A_i \times \rho \times p_{rcv_i} \times (1-\rho \times p_{rcv_i}) \quad (28)$$

Similarly, the mean and variance of the number of false positives in the area $A_{i+1} - A_i$ which are received by the aggregator, respectively $E(fp_i)$ and $V(fp_i)$, are:

$$E(fp_i) = \rho_{i(i+1)} \times (A_{i+1} - A_i) \times p_{fp} \times p_{rcv_i} \quad (29)$$

$$V(fp_i) = E(fp_i) \times (1-\rho_{i(i+1)} \times p_{fp} \times p_{rcv_i}) \quad (30)$$

From Eqns. 27, 28, 29 and 30, we obtain:

$$E(rcv_i)+3\times\sqrt{V(rcv_i)}+Efp_i+3\times\sqrt{V(fp_i)} < E(rcv_{i+1})-3\times\sqrt{V(rcv_{i+1})} \quad (31)$$

Let $\rho_{net_{i(i+1)}}$ be the minimum density required to maintain separation between the estimated influence fields of objects i and $i+1$, obtained using Eq. 31.

Theorem 8*: The minimum network density ρ_{net} required to maintain separation between all object types whp in the presence of false positives and end-to-end unreliability in the network is the maximum of all pairwise densities $\rho_{net_{i(i+1)}}$.*

4.2 Identifying and isolating objects

The analysis presented earlier describes how to choose network density and/or fault compensation techniques so that the estimated influence fields of all objects can be distinguished. Assuming these requirements are satisfied, we now describe a procedure to filter false positives from a set of detections received for one or more objects so that these objects can then be classified and tracked.

Let ρ_{net} be the network density which, as described in the previous subsection, suffices to distinguish between all object types in the presence of faults. We now define the *active density* ad_i for object i as the ratio of the minimum number of detection messages that may be received for object i to the influence field area of object i.

$$ad_i = \frac{(E(rcv_i) - 3 \times \sqrt{V(rcv_i)})}{A_i} \quad (32)$$

Let ad_{min} be the minimum active density among all objects and let i_{min} be the object with this minimum active

density. We then apply the following spatial filtering algorithm to isolate object detections from false positives:

1. For each received detection m, apply filtering windows of size A_1 around it.
2. If a window contains more than A_1 x ad_{min} detections, mark all these as object detections, else mark m as a false positive.
3. Repeat steps 1 and 2 till no more detections are unmarked.

It can be shown that the spatial filtering window size in Step 1 should be the influence field area of the smallest object, A_1. The uniform reliability property of the routing protocol in our model guarantees that if the minimum active density threshold is exceeded in a window, then all detections must belong to some object type. The above procedure thus identifies and isolates object boundaries.

4.3 Tracking using shape estimation

Having derived the minimum network density and/or compensation technique to achieve separation between estimated influence fields of detected objects, we can use the estimated shape of the influence field to track the object location. For instance, a metallic object, which generates a uniform, circular influence field around it can be tracked at the centroid of the locations of magnetometers that detect it, while a light source producing a conical beam can be tracked at the vertex of the photosensors detecting it.

To achieve tracking, we require that nodes from which detections used in shape estimation are received, should be distributed uniformly across the influence field, otherwise the estimated shape of the influence field may be distorted. False negatives and false positives occur independently at nodes hence their distribution is uniform across the influence field. The one-hop contention model for faults is based on nodes randomly selecting the same slot for transmission, hence the distribution of these faults is uniform. Also, recall from Sec. 2 that the GridRouting protocol has the property of uniform reception probability for nodes equidistant from the aggregator. In the multi-hop case though, the probability of failure is non-uniform because farther nodes are subject to a higher loss rate. However, as shown in Sec. 3, the techniques of distance dependent probabilistic reporting and spatial reconstruction compensate for this non-uniformity of network failures.

5 Case study: A Line In The Sand

In this section, we describe the design and implementation of a distributed classification and tracking system which we called *A Line In The Sand*. This system consists of 90 Mica2 motes deployed in a 1.5m spaced grid to cover a 18m x 7m area. *A Line In The Sand* has been deployed in several outdoor settings to accurately distinguish between civilians, soldiers and vehicles by estimating their influence fields based on magnetometer and micro-power impulse radar sensors. For simplicity of presentation, we only describe classification between a soldier and a car using magnetometer based influence fields. We first describe how we derived system parameters like density using the theorems presented earlier. We then validate, both theoretically and experimentally, that accurate classification and tracking can be achieved whp (99%) in this network.

5.1 Experimental measurements

We first describe the experimental setup to measure key system parameters. To measure the magnetic influence fields of a soldiers and a car, a dense, regular grid of Mica2 motes with magnetometers was deployed. These objects were then made to traverse this network at different speeds and orientations. We then averaged the observations from over 100 such trials. The influence field areas A_1 and A_2, for a soldier and a car, were thus measured to be $12m^2$ and $63m^2$ respectively. The probability of node faults, $1 - p_n$, which included node failures and false negatives was measured to be 10%. By observing the number and distribution of false positives in the network over time and location, we determined the probability of false positives, p_{fp} to be 2%.

5.2 Determining network density

Substituting for the experimentally measured values of A_1, A_2, p_n and p_{fp} in Theorem 7, we obtained the following conditions for minimum network density:

$$\rho_{01} > 0.6 \quad , \quad \rho_{12} > 0.65$$

where ρ_{01} is the minimum density needed to distinguish false positives from a soldier and ρ_{12} is the minimum density needed to distinguish a soldier from a car. The minimum density needed to distinguish between all three, $\hat{\rho}$ was thus 0.65. Based on the communication radius of these nodes, we deployed 90 nodes in a 1.5m spaced grid to cover the 18m x 7m area with a network density of 0.7.

5.3 Effect of network unreliability

As mentioned in Sec. 4, we characterized the end-to-end network reliability function empirically. To do so, we deployed the network with the desired density of 0.7 and repeated the earlier trials with a soldier and car traversing the network and measured the network reliability at the aggregator. By averaging over more than 100 such trials, we obtained the values of p_{rcv_1} and p_{rcv_2} as 0.85 and 0.55 respectively. Substituting for p_{rcv_1},p_{rcv_2} and p_{fp} in Eqns. 27, 28, 29 and 30, we obtain:

$$E(rcv_1) = 7.1, V(rcv_1) = 2.9$$
$$E(rcv_2) = 24.25, V(rcv_2) = 14.9$$
$$E(fp_1) = 0.6, V(fp_1) = 0.6$$

It can be seen that these values satisfy the inequality in Eq. 31, thereby validating that the estimated influence fields of a soldier and a car can be separated in the presence of end-to-end unreliability and false positives.

5.4 Experimental validation

We now present experimental data to demonstrate that we were able to distinguish between the given object types in the presence of node faults and network unreliability.

Fig. 4(a) shows the probability distribution function for the influence fields of a soldier and a car as measured at the aggregator. In this outdoor experiment, timestamped detections were recorded in the non-volatile memory at each mote during 100 runs of a soldier and vehicle each, moving through the network. These detections were then downloaded and time-correlated to recreate the measured influence fields and their probability distribution. It can be seen

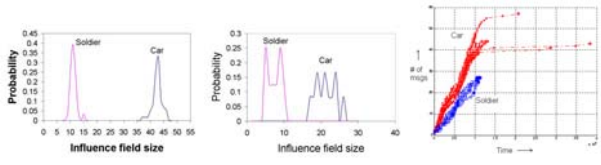

(a) Measured at motes (b) Measured at aggregator (c) Temporal evolution

Figure 4. Impact of network reliability on influence fields in *A Line In The Sand*

from Fig. 4(a) that these measured influence fields are indeed clearly separable. Fig. 4(b) shows the probability distribution of the influence fields estimated by the aggregator based on detections received using the GridRouting protocol in 100 runs of a soldier and a car each. It can be seen that in the presence of false positives and network unreliability, separation between the estimated influence fields of a soldier and a car is lower than in Fig. 4(a). However, as calculated earlier, we see that these distributions are non-overlapping meaning that the number of detections received for a vehicle is always greater than that for a soldier. From this data, we also calculate that there exists little variability (7.2%) in network reliability across these runs. Fig. 4(c) shows the temporal evolution of influence field estimation at the aggregator. The proximity of traces for individual runs indicates that in addition to reliability, network delays are also quite predictable, allowing us to choose tight latency bounds at the aggregator. This serves to validate our model that uniform, predictable end-to-end reliability can be obtained by careful design of the routing protocol.

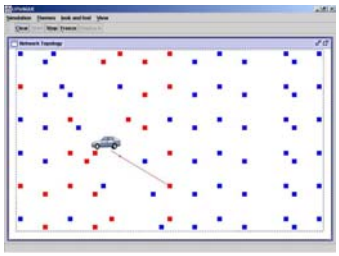

Figure 5. Classification and tracking of a car in *A Line In The Sand*

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

6 Extensions to the Influence Field Approach

In this section, we describe extensions to the influence field to increase the confidence in classification and also decrease the required density of deployment. One such extension is the use of multiple sensing modalities at each node to estimate multiple influence fields. As an example, consider the case where we wish to distinguish between a motorcycle, a SUV and a truck. The influence field areas for these objects using a magnetometer are $28m^2$, $78m^2$ and $150m^2$ respectively. The minimum density required to distinguish all three objects in the presence of faults, as calculated using the theorems presented earlier, is quite high (2-3m grid spacing). However, the influence field areas of the same objects with respect to an acoustic sensor are $1250m^2$, $700m^2$ and $1250m^2$ respectively. By estimating the influence fields for both these modalities, it is possible to distinguish between these objects with a much lower network density (8-10m grid spacing). The basic idea is to use the estimated acoustic influence field to distinguish a SUV from the other two object types and then use the magnetic influence field to distinguish between a motorcycle and a truck. Thus, using multiple sensing modalities reduces the network density to 6-10% of what was originally required. This simple example serves to demonstrate that estimating multiple influence fields can be used to classify multiple types of objects with higher confidence and lower network density. Multimodal influence field estimation was successfully demonstrated in the *ExScal* [3] network, which is one of the largest sensor deployments to date. *ExScal* used three types of sensors – magnetometers, acoustic and motion sensors – to estimate the influence fields of a person, a car and an ATV and classify and track them accurately.

Another extension to the influence field concept involves communicating not only presence outputs from each node, but some vector of features about the object like peak amplitude, total energy or distance from the object. Such an enhanced influence field can be used to improve confidence in the estimation output.

Regardless of whether the basis chosen for determining the influence field is merely presence as in the case of *A Line In The Sand*, a combination of multiple modalities, or a vector of features, reliable estimation is still an important problem and the same techniques described in this paper can be applied to each of these extensions.

7 Related Work

The notion of influence of an energy source is used in other science and engineering applications. In some formulations, the distribution of the intensity of the source at various points is considered while modelling its influence. For example, Kellogg et al [7] 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 [8] and ours, the distribution of the intensity is not modelled. Zhao et al [8] define an *influence area* as the number of sensors that detect an object. Our definition of the influence field also captures the shape of the influence field. We are unaware of previous work that has used influence field estimation as a basis for classification and tracking of objects.

The influence field approach should be contrasted to traditional approaches for classification and tracking using Unattended Ground Sensors [9]. The Remotely Monitored Battlefield Sensor System (REMBASS) is a representative example. The approach is centralized, requires complex pattern matching [10] and the sensing and processing devices are expensive, require careful and precise deployment as well as frequent remote monitoring. Meesookho, et al [11] describe a collaborative classification scheme based on exchanging local feature vectors, which imposes a high load

on the network. By way of contrast, most of the work on distributed tracking that decreases the load on the network is based on collaborative signal and information processing, sequential Bayesian filtering, and extended Kalman filtering [8, 12–17], that require significant node computation.

For the case of node faults, Krishnamachari et al [18] 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 and we have also presented 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. 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.

Our work also relates influence field to sensor coverage [19, 20] and highlights the important similarities and distinctions between the two concepts. Existing work on sensor coverage gives necessary conditions for detecting an object moving through a network. We are not aware of extensions that deal with the problem of classification.

8 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, 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 multiple concurrent targets that we dealt with experimentally, towards addressing the gap between existing theory and practice.

Acknowledgements. This work was sponsored by Defense Advanced Research Projects Agency (DARPA) contract OSU-RF #F33615-01-C-1901. We thank Santosh Kumar, Vinayak Naik, Hongwei Zhang and Hui Cao as well as the referees for their helpful comments and suggestions.

References

[1] A. Arora, P. Dutta, S. Bapat, and V. Kulathumani et al. A line in the sand: A wireless sensor network for target detection, classification, and tracking. *Computer Networks, Special Issue on Military Communications Systems and Technologies*, 46(5):605–634, 2004.

[2] Crossbow Technology Inc. MICA2 Mote. `http://www.xbow.com/Products/productsdetails.aspx?sid=72`.

[3] DARPA-NEST Extreme Scaling Project. `http://www.cse.ohio-state.edu/exscal`.

[4] Y.Choi, M.Gouda, H.Zhang, and A.Arora. Routing on a logical grid in sensor networks. Technical Report TR04-49, The University of Texas at Austin, 2004.

[5] S. Bapat, V. Kulathumani, and A. Arora. Reliable estimation of influence fields in unreliable sensor networks. *OSU Technical Report OSU-CISRC-8/04-TR49*, 2004.

[6] J. Hill, R. Szewczyk, A. Woo, S. Hollar, D. Culler, and K. Pister. System architecture directions for networked sensors. In *Architectural Support for Programming Languages and Operating Systems*, pages 93–104, 2000.

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

[8] F. Zhao, J. Liu, J. Liu, L. Guibas, and J. Reich. Collaborative signal and information processing: An information directed approach, 2003.

[9] M. Hewish. Reformatting fighter tactics. *Jane's International Defense Review*, June 2001.

[10] M. Caruso and L. Withanawasam. Vehicle detection and compass applications using AMR magnetic sensors, AMR sensor documentation. `http://www.magneticsensors.com/datasheets/amr.pdf`.

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

[12] A. D'Costa and A. 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.

[13] D. Li, K. Wong, Y. Hu, and A. Sayeed. Detection, classification and tracking of targets in distributed sensor networks. *IEEE Signal Processing Magazine*, 19(2):17–29, 2002.

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

[15] D. McErlean and S. Narayanan. Distributed detection and tracking in sensor networks. *36th Asilomar Conf. Signals, Systems and Computers*, 2002.

[16] F. Zhao, J. Shin, and J. Reich. Information-driven dynamic sensor collaboration for tracking applications. *IEEE Signal Processing Magazine*, 2002.

[17] M. Duarte and Y. Hu. Vehicle classification in distributed sensor networks. *Journal of Parallel and Distributed Computing*, 2004.

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

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

[20] H. Zhang and J. Hou. On deriving the upper bound of alpha-lifetime for large sensor networks. *Proceedings of ACM Mobihoc 2004*, 2004.