Improving System-wide Detection Performance for Sonar Buoy Networks using In-Network Fusion

Anshu Saksena, Lotfi Benmohamed, Jeffrey Dunne, Dennis Lucarelli, I-Jeng Wang

Johns Hopkins University Applied Physics Laboratory

Laurel, Maryland, USA

ABSTRACT

The problem of optimized distributed detection in a system of networked sensors involves a number of design aspects, including balancing probabilities of missed detection and false alarm as well as managing the communication resources through proper in-network information fusion. Moreover, a number of tradeoffs must be exercised, such as the one between the computational requirements for information fusion and sensor control and the communication requirements for information exchange. Therefore, overall system design decisions are best made by jointly considering the impact of design aspects and tradeoffs on the overall system performance. This paper addresses in-network fusion and associated networking algorithms that improve detection performance and energy efficiency for a multistatic sonar application. This is achieved by exchanging and fusing contacts among sonar buoys before transmission out of field. In-network fusion utilizes lower cost buoy-to-buoy communication for the majority of the data communication and enables a reduction in random uncorrelated false alarms by only reporting detections from multiple buoys that present sufficient correlation. The reduction of out-of-field contact transmissions allows a lower signal excess threshold for each buoy, corresponding to an increased probability of detection. We demonstrate the effectiveness of our distributed in-network fusion through both analysis and high fidelity sonar simulations.

I. INTRODUCTION

In this paper we study a number of design aspects for a large sensor field whose mission is to detect under-water objects like submarines. The system is buoy-based and distributed over a littoral area. The field has a number of gateway nodes that have the capability to transmit off-field through a satellite or an airplane. The system uses active sonar with acoustic sources sending acoustic energy (pings) and a number of sensor nodes receiving energy reflected by potential targets as well as non-target objects that may be present in the field. False alarms due to this clutter is controlled through proper fusion of sensor detections.

Our approach for distributed fusion exploits the peer-to-peer communications among the sensor nodes to perform *in-network fusion* before reporting the filtered detection information to an external decision center. The in-network fusion will take into account the bandwidth constraint for communicating off-field to the decision center and optimize the network-wide detection performance. There has been a significant amount of work in designing collaborative signal and information processing algorithms that minimize peer-to-peer communications while achieving satisfactory detection, classification, and tracking performance in generic sensor networks (see, for example, [1]-[5] and the references therein). Our results presented here focus on characterizing the benefits of in-network processing relative to the traditional centralized approach explicitly taking into account the communication constraint for a specific application with realistic system assumptions.

The paper is organized as follows. System model, performance metrics, and some analytical results are introduced in Section 2. Two correlation approaches are described in Section 3. Simulation results are discussed in Section 4 and network aggregation techniques are illustrated in Section 5.

II. SYSTEM MODELS AND PERFORMANCE METRICS

As mentioned above, our approach is based on distributed processing with in-network fusion. To compare and contrast our approach to that of the centralized fusion architecture, we start by briefly describing both approaches.

Centralized fusion: Under the traditional centralized fusion architecture each sensor buoy processes its locally received signal independently. A detection algorithm is implemented at each buoy to determine whether a target is detected. Based on the detection results, each buoy reports the positive detection with other important features associated with the detection. Driven by the bandwidth constraint on off-field communications, the local detection algorithm needs to be designed to manage the expected number of detection reports. This often entails setting the detection threshold high enough to limit the number of false alarms in a highly cluttered environment.

Distributed processing with in-network fusion: We propose to investigate an alternative approach assuming a realistic capacity for peer-to-peer communications among the buoys. We assume that the buoys will first communicate simple features of the local detection results in response to each ping to enable in-network cross-sensor correlation analysis. The set of features communicated include the sensor ID and the area of uncertainty for the target location based on locally received acoustic signals. Only detections from multiple buoys that present high enough spatial correlation are reported outside the sensor field. The detections from multiple buoys in response to a single ping are considered to be spatially correlated if they localize the target to the same grid point per our correlation approach described in section 3. The correlation is achieved through distributed in-network processing as well as efficient network aggregation described in section 5. Once detections are correlated, our fusion approach [6] consists of reporting detections that are supported by at least T spatially correlated detections. Note that in the centralized case all sensor detections are reported, which is effectively equivalent to operating with $T = 1$.

The correlation analysis can be performed at the gateway nodes or selected fusion nodes. In this paper we describe how to optimize the design of such in-network fusion to achieve higher network-wide detection performance.

A. Models

To characterize the detection performance of the proposed distributed fusion approach, we first describe our models and assumptions:

- We assume that there are N sensor nodes deployed in the field assumed to be an active sonar environment. Additionally, there are some acoustic sources capable of transmitting an impulsive signal (a ping).
- Following a ping from an acoustic source, each receiver processes the signal corresponding to the acoustic reflections off potential targets as well as clutter objects that may be present in the area (the sensor would look for returns which are the peaks in the received signal strength).
- Assume that a threshold-based detection algorithm is adopted for local detection at each buoy with some threshold τ applied to the received signal.
- The specific value of τ determines the associated probabilities of detection $p_d(\tau)$ and false alarm $p_{fa}(\tau)$ for each local detector.
- We assume that the local detection algorithm has independent performance across the buoy field and in response to pings from different sources given the hypothesis.
- The sensor is capable of deriving range (from time of arrival and acoustic signal speed) and bearing (from its sensor array) so that a spatial position is associated with each return (the number of returns is a function of the number of targets and how heavily cluttered the field is).
- Based on the estimated propagation time and bearing of the received signal, the target location can be estimated locally with some uncertainty for every local detection decision and an area of uncertainty (AOU) is associated with each detection.
- To approximate the effect of target localization uncertainty, we discretize the field into a set of grid points and associate each AOU with one or more grid points as we describe below.

B. Performance Metrics

We define the requirements of a performance metric for use in comparing the centralized and distributed fusion models. We present the obvious choice of metric and why it is not sufficient for this application. We then define the metric that we will use and provide analytical results that indicate that the distributed fusion approach would likely perform better than the centralized approach.

The quantities of interest to us are the detection performance of the sensor network and the out-of-field network communications of an approach. Therefore, our metric will be based on the network probabilities of detection, p_D, and false alarm, p_{FA}, and the number of returns transmitted out of the network, making the approximation that all transmitted returns consist of approximately the same number of bytes. In determining the p_D and p_{FA}, we must define a detection opportunity. In the traditional detection problem, a detection opportunity is defined as a measurement taken by a sensor. Since we are interested in the performance of a network of sensors, it may seem natural to define the network of sensors as a single sensor, and the collection of measurements made from all sensors for each ping as a single detection opportunity. As discussed above, the network responds with a detection if at least T sensors have detections, where the two models differ in what value of T they use and in how they determine the number of sensors with detections. In particular, in the centralized approach $T = 1$, resulting in the following expressions for the p_D and p_{FA} of the network:

$$p_{FA}^{cent} = 1 - \left(1 - p_{fa}\right)^N$$

$$p_D^{cent} = 1 - \left(1 - p_d\right)^N,$$

where p_{fa} and p_d are the probabilities of false alarm and detection at each sensor, respectively, and N is the number of sensors. Here we assume that the detection probabilities are identical for each sensor. These expressions approach 1 very quickly as N grows, rendering them virtually meaningless for large N. Additionally, comparisons of the models based on metrics defined this way would be dominated by the *1-of-N* processing of the centralized approach as opposed to the *M-of-N* processing of the distributed approach and would not consider other merits of each approach. Since these values are for detections from *at least 1* sensor from the network, they also do not indicate how many returns will be transmitted out of the network as a result of the detections, shedding no light on the communication requirement. Therefore, we will define these probabilities in a different way.

Instead of defining the detection opportunity based on the network of sensors acting as a single sensor, we retain each individual sensor's detection opportunities and define the probabilities accordingly. Again treating the sensors as identical and seeking to summarize their performance by a single quantity, the probability of false alarm will represent the chance that any non-target detection opportunity is considered a detection. Similarly, the probability of detection will approximate the chance that any target detection opportunity is detected. In practice, these will be calculated by looking at all target and clutter returns from each sensor for all pings and to calculate the following expressions:

$$p_{FA} = E\left[\frac{\#\,clutter\ reported}{\#\,clutter}\right]$$

$$p_D = E\left[\frac{\#\,target\ reported}{\#\,target}\right],$$

where *# clutter reported* and *# target reported* are the total number of clutter returns and target returns, respectively, that are reported out of the network, *# clutter* and *# target* are the total number of clutter returns and target returns, respectively, and $E[\ ]$ is the expected value of the argument over returns from all sensors and due to all pings. Here p_D only makes sense if a target is present. Intuitively, this definition of p_{FA} (p_D) characterizes the averaged percentage of clutter (true target) returns that the sensor network reports out of the entire field. Although these metrics do not consider localization, they do allow us to compare the number of target detections for each ping, which the network-level metrics did not, and the more target detections there are for each ping, the more information will be available for localization.

One key benefit of the metrics defined here is that the expected communication requirement is linearly proportional to p_{FA} and p_D. If we assume that the numbers of clutter and target returns are deterministic, then the expected communication requirement can be approximated by $(p_{FA}m_{FA} + p_D m_D)B$, where m_{FA} and m_D are the expected number of clutter and target returns, respectively, and B is the number of bytes required for each return reported out of the field. It is expected that m_{FA} grows linearly with the number of receivers N.

To derive analytical bounds for the presented definitions of p_{FA} and p_D requires stronger assumptions than the ones made in [6]. For the centralized approach, we can show that $p_{FA}^{cent} = p_{fa}$ and $p_D^{cent} = p_d$, assuming that the numbers of clutter and targets returns are deterministic and the clutter (target) returns are independent and identically distributed. Hence the communication requirement would grow linearly with N for the centralized approach.

For the distributed case, we can derive the following analytical approximations

$$p_{FA}^{dist} = p_{fa} \cdot GE\left(p_{cor}^f, N-1, T-1\right),$$

$$p_D^{dist} = p_d \cdot GE\left(p_d, N-1, T-1\right),$$

where $GE(p,N,k)$ is the probability of having at least k successes out of N independent Bernoulli trials with p as the probability of success, and p^f_{cor} is the probability that a clutter return exceeding threshold τ is spatially correlated with a clutter return exceeding threshold τ from another receiver. Note that p^f_{cor} is not a function of N and is expected to be much smaller than p_{fa} given the same local threshold τ. By applying the Hoeffding inequality we can further show that

$$p_{FA}^{dist} \le p_{fa} \cdot e^{-2(T-1-p_{cor}^f)^2 (N-1)},$$

$$p_D^{dist} \le p_d \cdot e^{-2(T-1-p_d)^2 (N-1)}.$$

Note that the communication requirement contributed by false alarms does not grow linearly with N as in the centralized case since p_{FA}^{dist} decreases exponentially with N given a fixed T. Furthermore, p_{FA}^{dist} decreases faster than p_D^{dist} as T increases since $p_{col}^f << p_d$ with appropriate choice of τ. These properties enable us to select appropriate parameters (τ and T) to effectively trade-off the system-wide detection performance.

III. CORRELATION APPROACHES

As an initial approach to fusion, we focus on the spatial correlation of returns from multiple sonar buoys. We discuss two approaches to correlate potential contacts: Coon's approach [8] for spatial correlation, and an approach that provides maximum likelihood soft associations. The latter uses two concepts from Coon's approach: correlation to grid points, and a discrete area of uncertainty (AOU) defined by bistatic time and bearing errors.

The problem of correlating contacts to one another has complexity $O(n\^2)$, where n is the number of contacts. In order to make contact correlation scale well to dense contact situations and noisy environments, the approaches we investigate define a fixed grid over the operating region. Contacts are correlated to grid points rather than to each other. Making the approximation that correlation follows the transitive property, contact-to-contact correlation is inferred between contacts that correlate to the same grid point. This approach scales as $O(n)$, making it more robust to large numbers of contacts. The two approaches we investigate differ in how they determine the correlation between returns and grid points and in how to disambiguate correlations of returns to multiple grid points. However, in both algorithms, correlations are determined based on defining a discrete AOU for returns.

Since contact localization is primarily based on two measurements, namely the bistatic return time and the bearing from the receiver to the reflector, and marginal error estimates are available for each of these measurements, the errors can be used independently to determine an AOU as follows. An acceptance region is determined for each measurement based on an assumed error distribution and a desired level of confidence. The acceptance region in bistatic time, when mapped back to spatial coordinates, forms an elliptical annulus. The acceptance region in bearing cuts out a radial slice of this annulus from the position of the sonobuoy. This slice is the AOU that forms the basis of the two correlation approaches investigated.

In Coon's algorithm, a return is considered to be spatially correlated to a grid point if the grid point lies within the AOU of the return. In this way, returns can be correlated to more than one grid point. In order to determine with which single grid point the return should be associated, Coon's algorithm maps returns to the grid point that has the most other returns that also can be mapped to it. Since each contact mapped to a grid point is independent evidence that there is a reflector near the grid point, this heuristic is justified.

In our second approach for spatial correlation, we assign a uniform distribution over the AOU of each return. Thus, smaller AOUs have higher probability densities than larger AOUs. We then define a region of space for each grid point that is the point's cell in a Voronoi diagram. For each return, we integrate the return's AOU density function in each grid point's cell. For each grid point, we maintain the sum of these integrals for all returns. In this approach, disambiguation of the grid point to which returns are mapped is accomplished by mapping returns to the point with the highest total integral. Since this approach takes into consideration size of AOUs and to some extent closeness of AOU to grid points, it is expected that it would give better localization results than Coon's algorithm.

IV. SIMULATION RESULTS

A. Sonar Simulation Models

In order to test performance over a range of environmental settings and conditions, numerical modeling was used to simulate the motion of a target through a distributed sensor network. The signal excess for echoes associated with source-receiver pairs was determined for each event in a ping schedule using a physics-based model. Clutter was established statistically from a measured characteristic distribution.

Target echo energy levels were established using the sonar equation:

$$TE = SEL - TL_{s\text{-}t} + TS - TL_{t\text{-}r} - Loss,$$

where TE is the target echo energy level, SEL is the signal energy level, $TL_{s\text{-}t}$ is the transmission loss from source to target, TS is the total energy target strength, $TL_{t\text{-}r}$ is the transmission loss from target to receiver, and Loss is a factor to account for additional losses, such as echo spreading, target splitting, processor losses, etc.

Transmission loss calculations are performed using APLNM, a range-dependent coupled normal mode code. Sound velocity profiles, bathymetry, and bottom properties are extracted from standard databases.

Background interference was calculated using a power sum of an assumed spectral ambient noise level (NL) and a calculated spectral reverberation level (RL). The RL was calculated by assuming a constant sound speed, with an equitime ellipse of reverberant patches constrained by the signal bandwidth. The receiver beam pattern is included directly in establishing RL, with a single directivity index (DI) applied to the NL. Both NL and RVB are determined in units of intensity, and are assumed constant over the bandwidth of the signal.

Potential clutter events are established using a uniform density of events for each receiver beam (e.g. one event per second). These events are drawn randomly from a generalized gamma distribution having mean, standard deviation, and skewness parameters fit to represent measured clutter statistics from at-sea experiments. The same detection threshold is then used on both clutter and target echoes to determine those events that are identified by the processor as “reportable” (i.e. a reported echo or a clutter spike).

B. Evaluation with Simulated Data

The simulation was used to generate signal excess, bistatic time and bearing for each receiver for many pings. Every receiver has exactly one return for the target and 6 clutter returns for every ping. The distribution of the signal excess of returns is plotted for clutter returns and target returns in Figure 1. Most target returns have lower signal excess than any clutter return.

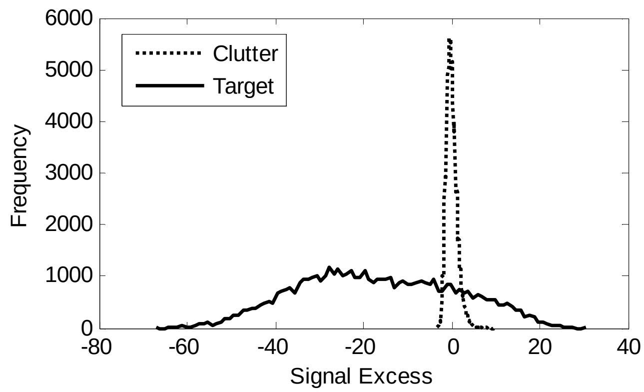

Figure 1 Distributions of the signal excess of target and clutter returns.

The system-wide detection performance of two approaches were compared using the simulated data: the centralized approach, and Coon’s approach for spatial correlation. Signal excess is the sole basis on which the centralized approach performs detection: each receiver determines whether it received any returns with a signal excess greater than the predetermined threshold and if so labels each as a detection. Given the unfavorable signal excess distributions, this approach performs poorly. The ROC curve is shown in Figure 2. As mentioned above, the communications requirements are linearly dependent on the probabilities of detection and false alarm and therefore can be ascertained from the ROC curve for each operating point.

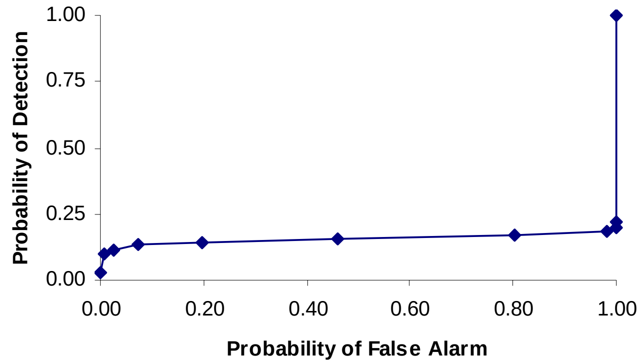

Figure 2 ROC curve for the centralized approach.

Detection performance may not be as hindered by the signal excess distributions when spatial correlation is taken into account. Care must be taken in setting the resolution of the spatial discretization in Coon’s approach. If the resolution is too coarse, returns from distinct objects may be associated with one another. On the other hand, if the resolution is too fine, then noisy returns from the same object may be associated with differing grid points. The typical size of the AOU is relevant in this manner. However, as AOU size is dependent on the range to the apparent reflector, some AOUs can be exceedingly small, making it difficult for the returns to be considered to correlate with any grid point. For these returns, the correlation is set to the grid point that is closest to the position of the apparent reflector.

In the distributed fusion approach, returns from each receiver are spatially correlated with returns from other receivers. For returns that are highly correlated, many receivers will have correlated detections. A histogram of the number of receivers with correlated detections is shown in Figure 3. The dashed lines show histograms for a local signal excess threshold that would be used in the centralized algorithm to limit off-field communications. The heavy dashed line is the histogram for true detections and the light dashed line is the histogram for false detections. The solid lines are the same for a much lower local threshold. These histograms show that in the space of spatial correlations, the distributions of true and false returns are more separable from a detection standpoint than the distributions in purely signal excess shown in Figure 1.

The detection performance using in-network fusion is summarized in the ROC curve in Figure 4. This curve was generated by using a low local threshold in signal excess, corresponding to the point at the top right of the centralized ROC curve and the solid lines of Figure 3. Using such a low threshold is traditionally not possible because of the ensuing flood of mostly false positives that would all be sent off-field in the centralized approach. However, with our additional layer of distributed processing to perform spatial correlation of these returns, we eliminate many of the local detections that are due to spatially uncorrelated clutter.

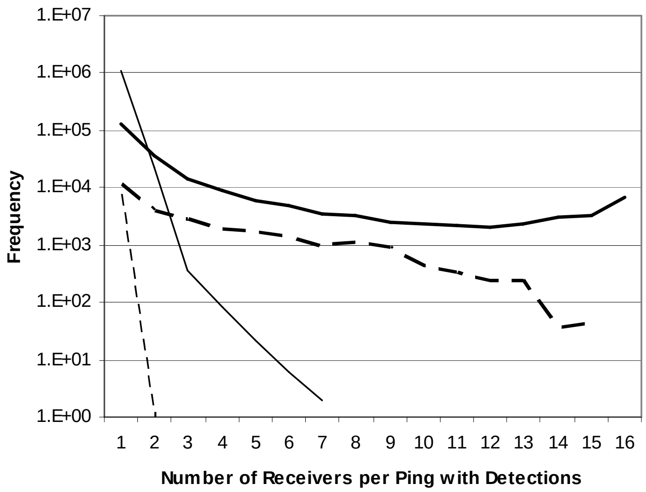

Figure 3 Histograms of the number of receivers with true (heavy lines) and false (light lines) detections for a low (solid lines) and intermediate (dashed lines) local signal excess threshold.

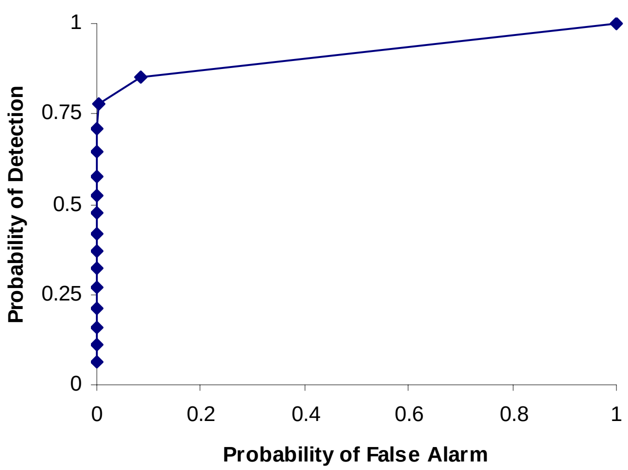

Figure 4 ROC curve for the in-network fusion approach.

V. NETWORK AGGREGATION

In this section we describe how network aggregation is used to efficiently implement the correlation of detections as described in Section 3. The brute force method of having correlated detections available at a gateway node would be to have every node send its detection information to the gateway where it gets correlated. This approach is very inefficient in energy and network resource consumptions. Indeed, consider the tree made up of all the paths between the gateway and all nodes in the sensor network. If each node needs to send a report to the gateway following a sonar ping, then each node in the tree would have to forward as many messages as nodes in the sub-tree rooted at this node. However, in our network aggregation approach described below, only one single message is forwarded by each node in the tree, thus yielding significant resource savings.

The example in Figure 5 shows a sensor field with four clusters of sonar returns each corresponding to a potential target. This example shows the AOUs of all receivers following a sonar ping, with each AOU belonging to one of the four clusters (the AOU in the middle of the grid is for illustration purposes only). Each AOU contributes a piece of information to the report generated by the node that owns the AOU (the node that received sonar reflections and estimated it to be from a reflector in that AOU). As an example, under Coon's algorithm the AOU in the middle of Figure 5 contributes the list of grid points it covers (g_1,g_2,g_3), and under soft association it contributes the list of grid points along with the integral of the AOU density in the grid point's cell in a Voronoi diagram (these cells are represented by the squares in this regular grid example). If a_i is the integral associated with grid point i the report contributes $\{(g_1,a_1),(g_2,a_2),(g_3,a_3),(g_4,a_4)\}$. Since a report in Coon's algorithm is a special case of a soft association report (covered grid points with the a_i weights being equal), we present below our network aggregation method in more detail for the correlation approach based on soft association.

Our network aggregation approach requires setting up a spanning tree rooted at the gateway node using a spanning tree establishment protocol. Once the tree is set up, and following a sonar source ping with each receiver node generating its detection information, a node which is a leaf node in the tree forwards its detections in a report message to its parent node. Each parent node collects detection information from all of its children and aggregates them along with its own before forwarding to its parent. In fact a parent node knows the set of its children and would perform the aggregation once it has heard from all of them (nodes that have no detections would still send a short "no-detection" message).

At any point in the tree, a report message contains the list of grid point clusters and the identifiers of the nodes that contribute to each cluster (i.e., nodes that have positive detections at some grid points in the cluster). A cluster C is described by the triple (G,A,I) where G is a set of identifiers of the grid points in the cluster, A is the set of weights a_i corresponding to each grid point g_i in G, and I is the set of identifiers of those nodes that made detections at one or more grid points in G. Once a node in the tree receives cluster information from its children it aggregates them along with its own generated clusters as follows: Let C_1, C_2, …, C_k denote the set of clusters that need to be aggregated. Two clusters C_i =(G_i, A_i, I_i) and C_j =(G_j, A_j, I_j) are said to overlap if the intersection of G_i and G_j is non-empty. These clusters get merged into one cluster C with G= G_i U G_j, I= I_i U I_j, and for each g in G its weight a is given by the sum of its weights in A_i and A_j.

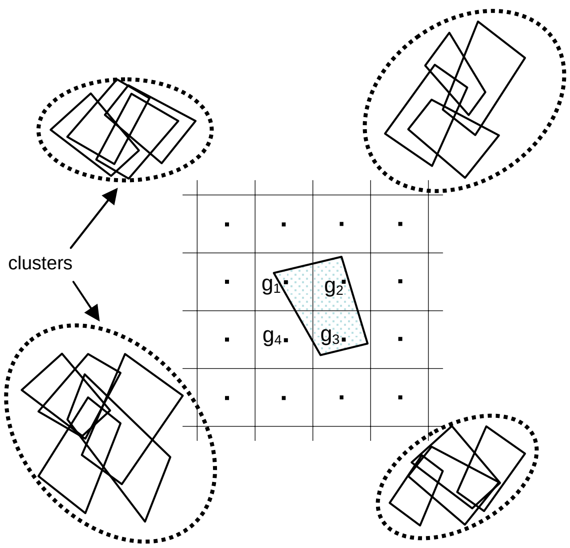

Figure 5 Sensor field with four clusters of sensor detection AOUs

As this aggregation process takes place starting from leaf nodes all the way to the gateway node along the tree, the set of network-wide aggregated clusters will be available at the gateway when it receives report messages from its children and aggregates them (with its own if any). At this stage, the detections within each cluster C get localized to the grid point in the cluster having the largest weight, with the new weight for this grid point being the sum of the weights a_i of all grid points g_i in the cluster C, which is equal to the number of nodes that have detections in the cluster (equal to the cardinality of the set I).

In the example of Figure 5, the gateway will have four clusters with detections localized to one grid point within each cluster. At the final stage in the fusion process, when the gateway has localized the detection to specific grid points g_i each with an associated number N_i of receivers that made the detection, those grid points with N_i larger than T will be reported to the external decision center. Moreover, if raw acoustic signal data needs to be sent to the decision center from some or all receivers that made detections, the identifiers of these receivers are available at the gateway since they are captured by the set I associated with each cluster.

VI. CONCLUSION

We present an approach to improving system-wide detection performance using in-network fusion. We illustrate the approach for a sonar buoy sensor network but it is applicable to other applications as well. In particular, most aspects of correlation, network aggregation, and distributed fusion apply in other domains. We show the improvements in network-wide detection performance for given network bandwidth constraints. The results are supported by high-fidelity sonar simulation data. Future availability of actual field data measurements can be used to further validate our approach.

VII. REFERENCES

1. F. Zhao, J. Liu, J. Liu, L. Guibas, and J. Reich, “Collaborative signal and information processing: and information-directed approach,” Proceedings of IEEE, 2003.
2. D. Estrin, R. Govindan, J. Heidemann, S. Kumar, "Next Century Challenges: scalable coordination in sensor networks," In Proceedings of the Fifth Annual International Conference on Mobile Computing and Networks (MobiCOM '99), Seattle, Washington, August 1999.
3. J. Hill, R. Szewczyk, A. Woo, S. Hollar, D.E. Culler, K.S.J. Pister. "System Architecture Directions for Networked Sensors," ASPLOS 2000.
4. D. Li, K. Wong, Y.H. Hu, and A. Sayeed. "Detection, Classification and Tracking of Targets in Distributed Sensor Networks", IEEE Signal Processing Magazine, 19(2), March, 2002
5. J.J. Liu, J. Reich, and F. Zhao, "Collaborative In-Network Processing for Target Tracking," J. of Applied Signal Processing, April 2003
6. L. Benmohamed, P. Chimento, B. Doshi, B. Henrick, and I-J. Wang, “Sensor Network Design for Underwater Surveillance,” in Proc. of IEEE Milcom, Washington, DC, Oct. 2006.
7. W. Hoeffding, “Probability inequalities for sums of bounded random variables,” Journal of American Statistical Association, vol. 58, p. 13-30, 1963.
8. A. Coon, "Spatial Correlation of Detections for Impulsive Echo Ranging Sonar," JOHNS HOPKINS APL Technical Digest, Vol. 18, No. 1, 1997.