

01 December 2005

Progress Report

CDRL A001 No. 03

**Third Quarterly Progress Report for Period of Performance
1 September 2005 – 30 November 2005**

**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
Professor Douglas Cochran/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.

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

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 21 NOV 2005		2. REPORT TYPE		3. DATES COVERED	
4. TITLE AND SUBTITLE Integrated Sensing Processor Phase 2 Third Quarterly Progress Report for Period of Performance			5a. CONTRACT NUMBER N00014-04-C-0437		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S) Harry Schmitt			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Raytheon Missile System,1151 Herman Rd ,Tucson,AZ,85706			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					
14. ABSTRACT 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.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT 1	18. NUMBER OF PAGES 36	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

TABLE OF CONTENTS

0. Technical Abstract	3
1.0. Management Overview and Summary	3
1. A. Program Summary	3
1. B. Program Status	3
1. C. Personnel Associated/Supported	3
1. D. Recent Accomplishments and Events	4
1. E. Near Term Events	5
2.0. Technical Progress and Accomplishments	6
2. A. Technical Progress	7
2.A.1. Raytheon Technical Progress	7
2.A.2. ASU Technical Progress	19
2.A.3. UM Technical Progress	20
2.A.4. FMAH Technical Progress	21
2.A.5. UniMelb Technical Progress	24
2.A.6. Georgia Tech Technical Progress	33
2. B. Publications	34
2. C. Conference Proceedings	34
2. D. Consultative and Advisor Functions	35
2. E. New Discoveries, Inventions or Patent Disclosures	36
2. F. Honors/Awards	36
2. G. Transitions	36
2. H. References	36
2. I. Acronyms	36

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 Program status can be summarized as “on track.” All of the negotiations have been completed and all of the subcontractors are now under subcontract. Raytheon is still running below its spending plan to better align with the subcontractor schedules; however, we expect to complete the contract on time and budget. As of 22 October 2005, 22% of contract funds had been expended with 31% of the program complete. The contract also reflects substantial under-runs due to delays in receiving the initial invoices from the all of the university subcontractors. We should recover from the spending profile deviation and expect to finish 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

ASU

Professor Darryl Morrell
Professor Antonia Papandreou-Suppappola

Georgia Tech

Professor David Anderson
Professor Paul Hasler

UniMelb

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

UM

Professor Al Hero

Significant Personnel Actions: There were two significant personnel changes during the current PoP. Mr. Thomas Stevens was added to the technical staff to support the distributed sensing and tracking demonstration. Thom is familiar with the MICA-2/Z sensors and supports Raytheon's DARPA NEST Shooter Localization program. Dr. Nicholas Coult was added to the FMAH technical staff to support their research on a diffusion based approach to high dimensional data processing.

1. D. Recent Accomplishments and Events

Two of the Algorithms Verification Units (AVU), comprised of the Crossbow wireless low-power sensor nodes and their associated sensors, needed for distributed processing demonstrations and evaluations were shipped to UniMelb on 15 August 2005. The remaining AVUs are available for distribution to other university personnel when needed. In support of the distributed processing demonstrations and evaluations, Raytheon (Harry Schmitt and Sal Bellofiore) spent 27 August 2005 to 4 September 2005 working with UniMelb personnel on the development and implementation of sensor scheduling and distributed tracking algorithms on the AVUs. Sal Bellofiore presented an overview of the mechanics of working with the AVUs and provided a CD to UniMelb with the AVU control and support software.

An amended Technical Assistance Agreement (TAA) was approved by the U.S. State Department on 6 October 2005. The amended TAA expands the technical scope to cover the research areas added under the ISP Phase II program as well as covering two additional UniMelb personnel, Darko Musicki and Sofia Suvorova, who are dual citizens of Serbia and Montenegro, and Russia, respectively. The amended TAA also covers Raytheon, Australia. The amended TAA is out for signature by the parties and will then be effective.

Other Accomplishments and Events:

- The annual Principal Investigator (PI) meeting was held in Washington, D.C. on 21-23 September 2005. Assigned Action Items are discussed in Section 1.E.

**ISP Phase II (Contract N00014-04-C-0437)
Quarterly Progress Report (CDRL A001 No. 3)**

- Acoustic Sensor characterization data (ROC Curves) were delivered by ASU to Raytheon and provided to UniMelb for tracker algorithm development. A more detailed technical discussion is provided in Section 2.A.2.
- A substantial set of publicly releasable IR imagery was provided by the US Army (Richard Sims) to Raytheon and has been distributed to FMAH for use in their diffusion algorithm research. A more detailed technical discussion is provided in Section 2.A.6. Raytheon has also provided this IR imagery to UM for evaluation.

1. E. Near Term Events

- A technical report on the mathematical foundation of their waveform family design is being developed by FMAH. This report is a contract deliverable and is due in conjunction with this Quarterly Report.
- Raytheon personnel (Waagen and Schmitt) will visit Georgia Tech 16 November to discuss current hardware and algorithm status.
- Raytheon personnel (Waagen, Schmitt and Samuel) will visit University of Maryland 15 November to discuss waveform design research being conducted by Professor John Benedetto.
- Dr. T.J. Klausutis (AFRL, Eglin) will visit Georgia Tech in November to discuss possible collaborative opportunities. In particular, there may be an opportunity to get the Georgia Tech CADSP imager included in AFRL test plans. AFRL, Eglin is particularly interested in learning more about the capabilities and maturity of the CADSP imager
- Several Action Items from the PI meeting were assigned and are being addressed. The Action Items that were assigned to the Raytheon Hub, as well as the responsible parties, are:

Notes and Actions from ISP II PI Meeting

1. [Georgia Tech] How is the CADSP Imager different from the ONR-funded CNN work? Should provide a chart specifically showing CNN *versus* CADSP technique.
2. [Raytheon/Georgia Tech] Need to define a test for the CADSP Imager that everyone (Raytheon ISP II Hub, ONR and DARPA DSO) agrees can be used to determine whether the Imager “works.” Specifically, we should address:
 - a. What are pass/fail metrics for the CADSP chip?
 - b. What constitutes good enough for the program to be a success?
 - c. DARPA DSO suggested a test consisting of a side-by-side comparison of JPEG compression using the CADSP and standard DSP. This was also earlier at our ONR Technical Interchange Meeting.
3. [Raytheon/Georgia Tech] Need to address how one “calibrates” the Georgia Tech CADSP. For example NUC:
 - a. Apply a uniform source and correct it for a non-uniform response of various grey levels.
 - b. Is this necessary to post correct the chip/process for it to be useful?

4. [Raytheon/FMAH] Need to more clearly address how FMAH implements the “ISP waveform post-processing.” Discussion and report should address the following concerns:
 - a. Provide a labeling of the axis on each chart. It was difficult to understand what was being measured.
 - b. If the targets were in fact stationary tetrahedrals (corner reflectors), what role did Doppler play in the measurements.
 - c. Other questions to consider: what are the benefits/impacts on Spur Free Dynamic Range (SFDR), bandwidth, resolution, clutter rejection, energy needed on target, *etc.* These are all metrics to consider and they seemed unclear in the presentation.
 - d. What are the implications of the results and how does one compare them to some standard or metric?

We are in the process of developing responses to these actions.

2.0. Technical Progress and Accomplishments

Probably the most effort during the current PoP has been spent on mathematical analyses and algorithm development for the distributed sensing demonstration. Key technical focus areas include:

- The development, implementation and evaluation of accurate and scaleable sensor self-localization approaches.
- Data collection for MICA-2/Z sensors characterization. The focus was on characterizing the acoustic sensors. Preliminary characterization of the vibration sensor was begun.
- The development of distributed tracking algorithms at UniMelb and ASU. The final tracking demonstration is still being refined.

Significant effort has gone into the development of a detailed test plan. While we expect this to evolve somewhat over time, the schedule shown in Table 1 is a solid baseline. These technical focus areas are discussed in significantly more detail in Subsection 2.A, where preliminary experimental results are also summarized.

We have also expended significant effort in the following technical evaluation areas.

- The evaluation of High Dimensional Data Processing algorithms on dual polarization radar field data.
- Stochastic approaches for Unmanned Aerial Vehicle (UAV) control and passive geolocation.
- Algorithm definition and hardware development/test for the implementation and demonstration of the Georgia Tech CADSP imager
- Waveform library selection using a policy-finding algorithm via T-step reinforcement learning.

ISP Phase II (Contract N00014-04-C-0437)
Quarterly Progress Report (CDRL A001 No. 3)

The next several subsections describe the technical approaches for Raytheon and for each subcontractor in greater detail.

Performer	Task Description (<i>Completed</i>)	Start	End
Raytheon	<i>Install updated VU MICA-Z Shooter Localization code</i>	9/19/05	10/14/05
Raytheon	Evaluate Vanderbilt Self Localization algorithm performance	10/17/05	11/18/05
Raytheon	Develop MICA-Z/MATLAB interface specification	11/28/05	12/23/05
ASU	<i>Characterize Acoustic Sensor</i>	9/15/05	11/11/05
ASU	<i>Provide results to tracker algorithm developers</i>	11/14/05	
Raytheon	<i>Modify Vibration Sensor</i>	10/3/05	10/21/05
Raytheon	Characterize Vibration Sensor	10/24/05	11/30/05
Raytheon	Provide results to tracker algorithm developers	12/5/05	
UniMelb	Develop GEN-I Phase II MICA-Z	9/6/05	11/30/05
UniMelb	Incorporate Sensor Specifications into GEN-II Tracker	12/5/05	2/24/06
Raytheon	Verification and Validation UniMelb GEN-II Tracker	2/27/06	4/28/06
ASU	Implement ISP Phase I Tracker Algorithm on MICA-Z	10/3/05	2/11/06
ASU	Develop GEN-I Phase II MICA-Z	10/17/05	12/30/05
ASU	Incorporate Sensor Specifications into GEN-II Tracker	1/9/06	3/10/06
Raytheon	Verification and Validation ASU GEN-II Tracker	3/13/06	5/12/06
Raytheon	Evaluate Michigan Self Localization Algorithm	10/12/05	11/11/05
Raytheon	Code/evaluate Modified Michigan Self Localization Algorithm	10/26/05	12/2/05
*Raytheon	Select Self Localization Algorithm for Demonstration	12/16/05	
Raytheon	Implement Self Localization Algorithm on MICA-Z	1/2/06	2/24/06
Raytheon	Evaluate Self Localization Algorithm Performance	2/27/06	3/17/06
Raytheon	Provide Localization results to tracker algorithm developers	3/31/06	
ASU	Incorporate Self Localization into GEN-III Tracker	4/3/06	6/2/06
UniMelb	Incorporate Self Localization into GEN-III Tracker	4/3/06	6/2/06
Raytheon	Verification and Validation GEN-III Trackers	6/5/06	7/21/06
All	Refine/Merge GEN-III Trackers to produce GEN-IV Tracker	7/24/06	9/8/06
All	Demonstration Dry Runs	9/11/06	10/9/06
All	Perform demonstration	10/12/06	10/31/06
Raytheon	Transition Self Localization Algorithm to NEST	3/31/06	
Raytheon	Transition Self Localization Algorithm to NICTA (ITAR?)	TBD	

Table 1: MICA-Z Tracking Demonstration Schedule

2. A. Technical Progress

2.A.1. Raytheon Technical Progress

2.A.1.a. Distributed Sensor Demonstration

Wireless low-power sensor networks have gained much deserved attention in many research fields. With the advent of low-cost digital signal processors, wireless sensor networks have begun to emerge in many applications. Many military applications, such as the use of wireless networks perimeter monitoring, have self-localization as a critical One very compelling application being pursued by Raytheon under the DARPA Information Exploitation Office (IXO) Networked Embedded System Technology (NEST) program is that of shooter localization through acoustic ranging. To locate the shooter more accurately, several wireless nodes are deployed within an area; thus, these nodes are self-localized first before the localizing the shooter. NEST has identified accurate and scalable localization algorithms as a critical program need.

In fact, self-localization is a key component of a wide variety of distributed

wireless sensing applications, including perimeter monitor and the detection and tracking of targets. Because such sensor networks will be laid down in an *ad hoc* configuration consisting of thousands of sensor nodes, accurate and scalable localization algorithms are critical to many, if not most, defense or homeland security applications.

The current generation of shooter localization algorithm is an acoustic ranging algorithm by Vanderbilt University (VU). The concept of this algorithm is based on measuring the time of arrival (TOA) of the sound signal between the signal source (actuator) and the acoustic sensor. The acoustic ranging algorithm has demonstrated localization accuracy sufficient for a proof-of-principle, and VU is developing an approach that should significantly improve localization accuracy. This new approach uses radio frequency instead of acoustic frequency for the ranging algorithm [Maroti2005]. Thus, this new approach provides more accurate localization with larger networks since radio waves propagate further than acoustic waves. However, both the baseline and improve VU self-localization algorithms rely on a genetic algorithm-based optimization approach which scales very poorly with the number of sensor nodes. As an alternative to the VU self-localization approach, we consider an algorithm based on concepts we are exploring for processing of high dimensional data. .

We began by evaluating the acoustic ranging and self-localization algorithm developed at Vanderbilt University (VU). In September, we tested the VU acoustic ranging algorithm behind the M09 building of Raytheon Missile Systems, Tucson, Arizona. Two sets of tests were performed; the first set comprising thirteen experiments with horizontal spacing between the notes of 18 ft (see blue dots on Figure 1 for geometry), and the second set comprising eleven experiments with horizontal spacing between the notes of 9 ft (see blue dots on Figure 2 for geometry).




Figure 1. Measured Notes' Positions for an horizontal spacing of 18 ft

ISP Phase II (Contract N00014-04-C-0437)
Quarterly Progress Report (CDRL A001 No. 3)




Figure 2. Measured Motes' Positions for an horizontal spacing of 9 ft

After the ranging data was collected, it was post-processed with the VU self-localization algorithm, which is a genetic algorithm, to estimate the position of each mote. The position estimates are shown in Figure 1 for the 18 ft spacing geometry and in Figure 2 for the 9 ft spacing geometry. Both figures show the average positions (red dots) and the median positions (green dots) for each set of experiments. Note that for the case of the 18 ft spacing, the measured positions deviate too much from the true positions to be of any use. A list of standard deviations of each mote's position is given on Table 2. The reason for such poor results is that the acoustic waves are attenuated to the extent that it is difficult to estimate the peak of the signal that determines an accurate TOA. For such large spacing, it appears that acoustic waves are inappropriate and, therefore, in the near future, we will be evaluating the VU Radio Interferometric Positioning (RIPS) algorithm which is based solely on radio frequency [1].

Table 2. Standard deviations for the 18 ft spacing experiments

Mote ID	Stdev x (cm)	Stdev y (cm)
50	426.36	324.66
51	449.51	222.20
52	34.54	407.02
53	236.97	343.99
54	204.12	213.36
57	N/A	N/A
58	200.28	255.75
59	441.89	247.39
60	311.83	170.99
61	80.62	103.92
62	387.07	267.00
63	265.76	206.18
64	583.26	285.67

RIPS exploits interfering radio waves emitted from two transmitting nodes at slightly different frequencies. Two receiving nodes measure the interference, which determines the relative phase offset, and this offset is a function of the relative positions of the four nodes, i.e. two transmitters and two receivers, and on the wavelength of the carrier frequency. By taking this measurement over multiple sets of four within the network, it is possible to collect enough information to uniquely locate each node. Once all of the information has been collected, the positions of the nodes are computed using a genetic algorithm. The advantage of this approach is that it doesn't require any sensors except the radio, and it is very accurate. One of the disadvantages is that it requires a long time to collect all the measurements. For example, for a 16-node network, it requires 10,920 measurements which (according to VU) would take several minutes. However, this data is highly redundant and it should be possible to reduce the set of measurements considerably without sacrificing accuracy, though not without redesigning the VU algorithm. Another disadvantage of the RIPS approach, as it currently stands, is the genetic algorithm used to post-process the data in order to compute the positions of the network nodes: genetic algorithms are known to converge slowly at best and there is no guarantee of convergence.

For the case of the 9 ft spacing network, the measured positions are within an acceptable range. A list of standard deviations of each mote's position is given in Table 3. Although the measured positions are within acceptable range, it should be possible to improve these measurements (once again) using the RIPS scheme. Also, we plan to replace the genetic algorithm currently used to estimate the positions of the nodes with a relatively new algorithm developed by Hero, et al, at the University of Michigan [2], named "distributed, weighted, multi-dimensional scaling," or dwMDS. This is an iterative algorithm which essentially solves a non-linear, weighted least-squares problem by minimization. The weights determine to what degree each measurement in the given set of ranging data influences the final position estimate for each node, and can be chosen to deemphasize the influence of measurements in which we have a low level of confidence. In addition, the least squares criterion is easily expressed as a sum of smaller problems, one for each node in the network, thus the computation is readily distributed by

allowing each node to compute the solution at that node. Furthermore, the weights can be used to completely remove some measurements from the local computations, to ensure that each uses only local information. For this reason the computation scales as $O(L)$ per iteration at each node, where L is the number of measurements included in each of the local computations, and thus the overall computation scales as $O(LN)$ per iteration, where N is the total number of nodes in the network. If ranging data between each pair of nodes were available and all were to be used at each node then the computational cost would be $O(N^2)$ per iteration, which is likely to be prohibitive if N is large. On the other hand, ranging data is likely to be absent for some pairs of nodes. Thus, in any case, the ability to use only local information at each node is important. It should also be remarked that dwMDS is formulated to include some prior information. In particular, if the positions of a few nodes are known in advance, then this information can readily be included as part of the least-squares criterion.

Table 3. Standard deviations for the 9 ft spacing experiments

Mote ID	Stdev x (cm)	Stdev y (cm)
50	10.54	50.53
51	11.97	18.19
53	9.82	9.35
54	1.42	11.55
57	N/A	N/A
58	15.88	4.71
59	41.54	83.51
60	15.64	35.40
61	13.21	14.52
62	21.46	23.77
63	30.58	63.34
64	17.12	50.04

We have obtained encouraging results with the dwMDS algorithm so far. Working initially from a “demo” provided by the University of Michigan team, a Matlab version of dwMDS has been programmed and is undergoing testing at Raytheon, and we show some typical results in the Figures 3 – 6. To initialize the dwMDS iterations for these experiments we generated random, uniformly distributed starting values for each node in the rectangular region $\{0 < x < 1800, 0 < y < 900\}$, however, with this choice of initial values sometimes several tries were needed to get a good solution. For these experiments, the dwMDS algorithm was not allowed to use any prior knowledge of actual positions, but computed the solution using ranging data only, after which the solution was rigidly translated/rotated/reflected onto the best fit with the actual data. Alternatively, a few known positions could have been provided to the algorithm as part of the initialization, and incorporated into the least-squares criterion, but this was not done for these experiments.




Figure 3: result of dwMDS algorithm on one of the 9 foot spacing data sets. Red circles are actual positions, black dots are estimates. A red circle with no black dot associated to it indicates a node for which we got no measurements during the data collection.

We emphasize that these are preliminary results only, and that some work is yet required to fine-tune the dwMDS algorithm to our specific application and ranging data. In particular, it may be useful to investigate different choices for the weights, as well as try to find a more optimal initialization for the iterations. And, of course, we plan to investigate different methods for measuring the ranges between nodes, as was mentioned above. However, it is our judgment, based on the preliminary results, that this algorithm, or some slight variant of it, will be satisfactory for solving the self-localization problem.




Figure 4: result of dwMDS algorithm for one of the 9 foot spacing data sets. Red circles are actual positions, black dots are estimates. A red circle with no black dot associated to it indicates a node for which we got no measurements during the data collection. This is a different data set than that shown in Figure 3, note the differences in locations of non-reporting nodes.




Figure 5: result of dwMDS algorithm for one of the 18 foot spacing data sets. Red circles are actual positions, black dots are estimates. A red circle with no black dot associated to it indicates a node for which we got no measurements during the data collection. Though there are large errors at some of

the nodes, the algorithm performed well on most of the nodes. We believe that with better ranging data the algorithm will perform well on all nodes.




Figure 6: This is the result of running dwMDS on only a selected subset of the data set in Figure 5. Note that when some “bad” nodes are removed, the algorithm performs well. As stated, we believe that a better method of collecting ranging data will solve this problem, which appears to be caused by a lack of sufficient connectivity throughout the data set.

2.A.1.b. Polarization-Resolution Trade Study for SAR Imagery

In the SAR ATR community, it has been noted that increasing polarization diversity can effectively serve as a surrogate for increasing sensor resolution. We have quantified this observation for a two-target case at the pre-classifier stage, in terms of information divergence between the SAR images corresponding to the two targets. The information divergence between distributions f and g , sampled with m and n points respectively, can be calculated using any of a number of methods. One approach is to use the Henze-Penrose Divergence $d(f,g)$, asymptotically approximated using the Friedman-Rafsky Multivariate Runs Test. By using a graph-theoretic method based on the minimal spanning tree for the joint data set of $n + m$ samples, this approach avoids curse-of-dimensionality problems associated with density-fitting or parameter estimation in high-dimensional spaces. The Henze-Penrose Divergence is given by:

$$d(f,g) = \int \frac{p^2 f^2(x) + q^2 g^2(x)}{pf(x) + qg(x)} dx, \quad p = \frac{m}{m+n}, \quad q = \frac{n}{m+n}.$$

We note that $d(f,f)=0.5$, and for f and g completely disjoint, $d(f,g)=1$. An asymptotic approximation for the Henze-Penrose Divergence is given by the following function of the test statistic R found in the Friedman-Rafsky Multivariate Runs Test:

$$d(f,g) \approx 1 - \frac{R}{m+n}.$$

For this study, we use the public domain portion of the Advanced Detection Technology

Sensor (ADTS) dataset¹. The ADTS is a fully polarimetric, air-to-ground SAR sensor operated by MIT Lincoln Laboratory. The ADTS operates at Ka band (32.6 GHz to 37 GHz). In 1989-1990, an M48 tank and an M55 self-propelled howitzer were measured at a fixed depression angle on 22.5°. Approximately 230 samples of each were collected over 360° in azimuth in 1.5° – 2° steps with several broad gaps in azimuth coverage. The pixel spacing is 9” in the object plane, and the images have 12” resolution in both down-range and cross-range.



Figure 7: ADTS radome and targets used in study.

The data consists of 128 pixel x 128 pixel SAR images in HH, VV and HV polarizations. These images are polarization-calibrated and for each scene, the three SAR images are registered. We first extract the central 36 pixel x 36 pixel region, assuming the targets are accurately centered in the images. We then simulate detectors of coarser resolution by block-averaging the complex data from the 12"-resolution images. The extrinsic dimensionality remains $36 \times 36 = 1296$, however the intrinsic dimensionality drops as the blocksize increases.




Figure 8: SAR images (M48 tank) with block-averaged regions, labeled by approximate spatial resolution / intrinsic dimensionality.

Given registered HH, VV and HV polarization images, we can construct images that combine any two polarizations or all three polarizations. We use the MIT Lincoln

¹ <https://www.sdms.afrl.af.mil/datasets/adts/>

$$\sigma_{\text{III}} = \sqrt{\frac{\sigma^2}{\sigma^2 + \sigma^2}} = \frac{\sigma}{\sqrt{2}} \quad \Sigma = \sigma_{\text{III}} \begin{pmatrix} 1 & 0 & \rho\sqrt{2} \\ 0 & \epsilon & 0 \\ \rho\sqrt{2} & 0 & \gamma \end{pmatrix} \quad , \quad \mathbf{X} = \begin{bmatrix} \text{HV} \\ \text{HV} \\ \text{VV} \end{bmatrix} \rightarrow \begin{bmatrix} \text{HV} \\ \text{HV} \\ \text{VV} \end{bmatrix} \xrightarrow{\text{"Whitening Filter"}} \mathbf{W} = \begin{bmatrix} \text{HV} \\ \text{HV} \\ \text{VV} \end{bmatrix} \xrightarrow{\text{"Whitening Filter"}} \begin{bmatrix} \text{HV} \\ \text{HV} \\ \text{VV} \end{bmatrix} \xrightarrow{\text{"Whitening Filter"}} \begin{bmatrix} \text{HV} \\ \text{HV} \\ \text{VV} \end{bmatrix}$$

Laboratory ‘polarization whitening filter’ approach, designed to minimize speckle. The $\{HH, VV, HV\}$ basis is whitened such that the resulting components are uncorrelated and have equal mean power.

The combined polarization images are formed by taking the non-coherent root-sum-square of the desired components, normalized by the number of components used:

$$PWF(HH, HV, VV) = \sqrt{\frac{|HH|^2 + \left|\frac{HV}{\sqrt{e}}\right|^2 + \left|\frac{VV - r^* \sqrt{g} HH}{\sqrt{g(1-|r|)^2}}\right|^2}{3}}, \quad PWF(HH, HV) = \sqrt{\frac{|HH|^2 + \left|\frac{HV}{\sqrt{e}}\right|^2}{2}}$$

Thus, we can form SAR images with different spatial resolutions and composed of one, two or three polarization components.

We apply two methods of dimensionality reduction to vectorized forms of the SAR images in these data sets. ISOMAP is used for nonlinear dimensionality reduction, with the first 10 embedding components retained. Metric MDS (essentially PCA in this context) is used for linear dimensionality reduction, with either the first 10 or the first 150 embedding coordinates retained. Divergence between data sets is measured for the original data, the ISOMAP embedding and the MDS embeddings. For ease of display, the three single-polarization results are averaged, as are the three two-polarization results.




Figure 9: Henze-Penrose Divergence in high-dimensional data space, as a function of spatial resolution and polarization diversity.

Our original goal was to quantify, at the pre-classification stage, the two-target (M48, M55) divergence as a function of spatial resolution and polarization diversity in the ADTS SAR imagery. This is shown in Figure 9. For example, having a single polarization at a spatial resolution of approximately 3' provides divergence similar to that provided by two polarizations at a spatial resolution of approximately 5'. As expected, as the spatial resolution coarsens, the divergence drops to 0.5 (sample distributions drawn from same parent distribution). However, there is unexpected behavior at the sharpest

resolutions: divergence decreases as resolution sharpens from 4-pixel blocking to the original 1-pixel blocking. We take this as an indication that there is a substantial amount of noise on spatial scales of 1 – 3 pixels.

Consider now the divergence behavior for the low-dimensional embeddings (Figure 10). ISOMAP (nonlinear approach), with 10 retained embedded coordinates, is fairly effective in capturing divergence (figure 10.b), and reproduces the divergence in the data vectors fairly well (figure 10 a). Metric MDS (linear approach) with 150 retained embedded coordinates does very well in capturing divergence (figure 10.d), and reproduces the divergence in the data vectors very well. Perhaps most interesting, Metric MDS with only 10 retained embedding coordinates (figure 10.c) actually outperforms the original data vectors. This highlights the effects of noise on the scale of 1 – 3 pixels. Such noise is captured in the nonlinear dimensionality reduction approach with 10 retained embedding coordinates, as well as in the linear dimensionality reduction approach with 150 retained embedding coordinates. However, the linear dimensionality reduction approach with 10 retained embedding coordinates does not have sufficient capacity to effectively embed the noise. Divergence is increased due to denoising at the sharpest resolution scales.




Figure 10: Henze-Penrose Divergence for the different representations, as a function of spatial resolution and polarization diversity.

One source of spatial noise is registration/centering errors. The original 128 pixel x 128 pixel images were manually cropped from larger images. It is possible that the

registration errors are predominantly on the scale of 1–3 pixels. To test this hypothesis, we used a cross-correlation filter to better register / center the images, and then we repeated the analysis. As seen in Figure 11, the data vectors no longer exhibit a decrease in divergence as the resolution sharpens to 12"/1 pixel. The images are still imperfectly registered/centered; with more accurate registration, particularly at the sub-pixel level, we expect the divergence of the data vectors to not flatten out at the sharpest resolutions (Figure 11.a), but rather to show the monotonic behavior seen in Figure 11.c.

In summary, we demonstrated a quantification of the resolution/polarization trade, using divergence at the pre-classifier stage. We also showed that in this particular example, a low-capacity linear dimensionality reduction approach can have utility in denoising (regularising) data. We also showed that for techniques that rely on image registration, going to sharper resolution places demands on registration accuracy.




Figure 11: Henze-Penrose Divergence for the different representations, as a function of spatial resolution and polarization diversity, for images recut with reduced registration / centering errors.

2.A.1.d. CADSP UCIR Evaluation Technical Support

There is currently a great deal of interest in UCIR sensors within the Department of Defense community in Automatic Target Acquisition (ATA) for smart munitions. A prime example of such a weapon system is the NetFires NLOS PAM that is being developed jointly by Raytheon and Lockheed Martin under U.S. Army sponsorship. We believe that the Georgia Tech CADSP imager has potential for improving the performance of the PAM UCIR imager. We are investigating its application for *on-Focal*

Plane Array (FPA) pre-processing operations; these include: Non-Uniformity Compensation (NUC) and non-linear/non-local pixel equalization. Traditional equalization approaches (*e.g.*, histogram equalization) tend to perform very poorly and it is likely that a localized, non-linear equalization approach is needed. Given ISP Phase II funding constraints, we will limit these pre-processing investigations to an evaluation of their implementation on the Georgia Tech CADSP imager. We have also been in preliminary discussions with Eglin, Air Force Base about using their optical flow test facilities.

2.A.2. ASU Technical Progress

2.A.2.a. Tracking and Sensor Scheduling with Motes Demonstration

We continue to investigate sensor configuration for the ISP demonstration of tracking a target moving through a network of sensor motes. We are investigating myopic sensor scheduling algorithms to minimize the network energy consumption while maintaining a desired tracking accuracy. The network comprises of acoustic-energy sensors on MICA-2/Z motes that provide noisy measurements of the target. Ultimately, the target will be tracked using a belief propagation algorithm, in which the target belief is transferred from one sensor to another as the target moves in the sensor network; in this algorithm, the sensor holding the current target belief at time k performs sensor scheduling to allocate sensors for the next time-step in order to minimize the energy consumption of the network while meeting the desired tracking accuracy at time-step $k+1$. In our initial implementation, the sensor observations will be communicated to a workstation with significant computational capacity that will perform both the tracking and scheduling operations.

We have formulated the sensor scheduling problem as a nonlinear integer programming (NLIP) problem in which the objective is to minimize the sensor activation cost subject to the accuracy constraint. We divide the NLIP problem into a sequence of linear problems, each of which can be posed as a linear integer programming (LIP) problem; for each LIP, we determine whether its solution meets the desired tracking accuracy constraint. The sequence of linear problems is determined by a greedy search; as it proceeds, the search provides a set of constraints on the sensor activations and reduces the dimension of the linear programs. In this formulation, we approximate the effect of each sensor on the target estimate error covariance matrix using a linear approximation and the information formulation of the filter step of the Kalman filter. The information formulation leads to a closed-form expression for a necessary constraint on the trace of the sensor information matrix to meet the desired tracking accuracy; this constraint is used by the greedy search algorithm. The advantage of the proposed method is that the sequence of LIPs performs intelligent enumeration in the search space and thus eliminates many sensing options that are not feasible solutions. Furthermore, the solutions of lower dimensional LIP problems provide linear sensor-usage constraints, which help in eliminating many infeasible sensing options.

In parallel with the development of the tracking and scheduling algorithms for the demonstration, we are characterizing the Mica-2/Z acoustic sensor response to a walking human target. We collected acoustic data for a target (person) walking in a circle around the sensor at various distances; we also collected acoustic data when no target was present. We divided the data into one second segments and computed the energy in each

segment. We developed receiver operating curves empirically by varying a detection threshold. We provided receiver operating curves and probability of detection versus threshold curves to the Raytheon and University of Melbourne ISP team members for use in the development of tracking algorithms. We are also formulating an acoustics observation model that is appropriate for our tracking and scheduling algorithms described above.

2.A.2.b. Multiple Target Tracking using the Configurable CADSP Imager

The objective of this demonstration is to track multiple targets in a surveillance region using image data from the CADSP imager, and specifically to configure the imager to provide only needed data. The imager can be configured to compute optical flow or image selective sub-areas of the field of view. Our emphasis to date has been to develop a measurement model for the image generation and processing of the CADSP imager; this model is necessary to integrate the imager with our existing multiple target tracker and configuration algorithms for the demonstration. We have used the POV ray-tracing program to generate video sequences for a generic moving object to be tracked. We have applied Gaussian filters and wavelet (Mexican hat) filters to process regions of the image; these filters are useful in classifying background and foreground image regions as well as detecting edges of moving objects. In the demonstration, the filters will be loaded onto the CADSP chip and applied to the image data as it is acquired. In addition, we have empirically estimated distributions of filtered image data from background and target regions of the images; we have used Gaussian mixtures to represent the distribution estimates.

We are currently developing processes to manipulate data over multiple frames, to better estimate foreground and background distributions, and to use the processed image data as observations in a particle filter to track objects. Once the particle filters are initially implemented, we will investigate methods to configure the CADSP chip by selecting between available CADSP configuration options.

2.A.3. UM Technical Progress

We evaluated the capabilities of classification constrained dimensionality reduction (CCDR) algorithm to improve performance of classifiers for ISP problems. In order to apply CCDR to ISP we had to generalize and extend the framework a little bit. The CCDR algorithm was introduced for binary classification. We developed an extension to multiple hypotheses that allowed us investigate classification performance using the CCDR algorithm on hyperspectral satellite imagery data (Landsat) for terrain classification. In our ICASSP paper [9] we did an extensive performance analysis for several baseline classifiers including neural nets, linear classifiers, and kNN classifiers. Use of CCDR as a preprocessing step prior to classification allowed a simple linear classifier to achieve over 200% reduction in misclassification probability of error and perform at or below the level of the benchmark non-linear classifier for this problem (kNN). Furthermore, the CCDR preconditioned kNN achieved a 10% improvement over the benchmark kNN without CCDR. Finally, we found an important connection between intrinsic dimension estimation via entropic graphs and the optimal embedding dimension for the CCDR algorithm.

We also demonstrated that classification reduction optimal policy search (CROPS) strategies can be applied to perform optimal online waveform selection with substantial improvements as compared to myopic (offline) methods such as DOE. For a fixed energy (average number of waveform bands transmitted) constraint, offline optimal scheduling simply chooses to simultaneously transmit the best set of waveform bands satisfying the constraint. Multi-stage online scheduling permits specifying an optimal sequence of waveform bands, again satisfying the constraint, where the next waveform band in the sequence is selected as a function of previous measurements. Online scheduling is an optimal sequential design problem while offline is an optimal fixed design problem. We applied these methods to the Landsat dataset and demonstrated that the use of online waveform scheduling significantly enhances performance on the Landsat test set. Since no statistical models are available for the Landsat dataset, these designs used non-parametric machine learning methods to capture the optimal scheduling strategy from the training set alone. In particular, we established that CROPS can come within less than 1% of the misclassification error performance of the optimal non-constrained strategy (use all waveform bands simultaneously) but using less than 3/4 of the transmitted energy of the non-constrained strategy.

2.A.4. FMAH Technical Progress

2.A.4.a. Synthetic Waveform Diversity

Note: The mathematical basis for the FMAH waveform design is provided as separate report.

In the work done over the last quarter we have established a solid foundation for the use of a new Processing/Modulation technique called Synthetic Waveform Diversity. The Code-generation schemes for this technology were introduced a in the six-month progress report presented at the ISP Phase II review. New suitable Radar applications for these mathematical Waveform Design and Processing methodologies are currently being identified.

Side-lobe Suppression in Radar

The side-lobe reduction problem is of fundamental importance in a variety of radar applications [5]. In the last quarter, we continued the development of a technique capable of reducing the Integrated Side-lobe Level (ISL) of a mixed signal r , consisting of a target return, multi-path (attenuated and delayed) and noise clutter returns:

$$\hat{r}(t) = \sum_{\lambda \in \mathbb{K} \cup N} a_{\lambda} p(t) + b_{\lambda} s(t - t_{\lambda}) + s(t) + n(t) \quad (1)$$

The assumptions usually are that the target components in Equation (1) are stationary over a few duty cycles and that the noise component is Gaussian. Detection is performed by simply thresholding the output of the correlator.

Since the classic work of Golomb, numerous mathematical techniques has been developed for the generation of code families with high combined auto-correlation performance. With few exceptions though, researchers in the field have been interested in the development of the algebraic aspect of the theory of code design. As for the code processing schemes, the conventional wisdom has always bees that linear techniques are to be used. In particular in most cases simple weighted averaging of the different waveform has been used to rains the Target's SNR. The goal of the present paper is to

demonstrate how analytical techniques can provide higher flexibility to build different requirements into the algorithms. In order to achieve this, we have proposed some techniques which make extensive use of tools from Multiresolution Harmonic Analysis as well as a general result of Benke regarding the fundamental construction of Rudin-Shapiro Polynomials. More specifically, we consider codes and code-families as finite approximants of bases in infinite-dimensional function spaces.

An approximation scheme is exhibited with the desired asymptotic properties. We then obtain a variety of new code generation algorithms and provide some explicit estimates for their performance.

Synthetic Waveform Diversity

The main idea for the algorithm is to combine cancellation properties of carefully designed waveforms with the advantages of statistically clever processing of the signal. The typical waveform we will use is chosen from the ones generated utilizing the new Benke-Scheme code generator. The fundamental innovation of our work is that the classic complementarity processing

$$r(t) = \sum_{\lambda} s_{\lambda}(t) \bullet \tilde{s}(t) \quad (2)$$

which can only be proven efficient in AWGN scenarios is completely supplanted².

A new, multi-dimensional non-linear functional is introduced. The approach we implemented is a stochastic learning strategy to identify the optimal parameters to deal with the real noise statistics. We write:

$$r(t) = F(s_1(t)m_1(t), K, s_N(t)m_N(t)) \quad (3)$$

Several different models are used for the Mapping F, as well as for the mismatched filters $\{m_1, \dots, m_N\}$. The burden of choosing their exact shape is entirely shifted onto an agile stochastic learning strategy. In doing so, we are in fact able to determine in an application specific manner the number and characteristics of the mismatched filters being used. To test the performance of the algorithm we simulate a highly noisy, multi-path return. The template waveform is chosen to have maximum (=near optimal) resolution around the main lobe a noise-like characteristics away from the latter.

2.A.4.b. Diffusion based high dimensional data processing

The goal of this portion of the project is to apply and adapt geometric diffusion methods of Coifman *et al.* to IR video data. As a simplification, we consider direct application of diffusion operators to the data itself.

Consider an image $\mathbf{u}$ consisting of a rectangular array of pixels. We construct a diffusion filter $\mathbf{K}$, such that repeated application of $\mathbf{K}$ to $\mathbf{u}$ suppresses the background while enhancing or preserving regions of interest in the image.

Construction of $\mathbf{K}$ proceeds as follows. For each pixel j in $\mathbf{u}$, a group of neighboring pixels of size $n \times n$, denoted $\mathbf{x}(j)$ and center on the pixel j , is selected. Next, a

² In the absence of colored noise Equation (2) will provide optimal range resolution. It is immediate to show that this is far from being the case with a simple noise statistics.

scalar non-negative kernel function k is selected to measure the degree of similarity or difference of two groups $\mathbf{x}(i)$ and $\mathbf{x}(j)$. The function k is therefore a bivariate function of vectors of length $n \times n$. If $\mathbf{x}(i) = \mathbf{x}(j)$, then $k(\mathbf{x}(i), \mathbf{x}(j)) = 1$, and if $\mathbf{x}(i)$ and $\mathbf{x}(j)$ are “dissimilar” then $k(\mathbf{x}(i), \mathbf{x}(j))$ should be close to zero. Usually, k is designed to be a symmetric function of its arguments, so that the degree of similarity between $\mathbf{x}(i)$ and $\mathbf{x}(j)$ is the same as the similarity between $\mathbf{x}(j)$ and $\mathbf{x}(i)$.

From this kernel function, together with an integer N , a matrix $\mathbf{K}'$ is constructed so that

$$\mathbf{K}'_{ij} = \begin{cases} k(\mathbf{x}(i), \mathbf{x}(j)) & \text{if } |i - j| \leq N \\ 0 & \text{otherwise} \end{cases}.$$

Then, $\mathbf{K}$ is constructed from $\mathbf{K}'$ by normalizing the row sums to one; that is,

$$\mathbf{K}_{ij} = \mathbf{K}'_{ij} / \sum_l \mathbf{K}'_{il}$$

where the sum is over the columns. The matrix $\mathbf{K}$ may be viewed alternatively as a *diffusion operator*, since its action on $\mathbf{u}$ is identical to a forward time step in a discrete diffusion problem, or as a *Markov matrix*, since the sums of its rows are unity. Each row has $(2N+1)^2$ non-zero entries, and the total number of rows is equal to the number of pixels in $\mathbf{u}$.

We note that the data set $\mathbf{u}$ need not be a two-dimensional image. For a sequence of images, as in video data, we select groups of pixels of size $n \times n \times m$, for example. In all other aspects the matrix $\mathbf{K}$ is constructed as above.

We intend to study a variety of choices for the kernel function k . Our initial results utilize the square of the Euclidean distance to define

$$k(\mathbf{x}(i), \mathbf{x}(j)) = e^{-\|\mathbf{x}(i) - \mathbf{x}(j)\|^2 / e},$$

a choice that, though apparently simple, is remarkably powerful.

An example of IR video data processed with this filter is shown in Figure 12. In this example, a single frame was processed with a 2D filter, with the parameters $N=1$ and $n=2$. The filter $\mathbf{K}$ was iterated 10 times. The resulting image has preserved the detail around the target, while significantly smoothing the background of the image. Some details such as roads, etc. are still present in the denoised image, however.

A number of possibilities exist for choices of the kernel function k :

- First, compute the left and right principal vector vectors of the $n \times n$ matrices centered on i and j . Let θ_l and θ_r be the angles between the left principal vectors and the right principal vectors. Then, $k(\mathbf{x}(i), \mathbf{x}(j)) = e^{-(\tan^2 \theta_l + \tan^2 \theta_r) / e}$. The resulting filter might show improved results for denoising, and do a better job of separating targets from noisy backgrounds. This filter can only be constructed for 2D images, however.
- Compute the angle, θ , between local orientation vectors centered at i and j . Then, $k(\mathbf{x}(i), \mathbf{x}(j)) = e^{-(\tan^2 \theta) / e}$. The resulting filter should segment the data into regions of similar orientation, which may be useful for certain types of targets or backgrounds.

Additionally, for video data, it may be possible to design kernel functions which exploit or utilize inter-frame motion data.




Figure 12: Single frame of IR video processed with diffusion filter. Background regions in the image have been significantly denoised.




Figure 13: Zoom from Figure 12. The details of the image around the target have been preserved.

2.A.5. UniMelb Technical Progress

2.A.5.a. Tracking Using Motes: Unscented Kalman Filter Approach

A new technique for tracking a target moving through a field of motes has been developed. The problem is formulated as follows. It is assumed that the number of targets is known and that the target states are Markov processes which evolve independently in a known manner. Measurements from the mote network, using, for instance, accelerometers or acoustic sensors, are made available at the regular time instants at a central node. These measurements are generated independently at each mote and are assumed to have a known probabilistic characterization, i.e., the distribution of the measurements conditional on the target state is known. The measurements can be transmitted to the central node as thresholded or non-thresholded measurements.

Optimal tracking of the targets, in the minimum mean square error sense, requires computation of the posterior density. This is not possible so a sub-optimal approximation is required. We consider the use of the Kalman filter. The Kalman filter has the desirable properties that it is the linear estimator with the minimum variance and it requires only second-order statistics so is therefore computationally inexpensive. The main challenge in using the Kalman filter is the computation of the predicted measurement, innovation covariance matrix and gain matrix required in the correction step. These quantities are generally intractable and must be approximated. Currently the unscented transformation and its higher-order generalizations are being used for this purpose.

The algorithm performance is demonstrated using the scenario of Figure 14. The straight lines are the target trajectories and the mote locations are indicated by small solid discs. The circles surrounding the motes indicate the area within which a detection in the thresholded measurement model is produced with probability greater than 0.5 for a single target. Note that the motes do not provide a complete coverage as the detection probability is less than 0.5 in much of the surveillance region. The algorithm performance for both thresholded and non-thresholded measurements is given in Figure 15. The plots show the RMS error in the y-position estimates of target 1, computed over 500 Monte Carlo realizations, plotted against time. The Cramer-Rao bound (CRB) is also shown. The algorithm achieves an estimation accuracy close to the CRB with both thresholded and non-thresholded measurements. The loss in performance resulting from thresholding the measurements is clearly evident.




Figure 12: Mote Tracking Scenario




Figure 13: RMS error of y-position estimate of target 1 plotted against time. (a) thresholded measurements; (b) non-thresholded measurements. Solid line is estimation error; dashed line is CRB.

Future work will involve extending the algorithm to handle target maneuvers, estimate target signal strength and detect the presence of new targets

2.A.5.b. Tracking Using Motes: Virtual Measurements

Our sensor clustering based method described in the previous report now has a formal name – the Virtual Measurement (VM) approach. The central idea of this approach is to define a mapping between the space of binary sensor observations and a discrete vector space, called the virtual measurement space, such that, any point within the VM space is a linear transformation of target state, as if it were generated by an equivalent “large sensor”. With VMs, conventional multi-target tracking (MTT) algorithms (with clutter filtering built in) can be straightforwardly used for multi-target tracking in binary sensor networks.

A significant amount of work has been done on the VM approach in this report period. In particular, we have concentrated on:

1. Refinement of concepts, theorem, and algorithms for the VM method. These include
 - a. Refinement of the concepts:
 - i. Distinct Sensing Area (DSA) – that is the area that is uniquely assigned to one sensor and to no others.
 - ii. VM – the assignment of virtual target positions to a given sensor field.
 - iii. Independently Activated Sensor Group (IASG) (or cluster) – that is the largest set of activated sensors sharing a common DSA.
 - b. Establishment of the mapping theorem; that is, that a 1:1 and unique mapping can be found from spaces of sensor observations to VM.
 - c. Develop IASG identification algorithm and exact VM assignment (EVMA) algorithm. These two algorithms can convert binary sensor observations into VMs as the input of a conventional MTT tracker.
2. Study of practical issues and solutions. In particular, we used the Gaussian distribution to model uncertainties for sensor location, sensing range. We have developed a Mixed VM assignment (MVMA) algorithm, which may enhance tracking performance (in terms of track loss and RMS error) when diverse uncertainties (noisy sensing range, sensor location and low PD) are presented.
3. Implementation of a multi-target tracker. The linear multi-target integrated probabilistic data association (LMIPDA) algorithm is a recent development by Musicki et. al. (Fusion 2004) for conventional (linear Gaussian) applications to multi-target tracking in clutter. In the context of mote tracking, this algorithm has the following advantages:
 - a. Lower computational complexity (linear to number of targets)
 - b. Good for dense targets but low clutter (ground targets) tracking.
 - c. Probabilistic target confidence feature which can be conveniently used for track assessment in the multi-target tracker.

As we mentioned in the previous report, much of the computational load of the VM approach can be put into the network initialization phase, which is an off-line single pass process, if all sensors of the network are stationary. In network initialization, base station will compute all DSAs based on known sensor locations and sensing ranges and thus build up a VM space for the ideal case (i.e., assuming sensor $P_D = 1$).

Tracking performance assessment:

We have tested the following scenario.

Number of Motes:	100
Sensing Range is a Gaussian random variable:	
Mean:	10 m
Variance:	0/3 m
Distribution of sensors:	
Uniformly spread with minimum separation	2 m
Surveillance region:	120 x 120 m ²
Sensor detection probability (P_D)	0.9/0.5
Data sampling interval:	1 sec
Number of samples in each run:	40

Track loss: A track that has not been updated in consecutive 10 scans is deemed loss.
 EVMA: Exact VM assignment – as described above
 MVMA: Mixed VM assignment, which computes the EVMA if there is an entry in data base corresponding to the measurement, otherwise it takes a mean position for the measurement.

Single target tracking performance statistics (over 100 runs):

Individual sensor PD	Sensing Boundary noise (B)	Track loss	Assignment Algorithm
0.9	0	0	EVMA
0.9	0	0	MVMA
0.9	3 m	0	EVMA
0.9	3 m	0	MVMA
0.5	0	75 or 75%	EVMA
0.5	0	21 or 25%	MVMA
0.5	3 m	78 or 78%	EVMA
0.5	3 m	41 or 41%	MVMA

VM distribution versus ground truth

To cope with conventional tracking technique, this distribution is modeled as a Gaussian distribution. The graphs below show that the Gaussian assumption is reasonable.




Figure 14: VM Distribution versus Truth




Figure 15: RMS Error Performance

Sensor location uncertainty has the same impact as the variability of sensing range. The impact of introducing these uncertainties is similar to the case of decreasing B_0 of individual sensors, which increases RMS errors and track loss.

Future areas of further investigation involve: i) include extra information (signal amplitude, power) to reduce the uncertainties of VMs; and (2) modeling and tracking targets with diverse dynamic using IMM framework.

2.A.5.c. Motes Health and Scheduling

This quarterly effort investigated tracking a target through a field of motes, some of which may be inoperative. If a mote is operational, and it detects a target within its sensing range, then it transmits a unique identifying code to a central processor. It is assumed that the central processor knows the location of each mote.

The simplest approach to tracking with such binary proximity sensors is to record the location of each mote that reported a detection and fit a straight line through these data points. However, this does not make use of all the available information. The absence of a detection in the predicted target location can provide information on the track prediction accuracy. A more accurate track estimate can be obtained if information, both its presence and its absence, from all sensors is used to update the track estimates.

The Gaussian Mixture Model approach operates in this way. Particle filter based approaches operate in a similar fashion but are computationally more complex.

Existing methods assume all motes in the surveillance region are functioning. These trackers therefore assume that the lack of a target detection means that the target is not near the sensor. However, as these sensors are battery powered an alternative explanation is that the mote is no longer operational. To track robustly with such sensors it is necessary for the central processor to know which motes in the surveillance region are functioning. As the wireless networking employed by such sensor networks is *ad hoc* such status information is not necessarily available to the central node.

The tracker developed here, called the **Health Tracker**, simultaneously estimates the target track and also the probability each mote in the surveillance region is operational. This algorithm is robust to sensor failures. In addition, the central processor can use the estimates of the probability a mote is operational to schedule targeted queries to individual motes, rather than requiring regular status updates from all motes.

The mote detection probability was modelled as a decreasing exponential function of distance, i.e. $P_D = P_C \exp(-2\lambda |x-z|)$, where P_C is the probability of detection at zero range; x and z are the locations of the target and the sensor respectively; and 2λ is a scaling parameter that determines the maximum sensing range of the mote. It was assumed that only a relatively small number of motes may be inoperative. This is reasonable as once too many motes fail the network is no longer connected and tracking is not possible. It was also assumed that motes do not die during the lifetime of a particular tracker. That is, the mote scenario (i.e. which motes are alive and which are dead) is fixed but some of the motes may be inoperative prior to the target entering the surveillance region. This assumption is based on the expectation that mote failures will be relatively rare.

Using these assumptions the Gaussian Mixture Model tracking approach was extended to estimate both the target track and the health of each mote. At each time instant, the target state for a given mote scenario was approximated by a single weighted Gaussian distribution, in a manner similar to Probabilistic Data Association (PDA) tracking. These Gaussians can then be combined to produce a single track estimate at each sampling instant. The weights are also used to approximate the posterior probability that a mote is alive.

1000 Monte Carlo simulations were used to evaluate the performance of the **Health Tracker** for a single target moving through a 2D field of sensors, one of which is dead. Figure 16 shows the RMS position errors in the y coordinate. The results for three trackers are shown. The kPDA results are the output from the tracking algorithm when it is *known* which motes are operational; the ePDA results are from the **Health Tracker** which *estimates* which motes are alive; and the bPDA results are from the equivalent tracker which *believes* all motes are alive. The target passes through the sensing range of the dead mote during scans 7 to 14. From this figure it can be seen that estimating mote health improves tracking performance as it reduces the bias in the track estimates that is due to the absence of a detection from the inoperative mote. Figure 4 shows the estimated probability that the dead mote is operational as a function of time, which shows a significant drop when the target enters its detection range. The probabilities for all other motes remained over 98% at all times.




Figure 16: RMS Position Errors in y-Coordinate




Figure 17: Probability Mote 14 is Operations

Work with this tracker is currently being extended to investigate the scheduling of mote status checks by the central processor. In this framework, at each sampling instant the central processor can either process incoming detections from the motes or query a mote to determine if it is still functioning. The comparison between the kPDA and ePDA trackers shown above suggest that this may further improve tracking performance. Two scheduling algorithms are being considered. The first is a simple, *ad hoc* scheme that queries a mote when its estimated probability of being alive falls below a threshold. The second scheme is a more sophisticated algorithm that considers the one-step ahead change in track estimation error due to querying a mote. The status of a mote is only checked if this reduces the overall uncertainty in the track state estimate. Results for these two scheduling algorithms will be available in the next report.

2.A.5.d. Scheduling of Observations of Hidden Markov Processes

We study the problem of scheduling observations of a dynamical process modeled as a Markov chain, with state space, S , and transition matrix, P . It is usual to consider a partially observed system of this kind in which a given state gives rise to a given measurement with a certain probability. Here we consider a situation in which multiple

measurements are possible, but only one is used at each epoch. Thus, various possible observation probability matrices

T_k ($k = 1, 2, \dots, K$)

are available for observing this system. We wish to schedule these observations to maximize the information rate coming to the observer from the system. While the general problem of this kind is difficult, it is possible to consider scheduling where the policy is stationary; that is depends only on the current information state of the system. We seek, then, a stationary solution to this problem. For this, we assume that the information state of the system is calculable at each epoch. This is a probability p on the (finite) state space S . A stationary solution will then depend only on the value of p . Accordingly, we define a policy to be a finite partition

$t = \{B_1, B_2, \dots, B_M\}$

of subsets of the S . This policy is to be interpreted as meaning that when p is in B_m , then observation matrix T_m is used. The aim then is to find the optimal partition in the sense described earlier. We have developed a considerable amount of this theory and have shown that there is a simple description of the optimal policy in terms of a calculable function of the dynamics, provided that there is a unique probability distribution associated with this policy. This is an irreducibility assumption for the system.

Our future work will be on finding conditions for this irreducibility assumption to be met. We believe that a theoretical result of this kind will provide a significant impetus to the study of scheduling of sensors.

2.A.5.e. Coordinated UAV Evaluation

In the previous report we have described one-step ahead scheduling for multiple UAVs. Our new effort over the last quarter focuses on multi-step ahead scheduling of multiple UAVs with either bearings or time-difference-of-arrival (TDOA) tracking. As in the earlier effort, the task was to track a single slow moving target. No clutter is assumed and noise on both the trajectory of the target and the sensor measurements was assumed independent and Gaussian. The idea of the method is the following:

1. At each epoch a linear dynamical model for target was assumed for the target.
2. Each UAV maintains the status of all of the other UAVs.
2. Tracking used measurements by all of the UAVs; that is, tracking was done by passing plot data to a centralized tracker.
3. The cost is the sum of the logarithms of the determinants of the track covariance over N steps ahead.
4. Course and velocity of every UAV (constant velocities are straight lines) to give minimum cost over N steps ahead using the Nelder-Mead simplex algorithm.
5. The optimization is able to be made subject to a constraint on the UAV velocity, to bound from below the closest approach to the target, to limit the UAV maneuver angle, and to avoid specified parts of the terrain.
6. The procedure is re-iterated every time the measurements are acquired.




Figure 18: (a) Three UAVs using TDOA; (b) Two UAVs using bearing-measurements, show the results of simulations of such scheduling. In each case, the UAVs fly out a base and position themselves to achieve optimal tracking. The thick lines are their trajectories optimally scheduled, in each case, over 10 steps ahead.




Figure 19 illustrates just one epoch of such scheduling in the bearings measurement case for given positions of the UAVs and a target estimate. The optimal course for the next epoch is also shown .





Figure 20: 2-D cost function for the possible courses of the two UAVs, UAV1 and UAV2. It is the implementation of the minimum value of this that is illustrated in Figure 15.

2.A.6. Georgia Tech Technical Progress

Algorithms Status

Work is continuing on optical flow applications on the CADSP Transform Imager. The primary efforts here have been the development and improvement of noise-robust OFE for use on the CMOS visible and UCIR applications. Work also continues on developing a transform imager simulator to be used by other groups in the Raytheon ISP team. Finally, a new OFE calculation technique is under investigation that will take advantage of the CADSP transform imager. This technique, referred to as checkerboard OFE will facilitate rapid OFE calculation at reduced resolution while allowing further processing to utilize full resolution.

Matlab tools for interfacing with the imager have been developed to allow for easy programming of the CADSP transform imager. Actual control of the imager operations is performed by code on an FPGA board via a daughterboard the holds the



imager IC. Communication to/from the PC is over an Ethernet cable. The boards and control code are all operational but development continues to add precision voltage and current test measurement capabilities to a daughterboard to make the system portable outside the lab. Work also continues to improve the communication speed over the Ethernet connection.

CADSP Hardware Status

Prior to the start of ISP Phase II, Georgia Tech designed and fabricated a mega-pixel CMOS transform imager. The characteristics of this imager are summarized below:

- $A'PB=Y$
- Design allows for up to 4 8×8 A matrices to be stored on chip
 - This is scalable to more matrices and different block sizes
- One 8×8 B matrix
 - This is also scalable, allowing for more matrices and different size matrices

- Input Matrix *A* can be bypassed with an External set of voltages
- Possible to read analog image values after first transform, before second multiplication
- Input for gain control can be adjusted at every read
 - This could allow block-by-block light-level normalization via an external feedback mechanism
- 8 ADCs allow for 8 parallel digital outputs, again scalable

Since coming under contract this chip has been tested and is operational but in need of various improvements. The primary deficiencies were related to the speed of the imager which runs too slow for most imaging applications. The speed problems were traced to a mistake in the integrating amplifier implementation and to a manufacturing process step that increased the opacity of the IC. Both of these problems are easily corrected for the next imager. After consultation with others in the Raytheon ISP team, a new imager is being designed that will be more useful for ISP systems.

The design for the new imager is currently on schedule and the imager will be implemented in stages. The first stage will be a small imager that can be submitted for manufacturing at the end of November 2005. The second stage will be a larger imager to be fabricated in January/February 2006. Major innovations being designed into the new imager architecture include:

- Multiple transform matrices that can be selected
 - This will facilitate data-adaptive algorithms
 - The goal is to have the ability to program one matrix while capturing images with the other matrix
- The achievable algorithmic shift of the transform blocks is altered to facilitate overlapping transforms
 - This system makes it possible to implement convolution
- The transform matrices will be larger (32x32 pixels)

2. B. Publications

There were no refereed publications that occurred during the current PoP.

1. Craig O. Savage and Bill Moran, "Waveform Selection For Maneuvering Targets Within An IMM Framework," IEEE Trans AES, accepted for publication.
2. A. Chhetri, D. Morrell and A. Papandreou-Suppappola, "Non-myopic sensor scheduling and its efficient implementation for target tracking applications," accepted for publication, EURASIP Journal on Applied Signal Processing, 2005.

2. C. Conference Proceedings

There were two publications in conference proceedings during the current PoP.

1. C. O. Savage, H. A. Schmitt, R. Cramer, and W. Moran, "Positioning and Scheduling UAVs for Passive Geolocation," Infotech@Aerospace, 26-29 September 2005, Arlington, VA.
2. A. Chhetri, D. Morrell and A. Papandreou-Suppappola, "Energy efficient target tracking in a sensor network using non-myopic sensor scheduling," IEEE Information Fusion, Philadelphia, PA, July 2005.

3. I. Kyriakides, D. Morrell and A. Papandreou-Suppappola, "Multiple target tracking with constrained motion using particle filtering methods," accepted to the Asilomar Conference on Signals, Systems, and Computers, Pacific Grove, California, October 30 - November 2, 2005.
4. I. Kyriakides, D. Morrell and A. Papandreou-Suppappola, "Sequential Monte Carlo methods for tracking multiple targets with stochastic kinematic constraints," invited to the First IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing, Puerto Vallarta, Mexico, December 2005.
5. S. Sira, A. Papandreou-Suppappola and D. Morrell, "Characterization of waveform performance in dynamically configured sensor systems," invited to the International Waveform Diversity and Design Conference, Kauai, Hawaii, January 2006.
6. W. Moran, C. O. Savage, S. Suvorova, H. A. Schmitt, D. E. Waagen and R. Cramer, "Dynamic Positioning and Scheduling of UAVs for Passive Geolocation," Session on cooperative dynamic systems, 2006 IEEE International conference on Networking, Sensing and Control, Ft. Lauderdale, FL, April 2006, abstract submitted.
7. R. Raich, J. A. Costa, and A. O. Hero III, "On Dimensionality Reduction for Classification and Its Application," 2006 IEEE International Conference on Acoustics, Speech and Signal Processing, submitted.
8. B. F. La Scala, M. Morelande, C. O. Savage, "Robust Target Tracking with Unreliable Binary Proximity Sensors," IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2006) , submitted.
9. A. Chhetri, D. Morrell and A. Papandreou-Suppappola, "On the use of linear integer programming for sensor scheduling in sensor networks," submitted to the 5th International Conference on Information Processing in Sensor Networks, Nashville, TN, April 2006.
10. I. Kyriakides, D. Morrell and A. Papandreou-Suppappola, "Multiple target tracking with constrained motion using particle filtering methods," Asilomar Conference on Signals, Systems, and Computers, Pacific Grove, California, October 30 - November 2, 2005.
11. I. Kyriakides, D. Morrell and A. Papandreou-Suppappola, "Sequential Monte Carlo methods for tracking multiple targets with stochastic kinematic constraints," invited to the 1st IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing, Puerto Vallarta, Mexico, December 2005

2. D. Consultative and Advisor Functions

There were two consultative or advisory functions 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 kindly given permission for several of these algorithms to be used in our ISP Phase II 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. Indeed, the DARPA NEST program hopes to demonstrate a 10,000 node network. We have identified and are evaluating several promising mathematical approaches to sensor localization developed by Al Hero (UM); these 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.

The second function relates to optical flow test facility at Eglin, Air Force Base. Raytheon and Georgia Tech have had preliminary discussion with Dr. T.J. Klausutis of Eglin AFB about the possibility of using their facility to evaluate the Georgia Tech CADSP imager being investigated on our ISP Phase II program. While these discussions are preliminary, Dr. Klausutis was interested in learning more about the capabilities and maturity of the CADSP Imager and plans to visit Georgia Tech this year.

2. E. *New Discoveries, Inventions or Patent Disclosures*

There were no patent disclosures filed during the current PoP. We are discussing the viability of patenting the elements of distributed tracking being developed by UniMelb.

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*

- [1] M. Maroti, B. Kusy, G. Balogh, P. Volgyesi, K. Molnar, A. Nadas, S. Dora, A. Ledeczi, "Radio Interferometric Positioning," *Technical Report, TR# ISIS-05-602*, Vanderbilt University, Nashville, Tennessee, 2005.
- [2] Jose A. Costa, Neal Patwari, and Alfred O. Hero III, "Distributed, Weighted Multi-Dimensional Scaling for Node Localization in Sensor Networks," *ACM Journal on Sensor Networks*, to appear.

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
CRB	Cramér–Rao Bound
CROPS	Classification Reduction Optimal Policy Search
DARPA	Defense Advanced Research Projects Agency
DSA	Distinct Sensing Area
dwMDS	Distributed, weighted, multi-dimensional scaling
FPA	Focal Plane Array
FMAH	Fast Mathematical Algorithms and Hardware
Georgia Tech	Georgia Institute of Technology
GPS	Global Positioning System

ISP Phase II (Contract N00014-04-C-0437)
Quarterly Progress Report (CDRL A001 No. 3)

IASG	Independently Activated Sensor Group
ISP	Integrated Sensing and Processing
IXO	Information Exploitation Office
kNN	k-Nearest Neighbor
LIP	Linear Integer Programming
MC	Monte-Carlo
MTT	Multi-target tracking
NEST	Networked Embedded System Technology
NLIP	Nonlinear Integer Programming
NLOS	NetFires Non-Line of Sight
NUC	Non-Uniformity Compensation
ONR	Office of Naval Research
PAM	Precision Attack Munition
PDA	Probabilistic Data Association
PWF	Polarization Whitening Filter
PoP	Period of Performance
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
UM	University of Michigan
UniMelb	Melbourne University
VM	Virtual Measurement
VU	Vanderbilt University