Raytheon

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

Progress Report:

2nd Quarter Progress Report
1 Jun 2006 – 31 Aug 2006

Acknowledgment of Support
This material is based upon work supported by the United States Air Force under Contract No. N00014-04-C-0437.

Contract No.: N00014-04-C-0437
Contract Line Item Number 0001
Deliverable Item: Publications (A001-006)

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

Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED
22 AUG 2006		

4. TITLE AND SUBTITLE	5a. CONTRACT NUMBER
Integrated Sensing Processor - Phase II	**N00014-04-C-0437**
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)	5d. PROJECT NUMBER
Harry Schmitt	5e. TASK NUMBER
	5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	8. PERFORMING ORGANIZATION REPORT NUMBER
Raytheon Missile System,1151 Herman Rd,Tucson,AZ,85706	

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	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			
unclassified	**unclassified**	**unclassified**		**38**	

01 September 2006

Progress Report

CDRL A001 No. 6
Fifth Quarterly Progress Report for Period of Performance
1 June 2006 – 31 August 2006

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

0. Technical Abstract 3
1. A. Program Summary 4
1. B. Program Status 4
1. C. Personnel Associated/Supported 4
1. D. Recent Accomplishments and Events 5
1. E. Near Term Events 5
2. A. Technical Progress 6
2.A.1. Raytheon Technical Progress 6
2.A.2. ASU Technical Progress 10
2.A.3. UM Technical Progress 16
2.A.4. UniMelb Technical Progress 18
2. B. Publications 34
2. C. Conference Proceedings 35
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 37

INDEX OF FIGURES

Figure 1: Separability of two classes as a function of dimension for four different training set selection techniques and the baseline generated by training with the complete data set. Training for four different techniques is performed using only 25% of original data set. 7
Figure 2: Test Bed Block Diagram. 8
Figure 3: Virtual Measurement Tracker Results. 8
Figure 4: UKF Tracker Results. 9
Figure 5: Sample of the Buffer which includes the Time Stamp. 9
Figure 6: Spectrogram and time domain plot of the acoustic signal generated by a person wearing boots walking on sand at a distance of 12 feet from the mote. 10
Figure 7: Spectrogram and time domain plot of the highpass filtered footstep data 11
Figure 8: Spectrum of footstep data before and after filtering. The top plot also shows the highpass filter response 11
Figure 9: ROC curves for the energy detector using the highpass filtered footstep data at a distance of two to twelve feet from the mote 12
Figure 10: ROC curves for the energy detector using the highpass filtered footstep data at a distance of fourteen to twenty feet from the mote. 12
Figure 11: Transfer function of the highpass filter implemented in the mote program 13
Figure 12: Simulated CADSP image with two targets. 14
Figure 13: RMSE performance of the tracker using the simulated imager (ISP) and the complete image frame (No ISP). 14
Figure 14: Average number of blocks requested from the simulated CADSP imager (ISP) compared to the total number of blocks in the image (No ISP) 15
Figure 15: Average number of lost tracks versus the frame number. 15
Figure 16: Estimated probability of n=0, 1, and 2 targets. 16

Figure 17: Photo of deployment environment. Here, Experiment III is photographed. All tests were run with sensors placed on the ground, in the grass. A laptop listens to and records the sensor packet traffic for later study; but is not involved in the coordinate calculation. 17
Figure 18: Ideal anchor geometry. The anchor nodes are shown by black squares, green circles indicate locations were the node location is unique, red asterisks indicate an ambiguous solution and black dots indicate no solution is possible. 21
Figure 19: Proportion of a beacon's region that can be uniquely localized as a function of relative range and the angle between the two anchors. 21
Figure 20: Proportion of a beacon's region that can be uniquely localized as a function of relative range and the angle between the two anchors. The maximum is indicated by the red asterisk. 22
Figure 21: Example of a single run of the collaborative localization algorithm. Beacons are indicated by black squares, black circles show the true location of the standard nodes, while green asterisks show their estimated locations. When an unambiguous estimate could not be found, this is indicated by a red asterisk. 22
Figure 22: Histogram of the average position errors for each standard node. 23
Figure 23: Simulation scenario layout. The location of the node that could not be localized, node 15, is shown by the blue asterisk in the lower part of the figure. The red asterisks indicate nodes whose estimated locations had significant errors. 23
Figure 24: Self-localization scenario. The positions of GPS-equipped motes are indicated by circles. The positions of unknown motes are indicated by crosses 26
Figure 25: Estimator bias plotted against measurement noise variance for motes F (solid), G (dashed) and H (dotted) 26
Figure 26: Estimator standard deviation plotted against measurement noise variance for motes F (solid), G (dashed) and H (dotted). 27
Figure 27: RMS position error for the UKF (dashed) plotted against the mote grid size for (a) the non-thresholded measurement model and (b) the thresholded measurement model. The solid line is the posterior CRB. The SNR is 15 dB. 28
Figure 28: RMS position error for the UKF (dashed) plotted against the mote grid size for (a) the non-thresholded measurement model and (b) the thresholded measurement model. The solid line is the posterior CRB. The SNR is 10 dB. 28
Figure 29: RMS position error for the UKF (dashed) plotted against the mote grid size for (a) the non-thresholded measurement model and (b) the thresholded measurement model. The solid line is the posterior CRB. The SNR is 5 dB. 28

INDEX OF TABLES

Table 1: Deployment Test Results 17
Table 2: Utility Function for Piecewise Linear Waveform Library for various Covariance Distributions 34

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; however, the revised schedule still supports demonstration before 28 February. We still expect to complete the contract on time and budget. ***One area of significant concern remains the availability of a suitable radar test and integration engineer. We continue to work this issue, but have yet to find an acceptable solution. This problem remains one of the higher risks for our program.***

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
Mr. William Daniels	Radar Test and Integration Support

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 was one significant personnel change during the current PoP. Dr. Neal Patwari has left the University of Michigan and taken a position at the University of Utah.

1. D. Recent Accomplishments and Events

- Integrated MATLAB code for two UniMelb 1-Bit trackers into the Raytheon Test Bed and collected preliminary test data in July 2006. Provided results and feedback to tune UniMelb researchers.
- Received University of Michigan dwMDS self-localization algorithm in July 2006. Integrated the TinyOS code into the Raytheon Test Bed in August 2006.
- Presented "Comparison of Inter-class Divergence for Linear and Nonlinear Dimensionality Reduction, with and without Class Labels" and "Correlation of Inter-class Divergence and Classification Performance," at the *Combat Identification Systems Conference (19-22 June 2006, Orlando, Florida).*
- Released MATLAB simulation code used for the analysis in the "Cooperative Control of Multiple UAVs for Passive Geolocation," paper to Professor Daniel Pack, United States Air Force Academy. This code is understood to be at a research-level and used for that purpose.
- Supported the ISP-II PI meeting (13-15 June 2006 in Phoenix, Arizona). Provided detailed responses (7 July 2006) to action items arising from the PI meeting.
- UM (Patwari) presented our technical results in the area of indirect radio interferometric localization at the EmNets 2006 workshop.
- Raytheon (Harry Schmitt) attended Technical Interchange Meeting at the University of Melbourne 5-13 August 2006.

1. E. Near Term Events

- Collect preliminary data set for the University of Michigan dwMDS self-localization algorithm and provide feedback
- Present "Localization, Detection and Tracking for Wireless Sensor Networks" at the *MSS BAMS* Conference. Our goal is to bring to the attention of the military

sensing community some of the work that is currently being done under ISP-II. We have provided summaries of algorithms which have been (or are being) developed for use in our distributed tracking demonstration. We discuss the dwMDS algorithm for self-localization (UM), the energy detector algorithm (ASU) and the Virtual Measurement tracking algorithm (UniMelb).

2. A. Technical Progress

2.A.1. Raytheon Technical Progress

2.A.1.a. Sparse Manifold Learning

We are interested in performing dimensionality reduction for the purpose of classification of SAR images using only a limited number of training points. The reason for this is to reduce the storage and computational requirements when performing real-time low-dimension embeddings of newly acquired points. We have tested the efficacy of a number of different training point selection techniques in addition to proposing our own algorithm. We show that our technique outperforms the existing methods by providing more exact low-dimensional embeddings using fewer training points. In these experiments we make use of the ISOMAP algorithm for performing the dimensionality reduction. The data set comes from the MSTAR database of SAR image data. The data is pre-processed using a CFAR algorithm that finds the target on the original SAR image and focuses on them using a 50 pixel by 50 pixel window. Below we describe the different set selection techniques:

Random Selection of points:

The simplest landmark selection technique is a random selection of points from the training set. This technique works fairly well provided that the number of points is sufficient.

MaxMin Greedy Optimization:

For selecting only a few training points, [Silva&Tenenbaum] suggest using a MaxMin greedy optimization technique that maximizes the separabilty of the selected points. This technique tends to outperform the random selection technique (in terms of maintaining manifold structure) when the number of selected landmarks is small, however as the number of landmarks increases they work equally well.

Clustering:

Another method for the selection of landmarks is to perform logical groupings and use a few training samples from each group. These groupings can be based on a priori information (*i.e.* groupings by azimuth angle for SAR data) or they can be data driven groupings (*i.e.* obtained using a clustering algorithm).

Proposed technique:

We are interested in reducing the storage and computational requirements for performing a low dimensional embedding of newly acquired points. Since the actual training is performed offline, we can afford to train with the whole data set. Rather than storing all the parameters of the full training procedure, however, we reduce the parameters by averaging the parameters associated with neighboring points in the original high-dimensional space. The assumption is that the candidate point to be embedded is approximately equidistant from training points in the same neighborhood. This makes it possible to combine the parameters associated with the neighboring points in the Nystrom

approximation. The actual storage and complexity requirements for performing the out-of-sample extension then become the same as the techniques shown above. The only drawback to the technique is that a longer training time is required, however, since the training is performed offline, this is not problematic.

We compare the efficacy of these techniques using a divergence measure technique that provides a measure of separability between two classes (). The graph-based technique is a multivariate generalization of the Wald-Wolfowitz runs test. A minimal spanning tree is constructed from the union of the two samples with class (sample) labels. The test statistic R corresponds to the number of disjoint subgraphs generated by removing all edges of the tree that connect vertices (data points) with differing labels. In these results we show the separability as a function of reduced dimension for two classes of SAR images (an M-2 tank and a T-72 tank).

In Figure 1 below, we present the divergence of the two data sets at different dimensions. We compare the four different set selection techniques described above (with only 25% of the original data) to the baseline generated using the complete data set. The MaxMin technique performs equally well to the random selection technique and has therefore been omitted from the plot. As the plot shows, the proposed technique (parameter averaging) outperforms the other methods for all dimensions and approaches the baseline for higher dimensions. The storage and complexity requirements for out-of-sample extensions are identical for all techniques. A technical manuscript is currently in preparation that will discuss our research results related to landmark selection.

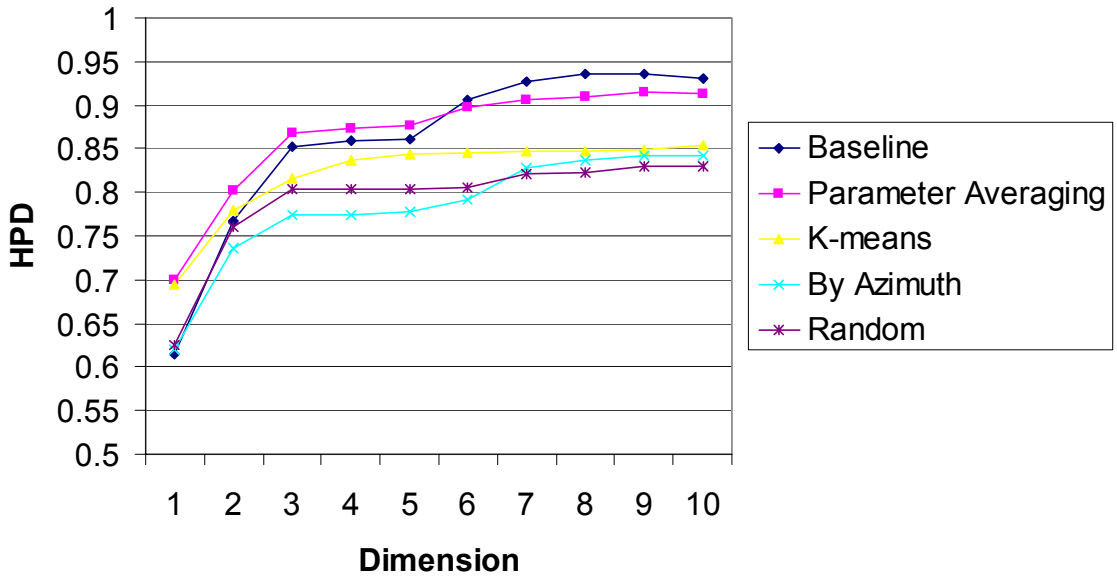

Figure 1: Separability of two classes as a function of dimension for four different training set selection techniques and the baseline generated by training with the complete data set. Training for four different techniques is performed using only 25% of original data set.

2.A.1.b. Test Bed Demonstration: Detector and Tracker Integration

At the Program Review held in Phoenix last June, we presented a test bed for the Distributed Tracking Demonstration. This test bed, shown in Figure 2, gives the capability to easily integrate and test a variety of wireless sensor detectors and trackers. In June, we had integrated the Particle Filter Tracker from Arizona State University and

only one of the University of Melbourne's trackers. At this time, we have completed the integration of the second tracker from the University of Melbourne. In addition to that, we have also created a Target Simulator for testing the test bed offline. The target simulator simulates the motion of a target that moves randomly through a field of motes, and the motes ID numbers are written to the Buffer (refer to Figure 2).

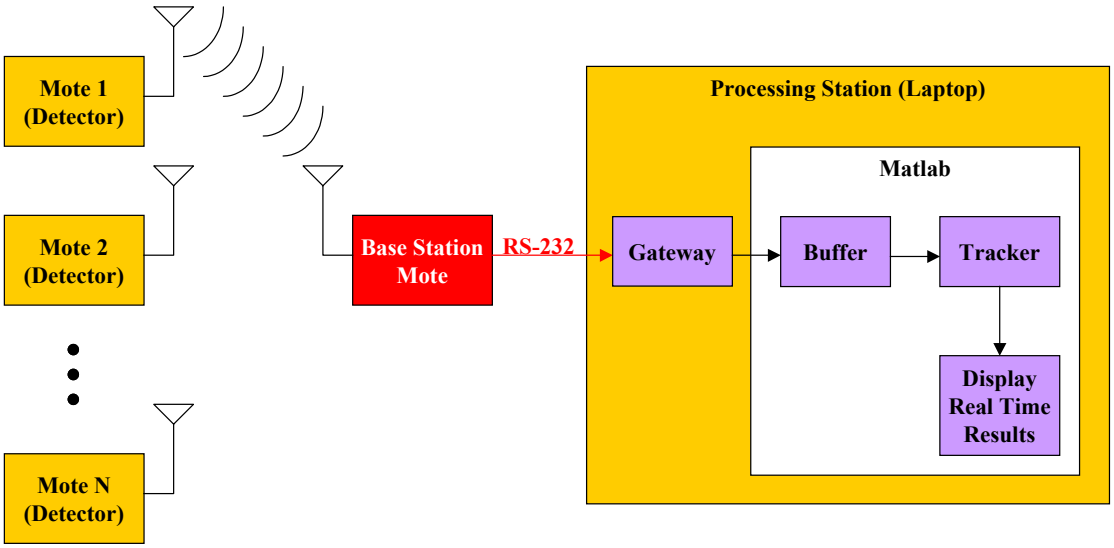

Figure 2: Test Bed Block Diagram.

The results using the target simulator are shown in Figure 3 and 4 for both of the University of Melbourne's trackers.

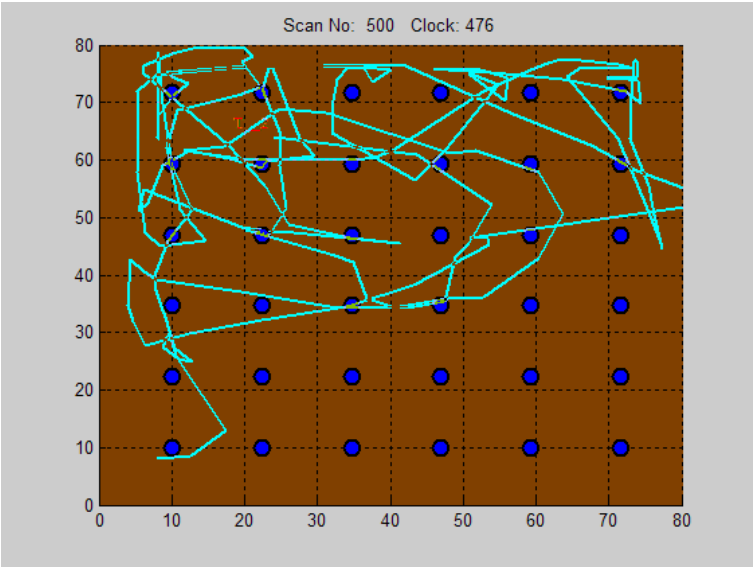

Figure 3: Virtual Measurement Tracker Results.

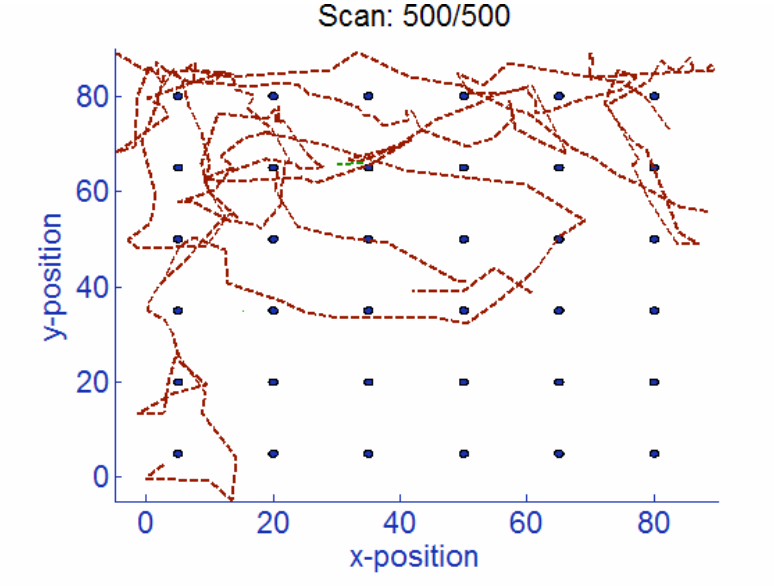

Figure 4: UKF Tracker Results.

Time-Stamping the detections was not available, and it is critical to playback the data offline to tune the trackers. Thus, we have used the provided by TinyOS method which we found to be inaccurate. The inaccuracy (i.e., bug) has been corrected and Time-Stamping the motes network detections is now possible. These Time Stamps are incorporated into the buffer of our test bed (refer to Figure 2). A sample of this buffer is shown on Figure 5 where its format is as follows. The first column consists of the Time Stamp, and its format is HH:MM:SS.SSS (i.e., Hours, Minutes, Seconds, and Thousand of a Second). The second column is an integer that represents the Mote ID that has detected the target.

```
10:29:05.940   12
10:29:05.960   12
10:29:05.980   19
10:29:05.990   2
10:29:06.010   19
10:29:09.204   2
10:29:09.224   19
10:29:09.234   19
10:29:09.244   2
10:29:12.468   19
10:29:12.508   2
10:29:15.732   2
10:29:15.772   19
10:29:18.976   19
10:29:18.986   8
```

Figure 5: Sample of the Buffer which includes the Time Stamp.

It has been noticed that as the motes network covers a large area, the reception of the detections degrades. Thus, some version of a network routing is needed to ensure reception from motes that are far away from the base station. The Vanderbilt FloodRouting is a fairly well qualified (not in the military sense) package for alleviating this problem. We have made progress in integrating the FloodRouting package into the

motes for the test bed demo. This task is not yet completed; however, its completion is estimated within 80 hours of effort.

2.A.1.c. Test Bed Demonstration: Self-Localization

We have made progress on integrating the Vanderbilt Radio Interferometer Positioning System (RIPS) code for self-localization on to the MICA2 900 MHz version. RIPS was originally written for the MICA2 433 MHz version, and it is highly dependent on radio frequency and phase differentiation between received signals. Operation of the Vanderbilt RIPS on the 900 MHz MICA2 would require significant effort to complete. Experiments with the Vanderbilt RIPS code will be conducted solely onto the 433 MHz MICA2. The completion of these experiments is estimated within 40 hours of effort.

We have received the self-localization work from University of Michigan. The code has been modified to comply with the TinyOS standard and TinyOS make system such that integration to the test bed demo is possible. Data has been collected for two topologies using the University of Michigan's code. Further testing will be conducted next week as well as the analysis of the data.

2.A.2. ASU Technical Progress

2.A.2.a. Mote Tracking Support: Detector Development

Data Analysis and Filter Design

The footstep data was analyzed using the spectrogram time-frequency representation (squared magnitude of short-time Fourier transform) to compute the frequency content of the footsteps as a function of time. Figure 6 shows the spectrogram of the acoustic signal generated by a person wearing boots walking on sand at a distance of 12 feet from the mote. The time-frequency plot shows that the footsteps have significant energy throughout the entire frequency band; it also shows that there is a significant low frequency noise component in the data.

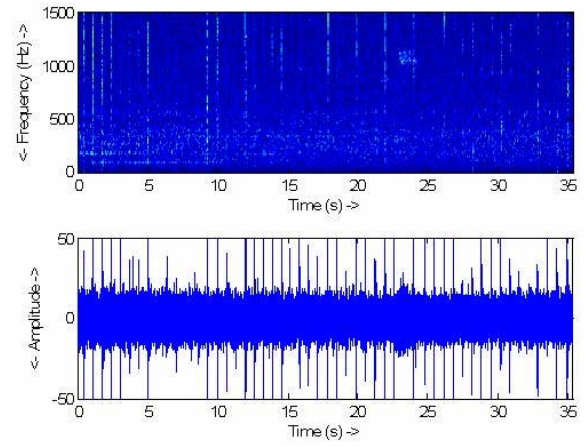

Figure 6: Spectrogram and time domain plot of the acoustic signal generated by a person wearing boots walking on sand at a distance of 12 feet from the mote.

This time-frequency analysis suggested the use of a highpass filter to eliminate the lower frequency noise while retaining the high frequency components of the footsteps. So we used the Remez algorithm in MATLAB to design a 21st order filter with a stopband edge of 606 Hz and cut off frequency of 1.0606 kHz. The resulting

filtered footstep data is shown in Figure 7. Figure 8 shows the spectrum of the original and filtered data as well as the frequency response of the filter.

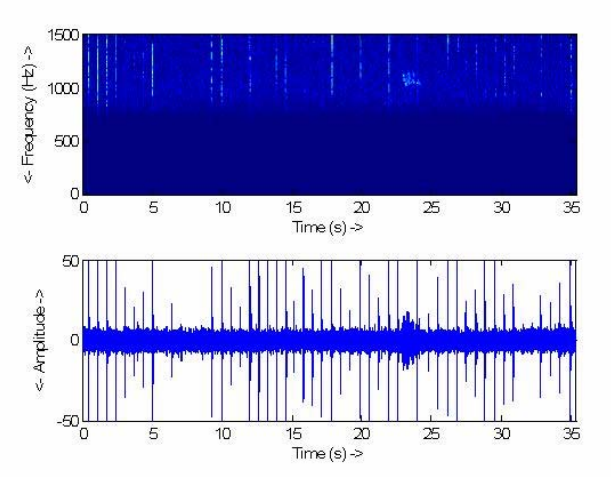

Figure 7: Spectrogram and time domain plot of the highpass filtered footstep data

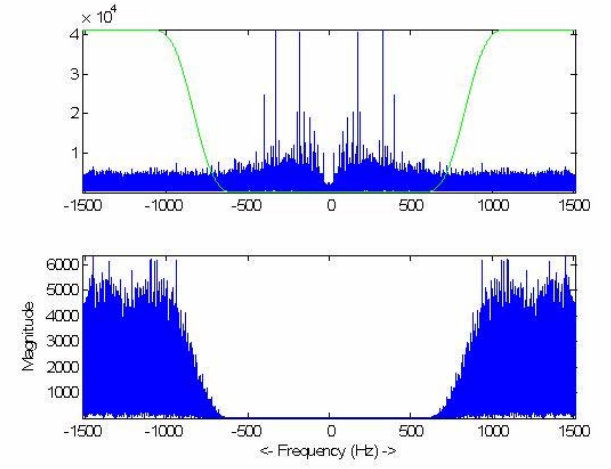

Figure 8: Spectrum of footstep data before and after filtering. The top plot also shows the highpass filter response

Data Collection

Using the modified mote collection system in which acoustic data is collected using a mote that communicates to the base station using a wired connection, we have collected a large database of footstep data. This database includes footsteps on several different surfaces (sand, brick, and cement) with several different shoe types. The data has been made available to the Raytheon team on the World Wide Web at http://www.fulton.asu.edu/~murishm/debejyo_web/motes/. For each combination of surface and shoe type, footstep data was collected at distances from two feet to twenty feet away from the mote. The sampling frequency was approximately 3.1 kHz with eight bits per sample.

Receiver Operating Characteristic (ROC) Curves of Energy Detector

In our last quarterly report, we discussed how we obtained ROC curves for the energy detector. With the highpass filtering operation, we now have improved ROC

curves. These are provided in Figure 9 (for footsteps at distances from two to twelve feet) and in Figure 10 (for footsteps at distances from fourteen to twenty feet).

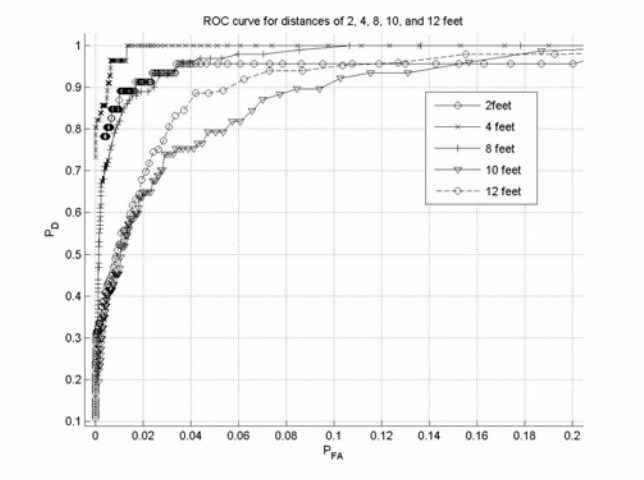

Figure 9: ROC curves for the energy detector using the highpass filtered footstep data at a distance of two to twelve feet from the mote

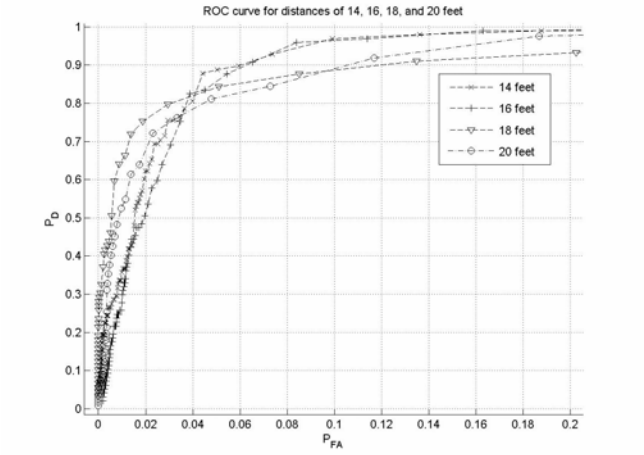

Figure 10: ROC curves for the energy detector using the highpass filtered footstep data at a distance of fourteen to twenty feet from the mote.

Mote Software Development

We have developed software for the motes that collects acoustic data at a sampling rate of 3.1KHz, filters the data with a highpass filter, and then computes the energy in the filtered data. Beginning with an initial implementation, we have significantly improved the speed and robustness of the code. The highpass filter is a 9-tap FIR filter with integer coefficients whose response is shown in Figure 11. The energy detector computes the energy in a 500 sample window of the filter output and compares the energy to a threshold. If the energy exceeds the threshold, the mote detects the presence of the footstep and transmits a packet that includes the mote ID number and the energy value. A video demonstration of the detection on the motes using two mote sensors can be found at
http://www.fulton.asu.edu/~murishm/debejyo_web/motes/videos/energy_detector.AVI

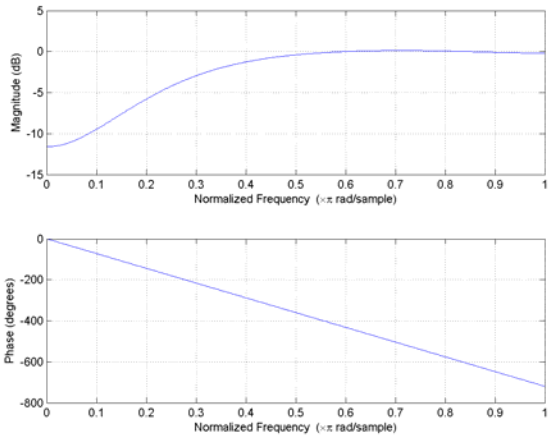

Figure 11: Transfer function of the highpass filter implemented in the mote program

2.A.2.b. Tracking Algorithms for the CADSP Configurable Imager

We have extended our previous single target tracker for the CADSP imager to track a varying number of targets; the tracker estimates the number of targets as well as the target state (target size and 3-dimensional position and velocity). In our current implementation, the tracker uses image data provided by a simulated CADSP imager whose configuration is under the control of the tracker; the imager is simulated by a web-camera whose output is filtered to simulate the CADSP imager. Our tracker is composed of the CADSP imager configuration strategy and the target tracking algorithm.

The imager configuration strategy determines which 8 × 8 pixel blocks will be requested from each image frame acquired by the imager as well as the type of filtering operation (Gaussian or Mexican hat filter) to be performed on each of these blocks. The tracking algorithm uses the output of the CADSP imager as measurements in the particle filter state estimator. It estimates the number of targets present in the scene and estimates the state of the targets. The tracking algorithm is adapted to the scene using training data to determine distributions of foreground and background image blocks.

The multiple target tracker has been evaluated for seven different video sequences. In three of the sequences the number of targets varies from zero to two; in the other four, the number of targets varies from zero to one. An example of an image with two targets is shown in Figure 12. The image blocks requested from the imager by the tracker at a given time are shown as white blocks. Which blocks are requested are determined by the tracker as follows. Each particle in the particle filter predicts the target location; for each predicted location, several image blocks in the center and on the edges of the predicted target location are requested from the imager.

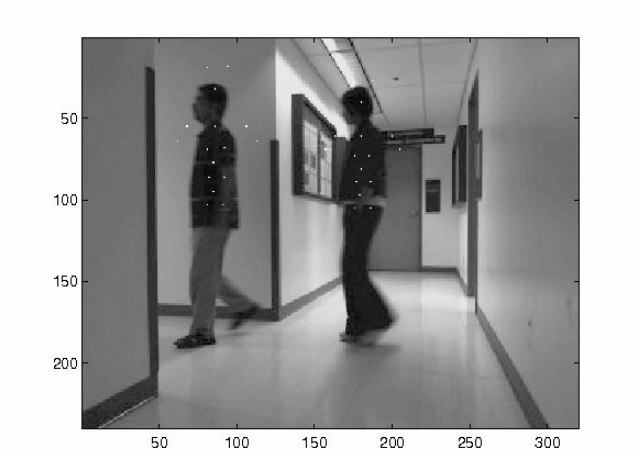

Figure 12: Simulated CADSP image with two targets.

Figure 13 shows the RMSE error of the tracker for the three video sequences with zero to two targets. The times of target arrival and departure, as well as the times where one target occludes another, are approximately the same in each of the three video sequences. The RMSE performance for the case where the simulated imager is used (labeled "ISP" in Figure 13) is comparable to the case where the complete image frame is used (labeled "No ISP"). This indicates that there is no performance degradation when the CADSP imager acquires only a subset of blocks in the image.

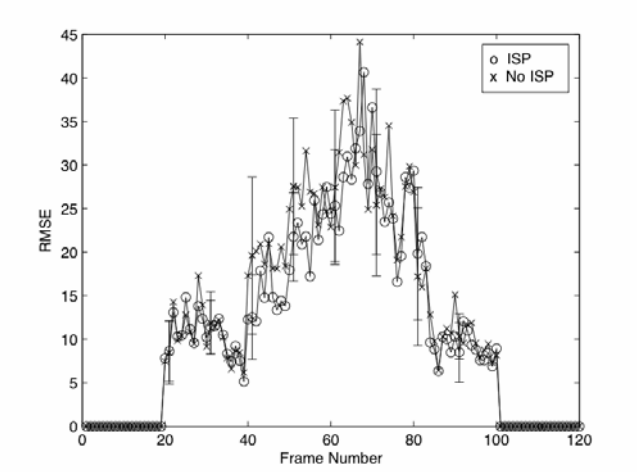

Figure 13: RMSE performance of the tracker using the simulated imager (ISP) and the complete image frame (No ISP).

Figure 14 shows the average number of blocks requested from the simulated CADSP imager (labeled "ISP") compared to the total number of blocks in the image (labeled "No ISP"); this figure shows that selectively requesting blocks that are important to the tracker can significantly reduce the number of blocks needed to be acquired by the imager.

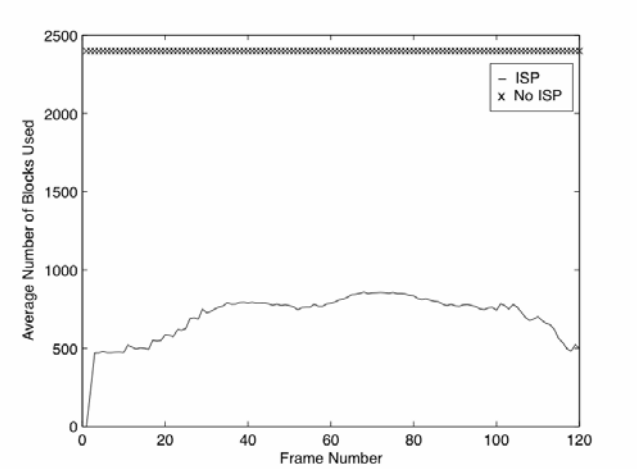

Figure 14: Average number of blocks requested from the simulated CADSP imager (ISP) compared to the total number of blocks in the image (No ISP).

Figure 15 shows the average number of lost tracks versus the frame number. The average number of lost tracks is close to 0.8 when one target is occluded by another target and is thus not observed by the camera.

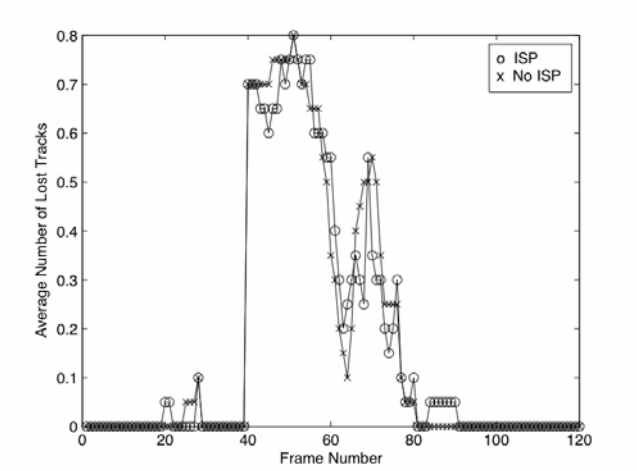

Figure 15: Average number of lost tracks versus the frame number.

Figure 16 shows the estimated probability of zero, one, and two targets; the true number of targets in the scene is also indicated. The results were averaged over 20 Monte Carlo simulations.

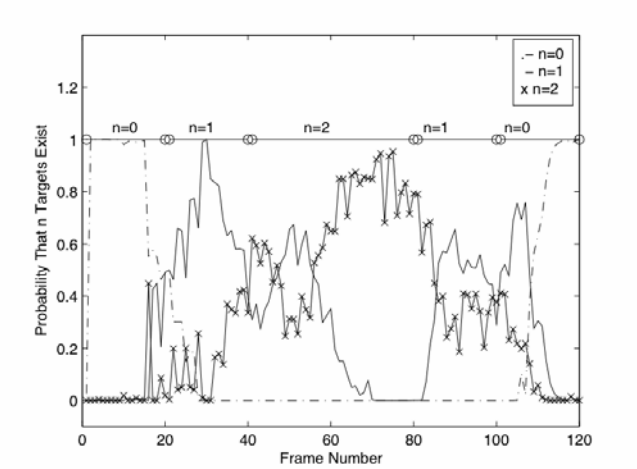

Figure 16: Estimated probability of n=0, 1, and 2 targets.

2.A.3. UM Technical Progress

In the three months since the last quarterly report we have made incremental progress on two fronts:

- We have made more progress on CCDR out of sample extension and will be submitting a paper to a journal in a couple of weeks. The progress is in applying the method to very large datasets where learning from the training sample is not practical given the size of the SVD. We are testing the training phase of CCDR with an OSE that allows labeled points to be included as OSE candidates. A copy of the paper and the related code will be delivered to Raytheon once we get it to a draft stage that is presentable.
- We are testing the GEM code and are trying to characterize when there are insufficient numbers of points to construct a reliable non-anomaly region. In conjunction with this, we are developing rules of thumb that can predict the number of points in training sample that are sufficient to detect an anomaly of given size - *i.e.*, distance from the non-anomaly region.

Most of the progress during the current PoP has been related to mote self-localization for the distributed tracking demonstration.

At the last quarterly report in May, we reported that the progress towards distributed localization algorithm implementation was at 75%. During the past quarter, we finished this implementation, and tested the distributed localization system in a deployment test in an outdoor field. This implementation used single-frequency measurements of received signal strength (RSS) between many pairs of neighboring sensors. In the deployment test (Experiment I) we placed 14 sensors in a 4x4 grid in a 4m by 4m square grassy area, allowed them to make measurements and then to perform the distributed coordinate calculation of the dwMDS algorithm described in [Costa2006]. In Experiment I, the root-mean squared error (RMSE) was 64 cm.

Figure 17: Photo of deployment environment. Here, Experiment III is photographed. All tests were run with sensors placed on the ground, in the grass. A laptop listens to and records the sensor packet traffic for later study; but is not involved in the coordinate calculation.

Following the successful deployment test using single-frequency RSS measurements, we upgraded the system to allow multiple-frequency RSS measurements. This required implementing a frequency-hopping (FH) protocol on the motes. In this implementation, motes measured the RSS of signals at 16 different center frequencies and averaged the result. The benefit of FH is that the variance of the RSS measurement due to frequency-selective fading can be sharply reduced. Using the new FH implementation, ran a new test (Experiment II) with 16 sensors in the same grid and 4m square area described above. In Experiment II, the RMSE reduced to 25.6 cm.

Finally, we made some adjustments to the implementation so that the system would allow for many more sensors, and then ran a deployment test (Experiment III) with 36 sensors in a 6x6 grid, in a 6.67m by 6.67m square area. This test ran successfully and recorded an RMSE of 55.3 cm.

Table 1: Deployment Test Results

Experiment	Sensor Number	Area Size	Frequency Hopping?	RMSE
I	14	4 x 4 m^2	No	64 cm
II	16	4 x 4 m^2	Yes	25.6 cm
III	36	6.67 x 6.67 m^2	Yes	55.3 cm

Ongoing work with the sensors will use the implementation experience gained from these deployment tests to further develop the dwMDS implementation. While the current dwMDS implementation separates measurement and calculation into two distinct stages, the next step will be to join them, so that movement can immediately impact position estimation, which is a key requirement for sensor tracking. We proposed and were accepted to conduct a distributed localization demonstration at the ACM SenSys 2006 conference in November [Patwari2006], which will incorporate this new update.

2.A.4. UniMelb Technical Progress

2.A.4.a Raytheon Technical Support

In the previous quarter, work has been performed in three ongoing research areas:

1. Theoretical scheduling
2. Geolocation of ground targets
3. Tracking smart targets

Additionally, through National Information and Communication Technology of Australia (NICTA), some gene expression data has been obtained for evaluation of some of the high-dimensional data processing algorithms developed under ISP.

Theoretical Scheduling

As reported previously, we have considered optimal scheduling of Gauss-Markov systems (GMS) under certain constraints for a terminal cost function. Optimality in this case refers to minimizing the sum of the estimated state error variance, or some function of the estimated state error covariance for vector-valued states. We are considering removing those constraints, in addition to using different cost functions (e.g. discounted cost). Along these lines, we have investigated rates of convergence of iterative equations [OrRh] to quantify reduction in the estimated state error covariance as measurements of each process are taken.

Research in this area is ongoing. The eventual goal is to determine an optimal scheduling methodology for general GMS, but the majority of the work to date has been on scalar systems.

Geolocation of Ground Targets

In previous reports, we have outlined the problem of geolocating a ground target utilizing a laser rangefinder and distributed passive sensors. Our scheduling solution addresses a number of practical concerns, such as

1. Scheduling passive sensors to conserve bandwidth and power
2. Dynamic measurement quantization based upon track information
3. Dynamic measurement rate scheduling to preserve track quality
4. Evaluations of multiple-step-ahead methods, such as the *rollout* algorithm [BeC99]

The work is being prepared for submission to *IEEE Transactions on Aerospace and Electronic Systems*.

Tracking Smart Targets

We continue research on tracking smart targets as described in our previous report. A paper is being prepared for the upcoming DASP conference, to be held in

December. Furthermore, we have begun research along the lines of Gittins [Gi79, GiRo79], which concerns itself with detection of an intelligent evader, as opposed to detection.

2.A.4.b Distributed Tracker Technical Support

Mote Self-localization in Distributed Networks

In the June 2006 report a computationally efficient method for mote self-localization using Radio Interferometric Positioning (RIPS) measurements was described. This algorithm assumed that all motes were in communication range of each other. However, typically nodes in a wireless sensor network are cheap and low-powered. Thus, their communication ranges are often significantly less than the area covered by the entire network. Therefore, practical self-localization algorithms must operate in a multi-hop fashion.

There are two main types of multi-hop localization — centralized localization and collaborative localization. In centralized localization all measurements are transmitted to a central processor which computes the locations of each node in the network. This is an inefficient method as it introduces bottlenecks in the communication channels from the nodes to the central processor and an undesirable power drain on the nodes one hop from it. In contrast, with collaborative localization the nodes in the network operate in a peer-to-peer manner to build up a map of the network.

Here we present an extension of the RIPS-based localization method described in the earlier report that is suitable for collaborative localization in a network where the surveillance region is larger than the communication range. In order to do this, we require that a small number of nodes, in known locations, are available. We will term such nodes beacons. Note, this requirement is not needed in other RIPS-based methods such as that of [Maroti2005]. However, use of beacons reduces the required number of RIPS measurements to a number that is linear in the number of nodes to be localized. This is a critical cost saving.

Collaborative Localization Algorithm

Consider the case when there is a sparse network of GPS-equipped nodes, i.e. beacons, spread across a surveillance region, whose size is significantly larger than the communication range of the nodes. In addition to the GPS nodes, there are many standard, non-GPS equipped nodes spread randomly through the area. The section of the entire surveillance region within the range of a GPS node will be referred to as its region. These GPS beacons are placed so that their regions intersect, so that some standard nodes may be in the region of more than one beacon.

In addition, we assume that an ad hoc network has been formed between all the nodes, so that each node knows the identity, but not necessarily the location, of its neighbors. That is, each node knows which other nodes are within one hop of itself.

Finally, we will assume that within the region of one beacon there are two other nodes in known locations. This is required to initialize the algorithm. An initialization method that does not have this requirement, but instead uses standard nodes in the intersection of three beacons, was discussed briefly in the previous report. Implementation of this method is currently being studied.

The localization algorithm is outlined below. Details of the **Select Anchors** step are discussed in detail below. The **Find Nodes** step uses the efficient localization algorithm described in the previous report. A successful localization attempt is one where a unique location for the node could be computed.

1) While there are unused GPS nodes (i.e. beacons) do
   a) **Select Beacon**: From the list of unused beacons, select the beacon, b, with the largest number of nodes in its region that are already localized. Create a list L of nodes whose locations are known and are in the region of b. Generate another list A of all possible pairs in L that are in range of each other and b.
   b) **Select Anchors**: Given the beacon b and the list A select the best pair of nodes to use as anchors along with b.
   c) **Find Nodes**: Given the anchor set, for each node n that is in range of all three anchors do
      i) Attempt to localize n, if the number of previous successful attempts to localize it is less than the allowed maximum.
      ii) If an unambiguous location is found (i.e. the localization attempt was successful) then
         (1) Update the estimated location of n by averaging the estimated position with the previous estimate.
         (2) If node n was not previously localized, add it to the list L. For every other node, m, that is in the list L and is in range of both the beacon and node n, add the pair (m, n) to the list A.
   d) Remove the current pair of nodes that were used as anchors from the list A.
   e) If the length of the list A is non-zero, return to the **Select Anchors** step.
   f) Remove the beacon from the list of unused beacons.

The indicative performance of this algorithm is examined using simulations and the results are discussed below.

Ideal Anchor Geometry

To localize standard nodes in the region about each beacon, two additional nodes, whose locations are already known, are chosen to be anchors. This pair is such that all three nodes are all in range of each other. From all such possible pairs, the pair that is chosen is such that the probability that an unknown node's position can be unambiguously calculated is maximized. That is, there is a valid solution to the closed-form RIPS measurement equation and this solution is unique.

The proportion of a region that can be uniquely localized with any given set of anchors cannot be computed exactly. However, numerical approximations indicate that the ideal choice of anchors is a pair that is equidistant from the beacon, at a range of 5% of the communication range of the nodes. The ideal angle between the anchors and the beacon is approximately 87^{o}. This results in just over 50% of the region being uniquely localizable. This geometry is illustrated in Figure 18.

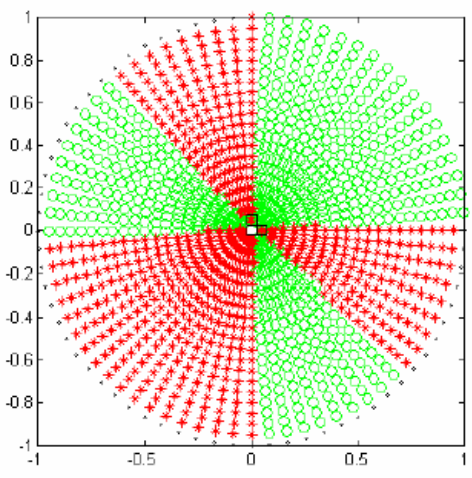

Figure 18: Ideal anchor geometry. The anchor nodes are shown by black squares, green circles indicate locations were the node location is unique, red asterisks indicate an ambiguous solution and black dots indicate no solution is possible.

Results of the numerical approximations for equidistant anchors are shown in Figures 19 and 20. The results shown in these figures were calculated assuming the two anchors were equidistant from the beacon and the range is given relative to the communication range of the nodes. The possible angle between these two anchors at the beacon was in the range $(0, \pi)$. Figure 19 shows the results as a three dimensional surface plot, while Figure 20 is the corresponding contour plot.

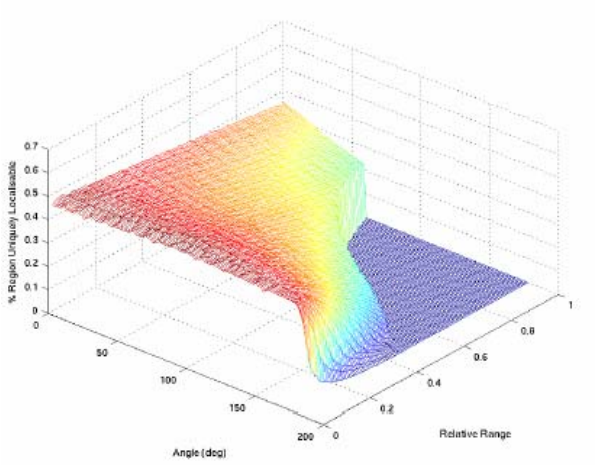

Figure 19: Proportion of a beacon's region that can be uniquely localized as a function of relative range and the angle between the two anchors.

The **Select Anchors** step in the algorithm uses linear interpolation over a pre-computed table of possible anchor locations to estimate the proportion of the region that is uniquely localizable. The pair with the highest proportion is selected to be the anchors, along with the beacon.

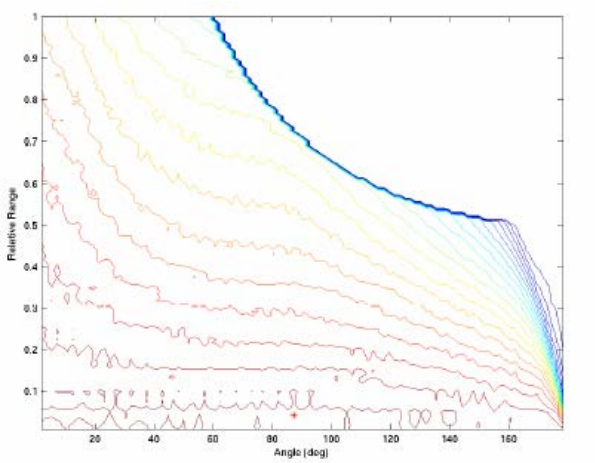

Figure 20: Proportion of a beacon's region that can be uniquely localized as a function of relative range and the angle between the two anchors. The maximum is indicated by the red asterisk.

Simulation Results

The indicative performance of the previously described algorithm described has been tested using simulations. The simulation scenario consisted of a surveillance region that was 300m square. There were 25 GPS-equipped nodes and 50 standard nodes in the region. The communication range of all nodes was 80m. The RIPS measurements were assumed to be corrupted by zero mean, white, Gaussian noise with a standard deviation of 0.04m. This is the accuracy that Maroti *et al.* state they were able to achieve in their field trials of the RIPS technique in [Maroti2005]. Ten Monte Carlo runs were performed. At most 5 attempts were made to localize a node during each run. An example of the output of the algorithm on a single run is shown in Figure 21.

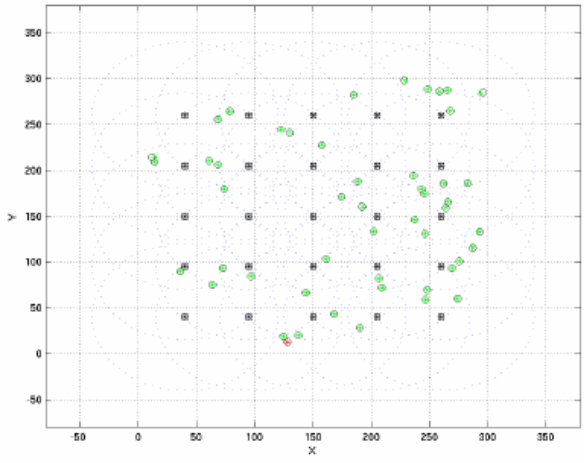

Figure 21: Example of a single run of the collaborative localization algorithm. Beacons are indicated by black squares, black circles show the true location of the standard nodes, while green asterisks show their estimated locations. When an unambiguous estimate could not be found, this is indicated by a red asterisk.

The average position accuracy was 42cm, while the median accuracy was 23cm. This pronounced difference is due to a number of outliers. A histogram of the average position errors for each standard node is shown in Figure 22 which illustrates this clearly. After excluding the outliers, the mean position accuracy was 28cm and the median was 21cm.

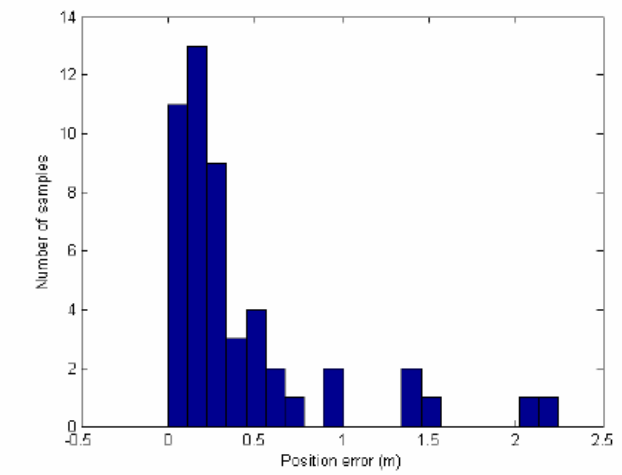

Figure 22: Histogram of the average position errors for each standard node.

The average number of successful localizations for a node was 4.5 and the average number of attempts per Monte Carlo run was 7.3. Thus, on average approximately 14 RIPS measurements were used to localize a node on each run. However, mote 15 was never successfully localized due to the geometry of the scenario. Figure 23 shows that location of mote 15 and also the four nodes that generated all the outliers in the position estimates.

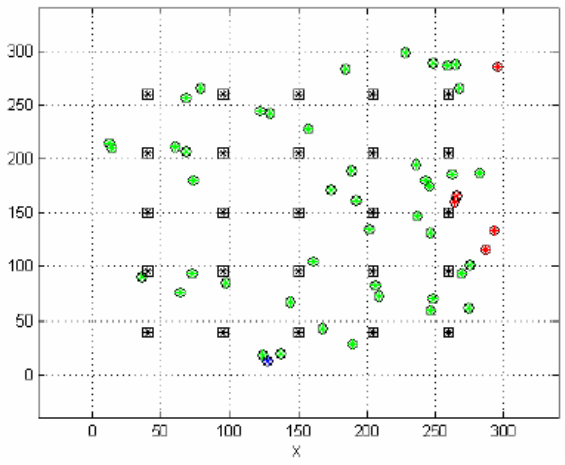

Figure 23: Simulation scenario layout. The location of the node that could not be localized, node 15, is shown by the blue asterisk in the lower part of the figure. The red asterisks indicate nodes whose estimated locations had significant errors.

Future Work

Missed Nodes: The current algorithm discards results when there is an ambiguous solution for the location of a node. This means useful information is lost and it is possible

that the node may never be localized. A simple extension that avoids this problem is to retain both possible locations for the node and seek to resolve the ambiguity when measurements become available from an alternative set of anchors.

Estimation Accuracy and Anchor Choice: The selection of nodes to use as anchors described in Section 1.3 aims to maximize the probability of finding an unambiguous solution. An alternative criterion is considered in [Savvides2003] when localization is performed using trilateration methods. There the ideal geometry is considered to be that which minimizes the Cramer-Rao bound on the estimation error. Calculation of the Cramer-Rao bound for this problem and the corresponding ideal anchor geometry is part of ongoing work.

Multi-Hop Error Propagation: The algorithm described here uses previously localized nodes as anchors to determine the location of new nodes. Errors in the estimated anchor locations will propagate and may potentially increase the location errors for neighboring nodes. The effect of this propagation of errors over multiple hops is being studied. The simulation results suggest that this effect can be significant at the perimeter of the surveillance region. Methods for detecting and, ideally, eliminating this are being investigated.

Self-Localization Using Noisy Measurements

The problem of self localization of motes using noisy RIPS measurements is posed as a Bayesian estimation problem. The problem is formulated as follows. Assume that there are K GPS nodes with the k^{th} node having position $\mathbf{y}_k$. The number of motes to be located is denoted as M and the unknown position of the m^{th} such motes is $\mathbf{x}_m$. The positions $\mathbf{x}_1,\ldots,\mathbf{x}_M$ are to be estimated using the efficient RIPS method proposed previously by Barbara La Scala and Xuezhi Wang. This involves selecting four motes which are all within Denote the collection of T RIPS measurements as $\mathbf{d}_1,\ldots,\mathbf{d}_T$. In a Bayesian framework, the problem is defined by the prior distribution of the unknown positions and the distribution of the measurements conditional on the unknown positions, *i.e.,* the likelihood of $\mathbf{x}_1,\ldots,\mathbf{x}_M$. The MMSE estimate of the unknown positions is the posterior expectation which can be computed from the posterior density:

$$\pi(x \mid d) \propto l(d \mid x)\psi)\psi(x)$$

where $\mathbf{x}=(\mathbf{x}_1,\ldots,\mathbf{x}_M)$, $\mathbf{d}=(\mathbf{d}_1,\ldots,\mathbf{d}_T)$, l is the pdf of the measurements conditional on the positions and ψ is the prior pdf of the positions. An intuitively reasonable prior pdf is to assume that the unknown motes are independent uniform random variables. The region over which a particular node is distributed can be limited by its proximity to the known locations of the GPS-equipped motes. We then have:

$$\psi(x) = \prod_{m=1}^{M} U_{R_m}(x_m)$$

The RIPS measurements are assumed to be subject to Gaussian noise which is independent for each measurement. This leads to a likelihood of the form:

$$l(d \mid x) = \prod_{t=1}^{T} N(d_t; h_t(x), \Sigma)$$

where the nonlinear function ht is determined by the distances between the motes used to generate the t^{th} RIPS measurement. The posterior density is then:

$$\pi(x \mid d) = c\prod_{t=1}^{T} N(d_t; h_t(x), \Sigma) \times \prod_{m=1}^{M} U_{R_m}(x_m)$$

where c is a normalizing constant. A closed-form expression for the posterior pdf π is unavailable due to the non-linearity of the RIPS measurements. At present we are pursuing a numerical solution based on important sampling. This works as follows. A collection of N samples is drawn from the prior distribution. These are denoted as $\mathbf{x}^1,\ldots,\mathbf{x}^N$. The posterior pdf is then approximated as:

$$\hat{\pi}(x \mid d) = \sum_{n=1}^{N} w^n \delta(x - x^n)$$

where the weights $w^1,\ldots,w^N$ are computed as:

$$w^n = l(d \mid x^n) / \sum_{r=1}^{N} l(d \mid x^r)$$

The posterior expectation is then approximated as:

$$\hat{x} = \int x\hat{\pi}(x \mid d)dx = \sum_{n=1}^{N} w^n x^n$$

The basic estimator described above will be close to optimal provided the sample size N is sufficiently large. In practice, the size of N is limited by the need for computational tractability. This is particularly the case if the measurement noise has small variance in which case the prior samples are distributed over a much larger area than the area of interest. To overcome this problem a kind of progressive correction is used. This involves a number of weighting steps in which the sample are initially weighted by a likelihood with a much larger spread than the actual likelihood, and then weighted using progressively tighter likelihoods in succeeding steps. The procedure can be summarized as follows:

1. Draw sample $\mathbf{x}^1,\ldots,\mathbf{x}^N$ from the prior distribution.
2. For S steps define the covariance matrices $\mathbf{\Sigma}_S > \mathbf{\Sigma}_{S-1} > \ldots > \mathbf{\Sigma}_1 = \mathbf{\Sigma}$.
3. For $s=S$:-1:1,
    a. For $n=1,\ldots,N$, compute weights: $w^n = c_1 \prod_{t=1}^{T} N(d_t; h_t(x), \Sigma_s)$ [c_1 is normalizing constant].
    b. Draw indices $j^1,\ldots,j^N$ such that $P(j^n=a) = w^a$.
    c. For $n=1,\ldots,N$, draw $\mathbf{x}^n$ from the distribution with pdf $k(x - x^{j^n})$ where k is a suitably chosen kernel density.

The method is demonstrated using an example with eight motes arranged on a regular 50m x 50m grid. The scenario, showing the known and unknown motes, is depicted in Figure 24. Three sets of RIPS measurements are used: one between motes A, B, C, and F, one between motes F, C, D and G and one between motes G, D, E and H. The aim of this quite simple example is to demonstrate the performance of the proposed approach and to

examine the effect of measurement noise on error propagation. It is expected that the estimates of the location of mote F should be more accurate than those of mote H.

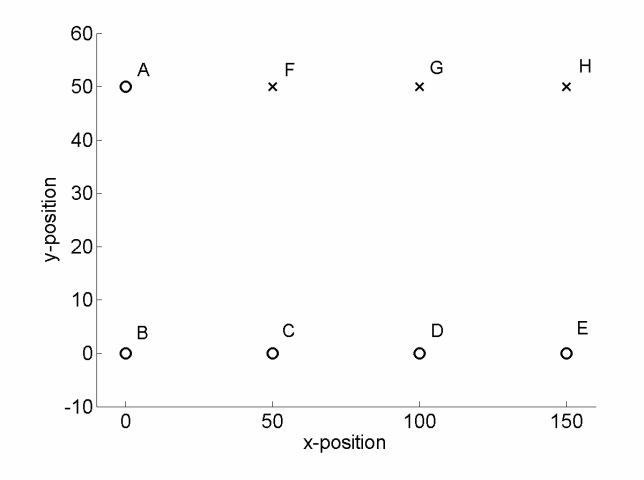

Figure 24: Self-localization scenario. The positions of GPS-equipped motes are indicated by circles. The positions of unknown motes are indicated by crosses.

The measurement noise covariance matrix is $\mathbf{\Sigma} = \sigma^2\mathbf{I}$, where $\mathbf{I}$ is the identity matrix. 100 Monte Carlo realizations are performed for several values of the variance σ^2 between 0.01 and 0.1. The bias and standard deviations of the x-location estimators for each unknown mote location are plotted against σ^2 in Figures 25 and 26. Error propagation is clearly evident in the standard deviations which are much greater for mote G than mote F, and slightly higher for mote H as compared to mote G. Interestingly, the ratio between these standard deviations seems to be unchanged by the variance of the measurement noise.

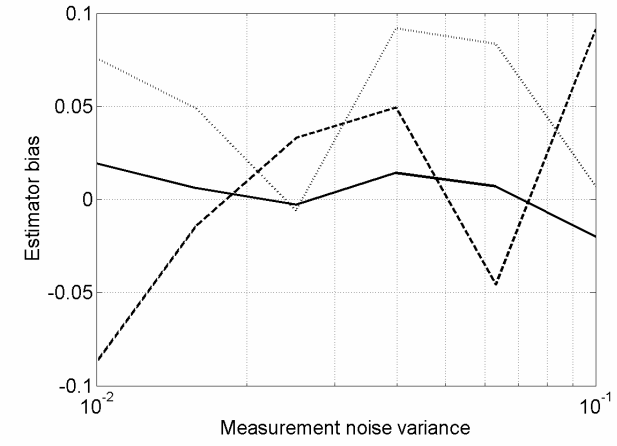

Figure 25: Estimator bias plotted against measurement noise variance for motes F (solid), G (dashed) and H (dotted).

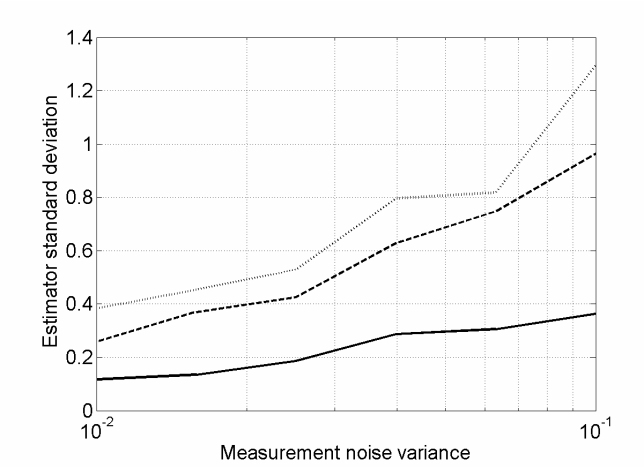

Figure 26: Estimator standard deviation plotted against measurement noise variance for motes F (solid), G (dashed) and H (dotted).

Mote tracking – the effect of mote density

In previous reports it was described how the unscented Kalman filter (UKF) can be used to track a target moving through a field of motes. It is of interest to examine how the density of the motes affects the accuracy and robustness of the tracking algorithm. This has been examined using simulations for the case of single target moving with a velocity subject to small random perturbations. The surveillance area is 150m x 150m and the the motes are arranged within this area in a regular grid. The half-energy sensing range of the motes is 12.5m, i.e., the signal from a target at this distance from the sensor is received with half the energy. Simulations are performed for m x m grids of motes with m= 6, …, 20. Two measurement models are considered: a non-thresholded model in which the intensities measured by the mote sensors are transmitted to the central processor and a thresholded model in which the intensities are thresholded and binary measurements are sent to the central processor. The sensor measurements are subject to a noise.

This noise will be varied to give signal-to-noise ratios (SNRs) of 15 dB, 10 dB and 5 dB. The results from 1000 Monte Carlo realizations are shown in Figures 27-29 for SNRs of 15, 10 and 5 dB respectively. As expected, the performance decreases as the number of motes decreases although this is considerably less so when the received intensities are available to the tracker. In this case accurate and reliable tracking can still be performed with rather sparse sensor configurations and a large amount of measurement noise. If the tracking algorithm is presented with binary measurements, even very dense sensor configurations do not result in reliable tracking if the measurement noise is large. However, reliable tracking with binary measurements in sparse sensor configurations is still possible provided there is not much measurement noise.

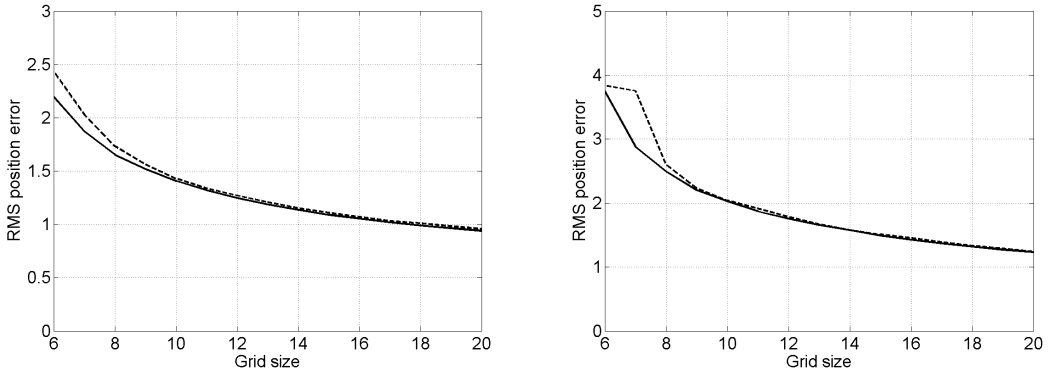

Figure 27: RMS position error for the UKF (dashed) plotted against the mote grid size for (a) the non-thresholded measurement model and (b) the thresholded measurement model. The solid line is the posterior CRB. The SNR is 15 dB.

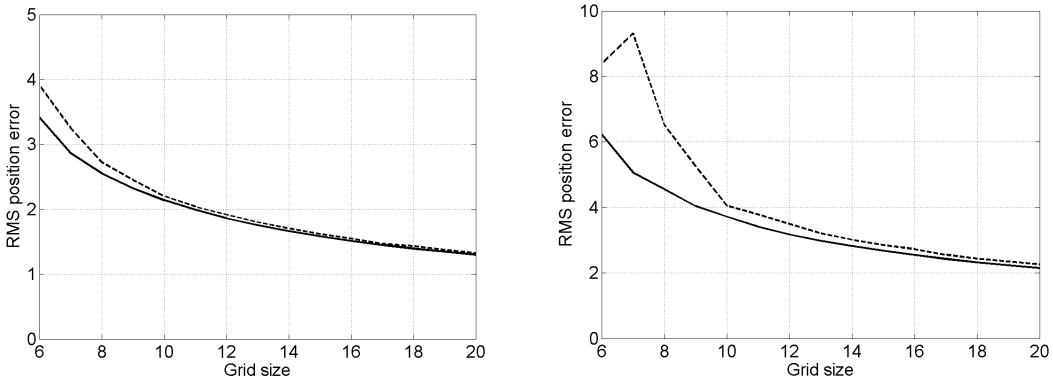

Figure 28: RMS position error for the UKF (dashed) plotted against the mote grid size for (a) the non-thresholded measurement model and (b) the thresholded measurement model. The solid line is the posterior CRB. The SNR is 10 dB.

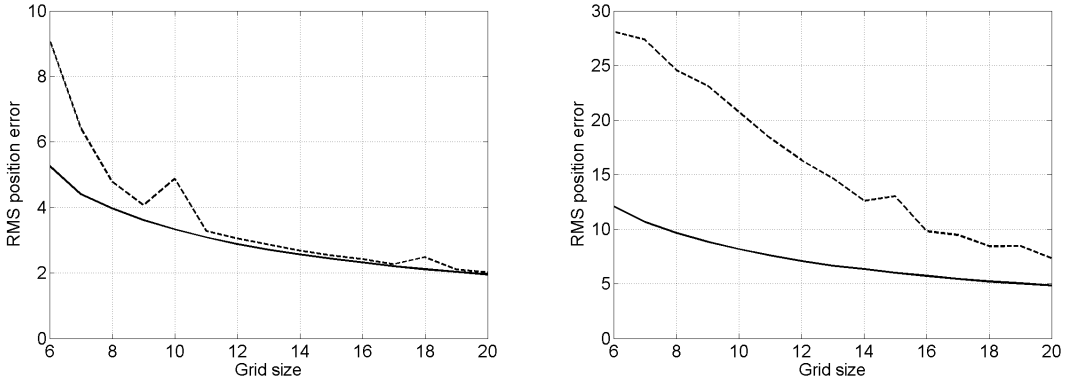

Figure 29: RMS position error for the UKF (dashed) plotted against the mote grid size for (a) the non-thresholded measurement model and (b) the thresholded measurement model. The solid line is the posterior CRB. The SNR is 5 dB.

2.A.4.c Generalized Frequency Modulated Waveform Libraries for Radar Tracking

Introduction

For radars capable of waveform agility the design of optimal waveform libraries comes into question. This piece of work considers the design of such waveform libraries for radar tracking applications from a generalized frequency modulated family of waveforms. Waveform libraries depend on the specific applications in which the systems are used. In designing or improving a waveform library certain questions arise. Firstly it is important to establish the measure of effectiveness for individual waveforms (cost function) and then to extend this to a measure of effectiveness (utility) for the library.

For the moment, we represent the measurement obtained using the waveform ϕ as a Gaussian measurement with covariance R_ϕ. The current state of the system is represented by the state covariance matrix P. The expected information obtained from a measurement with such a waveform, given the current state of knowledge of the target, is

$$I(X;Y) = \log\det\left(I - R_\phi^{-1}P\right). \quad (1)$$

This is the mutual information between the target state variable X and the radar return Y, resulting from the use of the waveform ϕ, I is the identity matrix. This cost function represents the immediate (one step ahead) gain of information with waveform ϕ. When we want to optimize the information gain, and to do this the dynamics of the target has to be taken into account.

Assuming two dimensional target state vector x (*i.e.*, range and velocity), we write

$$x_k = Fx_k + Q, \quad (2)$$

where $F = \begin{bmatrix} 1 & \delta \\ 0 & 1 \end{bmatrix}$, δ is time between the measurements and Q is dynamic noise covariance matrix, we write the mutual information between the state and the measurement, obtained with a waveform ϕ as

$$I(X;Y) = \log\det(P) - \log\det\left(Q + F\left(P^{-1} + R_\phi^{-1}\right)^{-1}F^T\right) \quad (3)$$

Minimizing this over some family of waveforms gives the optimum library of waveforms.

Suppose we know all the possible state covariances, P, generated by the tracking system. Next suppose this knowledge is statistical and is represented by a fixed probability distribution $F(P)$ over the space of all positive definite matrices. The utility function of a library $\Im$ of waveforms is given by

The utility function of a library $\Im$ of waveforms is given by

$$G_F = \int_{P>0} \max_{\phi \in \Im} \left[\log\det\left(I + R_\phi^{-1}P\right)\right] dF(P) \quad (4)$$

for one step ahead waveform scheduling and by

$$G_F = \int_{P>0} \min_{\phi \in \Im} \left[\log\det\left(Q + F\left(P^{-1} + R_\phi^{-1}\right)^{-1} F^T\right)\right] dF(P) \tag{5}$$

for long term scheduling.

In this work we assume that P belongs to a Wishart distribution, with the following definition and pdf. Suppose X_1 is a p column-vector-valued random variable that follows a p-variate normal distribution, with mean zero, and covariance nonnegative definite matrix V. Further suppose $X_1, \ldots X_n$ are independent and identically distributed (i.i.d.). Then the Wishart distribution is the probability distribution of the $p \times p$ random matrix

$$S = \sum_{i=1}^{m} X_i X_i^{'}. \tag{6}$$

The positive integer m is the number of degrees of freedom. The Wishart distribution is characterized by its probability density function. Random positive definite matrix P has a Wishart distribution with m degrees of freedom has a probability density function

The positive integer m is the number of degrees of freedom. The Wishart distribution is characterized by its probability density function. Random positive definite matrix P has a Wishart distribution with m degrees of freedom has a probability density function

$$f_{\mathbf{P}}(w) = \frac{|w|^{(m-p-1)/2} \exp\left[-\operatorname{trace}(\mathbf{V}^{-1}w/2)\right]}{2^{mp/2}\,|\mathbf{V}|^{m/2}\,\Gamma_p(m/2)} \tag{7}$$

where $\Gamma_p(\cdot)$ is the multivariate gamma function defined as

$$\Gamma_p(m/2) = \pi^{p(p-1)/4} \prod_{j=1}^{p} \Gamma\left[(m+1-j)/2\right]. \tag{8}$$

Generalized Frequency Modulated Waveform Library

In high SNR the waveform covariance matrix can be approximated by the CRLB on the estimation of delay and Doppler, which is obtained by inverting the Hessian of the Ambiguity Function, evaluated at the true target delay and Doppler [VanTrees71]. Computations (VanTrees71, loc.cit.) yield that within a constant factor the Fisher information matrix u can be calculated as,

$$\begin{aligned} u_{1,1} &= \int_{-\infty}^{\infty} \mathrm{s}(t)^{*''} \mathrm{s}(t)\, dt \\ u_{1,2} = J_{2,1} &= \int_{-\infty}^{\infty} t\mathrm{s}^{*'}(t)\mathrm{s}(t)\, dt. \\ u_{2,2} &= \int_{-\infty}^{\infty} t^2 |\mathrm{s}(t)|^2\, dt \end{aligned} \tag{9}$$

Each waveform in generalized frequency modulated library can be represented as

$$\mathrm{s}(t) = a(t) \exp(2\pi i \int f(t)\, dt), \tag{10}$$

where *a(t)* is a real valued even magnitude window function, b is a sweep constant and *f(t)* is a real valued frequency function. We shall assume that the window function has compact support, so that where integrals are over the whole real line, they are in fact only over a finite interval. Consider a library generated from a single window *a(t)* which satisfies the following conditions

$$\begin{aligned} & a(t) = 0, \text{ for } |t| > \Delta t \\ & \int_{-\infty}^{\infty} |a(t)|^2 \, dt \; = 1 \\ & \int_{-\infty}^{\infty} t|a(t)|^2 \, dt \; = 0 \end{aligned} \tag{11}$$

and *f(t)* which satisfies the following conditions

$$\begin{aligned} & \int_{-\infty}^{\infty} f(t)|a(t)|^2 \, dt \; = 0 \\ & |f(t)| \leq \Delta f, \text{ for } |t| \leq \Delta t \\ & |f(t_1) - f(t_2)| \leq M|t_1 - t_2|, \text{ for } t_1, t_2 \in [-\Delta t, \Delta t] \end{aligned} \tag{12}$$

These conditions give a convex compact set *C* of functions for the sup norm on the space of continuous functions with support a given interval. The Fisher information matrix for this type of waveforms is

$$\begin{aligned} u_{1,1} &= \int_{-\infty}^{\infty} \omega^2 |A(\omega)|^2 \frac{d\omega}{2\pi} + \int_{-\infty}^{\infty} |f(t)|^2 |a(t)|^2 \, dt \\ u_{1,2} = u_{2,1} &= \int_{-\infty}^{\infty} t f(t) |a(t)|^2 \, dt \\ u_{2,2} &= \int_{-\infty}^{\infty} t^2 |a(t)|^2 \, dt \end{aligned} \quad , \tag{13}$$

where *A(ω)* is the Fourier transform of the window function.

In short notation we write

$$\mathbf{U_s} = \begin{pmatrix} \rho_\omega + \overline{f^2} & \overline{ft} \\ \overline{ft} & \overline{t^2} \end{pmatrix}, \tag{14}$$

where

$$\begin{aligned} \rho_\omega &= \int_{-\infty}^{\infty} \omega^2 |A(\omega)|^2 \frac{d\omega}{2\pi}, \\ \overline{t^2} &= \int_{-\infty}^{\infty} t^2 |a(t)|^2 \, dt, \end{aligned}$$

$$\overline{f^2} = \int_{-\infty}^{\infty} |f(t)|^2 |a(t)|^2 \, dt,$$
$$\overline{ft} = \int_{-\infty}^{\infty} t f(t) |a(t)|^2 \, dt.$$

Note, that ρ_ω and $\bar{t}^2$ are identical for all waveforms in the library.

In this section we design library of waveforms which maximized instantaneous gain of information, when tracking with Kalman Filter. Given the frequency modulated waveform library, described above, we use the methods of calculus of variation to calculate the following expression

$$\max_{s \in \mathcal{L}} \log \det(\mathbf{I} + \mathbf{R}_\mathbf{s}^{-1} \mathbf{P}_0) \tag{15}$$

for some positive definite matrix P_0. Since $R_S^{-1} = 1/\eta T^{-1} U_S T$, we can rewrite this expression as

$$\max_{s \in \mathcal{L}} \log \det(\mathbf{I} + \mathbf{U}_\mathbf{s} \mathbf{P}), \tag{16}$$

where $P = 1/\eta T^{-1} P_0 T = \begin{bmatrix} p_{11} & p_{12} \\ p_{21} & p_{22} \end{bmatrix}$ is also positive definite. We interchange max and log and obtain the functional in frequency function *f(t)* which we want to maximize

$$\begin{aligned} \mathbf{J}[f] = \det(\mathbf{I} + \mathbf{U}_\mathbf{s}\mathbf{P}) &= 1 + \det(\mathbf{U}_\mathbf{s}) \det(\mathbf{P}) + \operatorname{tr}(\mathbf{U}_\mathbf{s}\mathbf{P}) \\ &= \overline{f^2 t^2} \det(\mathbf{P}) - (\overline{ft})^2 \det(\mathbf{P}) + p_{11}\overline{f^2} + 2p_{12}\overline{ft} + \\ &\quad 1 + \rho_\omega \overline{t^2} \det(\mathbf{P}) + p_{11}\rho_\omega + p_{22}\overline{t^2} \end{aligned} \tag{17}$$

We have

$$\mathbf{J}[f] = c + \overline{f^2} - \alpha(\overline{ft})^2 + 2\beta\overline{ft}, \tag{18}$$

where $c = \frac{1 + p_\omega \bar{t}^2 \det(P) + p_{11}p_\omega + p_{22}\bar{t}^2}{\bar{t}^2 \det(P) + p_{11}}$ is a constant, $\alpha = \frac{\det(P)}{\bar{t}^2 \det(P) + p_{11}}$ and $\beta = \frac{p_{12}}{\bar{t}^2 \det(P) + p_{11}}$

Denoting

$$F(t) = \int_{-\infty}^{t} f(\tau)\tau |a(\tau)|^2 \, d\tau, \tag{19}$$

we write

$$\mathbf{J}[f] = c + \int_{-\infty}^{\infty} \left(\frac{F'(t)^2}{t^2 |a(t)|^2} - 2\alpha F'(t) F(t) + 2\beta F'(t) \right) dt. \tag{20}$$

This is a quadratic functional in $F'(t)$ with gateaux second derivative positive. It is therefore a convex functional so that it achieves the maximum on the boundary ∂C, which is non-empty, since the set is compact and convex. One can show that the extreme points are functions which are all combinations of segments of the four straight lines:

$$\mathcal{B}_f = \{-\Delta f, \Delta f, Mt, -Mt\}. \tag{21}$$

The local minimum is given by the solution of the Euler-Lagrange equation. We have

$$-\frac{d}{dt}\left(\frac{2F'(t)}{t^2|a(t)|^2} - 2\alpha F(t) + 2\beta\right) - 2\alpha F'(t) \ = 0. \tag{22}$$

The solution to this equation is

$$f(t) = Ct. \tag{23}$$

Initial conditions give us $|C| \le \frac{\Delta f}{\Delta t}$. This result is remarkable, as it tells us that linear frequency modulated waveforms (chirps) are the minima of the short term cost.

If we impose the further constraint that $\Delta t = \frac{\Delta f}{M}$, the boundary is comprised of continuous piecewise linear functions constructed from up-sweep and down-sweep linear FMs.

Further constraining the waveforms to have a maximum number of linear pieces not exceeding two, gives a parametrized family of waveforms with one parameter $x \in \{-1,1\}; x\Delta t$ representing the point of connection between pieces.

The cost function in this case is 6th degree polynomial in x

$$f(x) = -1/4M\overline{t_2}(-3M\overline{t_2}x^4 + 3M\overline{t_2}x^2 - M\overline{t_2} + M\overline{t_2}x^6 - M\frac{p_{11}}{|\mathbf{P}|} + 3M\frac{p_{11}}{|\mathbf{P}|}x^4 - 6M\frac{p_{11}}{|\mathbf{P}|}x^2 - 12\frac{p12}{|\mathbf{P}|}x + 4\frac{p_{12}}{|\mathbf{P}|}x^3). \tag{24}$$

The zeros of the derivative of *f* with respect to the parameter x gives us at most three local maxima of *f(x)* in [−11], two of which are the end points x = −1 and x = 1, and the third point depends on *P*, i.e

$$\frac{df(x)}{dx} = -3/2M\overline{t_2}(x-1)(x+1)(Mx^3\overline{t_2} - Mx\overline{t_2} + 2xM\frac{p_{11}}{|\mathbf{P}|} + 2\frac{p_{12}}{|\mathbf{P}|}). \tag{25}$$

For example if $p_{12} = 0$ the local maxima is in $x = 0$ and it is greatest of the three local maxima when $p_{22} > \frac{3}{\bar{t}^2}$.

We evaluated the utility function (using Monte-Carlo integration) for Wishart Distribution of *P* with several degrees of freedom. The results are represented below. In these experiments we are looking at two Wishart distributions with 3 and 10 degrees of freedom, respectively, and with the same mean covariance $V = \begin{bmatrix} 1 & 0 \\ 0 & 10 \end{bmatrix}$. The state vector is [range velocity]^T and the units in the state vector are m in range and m/s in velocity. We compare the utilities of three waveform libraries: *L* consisting of down-sweep chirp, an up-sweep chirp, and a piecewise linear frequency waveform with connection point in the center of the time window, and four more piecewise linear frequency waveforms with

connection points not in the center of the time window; L_1 consisting of down-sweep chirp, up-sweep chirp, a piecewise linear frequency waveform with connection point in the center of the time window and L_2 consisting of down-sweep chirp and up-sweep chirp only. Of course, all pieces have maximum sweep rate M. For these experiments the window function is a truncated Gaussian window with duration $\Delta t = 100$ μs and frequency sweep Δf = 14 GHz. We give the actual value of utility function as well as the SNR improvement between the libraries. SNR improvement is given by the reduction in volume of expected posterior pdfs between the baseline library L_2 and each of L and L_1. The SNR gain in both cases is not null but it does not seem significant. For comparison the SNR gain between chirp library (2 waveforms) and chirp-rotation library (4 waveforms) for the same parameters as above is 5.2243 dB.

Table 2: Utility Function for Piecewise Linear Waveform Library for various Covariance Distributions

Degrees of Freedom	$G(L)$	$G(L_1)$	$G(L_2)$	SNR improvement $L_2 \rightarrow L_1$	SNR improvement $L_2 \rightarrow L$
3	4.6953	4.6953	4.6953	0.2418e-3 dB	0.2512e-3 dB
10	6.1190	6.1187	6.1034	0.0675 dB	0.0691 dB

2.A.5. Georgia Tech Technical Progress

Currently, the CADSP Imager IC testing is in progress utilizing newly developed custom PCB boards. These boards move towards a stable standalone system with on-board instrumentation needed to move the imager operation off the test bench. With the new boards, programming of floating gate transistors in the A and B matrices, which hold the coefficients, has now been performed. The imager pixel plane has be verified in a low resolution mode and nearly digital circuitry has now been tested. Still required is a full resolution test to verify and characterize the pixel level computation. Also remaining is an issue regarding the programming of a floating gate transistor on a trans-impedance amplifier which connects the pixel plane to the B matrix. A control line was mis-wired on chip, but possible workaround's have been developed and are being tested now. This will conclude basic level programming of all transistors on the IC. Following progressions will include optimizations for speed and accuracy. The last remaining subsystem testing is the verification of a final current-to-voltage conversion. Movement to full system-level operation will follow.

Work on the optical flow algorithms has been in two directions. First, we have been developing a system to generate synthetic image sequences that can then be projected onto the imager IC. These rendered sequences will have associated ground-truth motion vector information so that we can evaluate the performance of the system in a meaningful way. Second, we have been refining an RLS-based algorithm for efficiently and accurately extracting the motion vectors in the presence of imager noise.

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 at EURASIP Journal on Applied Signal Processing, 2006.
3. A. Chhetri, D. Morrell and A. Papandreou-Suppappola, "On the use of binary programming for sensor scheduling," under second review for IEEE Transactions on Signal Processing, (submitted February 2006).
4. Amit Chhetri, "Sensor Scheduling and Efficient Algorithm Implementation for Target Tracking," PhD Dissertation, Arizona State University, May 2006.
5. A. Chhetri, D. Morrell and A. Papandreou-Suppappola, "Sensor resource allocation for tracking using outer approximation," accepted for publication in IEEE Signal Processing Letters, 2006.
6. I. Kyriakides, D. Morrell and A. Papandreou-Suppappola, "Sequential Monte Carlo Methods for Tracking Multiple Targets with Deterministic and Stochastic Constraints," submitted to IEEE Transactions on Signal Processing, July 2006.

2. C. Conference Proceedings

[1] C. Savage, B. La Scala and B. Moran, "Optimal Scheduling for State Estimation Using a Terminal Cost Function," 9th International Conference on Information Fusion, Florence, Italy, July, 2006

[2] R. Cramer, S. Bellofiore, T. Stevens, H. A. Schmitt and D. Waagen "Localization, Detection and Tracking for Wireless Sensor Networks," MSS Specialty Group on Battlespace Acoustic and Seismic Sensing, Magnetic and Electric Field Sensors, JHU/APL, Baltimore, MD, August, 2006.

[3] V. Berisha, "Landmark Selection for Manifold Learning," in preparation.

[4] N. Patwari and A. O. Hero III, "Indirect Radio Interferometric Localization via Pairwise Distance," in *Proceedings of the 3rd Workshop on Embedded Networked Sensors (EmNets 2006)*, pp. 26-30, May 30, 2006, Cambridge, MA.

[5] A. Chhetri, D. Morrell and A. Papandreou-Suppappola, "Sensor scheduling using 0-1 mixed integer programming framework," IEEE Workshop on Sensor Array and Multi-channel Processing, July 2006.

[6] I. Kyriakides, A. Papandreou-Suppappola and D. Morrell, "Adapting Matching Pursuits Dictionaries to Signal Structure Using Particle Filtering," IEEE Workshop on Sensor Array and Multi-channel Processing, July 2006.

2. D. Consultative and Advisor Functions

The consultative or advisory functions that occurred during the current PoP have remained unchanged. 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 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 continue our investigation of several promising mathematical approaches to sensor localization; these will be made available to the Raytheon NEST

program if they are successful. Technical progress in these areas is conveyed to NEST regularly, more generally, the two programs have developed a strong working relation.

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. Dr. Klausutis has offered to make available a GPS-equipped truck capable of collecting ground-truth optical flow imagery should a "packaged" CADSP imager be ready in late fall 2006.

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.
MATLAB simulation code for cooperative control of UAVs for passive geolocation was released to Professor Daniel Pack of the US Air Force Academy.

2. H. References
[Silva&Tenenbaum] V. D. Silva, J. B. Tenenbaum, "Sparse Multidimensional scaling using Landmark Points," Technical Report, Stanford University, 2004.

[BeC99] D. Bertsekas and D. Castanon, "Rollout Algorithms for Stochastic Scheduling Problems," Heuristics, Vol. 5, pp. 89-105, 1999.

[Gi79] J. Gittins, "An Application of Control Theory to a Game of Hide and Seek," International Journal of Control, Vol. 30, No. 6, pp. 981-987, 1979.

[GiRo79] J. Gittins and D. Roberts, "The Search for an Intelligent Evader Concealed in one of an Arbitrary Number of Regions," Naval Research Logistics Quarterly, Vol. 26, pp. 651-666, 1979.

[Costa2006] J. A. Costa, N. Patwari, and A. O. Hero III, "Distributed Weighted Multidimensional Scaling for Node Localization in Sensor Networks", *ACM Transactions on Sensor Networks*, Feb. 2006, vol. 2, no. 1, pp. 39 – 64.

[Maroti2005] M. Maroti, P. Volgyesi, S. Dora, B. Kusy, A. Nadas, A. Ledeczi, G. Balogh, and K. Molnar, "Radio interferometric geolocation," 3rd International Conference on Embedded Networked Sensor Systems, pages 1–12, San Diego, CA, USA, 2005.

[Savvides2003] A. Savvides, W. Garber, S. Adlakha, R. Moses, and M. B. Srivastava, "On the error characteristics of multihop node localization in ad-hoc sensor networks," 2nd Workshop on Information Processing in Sensor Networks, pages 20–30, Palo Alto, CA, USA, April 2003.

[VanTrees71] H. L. van Trees, Detection, Estimation and Modulation Theory, Part III, Wiley, New York, 1971.

[Patwari2006] N. Patwari and A. O. Hero III, "Demonstrating Distributed Signal Strength Location Estimation", ACM SenSys 2006 (accepted), Nov 2006.

2. I. Acronyms

ADTS	Advanced Detection Technology Sensor
JHU/APL	Johns Hopkins University/Applied Physics Laboratory
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
DS	Danzig Selector
DSA	Distinct Sensing Area
dwMDS	Distributed, weighted, multi-dimensional scaling
FH	Frequency Hopping
FPA	Focal Plane Array
FMAH	Fast Mathematical Algorithms and Hardware
GEM	Geometric Entropy Maps
Georgia Tech	Georgia Institute of Technology
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
MSS	Military Sensing Symposia
MTT	Multi-target tracking
NEST	Networked Embedded System Technology
NDA	Non-disclosure Agreement
NICTA	National Information and Communication Technology of Australia
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
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