SensIT: Sensor Information Technology For the Warfighter

Srikanta Kumar, Ph.D.
DARPA – ITO
3701 North Fairfax Drive
Arlington, VA 22203-1714

David Shepherd
Strategic Analysis, Inc.
3601 Wilson Blvd Suite 500
Arlington, VA 22201

Abstract – *The Defense Advanced Research Projects Agency's Sensor Information Technology (SensIT) program is developing software for networks of distributed microsensors. This paper outlines the program goals, technical challenges, and some of the ongoing research. Specific ongoing work in collaborative signal and information processing and fusion is emphasized.*

Keywords: collaborative signal and information processing, distributed microsensors, fusion, query/ tasking

1 Introduction

Sensors have always been important to the military for use is reconnaisance, surveillance, and tracking and targeting. Recent advances in MEMS technology, embedded processing, and wireless communication, is enabling the development of microsensor devices that can be deployed in large numbers in the battlefield. A microsensor device will have multiple on board sensors (such as acoustic, seismic, infra-red, magnetic, etc.) embedded processing and storage, short range wireless links (ten to hundreds of meters), and positioning capability either through GPS or through a relative positioning mechanism. Continued advances in microtechnology will also enable "smart dust," i.e. microsensors, on the order of square millimeters or square centimeters in size, packaged with transceivers, having the capability to produce enormous quantities of data. In battlefield conditions, the large quantities of information available to the warfighter must be simplified, processed and presented appropriately and in a timely manner, due to the unique demands placed on warfighters and the penalties for cognitive overload on the battlefield. To accomplish these goals, methods of data fusion, collaboration and sensor networking are being developed by researchers in the Sensor Information Technology (SensIT) program, sponsored by the Information Technology Office (ITO) at the Defense Advanced Research Projects Agency (DARPA). SensIT researchers are also testing their software in the field, with the assistance of the United States Marine Corps and other service entities.

2 Program Goals

The SensIT program's primary goal is the creation of a new class of software for distributed microsensors. The program has two key thrusts: a) development of novel networking methods for ad hoc deployable microsensors, and b) leveraging the distributed computing paradigm for extraction of right and timely information from a sensor field, including detection, classification, and tracking of targets. The program has five tasks: Fixed Networking, Fixed-Mobile Networking, Collaborative Signal and Information Processing, Query/Tasking, and Integration.

2.1 Fixed Networking

Networked microsensors for the military must have many desirable characteristics. First, the networking algorithms and protocols should support ad hoc networking that can scale to large numbers. Second, networking must support rapid self-assembly without any manual intervention for configuration. Third, networking must be adaptive to the environment, as nodes or links may be dead on arrival or links may suffer outages due to fading. Fourth, networking must support the primary operational and technical requirements that drive system parameters, such as low latency and high energy efficiency. Sensor networking must be resource efficient, and support incremental deployment. Finally, networking must be survivable, secure, and have low probability of detection by others when the sensors are deployed.

2.2 Fixed-Mobile Inter-Networking

SensIT research aims to enable seamless interaction between forward-deployed, ad hoc sensors and fixed devices, and networks in rear areas. Sensors on mobile platforms might be located on UAVs, robots, ground vehicles, and even troop uniforms. Relays and other power-rich resources such as UAVs and mobile command posts will be integral elements of a fixed-mobile network.

2.3 Collaborative Signal and Information Processing; Fusion

A key feature of a sensor network is collaborative processing among nodes. Incorporating inputs from multiple sources increases the signal to noise ratio, i.e. the accuracy of useful target information. High node density enables dense spatial sampling and data sharing. This data must be fused, whereby the data from multiple nodes is combined to detect, classify, and track targets. The information must be exfiltrated to users or dynamic query points.

2.4 Querying and Tasking

User interaction with sensors is handled primarily through tasking of sensor operations and queries. In a sense a sensor field is a database, as sensor devices collect, process, and store data and information. Temporal and spatial data and information is stored in the distributed database of the sensor field. Queries and tasks must be simple for the user to input, sensor reports easy to understand. Dynamic querying and tasking should also be supported. The distributed nature of both users as well as data locations means information must be available at any location and recoverable with a minimum of latency.

2.5 Integration

A SensIT goal includes the integration of all the software from the above tasks, through well-defined APIs, and demonstration, in lab and in field, of new capabilities of microsensor networks. The software is being developed, tested, and integrated on hardware platforms developed by Sensoria [http://www.darpa.mil/ito/research/sensit/], with integration performed by BBN. Researchers have performed extensive integrated experiments in the summer of 2000 and 2001, and are planning additional experiments in the future.

3 Technical Challenges

The uncertain and harsh environmental conditions of sensor networks demand dynamic adjustment of mode of operations, including dynamic and adaptive routing, deployment of redundant nodes, and employment of robust communications protocols. Nodes and networks must be able to operate autonomously for extended periods of time. Resources such as transmission and processing power as well as bandwidth must be used only on a restricted or as-needed basis. Networks must be reliable, survivable and secure in the face of enemy jamming, friendly transmissions and incidental and spurious signals.

3.1 Fixed Networking

The primary challenge in operating a network of sensors is to ensure reliable operation in an uncertain environment. Traditional network performance metrics, such as latency, reliability, and energy performance, remain critical. Additional factors perhaps unique to sensor networks include scalability and device form factor. Sensors' small sizes limit battery size and performance. Given decreasing sensor sizes and intentions to deploy large quantities of sensors, how does network performance change as the numbers of nodes increase? How does network density, e.g. number of devices per square kilometer, affect performance? Yet throughout the requirements space, achieving reliable networking in an environment without always having to maintain end-to-end connections remains important. To achieve this goal, SensIT is examining alternatives to traditional Internet Protocol (IP) routing, such as diffusion routing [1].

3.2 Fixed – Mobile Networking

The challenges of fixed device networking are multiplied in the case of interaction between fixed and mobile devices. Challenges inherent in networking of both static and mobile devices include discovery of new devices, provisioning of services in a network of varying resource capacities, and handling of engagements between fixed and mobile devices. Intermittent connectivity as well as varying speeds of mobile devices must be handled. A mobile device such as a UAV might be in the range of several fixed devices; traditional protocols developed for cellular and PCS are not adequate, and as such new protocols, and new ways of handling handoffs, are required. Fixed networks with established user lookup tables must be able to handle incoming, new users. Finally, the user poses the ultimate requirement: does network operation appear seamless?

3.3 Collaborative Signal and Information Processing; Fusion

The high node densities found in sensor networks facilitate collaborative processing, yet truly collaborative behavior, combining efficient routing with the fusion of data from relevant nodes to determine consensus on targets, remains challenging. Processing, synthesizing and fusing inputs from multiple sources requires sophisticated distributed signal processing algorithms. Primary challenges include the processing of data asynchronously, as processing speeds may vary dramatically due to differing resource availabilities. When times of arrival appear irregularly or far apart in time, has more than one target appeared? Does a similar signal from multiple sensors spaced far apart mean one target is moving rapidly, or have several similar targets appeared simultaneously? How should processing be divided among multiple sensors to enable maximum power efficiency and minimum latency? Aside from latency, detection accuracy also plays a large role in this scenario. What is the fusion strategy for on-board multi-modal processing at one device? And what are good methods for cross-node fusion? Algorithms are needed that enable progressive accuracy, so processing can be stopped to conserve power when desired accuracy is reached. Yet throughout the process, energy and processing efficiency must remain critical.

3.4 Querying and Tasking

Challenges arising from distributed sensor networks' interaction with the environment include enabling users to simply and quickly query the networks, as well as to task them with requirements. Processing functions and database calls during data exfiltration must remain hidden to the user. Set up and maintenance of distributed databases and execution of remote queries requires coordination and fusion among nodes, while access to the database from potentially anywhere in the sensor field requires intelligent data and cache placement and servicing.

3.5 Integration

The primary challenge implicit in networking hundreds or potentially thousands of sensors is achieving a complete, end-to-end solution. Sensors come in a variety of shapes and sizes, with varying power, processing and connectivity requirements, varying task responsibilities, and varying reporting demands. Achieving timely interaction among the tremendous variety of sensors requires a program of iterative development and testing, particularly in military-relevant scenarios. SensIT will provide real-time proof of the functionality of a baseline system, and the collection of data to support laboratory experiments and ongoing algorithm development.

4 Collaborative Signal and Information Processing; Fusion

Collaborative behavior is superior to single processor behavior because nodes acting collaboratively can in effect perform computations in parallel, compare results to enable more intelligent decisions, minimize communication needs through local summarization and maximize data quality by selecting the highest value data sources. Primary collaborative signal processing (CSP) activities include

- The exchange of data among sensor nodes to enable better decisions and other high level data to be derived from raw sensor signals.
- Fusing data from multiple sensing modalities and irregular sensor placements.
- Minimizing power consumption on sensor nodes for communications, signal processing, and sensing.
- Reaching a consensus belief state among sensor nodes about what is occurring in the physical world and mapping sensor data to entities in this consensus.
- Creating timely reports in response to user needs.

4.1 SensIT CSP and Fusion Research

SensIT researchers are developing several CSP techniques for use in distributed networks. Collaborative signal processors being built by SensIT researchers have the ability to associate or correlate data sets for the same target from multiple nodes, selectively fuse data from multiple nodes and sensor types, select the geometrically optimal nodes for collaboration, iteratively update feature estimates based on asynchronous data received (possibly out-of-order) across the network, and distribute processing components across the sensor network to minimize power usage. Other methods under research include exploiting asynchronous feature update estimations to make on-the-fly decisions for localization and classification, and resource-bounded optimization techniques.

Analyzing data from multiple sources presents challenges not encountered in the case of single processor computation. Among these is threshold decisions: due to differences in background noise, the same signal may appear to be noise to one sensor while appearing distinct to others. When multiple sources appear in a sensor field simultaneously, how are the targets differentiated? Responses will vary among sensors, due to variations in sensor type or connectivity conditions.

One technically challenging area especially relevant to collaborative processing is finding a balance between the cost of communicating and the cost of computation. Traditional signal processing algorithms, which tend to rely on large matrix operations, require too much communication between nodes to be useful in a distributed signal processing environment. As such, it may be worth performing more than the minimum number of computations at the local level if the total cost of communications can be reduced [2].

Several SensIT projects are related to fusion and collaborative processing. First, multimodal onboard fusion of data from acoustic, seismic and infrared sensors for detection is being performed by Steven Beck and Joseph Reynolds of BAE Systems - Austin[3]. Second, the technique of Latent Semantic Analysis is being developed to convert large quantities of raw sensor data to useful information. This technique integrates signal data, in the form of time-series data, and semantic information, in the form of target attributes. This way, time series data from heterogeneous sensors can be compressed into semantic attributes. These attributes can processed locally at sensor nodes and, using pattern matching, compared to templates preloaded onto sensor platforms [4]. Third, Richard

Brooks and S.S. Iyengar are employing collaborative processing for target tracking using methods based on the Byzantine Generals based approach. In this approach, one node is chosen to combine the sensor readings and determine the target location most consistent with the information provided. Field tests showed that this local collaboration produced fused localization estimates robust to node errors and more accurate than individual readings [5, 13].

SensIT researchers from MIT LL have studied Time Difference of Arrival (TDOA) based approach for estimating a bearing to target. Whereas the Closest Point of Approach (CPA) method is useful when distances between nodes and targets are small, permitting reports from multiple nodes in close proximity, TDOA is most useful in estimating information when targets are located further from the sensors: TDOA works best when the distance between the targets and nodes is two to three times the distance between the nodes. It provides any-time information, unlike the tripwire approach of CPA, but it is more computation- and communication-intensive than CPA. On the other hand, compared to beamforming, TDOA can require less computation, makes fewer demands on precise node location information and time synchronization, and avoids a full exchange of time series data. TDOA produces a composite bearing estimate by choosing mutually consistent bearing solutions after comparing frequency tracking data across pairs of nodes Unneeded portions of targets' frequency profiles can be discarded, saving processor cycles but leading to a tradeoff between bandwidth and accuracy [6].

Another group of researchers is also improving multiple target tracking capabilities by combining classification algorithms with more traditional tracking algorithms. Researchers tracked targets using a two-step process: first, nodes surrounding a target detection share data to create a region of interest. The nodes in the vicinity of the detection select one node to be the manager for the detection process, and then report the energy detection and CPA information to the manager node. Second, nodes report energy levels and CPAs for successive time instants, and then use this information to predict the position of the target at later times. The information is also used to create and activate new regions for subsequent detection and tracking. Temporal processing is performed locally at each node in more sophisticated versions of this capability. To differentiate between targets that might appear closely in space and/or time, classifiers that contain data collected previously operate in parallel to recognize unfamiliar target [7].

Researchers from Xerox are studying techniques for information-directed sensor querying. These techniques allow the intelligent selection of sensors for collaboration based on information-theoretic and communication-based cost measures. This work may use a hybrid of traditional tracking techniques with a distributed particle filter. This nonparametric Bayesian approach to tracking and fusion of multiple sensing modalities allows incorporation of nonlinear dynamics and complex sensor noise models [8]

4.2 Experimentation

SensIT researchers have conducted two field experiments to test collaborative signal processing theories in real world environments. In these experiments, termed SensIT Experiment 2000 (SITEX00) and SensIT Experiment 2001 (SITEX01), researchers teamed with the U.S. Marine Corps to test SensIT CSP capabilities at the Marine Corps Air Ground Test Facility at Twentynine Palms, California.

For SITEX01, researchers conducted three experiments. In the first experiment, 10 nodes equipped with seismic sensors and separated by approximately 10 meters were used to track vehicle position, speed and direction of travel. Vehicles utilized included tanks, LAVs (light armored vehicles), HMMWVs (light to medium weight all purpose vehicles), and Dragon Wagons (heavy weight off-road tractor-trailer combinations). Data was remotely displayed. The second experiment used three passive infrared sensor nodes and one imager node to detect the presence and motion of ground vehicles and predict the vehicles' location so that the imager node, located further along the vehicles' direction of travel, was triggered to transmit the image of the vehicles, and then return to sleep mode. In the third demonstration, 15 nodes carrying magnetometers were deployed along a road, half of them hand emplaced and half dropped by a UAV. After circling back, the UAV queried the ground network for vehicle detection, time of occurrence and speed. The UAV then exfiltrated the responses to a remote base camp. The UAV also relayed images taken from cameras mounted in its nose and body [9].

Joe Reynolds and Steven Beck of BAE Systems - Austin have performed multi-node collaborative target detection and tracking. For the SITEX00 experiments, the group built 12 single sensor nodes containing omni-directional electret microphones, an 8-element linear array, a 2X4 element planar array, and two tetrahedron microphone arrays. They distributed eight of the single element microphone nodes along a road, and digitized the signals from different vehicles and troops using a 12-bit simultaneous sampling with 8 kilosamples per second per channel (20 kilosample/sec data at 16 bits is available for some of the nodes). The acoustic data was primarily used for target identification and for calculation of the time of CPA to the sensor based on Doppler line tracking. For the SITEX01 experiment, the group implemented a likelihood ratio vehicle detector for seismic and passive infrared sensing modalities, and combined the multiple detections using Bayesian techniques. The combined detections provided a false alarm rate less than 1 in a million and were input to a decentralized Kalman tracker used to trigger a wireless imager. On-node processing reduced the data rate to less than 100 bits of target state information per detection event, making low power wireless collaboration feasible. In a related project, this group recently demonstrated real-time vehicle and dismounted troop detection, discrimination, and tracking using four sensing modalities (acoustic, magnetic, seismic, and pulsed Doppler radar). The 8 KHz per channel data was collected, stored, and processed in real-time, and a Bayesian inference network was used to combine the detections from each sensor modality and discriminate pedestrian troop movement from vehicles and other clutter sources [10].

Similar experiments were conducted as part of SITEX00 in March, 2000. In 33 scripted runs, SensIT researchers gathered 13.5 hours of data on targets such as amphibious assault vehicles (AAVs), LAVs, Dragon Wagons, HMMWVs, 5-ton trucks and targets of opportunity. Forty nodes containing a total of 120 sensors (acoustic, seismic, and infrared) were used. Data on these experiments, available to SensIT researchers at http://dstl.bbn.com, includes information on each run, a narrative description of each run, run parameters, and node locations. Run data is listed by day and number of run and includes target data and node settings. Visitors may access a limited portion of the site by using the username of “visitor” and password of “visitor.”

4.3 CSP Canonical Scenarios [11]

To assist with the development of collaborative processing algorithms, SensIT researchers have written twelve “canonical scenarios.” These scenarios were written to test key capabilities of collaborative sensing systems. By utilizing combinations of these capabilities, it should be possible to demonstrate scenarios typical of those found in ground sensor deployments for military uses. Parameters employed in the scenarios include sensor field size, node type and deployment densities, initial states, trajectories and maneuvers of targets (vehicles only), desired accuracy, precision and resolution of target position, and network performance factors such as reliability, latency, and output rates. Targets included eight types of vehicles, ranging in size and acoustic signature from light wheeled trucks to heavily armored tanks. Sensors employed included directional and omnidirectional acoustic, seismic, passive infrared and magnetic sensors. Benchmarks measured for each scenario include energy consumption, detection accuracy, detection latency, tracking accuracy, and tracking latency. The scenarios, listed in rough order of processing difficulty, include:

1. Track Single Target: continually estimate target position versus time as the target moves along a road. Technical challenges include target localization, maintenance of estimate accuracy despite widely-varying temporal and physical distances between sensors, and the fusion of data from multiple sensor types.

2. Track Single Maneuvering Target (see figure 1): estimate target position versus time as a target maneuvers along an a priori unknown, offroad path. Challenges presented by this task include estimating position and time without prior knowledge of vehicle trajectory, and fusion of reports from many simultaneous observations in a 2D sensor network.

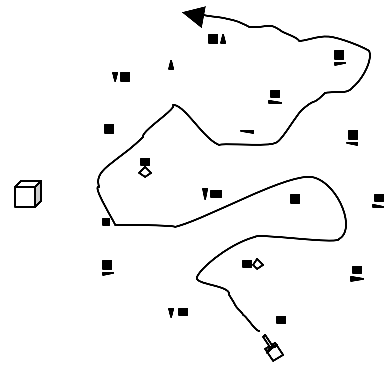

Fig. 1: Single Maneuvering Target scenario

3. Track One Accelerating/Decelerating Target: estimate target position versus time. In this scenario, target signatures vary with time, preventing the use of constant vehicle dynamics models. Shifts in vehicle gears may require maintenance of several discrete internal sensor states.

3. Count Stationary (Idling) Targets: counting and locating a number of pre-positioned targets. Challenges in accomplishing these tasks include source separation, achieving consensus on a distributed count of the vehicles, locating them without using closest point of approach (CPA) methods, and differentiating between types of vehicles.

4. Two-Way Traffic: Two-way or crossing traffic requires the sensor system to maintain vehicle identities through the time of crossover, when multiple vehicles of the same type produce simultaneous signals on the same sensors. This crossover time presents the primary technical challenge: maintaining accuracy of identity tracks by noting and tracking vehicle dynamics, and maintaining a running estimate of target crossing time. Useful parameters for these characteristics include determining time of crossing within 0.5 seconds and achieving 99% rate of correct target identification.

5. Convoy on a Road: Tracking multiple targets of various types along a road requires tracking an a priori unknown number of vehicles as well as determination of vehicle order. Technical challenges include initialization of new vehicle tracks, and handoff of vehicle tracks along a long road.

6. Track Multiple Maneuvering Targets: In this scenario three vehicles maneuver on the same layout as in the single target maneuvering scenario (#2). Two vehicles are of one type, the third is of another type. The task is to estimate the target positions versus time. The primary technical challenge presented by this scenario is the two-dimensional data association problem: how to map the raw sensor data to tracks and how to harmonize the beliefs of multiple nodes. No a priori knowledge of paths is provided.

7. Perimeter Violation Sensing: In this case, a perimeter has been set up; violators entering the perimeter must be identified and tracked, while activity outside of the perimeter ignored. The primary challenge arising in this scenario is filtering distracters from appropriate targets. Benchmarks used in this case include detection latency, power usage during periods when distracters are present but no violation occurs, and frequency of false positive reports.

8. Tracking in an Obstacle Field (see figure 2): track a vehicle amidst a field of obstacles of varying sizes and locations. Technical challenges include maintaining tracks despite loss of data due to obstacle interference, keeping sensors from locking on obstacles, and dealing with distinct obstacles affecting modalities in different manners. Obstacle configuration and density can be changed to vary the difficulty of the problem.

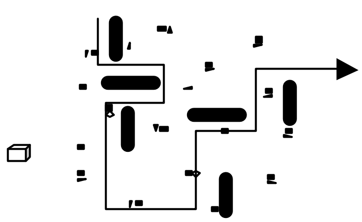

Fig. 2: Tracking an Obstacle in a Field

9. Identity Tracking and Mutual Exclusion using One-Time Sensors (see figure 3): When multiple targets are separated in space and time, a sensor establishes individual target identities, but when the targets pass together through a tunnel without sensors the distinct identities are lost. When the identity of one target is regained, the system must re-establish the connection between the original tracks and the identities of both vehicles. The primary technical challenge of this scenario is to reason about target identities despite periods of uncertainty. For example, if targets A and B merge then split and the target heading one direction is later identify as A, then the target headed the other direction must be B.

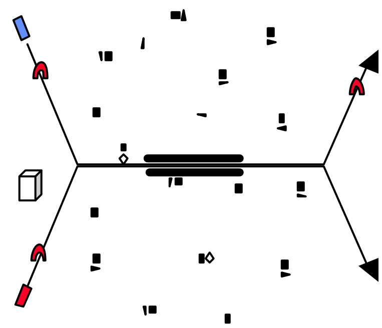

Fig. 3: Identity Tracking and Mutual Exclusion scenario

10. Cluster Behavior: When the separation between multiple vehicles becomes too small, tracking them individually becomes difficult if not impossible. This scenario demonstrates the ability to track the centroid of a group of vehicles and count the number of vehicles even while some cluster members arrive and depart. Challenges in this task include differentiating between multiple targets, coalescing multiple similar targets into a centroid, limiting exponential hypothesis blowup, and measuring the global properties of the cluster rather than properties of a single target.

11. Multiple Target Clusters: The most demanding of the scenarios presents multiple groups of many vehicles. The task is to determine the number of clusters and number of vehicles in each cluster, and to track the centroid of each cluster over time. Clusters may merge or split, and vehicles may depart or arrive. The amount of data involved might cause latency between cluster merge/split and notification to other clusters, and in the formation of cluster centroids. Cluster maintenance requires dynamic collaboration between nodes. Cluster size estimates and identities must be maintained over multiple queries. The size of the network and the

number of vehicles and clusters can be varied to stress the network's coordination capabilities.

5 Conclusion

Ongoing research in the SensIT program is developing collaborative processing and fusion techniques for distributed sensor networks. Further information on these and other topics is available at http://www.darpa.mil/ito/research/sensit/index.html.

References

[1] C. Intanagonwiwat, R. Govindan and D. Estrin, "Directed Diffusion: A Scalable and Robust Communication Paradigm for Sensor Networks," *Proceedings of the Sixth Annual International Conference on Mobile Computing and Networks (MobiCOM 2000),* August 2000, Boston, Massachusetts.

[2] A. Oppenheim and C. Sestok, presentation at SensIT Principal Investigators' conference, St. Petersburg, FL, April 19, 2001.

[3] S. Beck, J. Reynolds, and T. Leih, "Collaborative Signal Processing for Distributed Micro-Sensor Technology," forthcoming in *Proceedings of the BAE SYSTEMS Sensor Conference*, Dunchurch England, June 2001.

[4] D. S. Friedlander and S. Phoha, "Semantic Information Fusion for Coordinated Signal Processing in Mobile Sensor Networks," forthcoming in *Special Issue of the International Journal of High Performance Computing Applications*.

[5] R. R. Brooks, C. Griffin, and D. S. Friedlander, "Self-organized distributed sensor network entity tracking," forthcoming in *Special Issue of the International Journal of High Performance Computing Applications*.

[6] P. W. Boettcher and G. A. Shaw, "A Distributed Time-Difference of Arrival Algorithm for Acoustic Bearing Estimation," to appear in Proceedings of Fusion 2001 conference.

[7] D. Li, K. Wong, Y. Hu, and A. Sayeed, "Detection, Classification and Tracking in Distributed Sensor Networks," to appear in Proceedings of Fusion 2001 conference.

[8] M. Chu, H. Haussecker, F. Zhao, "Scalable Information-Driven Sensor Querying and Routing for ad hoc Heterogeneous Sensor Networks," forthcoming in *Special Issue of the International Journal of High Performance Computing Applications.*

[9] K. Pister, presentation at SensIT Principal Investigators' conference, St. Petersburg, FL, April 19, 2001. See also http://robotics.eecs.berkeley.edu/~pister/.

[10] S. Beck, presentation at SensIT Principal Investigators' conference, St. Petersburg, FL, April 19, 2001.

[11] J. Reich, "SensIT Collaborative Signal Processing Canonical Scenarios, v0.7," Xerox PARC, April 2001.

[12] S. Blatt, "Collaborative Signal Processing for Sensor Networks", *Proceedings of the SPIE*, vol. 4126, pp. 183-185, July 2000.

[13] R. R. Brooks and S. S. Iyengar, *Multi-Sensor Fusion: Fundamentals and Applications with Software*, Prentice Hall PTR, Upper Saddle River, NJ, 1998.