Raytheon

Integrated Sensing and Processing (ISP) Phase II: Demonstration and Evaluation for Distributed Sensor Networks and Missile Seeker Systems

Progress Report:

4th Quarter Progress Report
1 Dec 2006 – 28 Feb 2007

Contract No.: N00014-04-C-0437
Contract Line Item Number 0001
Deliverable Item: Publications (A001-008)

Raytheon Company
P.O. Box 11337
Tucson, AZ 85734-1337

Distribution Statement A.

Approved for public release.
Distribution is unlimited.

DESTRUCTION NOTICE

For classified documents, follow the procedures in DOD 5220.22M, National Industrial Security Program Operating Manual (NISPOM), Chapter 5, Section 7, or DOD 5200.1-R, Information Security Program Regulation, Chapter IX. For unclassified, limited documents, destroy by any method that will prevent disclosure of contents or reconstruction of the document.

Report Documentation Page		*Form Approved* *OMB No. 0704-0188*

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE **FEB 2007**	2. REPORT TYPE	3. DATES COVERED

4. TITLE AND SUBTITLE **Integrated Sensing and Processing (ISP) Phase II: Demonstration and Evaluation for Distributed Sensor Netowrks and Missile Seeker Systems**	5a. CONTRACT NUMBER
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Raytheon Company,PO Box 11337,Tucson,AZ,85734-1337**	8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT **Approved for public release; distribution unlimited.**

13. SUPPLEMENTARY NOTES **The original document contains color images.**

14. ABSTRACT

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**		**41**	

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std Z39-18

01 March 2007

Progress Report

CDRL A001 No. 8
Seventh Quarterly Progress Report for Period of Performance
1 December 2006 – 28 February 2007

Integrated Sensing Processor Phase 2
Program Manager: Dr. Harry A. Schmitt
Principal Investigator: Dr. Harry A. Schmitt

Sponsored By:

Defense Advanced Research Projects Agency/DSO
Dr. Carey Schwartz/DARPA DSO
Program Manager: Dr. Dan Purdy/ONR
Issued by ONR under Contract #N00014-04-C-0437

Prepared By:

Raytheon Systems Company
P.O. Box 11337
Tucson, AZ 85734

EXECUTIVE SUMMARY

The primary goal of this effort is to bring to maturity a select set of basic algorithms, hardware, and approaches developed under the Integrated Sensing and Processing (ISP) Phase I program, implement them on representative hardware, and demonstrate their performance in a realistic field environment. We have identified a few promising research thrusts investigated in ISP Phase I where field demonstrations are cost prohibitive but collected data sets are available. Here, we will conduct a thorough performance evaluation.

TABLE OF CONTENTS

INDEX OF TABLES

INDEX OF FIGURES

0. Technical Abstract

Advances in sensor technologies, computation devices, and algorithms have created enormous opportunities for significant performance improvements on the modern battlefield. Unfortunately, as information requirements grow, conventional network processing techniques require ever-increasing bandwidth between sensors and processors, as well as potentially exponentially complex methods for extracting information from the data To raise the quality of data and classification results, minimize computation, power consumption, and cost, future systems will require that the sensing and computation be jointly engineered. ISP is a philosophy/methodology that eliminates the traditional separation between physical and algorithmic design. By leveraging our experience with numerous sensing modalities, processing techniques, and data reduction networks, we will develop ISP into an extensible and widely applicable paradigm. The improvements we intend to demonstrate here are applicable in a general sense; however, this program will focus on distributed sensor networks and missile seeker systems.

1.0. Management Overview and Summary

1. A. Program Summary

The Raytheon Company, Missile Systems (Raytheon) ISP Phase II program is a twenty-four month contract with a Period of Performance (PoP) covering 1 March 2005 to 28 February 2007. Raytheon has four universities and one small business as ISP Phase II subcontractors: Arizona State University (ASU); Fast Mathematical Algorithms and Hardware (FMAH); Georgia Institute of Technology (Georgia Tech); Melbourne University (UniMelb) and the University of Michigan (UM).

1. B. Program Status

The Raytheon ISP Phase II Program status can be summarized as remaining "on track." All of the negotiations have been completed and all of the subcontractors are now under subcontract. We had incurred some minor schedule slips on both the distributed tracking and the Cooperative Analog Digital Signal Processing (CADSP) demonstrations during the PoP; the revised schedule supports demonstrations before 31 March 2007. This is a slip of approximately one month. We still expect to complete the contract on time and budget.

1. C. Personnel Associated/Supported

Raytheon

Dr. Harry A. Schmitt	Principal Investigator
Mr. Donald E. Waagen	Co-Principal Investigator
Dr. Sal Bellofiore	Distributed Sensing Lead
Mr. Thomas Stevens	Distributed Sensing Support
Dr. Robert Cramer	Mathematical Support
Mr. Craig Savage	Waveform Design and Control Lead
Dr. Nitesh Shah	High Dimensional Processing Data Lead

FMAH

Professor Paolo Barbano
Professor Ronald Coifman
Dr. Nicholas Coult

<u>**ASU**</u>

Professor Darryl Morrell
Professor Antonia Papandreou-Suppappola

<u>**Georgia Tech**</u>

Professor David Anderson
Professor Paul Hasler

<u>**UniMelb**</u>

Dr. Barbara LaScala
Professor William Moran
Dr. Darko Musicki
Dr. Sofia Suvorova

<u>**UM**</u>

Professor Al Hero
Dr. Raviv Raich

<u>Significant Personnel Actions:</u> There were no significant personnel changes during the current PoP.

1. D. Recent Events
The following events occurred during the current PoP

- Raytheon personnel (Waagen and Schmitt) visited Dr. Kirstie Bellman Aerospace Corporation 31 January 2007 to discuss possible collaboration and transition opportunities.

- Raytheon personnel (Stevens and Bellofiore) visited Professor Neal Patwari at the University of Utah on 30-31 January 2007 to discuss the University of Michigan RSS mote self-localization approach.

- Raytheon personnel (Bellofiore and Schmitt) visited Intelligent Automation Inc. (I-A-I) on 17 January 2007 to discuss various distributed mote algorithms and approaches being developed under a DARPA Phase II SBIR.

- Raytheon personnel (Schmitt) visited the University of Melbourne on 15-19 December 2006 to discuss work being performed.

- Raytheon personnel (Schmitt) visited DARPA DSO on 16 January 2007.

1. E. Near Term Events
- Conduct distributed tracking demonstration by 31 March.

- Conduct CADSP imager demonstration by 31 March

- Write and submit the final Technical Report by 30 April.

2. A. Technical Progress
    In this section we provide a more detailed discussion of the technical progress that occurred during the current PoP broken down by subcontractor.

<u>2.A.1. Raytheon Technical Progress</u>
    We now discuss the technical progress for the Raytheon team.

2.A.1.a. Identification of target support overlap in a low-complexity classifier

Target classification is an important step in Automatic Target Recognition (ATR). It is useful for classification algorithms to have low representation complexity, for purposes of meeting computational constraints in tactical systems and achieving generalization capabilities (avoiding over-training). Classifier complexity can be driven by several factors, including high-dimensional data structure, sparse sampling and complex decision boundary (possibly including regions of ambiguity) [Ho *et al.* 2000]. The Support Vector Machine approach can mitigate classifier complexity, however it does so at the cost of requiring substantial processing time for training [Burges 1998]. The Class Cover Catch Digraph approach produces a reduced-complexity partitioner/classifier without requiring extensive processing time for training [DeVinney *et al.* 2002], [Priebe *et al.* 2003].

We present here an extension to CCCD, in which we find regions of class support overlap in high-dimensional data while maintaining (or actually reducing) partitioning complexity with respect to that achieved by the standard CCCD approach. In the context of Integrated Sensing and Processing, there are three avenues for exploiting identified mixed regions: ATR – GNC feedback, ATR – sensor feedback, and decision trees. For platforms with reconfigurable sensors and/or constrained sensor suites, identified regions of target overlap can be used in developing management and scheduling algorithms for sensor configuration and sensor suite utilization by providing "block out" regions in the configuration / suite space. That is, if certain targets look the same when viewed by a particular sensor in a particular configuration, then it is better to avoid operating the sensor in that mode.

It is desirable for a tactical ATR system to be able to provide trajectory-shaping feedback to Guidance and Navigation Control (GNC). For example, it may be decided by a tactical missile's ATR algorithm that the current azimuth/elevation view (aspect) of the detected target provides insufficient information to perform the target classification task, and that changing the target azimuth by 15° would allow for high-confidence classification. The GNC algorithm could then adjust the platform motion towards the target to provide the appropriate target aspect. If there are a set of detected targets that require classification disambiguation, the GNC algorithm could optimize platform trajectory to best accommodate an overall improvement in the viewing aspects of the various targets.

The identified mixed regions can be used as part of an ATR decision tree. One can use feature-space regions that have significant target signature overlap, but little clutter signature overlap, as a robust mechanism for discriminating target detections from clutter detections. Then, given the detection is a valid target, one can use the results of another CCCD-mixed-region analysis that has been determined from only target signatures (no clutter included) to locate, for example, highly-overlapped "tracked targets" (with little "wheeled target" overlap) regions and highly overlapped "wheeled targets" (with little tracked targets) regions as robust tracked/wheeled recognizer. Then, given the target is tracked, one can use the results of another CCCD-mixed-region analysis that has been determined from only tracked-target signatures to know which highly-mixed regions to avoid when classifying the tracked target as a T72, BMP, SA12,

etc. So, in some cases, highly overlapped regions are to be used, in other cases highly overlapped regions are to be avoided – it depends on the question currently being asked.

The ATR – sensor and decision tree applications can certainly be combined. If a sensor is reconfigurable, and/or if multiple operationally-constrained sensors are available, the configuration / scheduling can be determined using decision-tree analysis. For discriminating targets from clutter, the choice of individual sensor parameters and/or which sensor is used can be guided by creating large volumes in hyper-feature space that have significant within-target-class overlap and significant within-clutter-class overlap, with little overlap between target and clutter classes. Similar consideration can be given for finer levels of the decision tree.

Generally ATR systems are *trained* on a set of data (measured or simulated) that spans the target set of interest, for example covering variations in target aspect. Given a set of classification features (either data-driven or defined by a domain expert) for target classes of interest, the high-dimensional feature space is partitioned into high-purity regions (data support limited to mostly one target class), and regions that are ambiguous (data support contains multiple target classes). In the context of the CCCD approach to partitioning high-dimensional feature spaces, we provide an extension to automatically identify regions of significant target support overlap, while maintaining low classifier complexity. We first provide an overview of the CCCD approach, then we discuss our extension for a two-class case (generalizable to n classes).

In the CCCD approach, class-conditional regions are modeled with a mixture of balls. The number, location and size of the balls are determined based on the proximity between training samples. The balls form a low-complexity representation of each class. In the standard version of (α,β)-CCCD, each class-cover ball contains a percentage β of out-of-class samples (purity factor / sensitivity to contamination) and the union of class-cover balls can neglect up to a percentage α of in-class samples (properness factor / sensitivity to outliers).

We are given d-dimensional training data sets $\mathbf{X}_{0\ldots n0}$ for class 0 and $\mathbf{X}_{1\ldots n1}$ for class 1, and a dissimilarity measure (*e.g.*, the standard Euclidean distance) over all pairwise combinations in the combined data set. Start with the samples of class 0 as the "base" set. Treat each sample as the center of a ball. For each ball, determine the maximum radius such that exactly $(\beta * n_1)$ class-1 samples are included in the ball. From the resulting n_0 balls, select a subset in the following greedy approach. First select the ball that contains the largest number of class-0 samples. If the number of covered class-0 elements is less than $(\alpha * n_0)$, select the ball that contains the largest number of as-yet-uncovered class-0 samples. Repeat this process until the union of balls covers at least $(\alpha * n_0)$ class-0 samples. This set of balls forms the cover for class 0, $\mathbf{C}_0$. Now, exchange the roles of class-0 and class-1 samples and repeat the process, resulting in a cover for the class-1 samples, $\mathbf{C}_1$. The covers $\mathbf{C}_0$ and $\mathbf{C}_1$ provide a compact description of the support of the two classes, with the problem complexity implicit in the number of balls in the cover, their relative positions and the distribution of their radii. The covers $\mathbf{C}_0$ and $\mathbf{C}_1$ found by the greedy approach are not necessarily optimized. Rather, the greedy approach is used to find an approximate dominating set in order to reduce the problem complexity from NP-hard to polynomial time.

We have developed the following extension. Define three parameters related to finding *mixed* regions: properness and purity parameters as before, α_M and β_M, and a fractional coverage threshold F_T. First compute the covers $\mathbf{C}_0$ and $\mathbf{C}_1$ with $\alpha = \beta = 0$ (proper, pure covers). For the example case shown in Figure 1a (2-dimensional, 2-class data with $n_0 = n_1 = 500$), the proper, pure covers are shown in Figure 1b.

Determine the fraction of in-class samples covered by each ball in $\mathbf{C}_0$ and $\mathbf{C}_1$. Identify the balls with membership falling below F_T, and label these as "potentially mixed" balls. The other balls are labeled as "large" balls. The result for our example (using $F_T = 0.005$) is shown in Figure 2a, with centers for "potentially mixed" balls highlighted in red. Note that some "potentially mixed" centers are in regions that are not mixed.
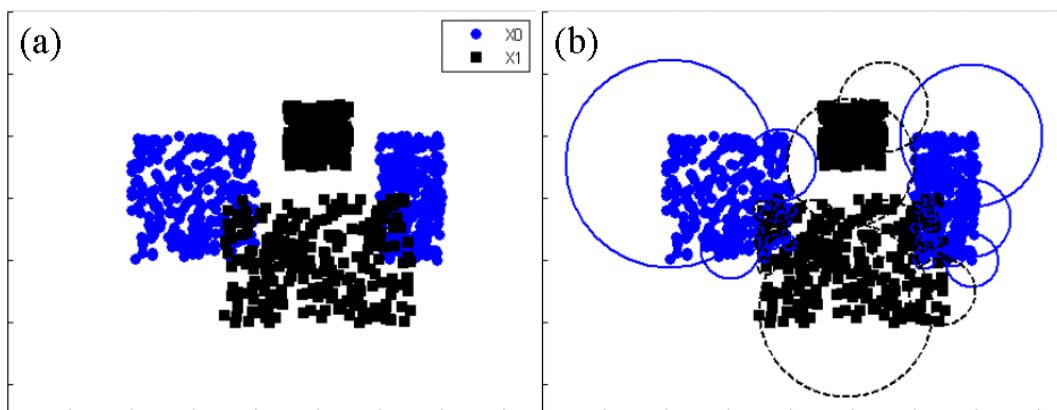
Figure 1: Overlapping 2-Class example, (a) original data; (b) pure, proper CCCD covers

We use the following clean up logic to remove these from further consideration. Remove from the "potentially mixed" list any ball whose center's closest neighbor in $(\mathbf{C}_0 \cup \mathbf{C}_1)$ is both a "large" ball and of the same class type. Remove from the "potentially mixed" list any ball whose center's (closest neighbor is both "potentially mixed" and of the same class type) AND (next closest neighbor is both "large" and of the same class type). The result of applying these rules is shown in Figure 2b.
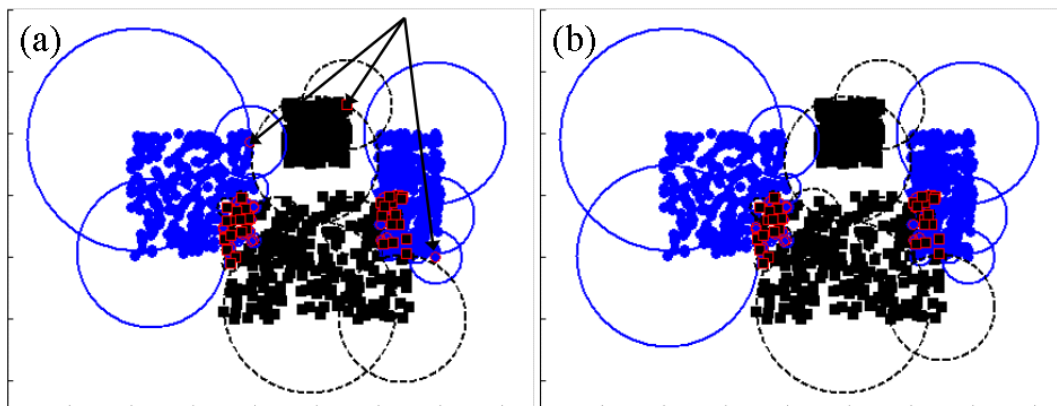
Figure 2: "Potentially mixed" centers highlighted in red. (a) Arrows point out undesired centers. (b) After applying cleanup logic

Apply a "potentially mixed" label to all the class-0 and class-1 samples that are covered by balls that remain in the "potentially mixed" list. The remaining samples are labeled "unmixed." Now apply the CCCD algorithm with α_M and β_M, with "potentially mixed" samples representing the new class 0 and "unmixed" samples representing the new class 1. One can then proceed with standard CCCD (or another partitioner) for the samples that are not in the "mixed" regions. The result for our example (with $\alpha_M = 0.05$, $\beta_M = 0.01$, $\alpha = 0.01$ and $\beta = 0.01$) is shown in Figure 3a, with the class-0 "mixed" region cover shown in red. For comparison, in Figure 3b we show the result of applying the standard CCCD (no extension for finding highly mixed regions) with $\alpha = 0.01$ and $\beta = 0.01$.
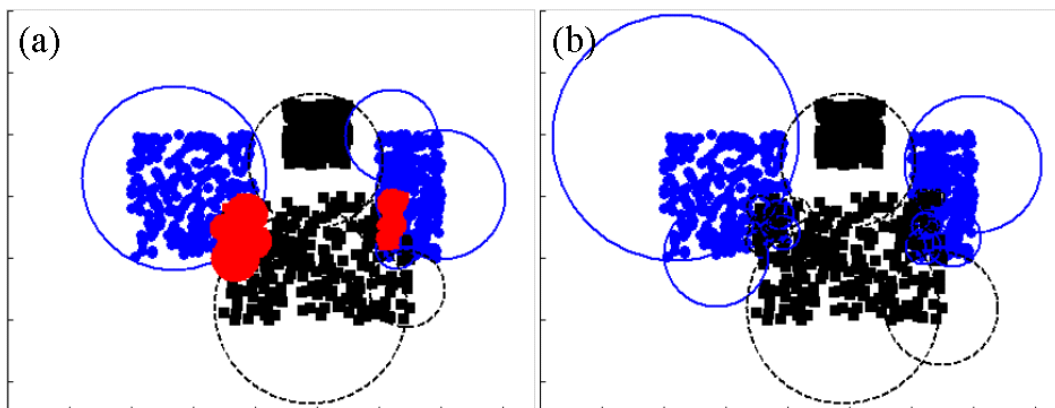
Figure 3: (a) Standard CCCD after determining mixed region (red). (b) Standard CCCD with no mixed regions identified

2.A.1.b. Distinguishing Between Direct Path and Terrain Bounce Jamming
<u>Introduction</u>

A radar system operates by radiating electromagnetic energy into space and detecting the echo signal reflected back to the radar from a target. The reflected energy not only indicates the presence of a target, but by comparing the received echo signal with the signal that was transmitted (matched filtering), the target location can be determined along with other target-related information [Skolnick 2001, p.1]. A hostile target would naturally like to deny the radar system access to this information, if possible, and thus may employ electronic countermeasures. In particular, noise jamming is the intentional transmission of energy in order to mask the target return and impair the effectiveness of the receiving radar. The receiving, or victim, radar may then employ a counter-countermeasure technique known as "home-on-jamming" designed to track the angle of the jamming signal and reveal the location of the jammer. To prevent discovery of his location, the jammer may employ "angle deception" or terrain-bounce jamming. This could be implemented, for example, by an aircraft flying at a relatively low altitude and transmitting a "noise waveform" toward the ground. This technique effectively presents a false targeting angle to the detecting radar system, thereby rendering ineffective the home-on-jamming counter-countermeasure. (See e.g. [Hsu]).

Obviously, it is necessary to be able to distinguish between direct path and terrain bounce jamming signals, since the detecting radar must know whether or not it can

Distribution Statement: Approved for public release; distribution is unlimited.

believe the target angle it is tracking. Under ISP, we have undertaken an information theoretic approach to this problem, which consists of the following two steps:

1. Process a collection of raw data collected by the receiving radar via the ISOMAP algorithm. The data is simulated and is known as to whether it represents a direct path or a terrain bounce jamming signal. The result is two distinct collections of points embedded in a linear space of lower dimensionality than the input samples, which are said to occupy the high dimensional space.

2. Compute the Henze-Penrose test statistic between the two collections of points in the low dimensional linear space, to determine whether the two collections are drawn from different underlying distributions.

If the Henze-Penrose test indicates that the underlying distributions from which the samples are drawn are indeed significantly different, then we have established the existence of information which could be further exploited, for example to design a method for automated classification of jamming signals, at least so far as terrain bounce vs. direct path is concerned.

Description of the Simulated Data

The data for these experiments were generated at Raytheon, and consists of several scenarios, six high altitude cases and six low altitude cases. In all cases the jammer is stationary and at a fixed altitude with a missile flying in from various azimuth and dive angles. The simulation time for all scenarios is 10 seconds. "High altitude" means the initial range separation is 11,000 meters, the missile velocity is 1000 meters per second, and the final range separation is 1000 meters. "Low altitude" means the initial range separation is 1100 meters, the velocity is 100 meters per second, and the final range separation is 100 meters. In some cases the missile flies directly at the jammer and in some cases the flight path is directed toward the bounce point in the ground plane.

This is a bi-static system, with the jammer transmitting continuously and the missile passively receiving. Noise energy is transmitted either through an isotropic, or spherical, antenna, or a directional antenna with very low side lobes pointed toward the ground. Each scenario consists of 10 coherent processing intervals (CPIs), and each CPI consists of 128 pulse repetition intervals (PRIs), 100 range gates, and 4 antenna channels.

The ISOMAP Algorithm

This algorithm is fairly well-known by now and a detailed description seems unnecessary, thus we make only some brief comments and refer the interested reader to the references, e.g. [Tenenbaum 2000]. The ISOMAP algorithm attempts to extract data which lie on a nonlinear manifold and place them on a linear one, such that the geodesic distances between points on the nonlinear manifold are closely approximated by Euclidean distances on the linear one. Geodesic distances are approximated in piecewise linear fashion, by added up a series of "short hops" between data points, which is the only manner in which this can reliably be done since outside of the data itself we have no knowledge of the underlying nonlinear manifold. It follows that the geodesic approximations can only be accurate if we have "lots" of data points. The aim is to faithfully represent the statistics of the data with a much smaller number of dimensions in the linear space, thus the notion of *dimensionality reduction*; however, we also

11

experiment below with embedding points in a linear space of higher dimensionality than what we started with. This seems odd in light of the manner in which ISOMAP is normally used, but the point is that it gives us *something,* and that this *something* may contain useful information.

Henze-Penrose Two Sample Test

This is a test for determining whether two sets of samples were drawn from the same, or different, underlying distributions. It is based on the earlier Friedman-Rafsky test [Friedman-Rafsky 1979], which begins by constructing a minimal spanning tree on the pooled, or combined, data. The spanning tree is then pruned, to remove all edges which connect points from different samples. The pruning procedure breaks the tree into a number of connected components which, if the data were statistically quite similar and thus heavily intermingled, will be large. A small number of connected components indicates the data are statistically dissimilar and allows us to conclude that the two samples were drawn from different underlying distributions. The contribution of Henze and Penrose was to provide proofs of Friedman and Rafsky's conjectures. See [Henze-Penrose 1999].

Results of Experiments

The raw signal data consists of complex voltages. Since we prefer to work with real numbers our first step was to replace each complex valued sample with the complex magnitudes, then normalize each data vector to have unit length. There are three different ways to extract data vectors, since the data is contained in a three-dimensional array each element of which corresponds to the following:

$$\text{data}(i, j, k) \quad \Leftrightarrow \quad \text{range gate} = i, \text{pri} = j, \text{antenna channel} = k.$$

The three methods for selecting data vectors are as follows:

1. All i for a fixed j and k.
2. All j for a fixed i and k.
3. All k for a fixed i and j.

Having extracted data samples by one of these three methods, then pre-processed it as described above, the collection of samples for either direct path or terrain bounce are processed through the ISOMAP algorithm, and embedding coordinates one through twenty for each sample are computed. The Henze-Penrose test statistic is then computed between these collections of embedding coordinates, for one embedding dimension, then two, then three, etc. up to twenty embedding dimensions. Finally we plot the test statistic as a function of the number of embedding dimensions.

For samples with an equal number of points in each, the test statistic should be near 0.5 if the two samples are very similar and near 1.0 if the two samples are very dissimilar. A value, say, above 0.8 might be considered a "good" indication, while a value above 0.9 might be considered a "terrific" indication of dissimilarity. Of course these thresholds are somewhat arbitrary. As mentioned above, what we are trying to do here is establish the existence of information that can be exploited, and this is indicated by high test scores. We present some results in the following figures. The caption below each figure gives the details.
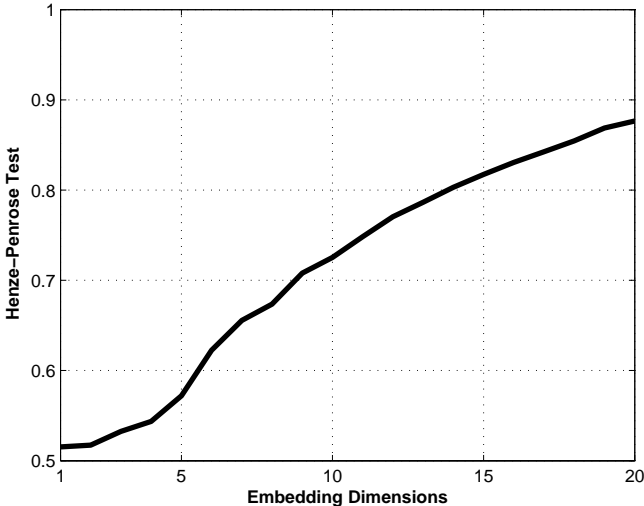
Figure 4: Data with a directional antenna, first CPI (furthest from the target), showing two-sample test statistic for one through twenty embedding dimensions. The underlying distributions become more dissimilar as the number of dimensions increases. Data vector were extracted by taking all range gates for a given PRI and antenna channel, which resulted in a total of 512 samples, each of length 100.
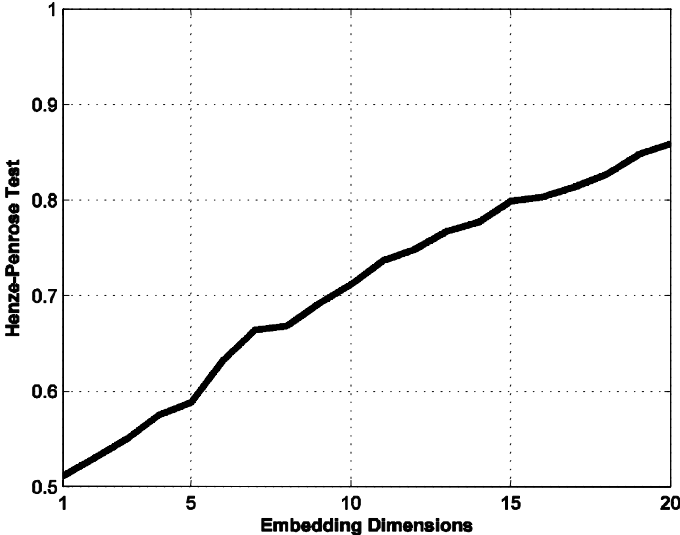
Figure 5: Data with a directional antenna, first CPI (furthest from the target), showing two-sample test statistic for one through twenty embedding dimensions. The underlying distributions become more dissimilar as the number of dimensions increases. Data vectors were extracted by taking all PRIs for a given range gate and antenna channel, which resulted in a total of 400 samples, each of length 128

Distribution Statement: Approved for public release; distribution is unlimited.
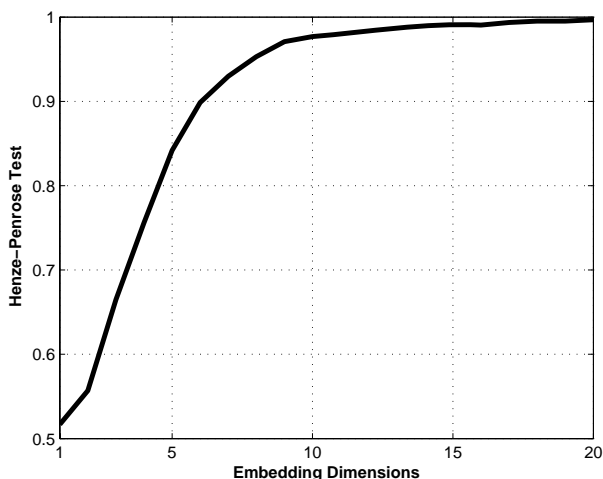
Figure 6: Data with a directional antenna, first CPI (furthest from the target), showing two-sample test statistic for one through twenty embedding dimensions. The underlying distributions become more dissimilar, in this case markedly so, as the number of dimensions increases. Data vectors were extracted by taking all antenna channels for a given range gate and PRI, however we used only the first five PRIs, since if we try to use all of them the number of data points becomes too large for the computer to handle. This resulted in a total number of 500 samples, each of length 4. Note that the number of embedding dimensions becomes larger than the number of dimensions in the "high" dimensional space. However, the Henze-Penrose test seems to be telling us that there is useful information here and, furthermore, that these two sample distributions are very dissimilar in these regions.
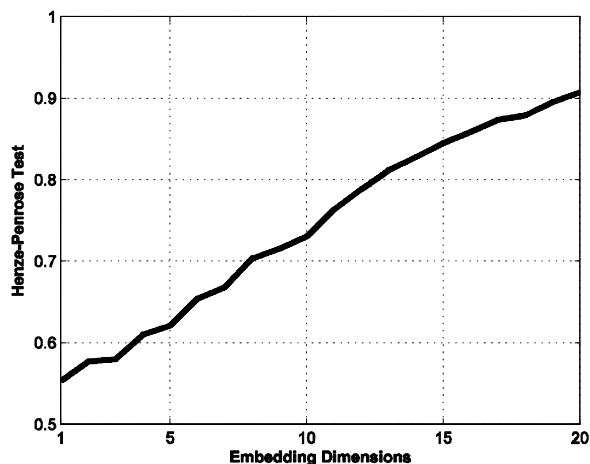
Figure 7: Data with a directional antenna, last CPI (nearest to the target), with data vectors extracted by taking all range gates for a given PRI and antenna channel. This figure should be compared with Figure 5, to get a feel for the variation with range of our results. It can be seen (empirically) that the test statistic for the last CPI increases a little faster than for the first CPI, but it appears that the statistic we are measuring is not a strong function of range to the target.

Conclusions

We have processed radar data, formed into one-dimensional vectors using one of three possible methods, through ISOMAP and then computed Henze-Penrose two-sample test statistics, in an effort to determine whether samples of direct path and terrain bounce jamming signals show significant statistical differences. The results show that in fact they do, which implies that there is information here which can be exploited. We noted that the test statistic is a strong function of embedding dimensions, but only a weak function of range to the target. Also, these results would seem to indicate that forming data vectors by taking all the antenna channels for a fixed PRI and range gate gives the greatest statistical separation, however we do not wish to commit ourselves to such a conclusion too early, so at this point we merely call it an "indication."

2.A.1.c. Distributed Tracking Demonstration

ISP Motes Demo Status

Since our last status report, we have made some progress on the ISP Motes Demo. In particular, we have completed the Graphical User Interface (GUI) of the demo, we are one week away from completing the integration of the Unscented Kalman Filter (UKF) tracker (from University of Melbourne), and we are collecting data on the DWMDS self localization to improve the estimated location results. Details on each area follow.

ISP Motes Demo's GUI

We have finally completed the ISP Motes Demo's GUI. A snapshot of the GUI is shown in Figure 8. The GUI contains several features to help the user design his test. For example, the user can if he/she wishes name the test, and this test name (or experiment name) is carried out throughout all the output files. To avoid overwriting output files, a date and time stamp is attached to the experiment name. In addition to the experiment name, the user can choose among three different trackers: Virtual Measurement (VM) tracker, UKF tracker, and Particle Filter (PF) tracker. Another feature to help the user is to choose among a live test versus a playback test.
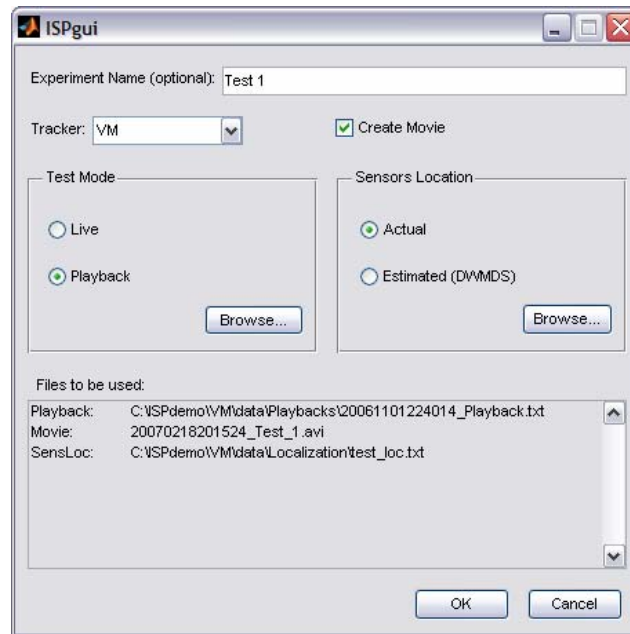
Figure 8: A snapshot of the ISP Motes Demo's GUI

A live test operates on real time detections while the playback operates on saved time stamped detections. This is good because the user can replay the same detections with a different tracker offline. For the nodes location, the user has a choice to either use the actual location or estimated location (produced by the DWMDS code from the University of Utah in Salt Lake City, Utah). Finally, all input and output files needed for the demo are summarized at the bottom of the GUI so that the user is aware of which files are being used, and all settings of the GUI are saved to minimize user's input. In a week, all trackers' codes will be tested to verify full compatibility with this new GUI.

UKF Tracker
The new version of the UKF tracker's code was received from the University of Melbourne. The new version includes a fix to take into account detections time stamps. In addition to this fix, the new version was built to be easily integrated with the new GUI. The UKF tracker's code will be verified by the end of the week.

DWMDS Self Localization
    Two weeks ago, we visited the University of Utah in Salt Lake City, Utah, to understand why the code was producing results outside of the anchor nodes. In other words, the algorithm estimates the relative distance of all the nodes within four known positioned nodes (anchor nodes). One of the reasons why the code was producing such results was that it was compiled with a different version of TinyOS, and that we were also using an outdated header file. We have been testing the self localization code for the past two weeks, and although our original results have improved, they are currently still poor. That is, the nodes are placed within the anchor nodes; however, the error on the estimated positions is too large. Figure 9 shows the graphical representation of the results from the DWMDS Self Localization code. The network tested was of 16 motes in a 4x4 configuration with inter-spacing of 1 ft. The error varied anywhere from half a foot to 3 ft. Figure 10 shows how each network node converges after several minutes.

16
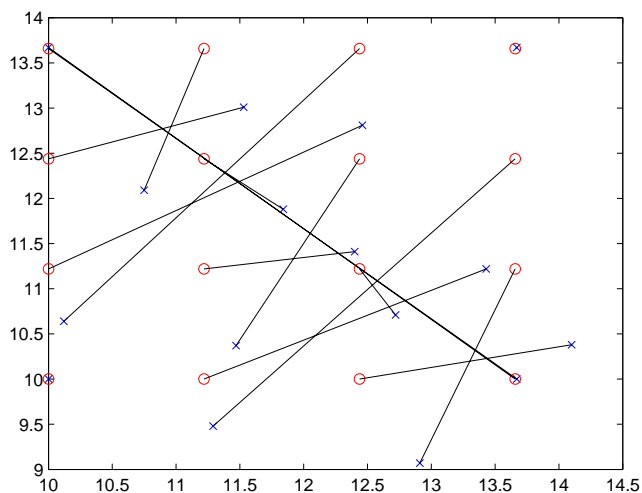
Figure 9. The red circles represent actual location of the motes, the x's represent the estimated locations, and the black lines are the errors.
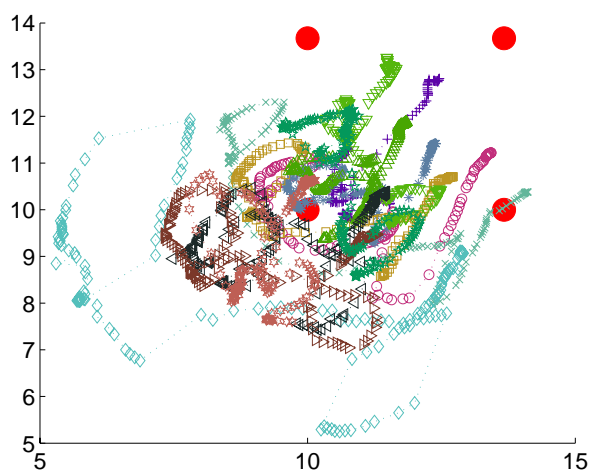
Figure 10. Convergence of the network nodes after several minutes

2.A.1.d. Heterogeneous Sensor Scheduling and Allocation

How does one make the best use of available resources? This central question or theme reverberates throughout virtually all Operations Research type problems; in particular, the following one from the ISP-II program. Given the limited nature of the resources at hand (*e.g.*, the number, type and mobility of sensors; communications bandwidth; and sensor computational, memory and power specs) and the tasks to be performed (*e.g.*, target location and motion estimation; sensor scheduling, allocation and movement), one must find optimal (or, more realistically) sub-optimal yet still acceptable solutions. Pragmatic approaches to this problem require that the sensing and computational demands be jointly engineered, to enhance the results of this and all other downstream phases of the analysis.

Heterogeneous sensor scheduling implies the use of more than one type of sensor. The two types of sensor considered below are (1) "event-timing" sensors (labeled with prefix "S") that accurately record the arrival time of an impulse signal and (2) sensors that measure the bearing angle of the impulse signal (prefix "A").
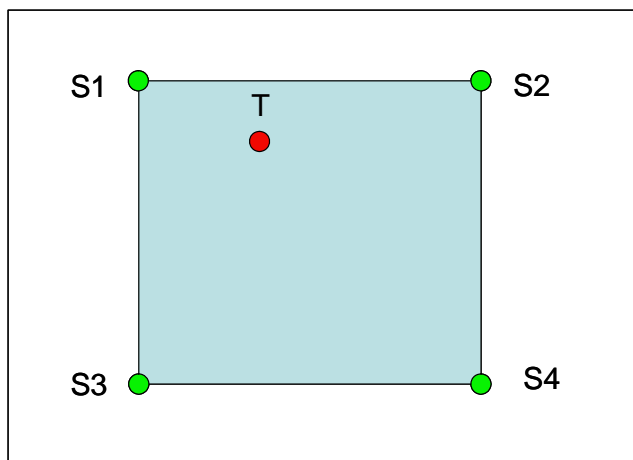
Figure 11: Blue-colored square region defined by four green-colored "arrival time" sensors (S1, S2, S3, S4) that measure an impulse sent out by red-colored target (T).

We look first at a very particular simpler sub-problem: determining the optimal use of measurements from a set of four TDOA sensors arranged in a square, with the source located within the square. The geometry of this layout is depicted in Figure 11. Due to the multiple symmetries of the square region (about vertical, horizontal, and diagonal lines, without any loss of generality the location of the target may, for simplified analysis purposes, be restricted to the triangular region shown in Figure 12.
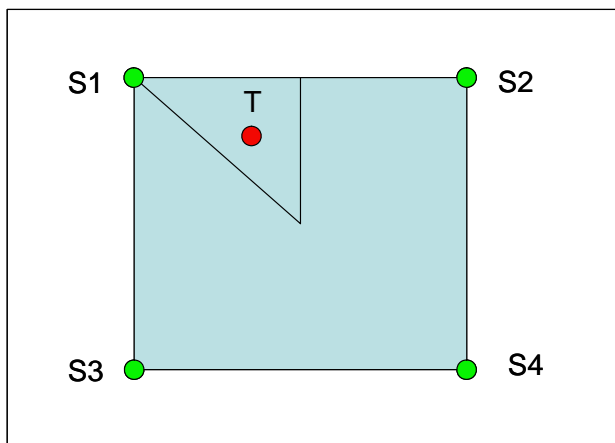
Figure 12: Blue-colored square region defined by four green-colored "arrival time" sensors (S1, S2, S3, S4) that measure an impulse sent out by red-colored target (T). Here, target (T) is assumed to lie within the triangle whose interior vertex is located at the center of the square. The triangle's area is one-eighth the square's area

Four sensors are used instead of only three sensors in order to remove potential ambiguities that arise when using only three TDOA sensors. This "3-sensor" scenario is illustrated in Figure 13. It shows the red, green, and blue "noise-added" hyperbolic

envelopes that are associated with the three possible ways of choosing two focal points (that define a hyperbola) from a set of three possible focal points. As shown in the figure, there are *two* solution regions associated with intersections of any of the three hyperbolas. This is the key reason why, in general, *four* TDOA sensors are required in order to accurately identify the location of a target emitter in the plane. Also, by using this approach (of intersecting hyperbolas), one may obtain a good measure of the geo-location error by counting the number of pixels within the black intersection regions of the red, green, and blue hyperbolas. In this figure, it is obvious that the location of the left intersection region is known with greater accuracy than the right intersection region.

Returning to the scenario of Figure 12, there are 16 different combinations of three pairs of TDOAs that utilize all four of the sensor time arrival measurements. These are displayed in Table 1. In the vast majority of published multi-lateration studies that use the TDOA method, a "fan-type" selection of TDOA measurements is used for geo-location. One special sensor is considered to be the basis point for time measurements and this time is subtracted from the time measurements of all other sensors. For example, if the special basis sensor is *S1*, then the "fan-type" TDOA measurements are given by the delta-times corresponding to: *S2-S1*, *S3-S1*, and *S4-S1*. This scenario is equivalent to Combination #1 in Table 1. The other three possible "fan-type" setups in Table 1 are Combination #6 (with basis sensor *S2*), Combination #12 (with basis sensor *S3*), and Combination #16 (with basis sensor *S4*).
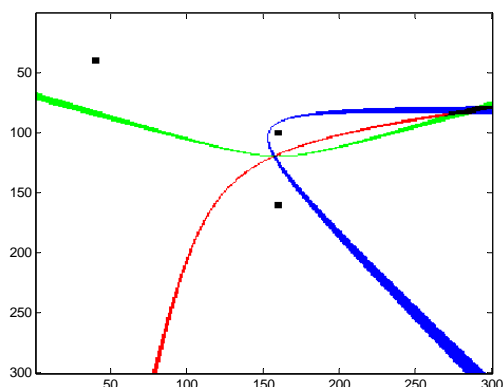
Figure 13: This shows the three "noisy" hyperbolic envelopes generated by the three TDOA sensors portrayed by the three black squares located at the upper-left (UL), the upper-middle (UM), and the lower-middle (LM). The green hyperbolic envelope is the hyperbola generated by focal points (UM) and (LM), the red hyperbolic envelope is generated by focal points (UL) and (LM), and the blue hyperbolic envelope is generated by focal points (UL) and (UM). A fixed error range corresponding to time arrival uncertainty is added to each of the TDOA measurements; note how the error envelope grows wider as one proceeds further out on the hyperbolic branches. (Slight differences in the horizontal and vertical scaling in the figure have introduced apparent distortions in the hyperbolas.)

Table 1: This table displays the sixteen combinations of three hyperbolas that use all four sensor measurements at least once. Combinations #1, #6, #12, and #16 are the "fan-type" sets of TDOA measurements that have been historically most commonly used

Combination	Hyperbola 1 Focal Pts.	Hyperbola 2 Focal Pts.	Hyperbola 3 Focal Pts.
#1	S1 & S2	S1 & S3	S1 & S4
#2	S1 & S2	S1 & S3	S2 & S4
#3	S1 & S2	S1 & S3	S3 & S4
#4	S1 & S2	S1 & S4	S2 & S3
#5	S1 & S2	S1 & S4	S3 & S4
#6	S1 & S2	S2 & S3	S2 & S4
#7	S1 & S2	S2 & S3	S3 & S4
#8	S1 & S2	S2 & S4	S3 & S4
#9	S1 & S3	S1 & S4	S2 & S3
#10	S1 & S3	S1 & S4	S2 & S4
#11	S1 & S3	S2 & S3	S2 & S4
#12	S1 & S3	S2 & S3	S3 & S4
#13	S1 & S3	S2 & S4	S3 & S4
#14	S1 & S4	S2 & S3	S2 & S4
#15	S1 & S4	S2 & S3	S3 & S4
#16	S1 & S4	S2 & S4	S3 & S4

Using MATLAB, a study was made of the performance of each of the above 16 combinations of TDOA measurements for locations of the target (T) within the triangular region shown in Figure 12. The figure of merit was related to the size of the intersection region that bounded the probable location of the target. This size, given by the number of pixels that occurred simultaneously in all three hyperbolic envelope regions, was deemed to be better, the smaller it was.

Somewhat surprisingly, a single combination (*viz.*, #4) of TDOA sensors produced the tightest localization boundary around the target T over almost the entire triangular region. Quite often the improvement over the "fan-type" combination was substantial, amounting to a localization area (in pixels) that was half of the "fan-type" result. The only exception to the superior performance by combination #4 was when target T was located quite close to sensor $S1$ (approximately within a distance corresponding to the input noise modeling factor associated with the thickness of the hyperbolic envelopes). When T was within tight proximity of $S1$, an assortment of several TDOA combinations produced the best (although substantially degraded) results: combinations #1, #2, #3, #7, #8, and #14.

Figure 14 shows the TDOA sensors that are affiliated with combination #4, the one that gives the best overall localization performance on target T. Inspecting the figure, the overall symmetry affiliated with this choice is evident and one can readily justify with physical reasoning why this choice is the best one.
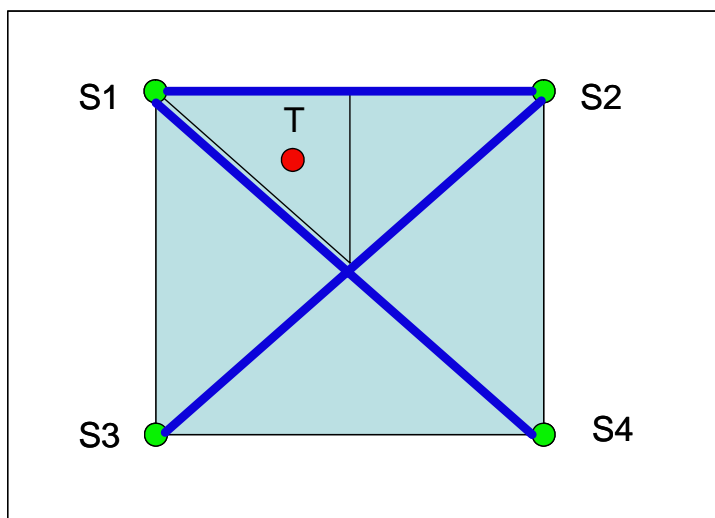
Figure 14: The diagram shows combination #4 of the TDOA sensors; the wide blue lines represent TDOA measurements between sensors S1 & S2, S1 & S4, and S2 & S3

Based on the above results, a heuristic approach to sensor scheduling and allocation presents itself quite naturally. Spread the sensors of type *S* (for TDOA analyses) in a square lattice array. In the initialization, or first measurements step, select four sensors at the corners of the largest available square. Any combination of two diagonals plus a side edge is adequate at this stage. Then perform a "hyperbolic envelopes" analysis; this should yield a smaller square which more tightly bounds the target's location and allows a better "square's edge" TDOA measurement selection to go with the two diagonal TDOA measurements. Repeat this procedure to get an initial estimate of the target's motion. At this stage, the "Unscented Kalman Filter" algorithm can be incorporated for following the target's motion. Based on predictions in this stage of the algorithm, an appropriate set of sensors can be queued up for the next set of measurements; also, instructions can be given to move an appropriate subset or "square lattice" constellation of sensors to best maintain precise and efficient tracking of the target. This modeling work will be done next.

Distribution Statement: Approved for public release; distribution is unlimited.
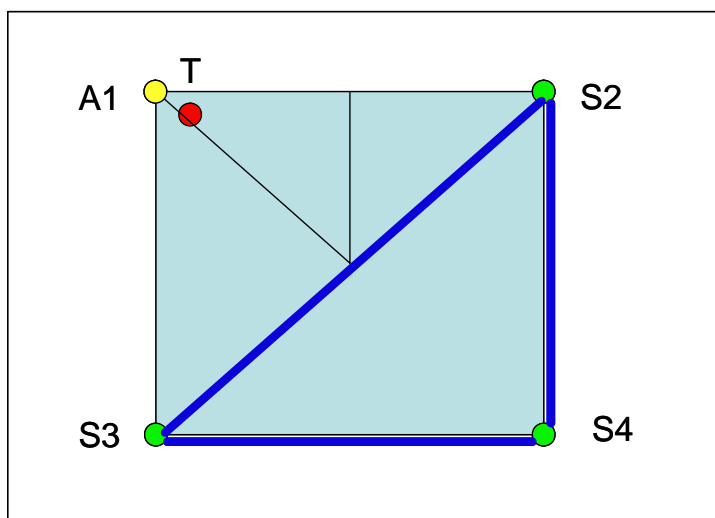
Figure 15: When target T is close to sensor 1, a combination of using TDOA (S2 & S3, S2 & S4, S3 & S4) plus a single AOA (A1) measurement works best

In the above procedure, whenever the target approaches a sensor location, it would be highly useful if one had dual measurement capability (*i.e.*, TDOA & AOA) sensors; then, for the sensor very close to the target, an AOA measurement would help to localize the target more efficiently (as in Figure 15). Although the above outlined procedure is essentially a greedy algorithm (*i.e.*, myopic) and only looks one time step ahead, its performance should be evaluated since it may be "good enough" and obviate the need for non-myopic algorithms (looking several time steps ahead) which entails considerably greater computational requirements.

Future considerations and algorithm enhancements include: multiple targets and the "de-interleaving" problem, multi-path signals, generalizing from 2-D to 3-D, and incorporating "elliptical" error bars from the intersecting hyperbolas into the UKF algorithm.

2.A.2. Melbourne Technical Progress

2.A.2.a. Multiple target tracking using motes

In previous work an algorithm based on the unscented Kalman filter has been developed for detection and tracking of multiple maneuvering targets using measurements from a collection of motes. A paper describing the algorithm and giving simulation results has been prepared for submission to the International Conference of Information Fusion. Matlab code has been prepared for the demonstration with real data and sent to Raytheon.

2.A.2.b. Particle filters for EKV tracking

The EKV tracking problem involves guiding a ballistic missile to the intended target in the presence of spurious targets, or decoys. This involves recursive state estimation to find the positions of the target and decoys, classification to distinguish the target from the decoys and stochastic scheduling to position the missile. In our approach joint classification and state estimation is performed using a particle filter. Positioning of the missile is achieved by selecting at each time step, from a finite number of candidates, the target move which minimizes a suitable criterion. The criterion used is the expected length of the line which projects the target position onto the direction of travel of the missile. The algorithm has been implemented and successfully tested using simulations.

2.A.2.c. Emitter Geolocation with Two UAVs

This subsection describes research investigating resource allocation with geolocation of emitters, using passive sensors located on airborne vehicles (UAVs or manned aircraft, in further text the generic term UAVs is being used). The chosen method of geolocation is the time difference of arrival (TDOA) of emitter pulses to individual sensors. Possible resources to allocate include

- Number of UAVs
- Communication Resources
- Computational Resources
- Choice of received pulses to process
- UAV Trajectories

Distribution Statement: Approved for public release; distribution is unlimited.

- etc

The resource allocation problems increase dramatically with the number of sensors, especially in the situation of limited communication bandwidth and limited computational resources available onboard the vehicles. A large number of emitter pulses arrive to each sensor, with a certain probability of detection. The number of possible combinations of received pulses, one per sensor, which need to be investigated to determine a correct combination, grows combinatorially with the number of the sensors. Required communication resource (transmission bandwidth) grows linearly with the number of sensors.

The approach investigated in published literature is to use at least three, preferably four sensors simultaneously for emitter location triangulation. To reduce complexity, previous research phase concentrates on emitter geolocation when only two passive sensors on two UAVs are deployed, without data association issues. The algorithm used is relatively simple, and can be implemented on one of the UAVs, called here the processor UAV. Only one communication stream, from the other UAV to the sensor UAV is required. This work resulted in optimization of the required communication resources, as well as the optimization of the required number of UAVs.

In this phase the data association issue was tackled. We assume that the emitter transmits pulses, so that multiple pulse combinations are feasible. One pulse transmitter position uncertainty in this case is shown in Figure 16, where the line widths indicate a one σ uncertainty.
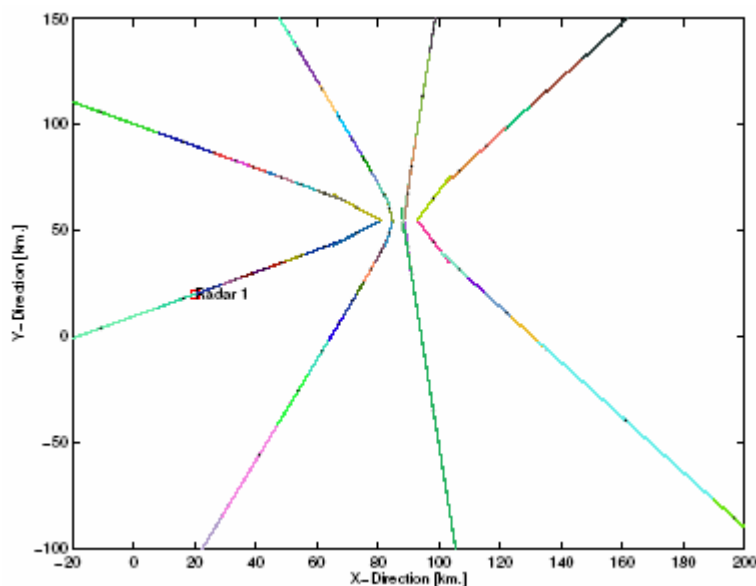
Figure 16: Two UAVs, four pulse / measurement uncertainty

As previously, the uncertainty (four hyperbolae) is represented as a Gaussian sum. Track state is represented also as a Gaussian sum, where each component of the sum represents a possible emitter location uncertainty based on used measurement component sequence

After 10 batches of 4 pulses have been processed, with all measurement batches arriving within 150 ms, all remaining track components are shown on Figure 17. After a

23

delay of 2.5 seconds with no measurements arriving, another 10 measurement batches have arrived, and all remaining track components are shown in Figure 18. The width of the area showing track components on Figure 18 indicate the one σ uncertainty. This work has minimized Data Association issues.
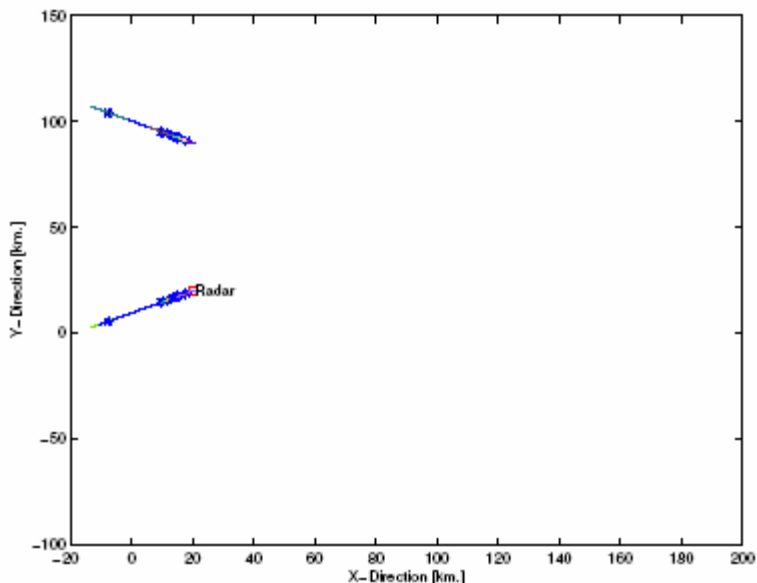
Figure 17: Remaining track components after 10 measurement (150 ms)
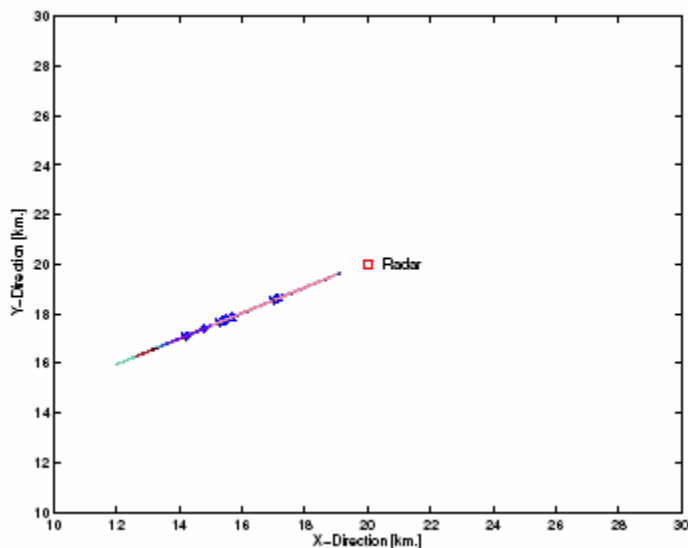
Figure 18: Remaining track components after 20 measurement batches (2.5s)

After first 10 batches of measurements have been processed, possible emitter location uncertainty is greatly reduced. It is now possible to choose another sensor UAV at the optimal location (to be determined), to aid in further speedy decrease of the target emitter uncertainty. TDOA measurement between the additional sensor UAV and the processor UAV will create another uncertainty hyperbola, with the emitter location being estimated at the intersection between the new hyperbola and the remnant uncertainty

depicted in Figure 18. Precise localization of the emitter is expected within a small number of seconds using this approach. Due to already roughly localized emitter, the additional UAV will not significantly increase computational, communication and data association usage. This is the topic of the next phase.

2.A.2.d. Scheduling for Distributed Passive Sensors

This work investigated sensor scheduling for target tracking with a network of distributed passive sensors.

Two aspects of this problem were investigated. These were:

• Scheduling sensor activation to minimize bandwidth and battery usage while maintaining track accuracy; and

• Sensor mode selection when tracking with a network of heterogeneous sensors.

In the first case, all sensors were simple acoustic, binary proximity sensors. That is, they would return a single bit of information if a sufficiently strong target was in their sensing range. In the second case, the sensors were assumed to be a mix of acoustic sensors and passive infra-red (PIR) sensors. The PIR sensors report detections in a quadrant and hence provide more information than the simple acoustic sensors.

In both cases the basic tracking algorithm employed was based on the Unscented Kalman Filter (UKF). A particle filter is normally employed in such problems however, as was reported previously the UKF outperforms a particle filter when tracking with binary proximity sensors. In addition, the computational costs are far lower and hence more suitable for distributed processing. Details of this study are presented in [Morelande 2006].

Scheduling Algorithms

A total of four scheduling algorithms were investigated. However, not all were used each of the two cases. Details of which algorithms were used in which case are given below. The first two scheduling methods were reported on previously. The two cost-based algorithms are new.

In Range Scheduler

This algorithm is based on the one used in [Stevens&Morrell 2003]. With this method, the all motes within a given range of the predicted target location are activated. That is, these motes return a measurement while all other motes in the region are switched off. The range around the predicted location is a multiple of the predicted position error.

MaxS Scheduler

This algorithm is a heuristic algorithm that is similar to the *In Range* scheduler. With this algorithm a fixed number, *maxS*, of motes are activated on each scan. The activated motes are those that are closest to the predicted target location.

Cost-Based Scheduler

This algorithm seeks to maximize track accuracy while simultaneously minimizing bandwidth usage. It makes use of the fact that the UKF filtered state error covariance, $\mathbf{P}_{k|k}$, at the end of scan k is not a function of the sensor measurements during

that scan. Thus it can be computed at the start of the scan. The following cost function was used to quantify the tracking accuracy improvement per bit

$$V(J_k) = \frac{tr\left(P_{k|k}\left(J^0\right)\right) - tr\left(P_{k|k}\left(J_k\right)\right)}{nbits(J_k)} \qquad (1)$$

where

- J_k is a set of sensors to activate at time k
- J^0 is the set of all inactive sensors
- nbits(J_k) is the number of bits needed to transmit all measurements from the set of active sensors J_k

This scheduler then selects the set of motes J_k that minimizes this cost without exceeding a upper limit on the available bandwidth of *Bmax*. Due to the properties of the UKF, the numerator of the cost function is guaranteed to be positive semidefinite. However, the cost does not change monotonically with the number of activated sensors.

TCost-Based Scheduler
An additional cost-based method was also considered. Here, the aim was to maximize track accuracy within a given bandwidth allowance. With this method the set of motes to activate at time k were those that minimized the cost

$$V(J_k) = tr\left(P_{k|k}\left(J^0\right)\right) - tr\left(P_{k|k}\left(J_k\right)\right) \qquad (2)$$

Test Scenario
The basic test scenario was identical in all the simulation studies carried out. The parameters are based on those used in [Morelande 2006] and [Stevens&Morrell 2003]. The surveillance region was $1000m2$ with sensors placed randomly within it. The sensing range of both the acoustic and PIR sensors was identical and set to $75m$. The probability of target detection when the target was on top of the sensor was unity and dropped off with the square of the distance between the target and sensor. The probability of a sensor reporting a false alarm was set to 0.001.

As the focus of this work was on the effect of sensor scheduling, only a single target was present in the region. It moved through the region with near-constant velocity for 25 scans, with measurements taken every second. The target SNR was 20dB. One hundred Monte Carlo simulations were performed for each study.

Sensor Activation
For these simulations, all motes were equipped with acoustic, binary proximity sensors. One hundred such motes were randomly distributed throughout the surveillance region. The first set of simulations compared the performance of all four scheduling methods with a baseline scenario when all motes were active. The parameters for each scheduling method were:

In Range - all motes within three times the predicted position error were activated;
MaxS - the four motes closest to the predicted target location were activated, i.e. *maxS* = 4; and
Cost-Based - for both methods the best set of motes were activated, given a maximum bandwidth allowance of *Bmax* = 4.

The tracking performance is illustrated in Figure 19 and Figure 20. Note, the *Baseline* tracker, which has all motes active on every scan, does not necessarily provide the best performance. This is due to false detections.
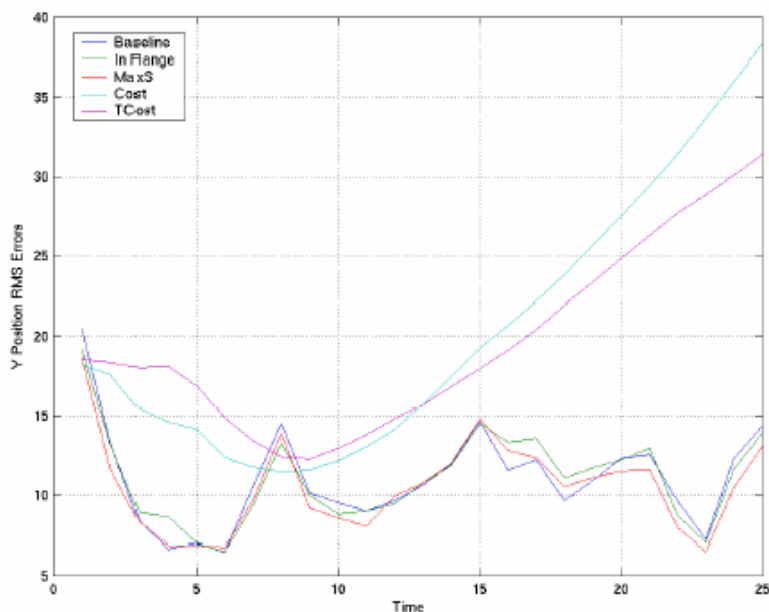
Figure 19: Y position RMS errors over time for each scheduling method. All motes have acoustic sensors
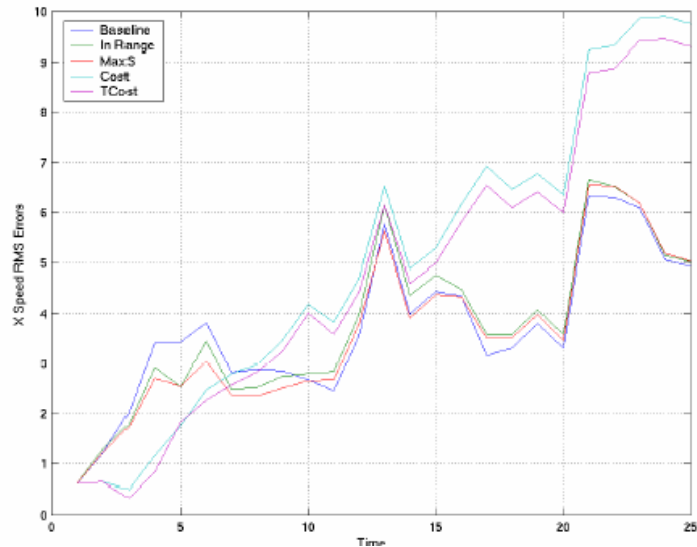
Figure 20: X speed RMS errors over time for each scheduling method. All motes have acoustic sensors

From these figures it can be seen that the two cost-based methods actually provide poorer performance than the two heuristic scheduling algorithms. The reason for this appears to be that the cost-based methods are too parsimonious with bandwidth. This is illustrated in Figure 21 which shows the average number of motes activated on each scan by each of the four scheduling methods. The two cost-based methods typically only

27

activated one mote per scan. In addition, often this mote was not the one closest to the predicted target location. The reason for this counter-intuitive behavior is the subject of current investigations.
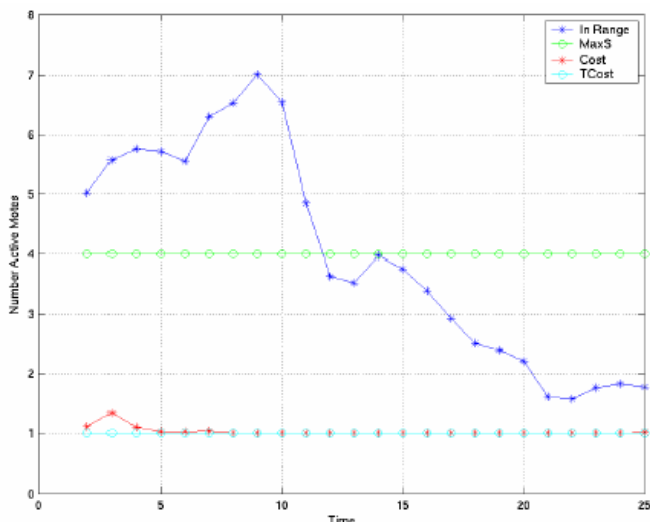
Figure 21: Average number of active motes over time. All motes have acoustic sensors.

An additional point to note is the drop off in the number of motes activated by the *In Range* scheduling method. As the track progresses the volume of the state error covariance matrix reduces. Consequently, fewer motes are activated as the track progresses. On occasion, this can lead to track loss if the tracker is seduced from the true target trajectory by false alarms. The *MaxS* scheduling method is the least susceptible to false alarms and, as a result, provides the best performance. This method also had the lowest computational cost. Both the *In Range* and *MaxS* trackers took approximately one fifth the runtime of the *Baseline* tracker, while the two cost-based methods had a runtime that was almost 2.5 times longer than the *Baseline*.

The second set of simulations was designed to investigate the effect of increasing the bandwidth on tracking accuracy. Only the two cost-based methods and the *MaxS* method were used, along with the standard baseline. In these simulations the available bandwidth was increased from 1 bit to 8 bits.

Figure 22 shows the X position RMS error averaged over time for each method, with the *Baseline* for comparison. As expected, the two-cost based trackers have identical performance to the *MaxS* method when only one mote can be activated at each scan. The performance of the *MaxS* tracker thresholds at approximately 4 bits of bandwidth. In contrast, the performance of the two cost-based methods actually worsens as the amount of available bandwidth increases. As shown above, these two methods typically do not use all the available bandwidth but select a single mote instead to activate. As the available bandwidth increases, the set of motes that can potentially be activated increases. The two cost-based methods most often select a mote that is not the closest to the predicted target location, which generally results in no detection being recorded.
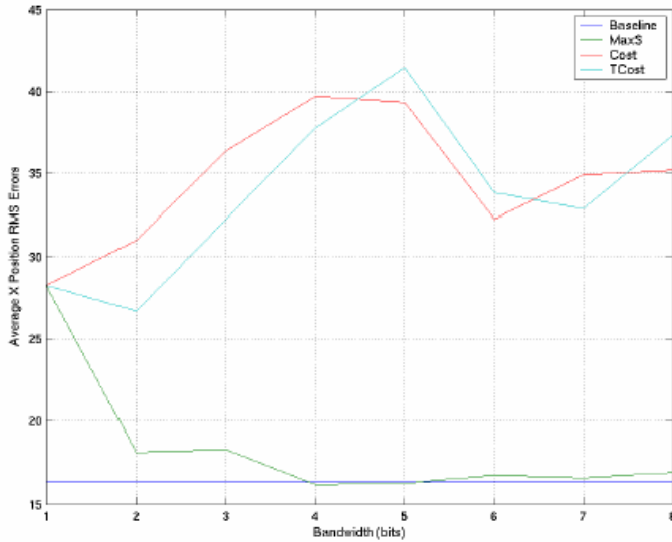
Figure 22: Average X position RMS as a function of bandwidth for bandwidth-limited scheduling methods. All motes have acoustic sensors.

Sensor Mode Selection

PIR sensors require four times the bandwidth of binary proximity sensors as they must report a detection or non-detection in each quadrant. However, they provide more information. These simulations were designed to quantify any performance improvement that could be obtained from using PIR sensors. The first set of tests had 50 acoustic sensors randomly located in the surveillance region. The number of PIR sensors was increased from 1 to 50, with 100 Monte Carlo runs performed at each step. Two tracking methods were compared for each set of sensors — the baseline of all motes active and the *MaxS* method. This scheduling method was used as it is computationally cheap while minimizing the effect of false alarms on track accuracy.

The performance change as the number of PIR sensors increases is illustrated in Figure 18. This shows the average Y position RMS over time as a function of the number of PIR sensors. The tracking performance when there are only 50 acoustic sensors and no PIR sensors is also shown as the *Baseline – All Acoustic* case. From this figure it can be seen that the extra information provided by the PIR sensors actually worsens performance, rather than improves it. Again, the reason for this unexpected outcome is being investigated.

Distribution Statement: Approved for public release; distribution is unlimited.
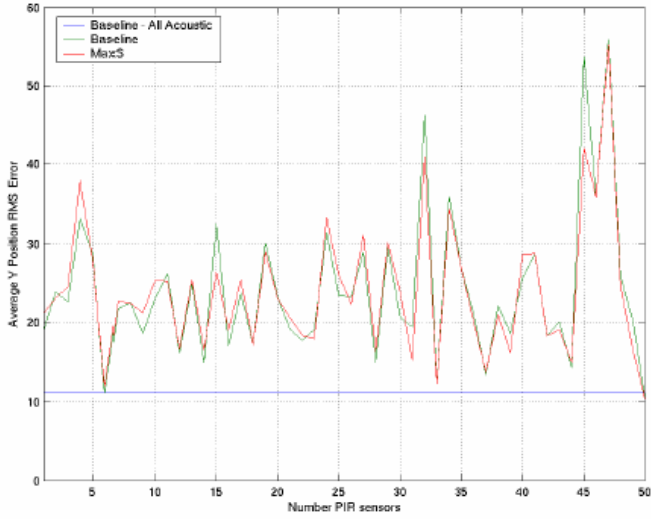
Figure 23: Average Y position RMS errors as a function of the number of PIR sensors.
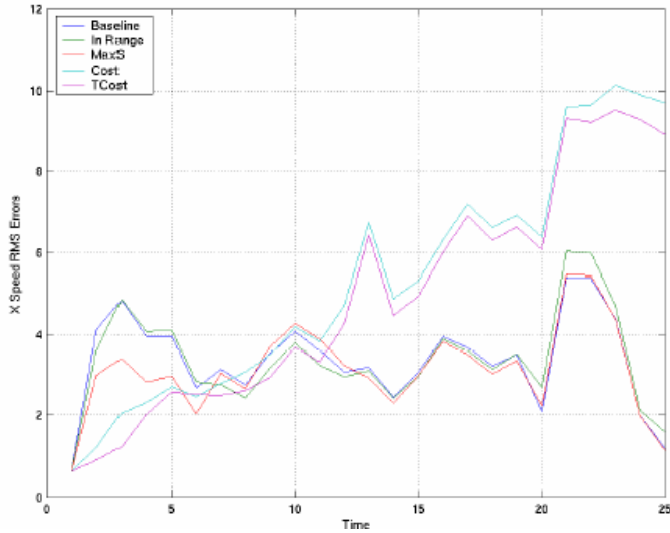
Figure 24: X speed RMS errors over time for each scheduling method. Half the motes have acoustic sensors, the rest PIR sensors.

The final set of simulations used a network of 50 acoustic sensors and 50 PIR sensors. All four scheduling methods were used, along with the baseline scenario. The parameters for the scheduling methods were the same as those used in the first set of simulations. The results of these simulations are shown in Figure 24 and Figure 25. These show similar behavior to the results when all the motes are equipped with acoustic sensors

Distribution Statement: Approved for public release; distribution is unlimited.
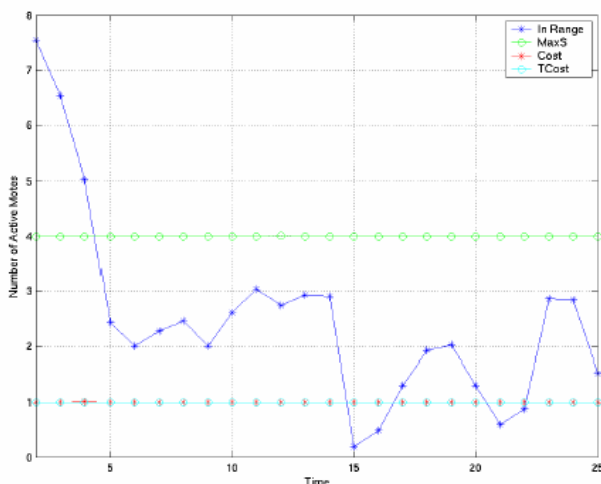
Figure 25: Average number of active motes over time. Half the motes have acoustic sensors, the rest PIR sensors.

Conclusions

These preliminary results suggest that the most robust and cost-efficient method for tracking with passive sensors is to make use of all the available bandwidth, i.e. use the *MaxS* scheduling method. However, the amount of bandwidth required for reasonable tracking accuracy does not seem to be high. An allowance of approximately 4 bits per track is adequate for this scenario.

2.A.2.e. Distributed Algorithms on Motes

Here we examined the problem of double-counting of data when performing distributed estimation with a sensor network. This is also sometimes termed *data incest*. The re-use of the same piece of information during data fusion can result in false levels of confidence and bias in the estimates. This problem is particularly acute in wireless sensor networks as network connectivity is not fixed and bandwidth is often limited.

A number of results have been obtained and have been submitted for publication. The first considers the idealized case of localization of a single, static source. Here it is assumed that the network is fully connected; any communication delays between network nodes are bounded and the bandwidth is not limited. An optimal algorithm for eliminating double-counting has been developed and submitted to *IEEE on Signal Transactions Processing* [McLaughlin_2 2007]. A companion paper considers the case of a dynamic source or target [McLaughlin_1 2007].

While the two previous results provide optimal algorithms for distributed data fusion, the conditions for optimality may often be breached in practical systems. Using the insight gained from the previous work, three novel, sub-optimal algorithms for data fusion in distributed networks have been developed. These algorithms are designed for use in wireless sensor network applications. Simulation studies show that these algorithms eliminate a significant proportion of the error due to data incest. An outline of these results has been submitted to the *10th International Conference on Information Fusion* to be held in Quebec City, Canada in July 2007 [McLaughlin_3 2007].

Distribution Statement: Approved for public release; distribution is unlimited.

2.A.2.f. Raytheon On-Site Support

Work during the previous quarter was focused on two main areas:

1. Optimal scheduling of Gauss-Markov systems
2. Cancer prognosis prediction via random projections

Additionally, joint work has been done for the integration of the UKF-based mote tracker for the demonstration. This has involved interfacing the UniMelb tracking routine with the Raytheon tracking GUI. A new version, updated to utilize time-stamped data, has been submitted to Raytheon.

Optimal Scheduling of Gauss-Markov Systems

We have been further investigating Gauss-Markov Systems (GMS), as identified in previous reports. During the past quarter, we have been investigating formulations of GMS for which we have provably optimal schedules. These are presented in Table 2; see also [SavLaSMor] and [HowSuvMor]. These results are being prepared for publication in the next few months.

Table 2: Summary of scheduling algorithms with optimality proofs

# Systems	# Measurements	Cost Type	f-value	Comments
1	N	Terminal	Nonzero	In preparation
$T < \infty$	$n_i \leq 1$	Terminal	1	[SavLaSMor]
$T < \infty$	$\sum_i n_i \leq B$	Terminal	Nonzero	In preparation
1	1	Cumulative	Nonzero	In preparation
2	N	Cumulative	1	[HowSuvMor]

Cancer Prognosis Prediction via Random Projections

As previously mentioned, through collaborations with National Information Communication Technology of Australia (NICTA), we have obtained genomic data for purposes of determining breast cancer prognosis. This type of data is an extreme case of sample starvation; we hope to gain some insight into problems such as the ballistic missile defense Kill Vehicle (KV) discrimination problem. In the BMD KV discrimination problem, the available collected training data is extremely sparse and the possible physical feature number on the order of tens. A plot of the log of gene concentrations is presented in Figure 26. The first task was to confirm that random projections (RP) would preserve distances on these data. The results for epsilon, defined by equation 1, are shown in Figure 27, considering all 97-choose-2 pairs.

Equation 1: Definition of the squared distance ratio. The matrix P is a random projection, while x and y are two non-identical gene vectors.

$$\varepsilon(x, y) = \left(\frac{\|P(x - y)\|}{\|x - y\|} \right)^2$$

Note that projection onto a larger number of dimensions results in tighter bounds of epsilon. From this, we are investigating adaptively selecting the number of dimensions to maintain epsilon and *still be useful for prognosis*. The aim here is to classify gene

concentrations into "good" or "bad" prognosis, *not* perform signal reconstruction. Based upon the results of Figure 27, it appears that the impact of taking a RP with "enough" dimension adds an amount of approximately Gaussian noise. The aim is to adaptively select a noise value that is still sufficient to perform classification in a Bayesian setting. The random projections have a benefit over more traditional feature selection algorithms in that the random projections have theoretical bounds as to distance preservation, while such bounds on traditional feature selection (*e.g.,* the best N genes) have no such bounds.
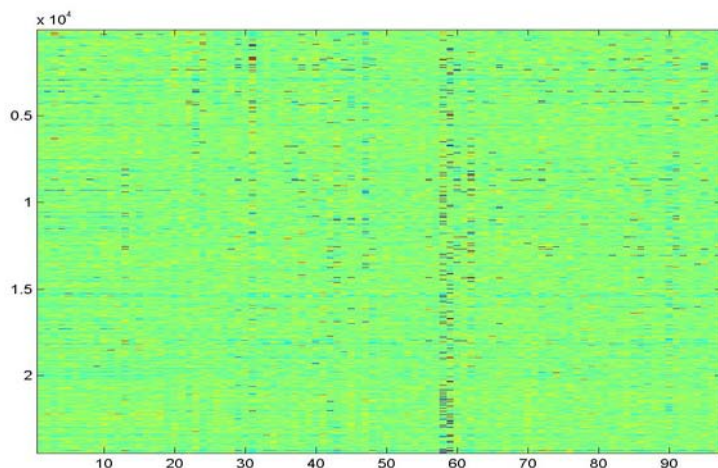
Figure 26: Plot of 24,481 selected genes for 97 patients
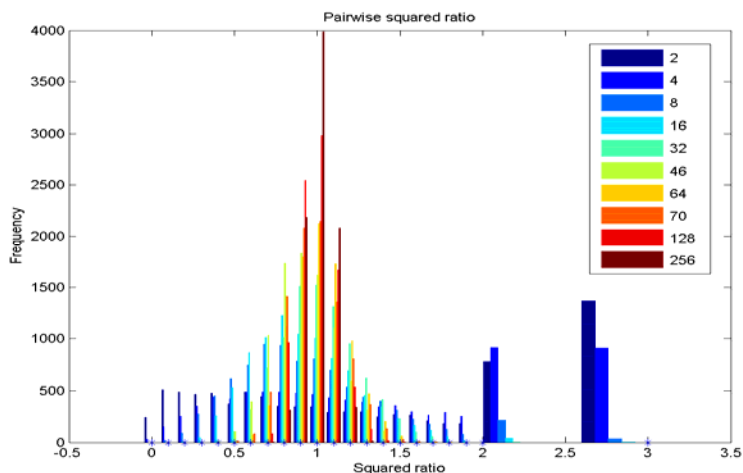
Figure 27: Squared distance ratio as a function of the projection dimension, as shown by the various colors. The original data are as shown in Figure 4, over all x, y pairs

2.A.3. ASU Technical Progress

Since the last progress report, we have modified the ASU person tracking algorithm to use the initial imager interface supplied by the Georgia Tech team. This interface is implemented as a MATLAB Application Program Interface (API). Currently the API is for a MATLAB emulator of the imager; the same API will be used to interface to the actual imager hardware. In addition to integrating the imager API into the tracker

Distribution Statement: Approved for public release; distribution is unlimited.

code, we have mapped the image filtering operations used by the tracker onto the imager hardware model.

The imager acquires a 256x256 image (pixel intensities only) and is capable of performing separable row/column-wise FIR filtering and separable row/column-wise block transforms on blocks of the acquired image. The imager has four programmable 8x8 matrices: $\mathbf{A}_1$ and $\mathbf{A}_2$ are used to perform column filtering/transforms and $\mathbf{B}_1$ and $\mathbf{B}_2$ are used to perform row filtering/transforms. In the imager simulation, these matrices, as well as each image block, are padded with zeros to create 16x16 matrices. Note that while the imager simulation internally uses 16x16 matrices, one can only obtain an 8x8 matrix from the imager simulation. This is not to trust explicitly in the discussion below.

The imager is capable of performing 1-D block transform, 2-D block transform, 1-D filtering and 2-D filtering. Note that 1-D row filtering/transform, followed by 1-D column filtering/transform (or vice-versa) performed sequentially, results in 2-D filtering. To perform a row or column filtering/transform, the appropriate $\mathbf{A}$ and $\mathbf{B}$ matrices are selected; in this report, we denote the selected column matrix as $\mathbf{A}$ and row matrix as $\mathbf{B}$.

The ASU tracker algorithm selects image blocks and correlates them with a Gaussian or a Mexican hat matrix; the correlation values for the selected blocks form the input to the tracker algorithm. Computing the correlation c requires component-by-component multiplication of each pixel block $\mathbf{P}$ with a Gaussian matrix $\mathbf{G}$ or a Mexican hat matrix $\mathbf{H}$, followed by a summation of the products.

Note to that the correlation c can be computed as

$$c = trace(\mathbf{PH}^\mathrm{T}).$$

To compute $\mathbf{PH}^\mathrm{T}$ using the imager, we load the $\mathbf{A}$ matrix with an identity and the $\mathbf{B}$ matrix with $\mathbf{H}^\mathrm{T}$. The trace operation must be performed in the tracker software.

The two-dimensional Gaussian matrix $\mathbf{G}$ can be de composed into a product of vectors as

$$\mathbf{G} = \mathbf{xx}^\mathrm{T}$$

where $\mathbf{x}$ is a Gaussian vector. In this case, the correlation c can be completely computed by the imager chip as follows. $\mathbf{A}$ is loaded with the FIR filter $\mathbf{x}$, $\mathbf{B}$ is loaded with the FIR filter $\mathbf{x}^\mathrm{T}$, and the resulting scalar c is the upper left element of $\mathbf{APB}^\mathrm{T}$.

The performance of the tracker/imager simulation has been evaluated using the same video sequences used to originally evaluate the tracker, and the performance is substantially the same as the original tracker performance.

2.A.4. Georgia Tech Technical Progress

Work continues on the imager simulator and associated API. This work is primarily concerned with testing and predicting algorithm performance for CADSP-type optical flow and compression algorithms. We have also continued to develop and test efficient digital algorithms for optical flow that complement the imager, focusing on a recursive least-squares filter architecture that is memory and computationally efficient and also robust to imager noise.

The CADSP Imager has demonstrated the ability to perform focal-plane processing. Images have been acquired and a DCT has been computed in the pixel plane.

To complete the full two-dimensional transformation, the backend vector-matrix multiplier (VMM) programming algorithm is being developed. For the demonstrated focal-plane DCT, it was configured as an identity matrix. We have successfully programmed the sense amplifiers which feed into the VMM and multiplier elements in the VMM. Developments now are concentrating on achieving target multiplications quickly and accurately so that arbitrary matrices may be programmed.

The image reading process remains slow, taking several minutes to extract an image. The communication from the system to the computer is now done using USB connections and the desired communications speeds have been achieved for video frame rates. Still remaining are unidentified system delays on the IC, requiring pauses for settling. Present component characterizations suggest faster times are possible, so the source of the delay is being tracked down. Once the discrepancy is resolved, which maybe a digital switching issue, faster reading should be possible.

We are packaging the hardware and software into a stand-alone system which is primarily composed of a computer, an FPGA board, an imager board, and an external current measurement device. The external measurement device is to be eliminated in the final version once system components are characterized. A new user is being trained to use the system; they will, in turn, help develop the training requirements and procedure for other new users to freely use the imaging system.

2.A.5. UM Technical Progress

In the last quarter, we have wrapped up our research in classification constrained dimensionality reduction; self localization; and geometric entropy minimization (GEM). The major advance that we report here is the development of an out-of-sample extension of classification constrained dimensionality reduction (CCDR) that includes simultaneous updates for both unlabeled and labeled data.

In our last review we presented results for the out-of-sample extension (OSE) for our CCDR that can only be applied to unlabeled data, *e.g.* test samples to be classified. This period we have concentrated on extension of CCDR to labeled as well as unlabeled data. An outline of the extension is given below.

Let $\{x_1, \quad x_2, \quad \ldots \quad x_n\}$ be high-dimensional samples, and $\{y_1, \quad y_2, \quad \ldots \quad y_n\}$ be their lower-dimensional (d-dimensional) embedding found by SVD. As usual, define Λ as the $d \times d$ diagonal matrix of the first d eigenvalues of the Graph Laplacian. For a new unlabeled data point, x_{n+1}, the out-of-sample-extension for an unlabeled point is

$$y_{n+1}(x_{n+1}) = \Lambda^{-1} \frac{\sum_{i=1}^{n} K(x_i, x_{n+1}) y_i}{\sum_{s=1}^{n} K(x_s, x_{n+1})} \tag{1}$$

For a new labeled data point x_{n+1} which belongs to class k we can show the modified out-of-sample-extension:

Distribution Statement: Approved for public release; distribution is unlimited.

$$y_{n+1} = \Lambda^{-1} \left(\beta \frac{\sum_{i=1}^{n} K(x_i, x_{n+1}) y_i}{\beta \sum_{s=1}^{n} K(x_s, x_{n+1}) + 1} + \frac{z_k}{\beta \sum_{s=1}^{n} K(x_s, x_{n+1}) + 1} \right) \tag{2}$$

where z_k is the centroid for class k, given by the SVD performed for the first n points, and β is the regularization parameter applied to the centroid points. A Matlab program is available to compute (1) and (2) and has been delivered to Raytheon for testing.

2.A.6. FMAH Technical Progress

In this last phase of the ISP Program we have integrated some new Signal Processing Algorithms from Raytheon into a Radar Hardware Demo. We have also created a simulated RF environment to test the basic capability of the Ent-To-End Design/Processing strategy based on ISP principles.

Experimental Results - Setup

The Radar System utilizes Chirp Waveforms at 1GHz, Analog Devices DDS card AD9854. The picture below sketches the position of four corner reflectors placed so that two lie at almost the exact same angle with respect to the radar antenna
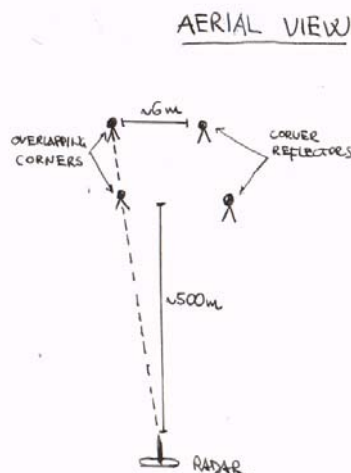
Figure 28: Waveform Experimental Set UP

Results - Detection

Abundant Experimental Evidence and extensive Simulations show that "Highly Textured Noise" (more formally called "Sparse") due to the presence of a large number of scattering surfaces, when superposed with hardware-induced non-linearities, heavily affect the performance of any deterministic receiver. We conclusively demonstrated in that ConvNets are capable of learning the extraordinary high number of variables, and thus they are natural candidates to break the complexity of a RF landscape data.

As the Receiver is training to detect corners we plot the "ROC" curves for the ConvNet: the Red Curve represents the performance of the ConvNet as it tests its performance just against corner reflectors:

$$D_t - p \cdot D_b > < 0$$

36
Distribution Statement: Approved for public release; distribution is unlimited.

Where the outputs of the target detector Dt and the one of the background Db are between -1 and 1 my knob "*a*" goes from 0 to 1. The Green curve represents the performance once the classifier is asked to decide on data containing returns from the multi-path signal of a pick-up truck parked sideways. The Following Picture illustrates the results:
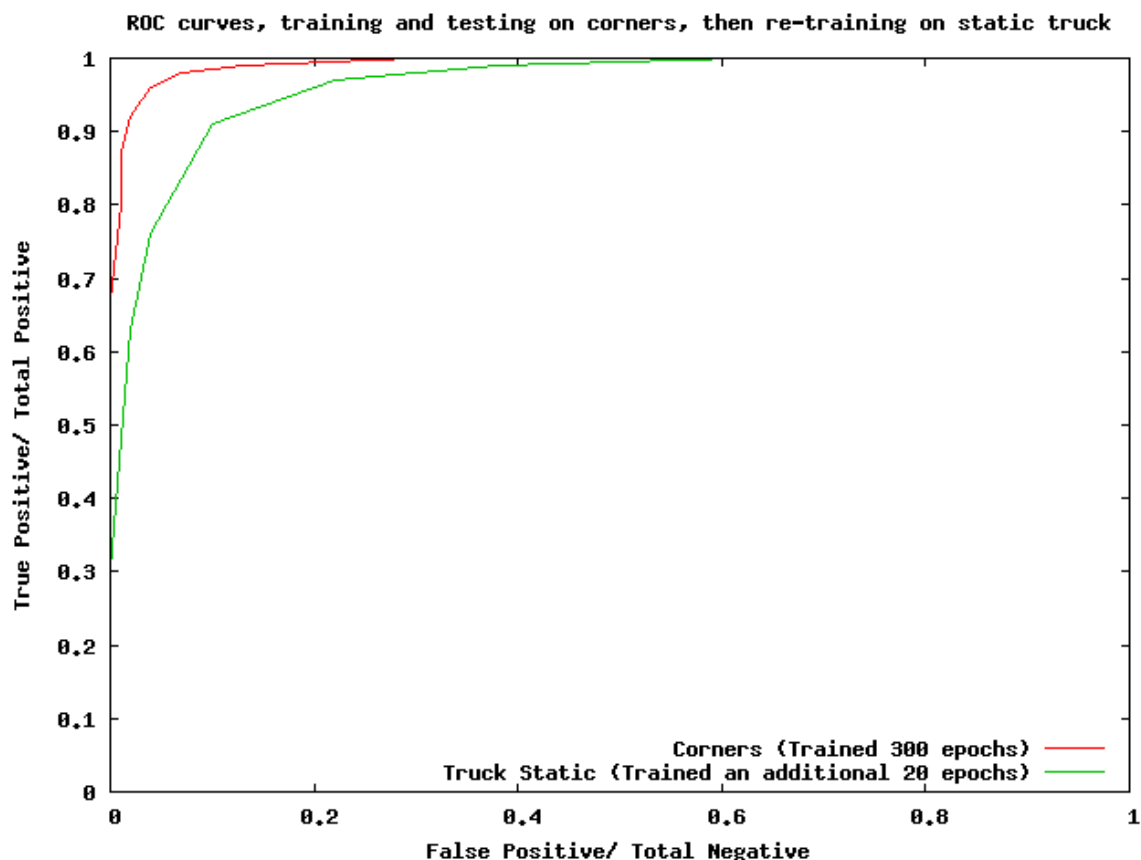

Figure 29: Waveform Processing Results

2. C. Conference Proceedings
There was one publication in conference proceedings during the current PoP.

1. "Integrated Sensing and Processing Phase II Demonstration Program Overview," H. Schmitt, D. Waagen, S. Bellofiore, N. Shah, R. Cramer, T. Stevens, C. Savage, and V. Berisha, Defense Applications of Signal Processing 2006, 10-14 December 2006, King Fisher Bay, Australia.

2. "Sparse Manifold Learning with Applications to SAR Image Classification," V. Berisha, N. Shah, D. Waagen, H. Schmitt, S. Bellofiore, A. Spanias, and D. Cochran, 32nd International Conference on Acoustics, Speech, and Signal Processing (ICASSP), Honolulu, HI, April 15-20, 2007, accepted.

2. D. Consultative and Advisor Functions
There was one consultative or advisory function that occurred during the current PoP. The first relates to a Raytheon Shooter Localization demonstration using the MICA-2/Z sensor nodes. This work is being funded under the DARPA IXO NEST Phase II

program. The Phase I shooter localization algorithms were developed by VU. Preliminary results indicated that the shooter localization algorithm has significant potential. The program was subsequently classified and was ultimately transitioned to Raytheon for demonstration and refinement under Phase II. The DARPA IXO Program Manager has given permission for several of these algorithms to be used in our program. The Raytheon NEST program has identified a critical need for the development of an accurate sensor localization algorithm that is scalable to hundreds or thousands of nodes. We have identified and are evaluating several promising mathematical approaches to sensor localization that will be made available to the Raytheon NEST program if they are successful. Thom Steven supports the DARPA ISP II and DARPA NEST programs, and, more generally, the two programs have developed a strong collaboration.

2. E. New Discoveries, Inventions or Patent Disclosures
There were no patent disclosures filed during the current PoP.

2. F. Honors/Awards
There were no honors or awards received during the current PoP.

2. G. Transitions.
There were no technology transitions achieved during the current PoP.

2. H. References
[HowSuvMor] Howard, S., Suvorova, S., and B. Moran, "Optimal policy for scheduling of Gauss-Markov systems," Proceedings of the 7th International Conference on Information Fusion, Stockholm, Sweden, June, 2004.

[SavLaSMor] Savage, C., La Scala, B., and B. Moran, "Optimal scheduling for state estimation using a terminal cost function," Proceedings of the 9th International Conference on Information Fusion, Florence, Italy, July, 2006.

[Burges 1998] Burges, C.J.C., "A Tutorial on Support Vector Machines for Pattern Recognition," *Data Mining and Knowledge Discovery*, **2**, pp. 121-167 (1998).

[DeVinney *et al.* 2002] DeVinney, J.G., Priebe, C.E., Marchette, D.J., and Socolinsky, D.A., "Random Walks and Catch Digraphs in Classification," *preprint*, 2002.

[Priebe *et al.* 2003] Priebe, C.E., Marchette, D.J., DeVinney, J.G. and Socolinsky, D.A., "Classification Using Class Cover Catch Digraphs," *Journal of Classification*, **20**, pp. 3-23, 2003.

[Ho *et al.* 2000] Ho, T.K. and Basu, M., "Measuring the Complexity of Classification Problems," *Proceedings of the 15th International Conference on Pattern Recognition*, **2**, pp. 43-47, 2000.

[Skolnick 2001]   Skolnick, *Introduction to Radar Systems*, McGraw-Hill, 2001.

[Hsu] Hsu, et. al., *Radar Terrain Bounce Jamming Detector*, software patent application, http://gauss.ffii.org/

[Tenenbaum 2000] Tenenbaum, et. al., "A Global Geometric Framework for Nonlinear Dimensionality Reduction," *Science*, v. 290, 2000.

Distribution Statement: Approved for public release; distribution is unlimited.

[Friedman-Rafsky 1979] Friedman, J. H. and Rafsky, L. C., "Multivariate Generalizations of the Wald-Wolfowitz and Smirnov Two-Sample Tests," *Annals of Statistics*, v. 7, n. 4, 1979.

[Henze-Penrose 1999] Henze, N. and Penrose M. D., "On the Multivariate Runs Test," *Annals of Statistics*, v. 27, n. 1, 1999.

[McLaughlin_1 2007] S. McLaughlin, "Optimal data incest management: Dynamic case. *IEEE Transactions on Signal Processing*," 2007, submitted for publication.

[McLaughlin_2 2007] S. McLaughlin, "Optimal data incest management: Static case. *IEEE Transactions on Signal Processing*," 2007, submitted for publication.

[McLaughlin_3 2007] S. McLaughlin, "Towards real-time track fusion for network centric warfare using wireless sensor networks," in *10th International Conference on Information Fusion*, Quebec, Canada, July 2007.

[Morelande 2006] M. M. Morelande, "Tracking multiple targets with a sensor network," in *9th International Conference on Information Fusion*, Florence, Italy, July 2006.

[Stevens&Morrell 2003] T. J. Stevens and D. Morrell, "Minimization of sensor usage for target tracking in a network of irregularly spaced sensors" in *IEEE Workshop on Statistical Signal Processing*, St Louis, MO, USA, September 2003.

2. I. Acronyms

ADTS	Advanced Detection Technology Sensor
ASU	Arizona State University
ATA	Automatic Target Acquisition
AVU	Algorithms Verification Units
CADSP	Cooperative Analog Digital Signal Processor
CCDR	Classification Constrained Dimensionality Reduction
CPI	Coherent Processing Interval
CRB	Cramér–Rao Bound
CROPS	Classification Reduction Optimal Policy Search
DARPA	Defense Advanced Research Projects Agency
DS	Danzig Selector
DSA	Distinct Sensing Area
dwMDS	Distributed, weighted, multi-dimensional scaling
FPA	Focal Plane Array
FMAH	Fast Mathematical Algorithms and Hardware
GEM	Geometric Entropy Maps
Georgia Tech	Georgia Institute of Technology
GMS	Gauss-Markov Systems
GPS	Global Positioning System
IASG	Independently Activated Sensor Group
ISP	Integrated Sensing and Processing
IXO	Information Exploitation Office
kNN	k-Nearest Neighbor
LEAN	Laplacian Eigenmap Adaptive Neighbor
LIP	Linear Integer Programming

M2M	Multipoint-to-multipoint
MC	Monte-Carlo
MTT	Multi-target tracking
NEST	Networked Embedded System Technology
NDA	Non-disclosure Agreement
NLIP	Nonlinear Integer Programming
NLOS	NetFires Non-Line of Sight
NUC	Non-Uniformity Compensation
ONR	Office of Naval Research
OSE	Out-of-sample extension
PAM	Precision Attack Munition
PDA	Probabilistic Data Association
PRI	Pulse Repetition Intervals
PWF	Polarization Whitening Filter
PoP	Period of Performance
RIM	Radio Interferometric Measurements
RIPS	Radio Interferometric Positioning
RISCO	Raytheon International Support Company
RSS	Received Signal Strength
TAA	Technical Assistance Agreement
TDOA	Time Difference of Arrival
TIM	Technical Interchange Meeting
UAV	Unmanned Aerial Vehicle
UCIR	Uncooled infrared imaging
UKF	Unscented Kalman filter
UM	University of Michigan
UniMelb	Melbourne University
VM	Virtual Measurement
VU	Vanderbilt University