

AFRL-IF-RS-TR-2004-217
Final Technical Report
July 2004



LAYERED FAULT MANAGEMENT ARCHITECTURE

Vanderbilt University

Sponsored by
Defense Advanced Research Projects Agency
DARPA Order No. J184

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK

STINFO FINAL REPORT

This report has been reviewed by the Air Force Research Laboratory, Information Directorate, Public Affairs Office (IFOIPA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

AFRL-IF-RS-TR-2004-217 has been reviewed and is approved for publication

APPROVED:

/s/
DANIEL E. DASKIEWICH
Project Engineer

FOR THE DIRECTOR:

/s/
JAMES A. COLLINS, Acting Chief
Information Technology Division
Information Directorate

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 074-0188

Public reporting burden for this 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 this 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, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE July 2004	3. REPORT TYPE AND DATES COVERED Final Aug 02 – Dec 03	
4. TITLE AND SUBTITLE LAYERED FAULT MANAGEMENT ARCHITECTURE			5. FUNDING NUMBERS C - F30602-02-2-0206 PE - 62301E PR - ANTS TA - 00 WU - 08	
6. AUTHOR(S) Janos Sztipanovits, Ted Bapty, Ben Abbott				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Vanderbilt University 2015 Terrace Place Nashville, TN 37203			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Defense Advanced Research Projects Agency AFRL/IFTB 3701 North Fairfax Drive 525 Brooks Rd Arlington, VA 22203-1714 Rome, NY 4505			10. SPONSORING / MONITORING AGENCY REPORT NUMBER AFRL-IF-RS-TR-2004-217	
11. SUPPLEMENTARY NOTES AFRL Project Engineer: Daniel E. Daskiewich, IFTB, 315-330-7731, Daniel.Daskiewich@rl.af.mil				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution unlimited.			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 Words) This report describes Vanderbilt's contribution to the ability to build systems that use decentralized control and fault tolerance techniques to support applications such as large clusters of Micro UAVs or Organic Air Vehicles. The approach of this effort was to analyze fault management requirements of formation flight for fleets of UAVs, and develop a layered fault management architecture which demonstrates significant improvement over current technology. The target demonstration was a radio-geolocation system, using 3-10 UAV mounted time-of-arrival measurement nodes and a single base station.				
14. SUBJECT TERMS Organic Air Vehicles, Fault Management, Time Difference of Arrival			15. NUMBER OF PAGES 36	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL	

Table of Contents

List of Figures.....	ii
Participants.....	iii
Executive Summary	1
Introduction.....	4
Methods, Assumptions, and Procedures	5
Demonstration Scenario/Mission CONOPS	6
Impact of Position on Accuracy.....	8
UAV Platforms	8
Challenges Associated with Low Cost TDOA Nodes	8
Hardware Architecture.....	9
Radio Sensor Units	11
Software Radio -- Signal Processing	14
Synchronization and Clock Correction.....	16
Real-Time Processing.....	19
User Interface and Communications Infrastructure.....	20
Results and Discussion	21
Basic Radio Performance.....	21
FRS Radio Attenuation Performance.....	21
Signal Processing Measurements.....	22
Flight Package:.....	25
UAV-based Antenna Arrays	26
Hardware Fault Tolerance.....	28
CONOPS.....	29
Conclusions.....	29
Recommendations.....	30

List of Figures:

Figure 1: Layered Fault Management Architecture.....	5
Figure 2: UAV-Based TDOA Scenario	6
Figure 3: UAV Positioning With Multipath	7
Figure 4: UAV Reposition to Avoid Multipath	7
Figure 5: TDOA Accuracy with Positioning of UAV	8
Figure 6: WASP Sensor Node Hardware Architecture.....	10
Figure 7: Sensor node Assembly & Carrier Board	11
Figure 8: Motorola FRS Block Diagram	12
Figure 9: Modified Radio Block Diagram, with SWR Front End Electronics	13
Figure 10: Photo of Modified Radio	14
Figure 11: Signal Processing Block Diagram.....	15
Figure 11a: Gross-level Signal Location	16
Figure 12: Example GPS Jitter	17
Figure 13: Schematic of 1 PPS Sampling	18
Figure 14: Structure of Node-based Real-time Processing.....	19
Figure 15: WASP User Interface	20
Figure 16: FRS Laboratory Experiment	21
Figure 17: SNR Measurements at Honeywell	22
Figure 18: Expected Accuracy vs SNR vs Average Time.....	22
Figure 19: Raw Input from Radio.....	23
Figure 20: Results of Demodulation.....	23
Figure 21: Cross Correlation Plot	24
Figure 22: Long Term TOA Results from Radio Pair	24
Figure 23: Incremental TOA Measurements	25
Figure 24: Mounting on eMAV	26
Figure 25: Planned Flight Package	26
Figure 26: Antenna Phase Correction	27
Figure 27: Hardware Fault Tolerance.....	28
Figure 28: Concepts of Operations for Mountainous Surveillance	29

Participants

1. *ISIS-Vanderbilt:*

Faculty:

Prof. Janos Sztipanovits

Senior Research Scientist:

Dr. Ted Bapty

Dr. Tim Holman

Research Assistants:

Dr. Jason Scott

Kumar Chhokra

Simon Winberg

2. *Cornell University*

Faculty:

Dr. Raffaello D'Andrea

3. *Southwest Research Institute*TM

Principal Engineer:

Dr. Ben Abbott

Executive Summary

The Layered Fault Management project was a joint effort under the WASP (Widely Adaptive Signal Processing) Seedling, between Vanderbilt, Southwest Research Institute (SwRI), and Cornell University. The primary focus was the development, analysis, and prototyping of techniques and technologies for deploying signal processing-intensive systems on highly mobile platforms, such as UAV's (Unmanned Air Vehicles) and OAV's (Organic Air Vehicles). The target demonstration was a radio-geolocation system, using 3-10 UAV-mounted time-of-arrival measurement nodes, and a single base station.

A novel low-cost, distributed TDOA (Time Difference Of Arrival) algorithm was developed and implemented on highly resource constrained sensor nodes. Constraints, driven by the concepts of operation, include: 1) low cost, for high expected attrition systems; 2) low power for battery operation; 3) low weight for UAV deployment; and 4) minimal communications for stealth and power considerations. The constraints implied local processing, using commodity/COTS (Commercial Off-The-Shelf) hardware. Local processing enables independent TOA (Time Of Arrival) measurements, with minimal information sent to the base station. The constraints mandated development of several accuracy enhancing techniques, including 1) demodulation via software radio processing; 2) local, template-based time of arrival estimates; 3) estimation and correction for local oscillator variations; and 4) accuracy enhancement of the global time-base.

Vanderbilt served as the system integrator, hardware developer, and embedded software developer. SwRI provided expertise on algorithms and data analysis. Cornell worked on the control and processing issues involved with forming UAV-based array antennas.

The following component technologies were demonstrated for feasibility:

1. A WASP sensor processing node was developed based on a high performance DSP (Digital Signal Processor) with attached FPGA (Field Programmable Gate Array) and data acquisition hardware.
2. A RF (Radio Frequency) receiver based on COTS hardware and additional custom preprocessing electronics was developed and tested.
3. A Software-Radio was developed and implemented to provide a low-phase-distortion output signal for high accuracy time estimation.
4. Time of Arrival Estimation algorithms were developed for distributed sensor nodes and implemented on the embedded DSP.
5. A communication infrastructure was developed to allow TDOA computations on a central station based on multiple remote TOA measurements.
6. GPS clock correction algorithms were developed and implemented for accuracy enhancement of the COTS GPS (Global Positioning System) timebase.
7. Mathematical accuracy estimation for Radiolocation process, based on arbitrary placement of sensors and targets.

8. Implementation of a phased-array antenna based on multiple coordinated UAV's. Transmission signals are phase-adjusted dynamically to compensate for measured node position errors.

In addition, a compelling concept of operations was developed to drive prototype and program requirements. The CONOPS (CONcept of OPERATIONs) describes the use of a mobile network of sensor nodes covering a mountainous region. Initial, low accuracy measurements are used to refine vehicle/sensor positions, with ultimate accuracy sufficient for targeting or tasking of video surveillance assets.

Results:

The seedling effort investigated the feasibility of using multiple, mobile unmanned sensor platforms for a range of potential CONOPS from geo-location of RF sources to SAR (Synthetic Aperture Radar). The high risk technologies in these CONOPS are

1. The feasibility of light weight sensor packages for these challenging applications,
2. Control technology for the mobile platforms,
3. Algorithms for highly accurate, distributed sensor processing, and
4. Platform technologies.

This project focused on the analysis and demonstration of the feasibility of UAV-hosted TDOA sensor technology. High risk aspects of the problem were addressed, and prototype solutions developed.

According to our results, the full development of low-cost, distributed, mobile TDOA concept is feasible. Our final demonstration has shown that, using novel architectural solutions, TDOA sensor:

1. Cost can be decreased by a factor of 30. (Prototype node was ~\$2K)
2. Power can be reduced to under 2 watts, (Prototype node was ~8W)
3. Weight can be reduced well under 1 pound (Prototype was ~1 lb)

These flight package parameters are suitable for light UAV-s.

Our demonstration showed that the accuracy of time of arrival measurements at a single node is less than 300ns. We expect that this accuracy can be doubled by pushing more processing into digital devices.

Recommendations

Work on the TDOA package has continued under the P&M (Perch & Move) seedling effort. Further development of software and hardware has proven the system performance in full-scale testing. To prove the system out for the WASP CONOPS, several optimizations and technology advances are needed:

1. Power: The initial prototype was done cheaply, leveraging existing compute and acquisition platforms. Power was not a consideration in these designs. A power

optimized design should lower consumption by 10x or more. This will allow a 10x improvement in runtime.

2. Cost: The current TDOA package costs approximately \$2000 to build. Design optimization can reduce this to 1/5th the cost.
3. Sensor Placement: The seedling work shows the expected dependency of SNR (Signal to Noise Ratio) on sensor location. To enable effective deployment, a sensor placement is critical. These placement decisions are multi-dimensional optimizations, planning for maximum accuracy while minimizing energy expenditures, threat to assets, speed of position fix, etc.
4. Accuracy enhancement: Initial correction algorithms show large potential for accuracy enhancement. This, combined with a dynamic calibration process will result in consistent, high quality results.
5. Sensor fusion: Radiolocation is just one part of the wide area surveillance CONOPS. The sensor node must be extended or combined with sensors for video, dismount detection, vehicle detection, etc. Optimum platform efficiency will be obtained if all sensor modalities are considered in the design of the sensor platform.
6. Multiple CONOPS are possible. Due to the software radio approach, the system is amenable to other radio waveforms and units. This would enable a TDOA system to be rapidly deployed using captured radio units. A lengthy design and characterization process could be avoided by treating radio electronics as a 'black box'.

Introduction

Over the past few years, the army's need for a lighter force that is rapidly deployable resulted in an increased need for command, control, reconnaissance and surveillance capabilities. While the armed forces recognize various means of gathering intelligence, such as HUMAN INTelligence (HUMINT), SIGnals INTelligence (SIGINT), measurement and signature intelligence (MASINT), Unattended Ground Sensors (UGSs) and Unmanned Aerial Vehicles (UAVs) have by far the most direct impact. Such automated sensor networks provide nearly instantaneous intelligence capabilities, with low risk to human life and high reliability. With proper design and understanding of the operating environment, such sensors can be tailored to have a long operating life, stealth and a low False-Alarm-Rate (FAR). Such sensor networks can also help discover adversary movement and activity patterns not exposed by other modalities of intelligence.

The Vanderbilt University effort under the Widely Adaptive Signal Processing (WASP) project was an effort to produce technology to enable a fleet of unmanned aerial vehicles and Organic Aerial Vehicles (OAVs) to triangulate a radio source in a hostile environment. While radio geolocation systems such as the ship-based units, already exist, there are few and expensive and also often used as shared resources. These long-stand-off systems are limited in accuracy and are not effective for low power radios and in complex terrains. The mobility of UAV/OAV's permit the sensors and processing to be moved closer to the signal source.

The project involved several main thrusts:

1. Development of a reliable sensor node for localized processing.
2. Development of receiver electronics for capturing and measuring target signals.
3. Development of algorithms for distributed geolocation.
4. Development of techniques for assessing the impact of geographical structures on TDOA results.
5. Demonstration of a prototype system for radio geolocation.

Several additional issues were studied:

1. The use of mobile communications nodes to enhance antenna capabilities (power and directivity).
2. Tools for composition of layered fault-tolerant embedded systems.

The results indicate that large gains can be achieved with relatively low cost techniques and equipment. The TDOA demonstration, along with the hardware and algorithms developed in implementing the demonstration, are discussed below.

Questions and technical challenges in UAV-based TDOA:

- 1) Accuracy of radios, 2) Accuracy of GPS, 3) Ability to do distributed, silent TDOA

Methods, Assumptions, and Procedures

The guiding architectural principals follow a layered structure. Figure 1 shows the general principles:

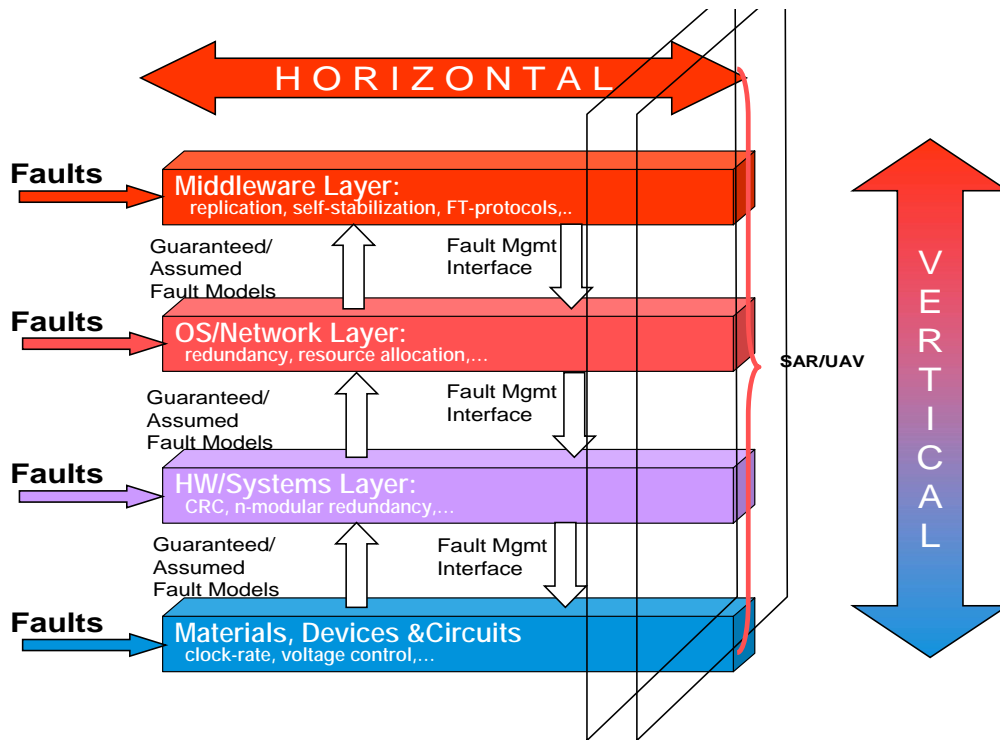


Figure 1: Layered Fault Management Architecture

Within each layer, fault management is composed of many tightly interlocking components. For example, middleware uses several facilities: Replication provides multiple copies and a voting or timeout failure sensing mechanism. This mechanism will closely link with self-stabilization and fault tolerant protocols. For instance, the replicated components will use fault-tolerant protocols to interact.

Vertical integration isolates fault tolerance mechanisms, providing a guaranteed level of reliability to higher levels. Likewise, it relies on a level of fault tolerance from lower levels. This layered approach is common in many fault tolerant infrastructures. The goal for WASP systems are to optimize for specific levels of fault tolerance by balancing the load of resilience across the vertical layers. Each application will result in a unique cross slice of fault tolerance mechanisms and parameters, taking resources and requirements into account.

Within this project, fault tolerance was explored at the materials, hardware, and application layers. These were integrated into the demonstration, as noted below.

Demonstration Scenario/Mission CONOPS

Figure 2 shows the scenario for UAV-based radio geolocation. Since a Time-Difference of Arrival technique was employed, we refer to the radio geolocation process as TDOA hereafter.

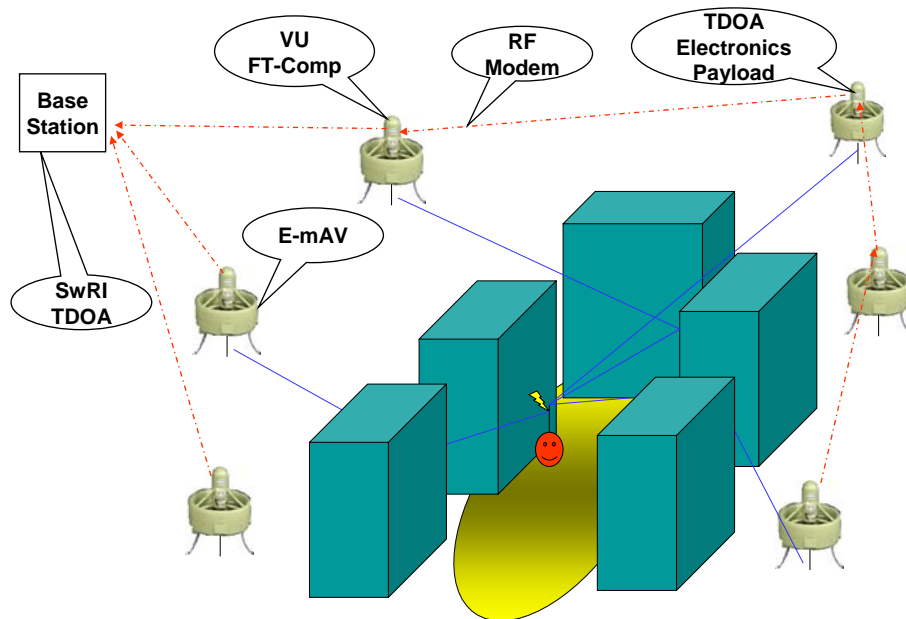


Figure 2: UAV-Based TDOA Scenario

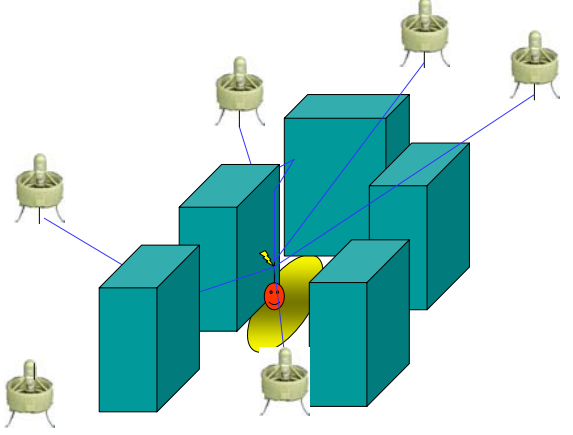
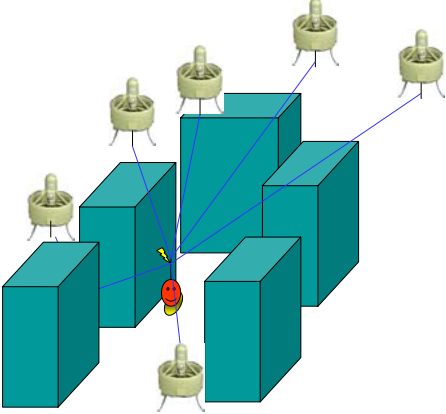
It is assumed that the target can be located in a geographically complex area. This can include city/urban areas, or mountainous regions. Radio signals in most commonly used systems are line-of-sight. If a building, or other obstruction, is between the transmitter and receiver, signals are greatly attenuated. (For example, this report contains a measurement set showing large attenuation caused by trees for the Family Radio Service (FRS)-band). Mountains can completely block most signals.

TDOA is also highly dependent on geometry. Simulation, underlying math, and experience show that the sensors should surround the target for maximum accuracy. The mobility of UAVs will permit the target geometry to be achieved regardless of the location of the transmitter.

UAV repositioning is an iterative process. Initial placements are necessarily flawed, since the initial position of the target is unknown. With each measurement, the target estimate is refined,

due to a more optimal sensor placement. Once the required accuracy is achieved, UAVs can be retasked to seek other targets, or kept to track target movements.

Note that WASP was not concerned with the planning of sensor allocation, placement, or path planning to move UAVs to the proper locations. These problems must be addressed in a later program.

First UAV Reposition, Multipath still a problem	2 nd UAV reposition. Accuracy improved
 <p data-bbox="248 1035 740 1066">Figure 3: UAV Positioning With Multipath</p>	 <p data-bbox="868 1041 1393 1073">Figure 4: UAV Reposition to Avoid Multipath</p>

The Figure 3 and Figure 4 show an example scenario where a target is located in an ‘Urban Canyon’. The first reposition (Figure 3) suffers from multipath and attenuation, as sensors on the right do not have Line-Of-Sight (LOS). The repositioning in Figure 4 moves all UAV’s to the LOS positions. The current geolocation uncertainty (yellow circle) is greatly reduced.

Impact of Position on Accuracy

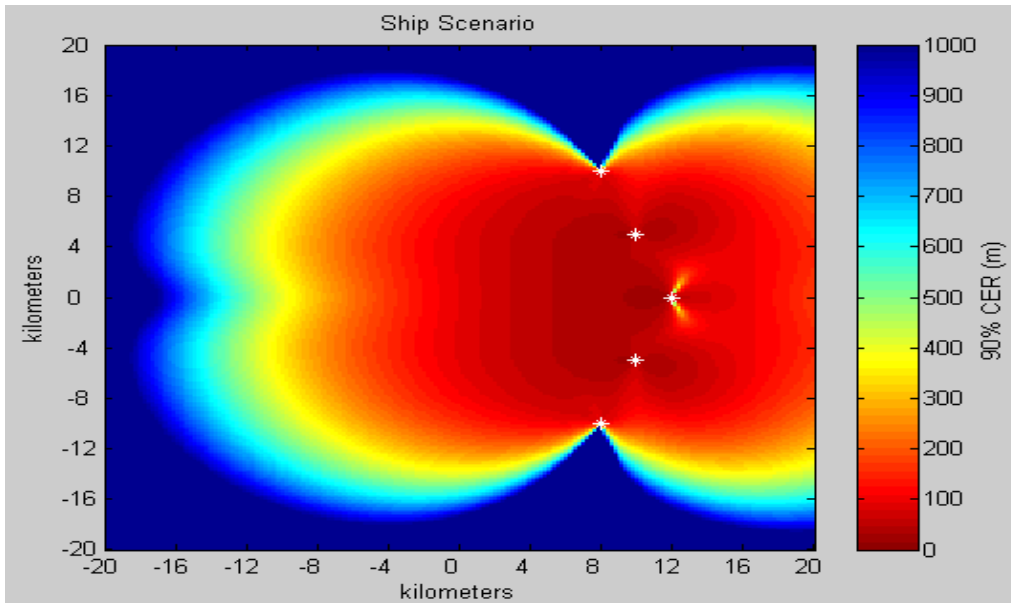


Figure 5: TDOA Accuracy with Positioning of UAV

Figure 5 shows SwRI's computations of error vs. position for a specific configuration of sensor locations. Darker (red) areas show higher precision. As can be seen, accuracy varies greatly with position and sensor placement.

UAV Platforms

The initial target platform for the research was the Allied eMAV, a 9" diameter UAV with an electrical motor. The payload size/weight/power of this vehicle were used as guiding constraints for the design of the TDOA sensor node. The basic constraints were: 0.5 lbs, 3" diameter, 2" tall. Midway through the project, it became apparent that the Allied vehicle would not be available in time for the planned demonstration.

The second UAV considered was the Honeywell/Allied 29" UAV. Size/Weight/Power constraints were greatly relaxed with this platform. This platform was not available for a TDOA demonstration. Consequently, demonstrations were performed with ground-based mobile sensor nodes. The constraints of UAV deployment were observed in the design.

Challenges Associated with Low Cost TDOA Nodes

The constraints of the concepts of operations are many fold:

1. **Cost:** The sensor node must be low cost. The sensors will be mounted on low-cost UAV's, with limited flight range. We anticipate many of the vehicles will be lost or destroyed, and therefore must not add significantly to the cost of the platform.
2. **Size:** Micro-UAV's are small, and thus the payload must be small. The Allied 9" eMAV UAV has an upper payload bay that is ~3" in diameter and 4" tall.
3. **Weight:** Every ounce of payload reduces the flight range and duration of the UAV. Max weight for the eMAV.
4. **Power:** Batteries are extremely heavy. Reducing power consumption of the payload is the best way to reduce the battery weight.

These payload constraints have a significant impact on the design of the sensor nodes and TDOA approach. Typical TDOA systems are hosted by a large structure (i.e. a ship, ground vehicle, airplane, or a ground station) with few limits on size/cost/power. This enables designers to use expensive, high-power GPS units with tightly synchronized clocks. Signals can be sampled as fast as desired. High quality (heavy, costly, high power) radios enable very high signal-to-noise ratio acquisition. High bandwidth connections allow all data to be exchanged and processing to be done at a centralized location. The UAV-based TDOA solution must take an entirely different approach:

1. **Time Reference:** Low power, low cost GPS units must be used. These typically have clock accuracies (jitter) of 90 ns or more, providing only a 1 pulse-per-second signal. Since clocks for GPS are ultimately derived from a cesium clock in orbit, the clock is long-range stable and accurate.
2. **Synchronization:** High-end GPS units provide a synchronized sampling clock. Since this is not available, we must use unsynchronized, low-quality crystal oscillators.
3. **Sampling:** Since each unit uses a different crystal, and crystals vary by ~20 parts-per-million, we can assume that each unit will have a separate clock. In addition, since we cannot afford the size/weight of an ovenized crystal, we must assume clock drift.
4. **Localized Processing:** The lack of a high-bandwidth network forces all processing to be done at the node. A very limited quantity of results can be transmitted back to the base station. Stealth in the RF spectrum also imposes this constraint for non-wired systems.

Achieving accuracy in TDOA results means overcoming each of these constraints. The design approaches are described below.

Hardware Architecture

Due to budget and time limitations, existing components were employed where possible. The computational platform consists of (see Figure 6):

1. A TMS320C6711 DSP. The DSP provides up to 800 MFLOPS, with 32 bit IEEE floating point.
2. Three banks of RAM. 256 KB EEROM is used to store bootstrap code and coefficients. 1 MB of SBSRAM provides fast scratchpad memory. 64 GB of SDRAM provides a large

buffer for data capture and intermediate processing results. Redundant banks of RAM provide hardware redundancy.

3. An on-board Field Programmable Gate Array (FPGA) provides the basic glue logic along with high-speed communications with adjacent processors and data acquisition devices. The FPGA also supports interfaces to asynchronous devices, including the GPS and radio modem. Finally, watchdog timers are implemented in the FPGA for breaking deadlock and recovering from software/hardware errors.

The Data Acquisition module contains 2 A/D converters, connected to an onboard FPGA via First-In-First-Out (FIFO) buffers. The FPGA performs all the control and synchronization of the A/D's using the local oscillator. The local oscillator is accurate to 20 parts per million. While accurate to commercial standards, this accuracy, if uncorrected, is too low for node-to-node synchronization.

The data acquisition module also samples the 1 Pulse-Per-Second signal. It is critical that this is sampled with the same clock as is used for data acquisition, so the GPS time base can be used to compute actual sample rate.

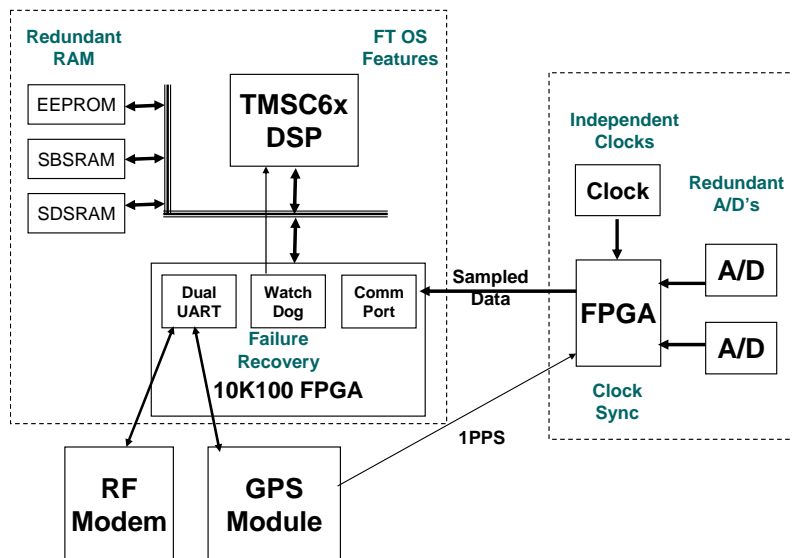


Figure 6: WASP Sensor Node Hardware Architecture

The sensor node uses a Trimble LassenSQ, low cost, lightweight GPS module for positioning and timing. Positions are accurate to 5 meters and clock pulses are specified accurate to 95 ns (although we have measured jitter of up to 250 ns).

A MicroHard CompactRF modem is used for communications to the base station and to other nodes.

VU FT Comp Implementation

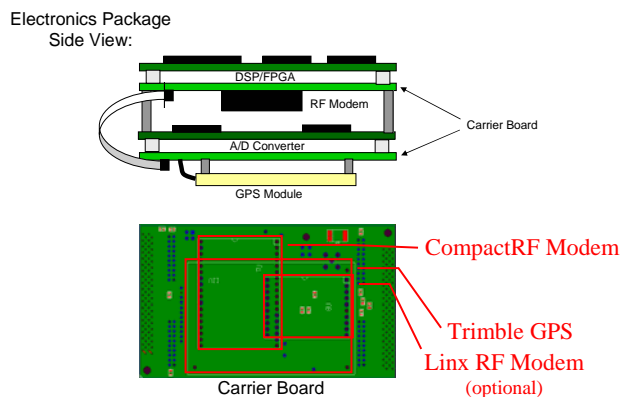


Figure 7: Sensor node Assembly & Carrier Board

The modules are physically assembled via a stack of printed circuit boards, shown in Figure 7. A special carrier board was designed and fabricated to hold either the computational module or the acquisition module.

The size/weight/power specifications for the processor stack are:

- Size: 6.3 x 3.15 x 3.35 in.
- Weight: 0.9 kg
- Power: 9.2W

Note that no active power management has been incorporated into the processor stack. Future redesigns could drastically reduce the power, weight and size of the sensor node.

Radio Sensor Units

Consistent with the concept of operation, a low cost, low power radio waveform was chosen. For the purposes of the project, the Family Radio Service (FRS) was used. Since the FRS radios were designed primarily as a cost-effective device for voice interactions, their inherent performance metrics are far below the necessary limits for an accurate geolocation system. In particular, the phase linearity across the audio spectrum is extremely poor. The radios also exhibit, for the current geolocation application, spurious, inadequate base-band performance and inconsistent near-field effects.

The Motorola TalkAbout T5320 radio was used for the prototype. Its design was reverse engineered. The block diagram is shown in Figure 8.

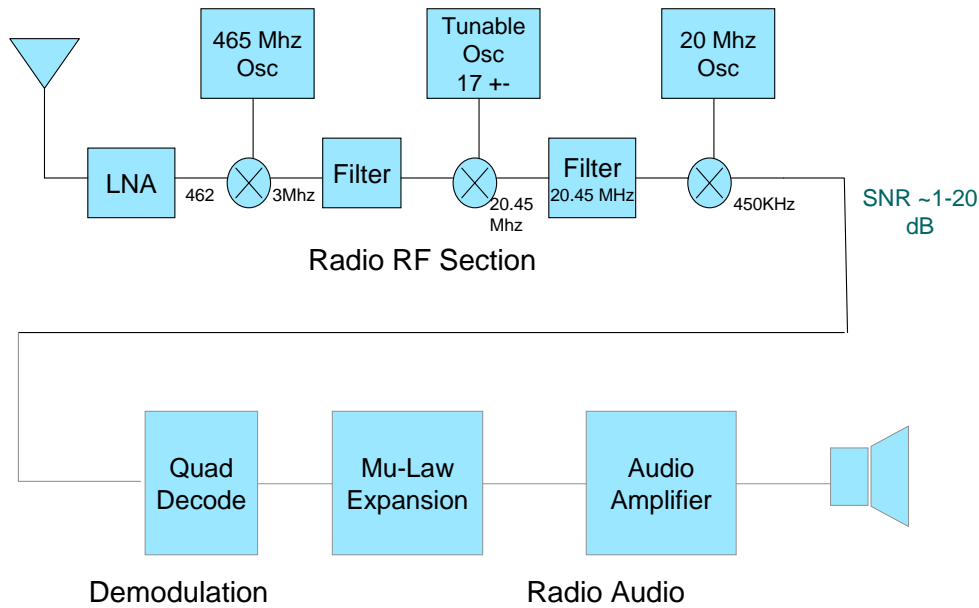


Figure 8: Motorola FRS Block Diagram

The signal received at the antenna is the input to the first stage buffer, a LNA (Low Noise Amplifier). A local oscillator is mixed with the signal, followed by a filter stage to remove sidebands. The tunable oscillator selects a FRS channel (25 KHz bandwidth) using the narrowband 20.45 MHz filter. Finally, a 20 MHz oscillator/mixer combination generates the FRS intermediate frequency of 450 KHz.

Demodulation is performed with an analog quadrature decoder. A Mu-law expansion, followed by a voice-quality audio-range power amplifier completes the process of recovering the audio signal.

The initial approach in using the FRS radio was to employ the entire audio recovery signal chain. Measurement of signals revealed a highly non-linear transfer function, in amplitude and phase. Phase nonlinearity is a serious drawback for TDOA use. Working back through the signal chain revealed that the analog demodulation and associated high-Q filters were the source of the distortion.

These findings necessitated the use of *software radio* (SWR) techniques to provide a stable, linear phase transfer characteristics. The software radio implementation will be described in a later section. In order to enable SWR computation, a small amount of analog preprocessing was

required. Vanderbilt designed and fabricated a Printed Circuit Board (PCB) to perform these functions. The block diagram in Figure 9 shows this circuit.

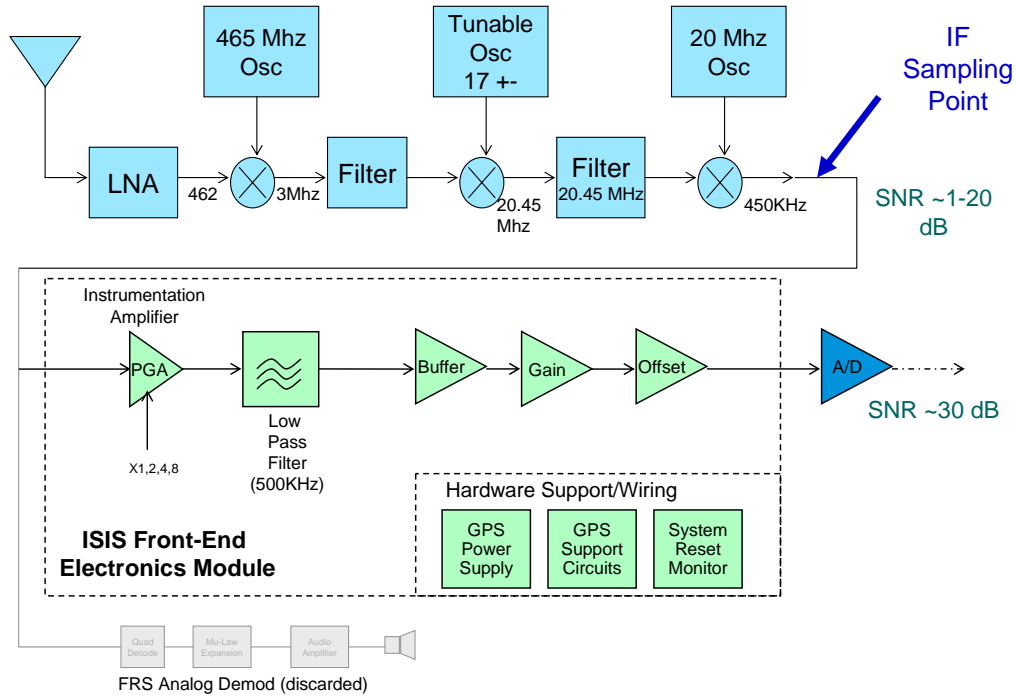


Figure 9: Modified Radio Block Diagram, with SWR Front End Electronics

The IF signal is amplified with a programmable gain amplifier to permit software-adjustable auto-ranging, for pickup of low level signals. A 1 MHz cut-off low-pass filter removes high-frequency signal components to prevent aliasing of the 2 MHz sampled data. Gain and offset stages condition the signal to the proper levels for the A/D converter.

This circuit was fabricated on a 2 layer PCB. The size of the board and mounting holes match that of the processing nodes for incorporation onto the processor stack.




Figure 10: Photo of Modified Radio

Figure 10 shows the modification to the Motorola radio. The demodulator chip was disconnected from the IF signal, along with its reactive components. The signal pick up wires were attached to the IF signal points. A buffer amplifier was added to reduce cable effects and the effect of loading of the programmable gain amplifier.

Software Radio -- Signal Processing

The first phase of SWR signal processing is in the detection of a transmission. Current DSP performance prohibits a constant demodulation approach. Instead, raw energy detection is performed on all input samples. When the energy passes above a fixed threshold, data capture and processing is triggered. This initial process is classically called “New Energy Arrival” (NEA) by the SIGINT community. NEAs are subsequently processed to see if they are a Signal Of Interest (SOI). In the case of the WASP system, the time of the SOI is particularly important since it (when combined with the TOA of geographically separated receivers) is used to determine the geolocation of the transmitter.

Upon detection of the NEA, approximately 1 second of data is acquired for SOI time evaluation. This data is stored in onboard RAM and then evaluation begins.

Signal Detection/Estimation (On-Sensor Platform)

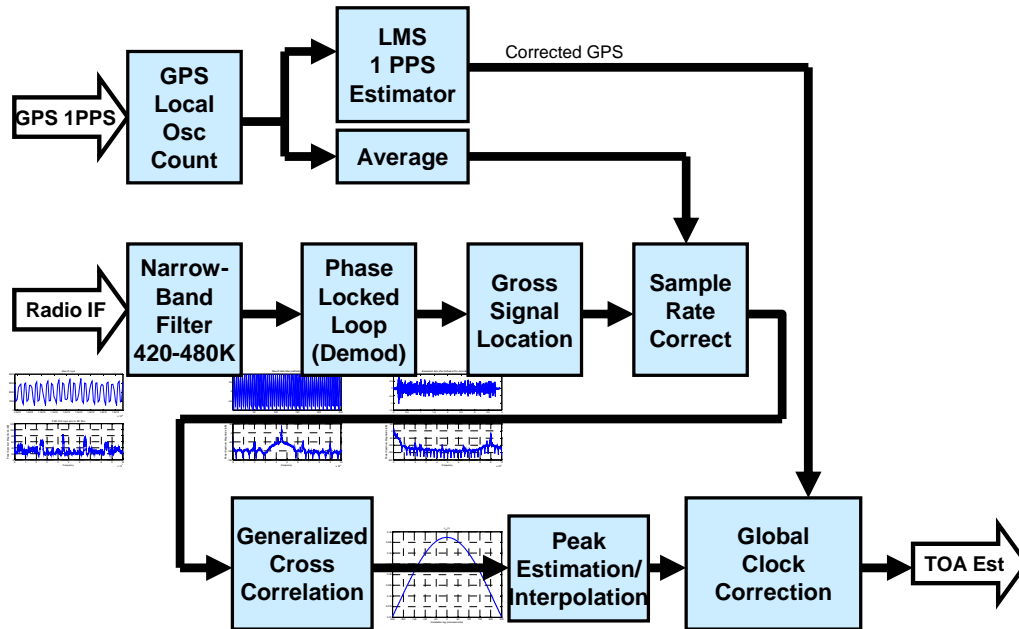


Figure 11: Signal Processing Block Diagram

Figure 11 provides details concerning the software processing associated with each NEA. The sampled IF signal is first filtered by a high quality FIR (Finite Impulse-Response) bandpass filter. A digital phase-locked-loop is used to perform the FM demodulation of the signal.

After demodulation, a rough signal detection is used to locate the target template within the 1 second of stored data. The signal search is shown in the diagram below. This process is basically a “matched filter”. For more exotic radios, it can be equated with watching for a particular phone, ID, phone number, launch code, etc. Pinpointing the transmission location of this type of information can often be very important.

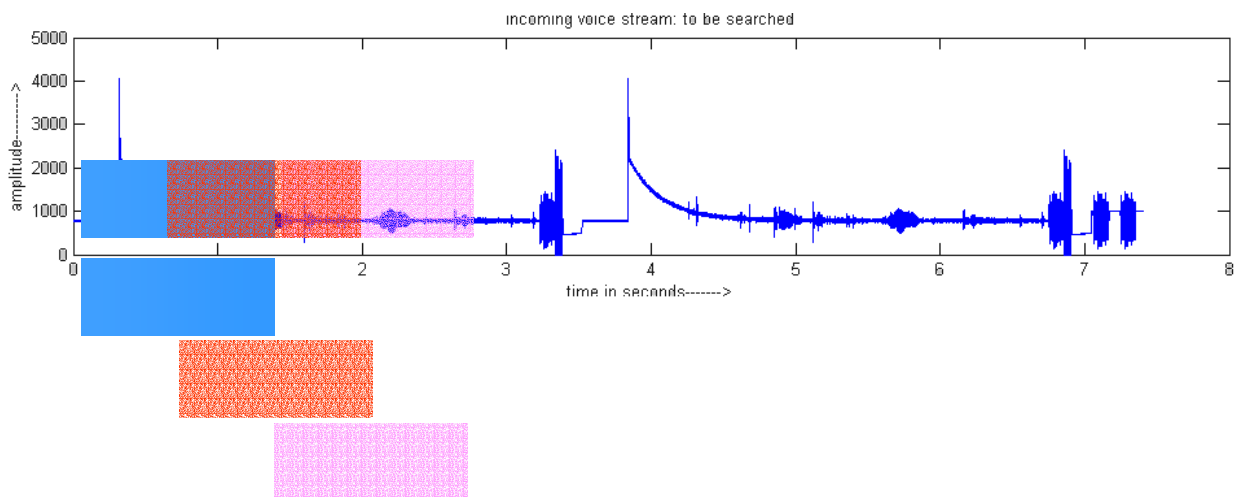


Figure 11a: Gross-level Signal Location

To perform this matched filter, the 2 MHz signal is decimated to an effective bandwidth of 32 KHz, via filtering and downsampling (Figure 11a). An FFT (Fast Fourier Transform)-based cross-correlation is used, with an overlap factor of $\frac{3}{4}$. The maximum of the cross-correlation is used to extract the block of data containing the signal of interest.

The next step is to find an extremely accurate position of the signal within the block of collected data. For this, the Generalized Cross Correlation (GCC) is used. The signal is processed in blocks of 8K samples, with an overlap of 75%. For each block, a complex FFT is computed. The complex spectrum is multiplied with the complex spectrum of the template signal. This complex product is accumulated, effectively implementing a Welch Transform. Several computational optimizations are implemented:

- Since the signal is band limited, only the first 40 frequency bins are within the signal bandwidth. Consequently, only the first bins are computed in the template. The same applies for the cross multiply and accumulation.
- The template is pre-computed at bootstrap.
- 8K FFT's are chosen to optimize the use of internal RAM on the DSP.

At this point, the fractional sample index of the TOA for the SOI has been evaluated. The index then needs to be rolled back into the precise time it occurred such that the arrivals at distinct nodes can be compared. The process of correlating the fractional index to a global clock standard GPS is described next.

Synchronization and Clock Correction

In signal time of arrival estimation, accurate time measurement and synchronization is paramount. In the low-cost approach, this must be accomplished with low quality hardware and

post-processing algorithms. Figure 12 shows an example of the jitter associated with low-cost GPS units.




Figure 12: Example GPS Jitter

Synchronization relies on several assumptions:

1. GPS clocks are highly accurate over the long term. This is due to the Cesium Atomic Clocks used to synchronize the satellites.
2. GPS clocks have high jitter for any one measurement. This is due to the variations in RF propagation, signal-to-noise ratio at the receiver, etc.
3. Local oscillators have high variation in base frequency. This is due to the sensitivity in the manufacturing process.
4. Local oscillators can have drift, due to temperature variations at the crystal package.
5. Local oscillators have low jitter, due to the stability of crystal oscillators.

The hardware collects the raw data for the corrections as follows (See Figure 13):

- A 48 MHz clock drives all circuits.
- The 2 MHz A/D sampling clock is directly derived from the 48 MHz clock with a division by 24.
- The 48 MHz clock drives a free-running 32 bit clock.
- The 1 PPS signal is sampled and used to snapshot the counter. These counter snapshots are reported to the DSP for further processing.

To provide data for the correction algorithms using the collected clock and A/D time data, the following are computed:

- Multiple deltas between readings of the Free-Running Counter (FRC) can be averaged to estimate the frequency of the 48 MHz crystal.
- In software, the 32 bit counter is extended to 64 bits. This forms an absolute time base.
- Signal samples are also counted exactly, (block # * number of blocks + offset within a block).
- Samples can be correlated to the extended 64 bit counter, since both sample count and FRC started at the same time and were derived from the same clock.

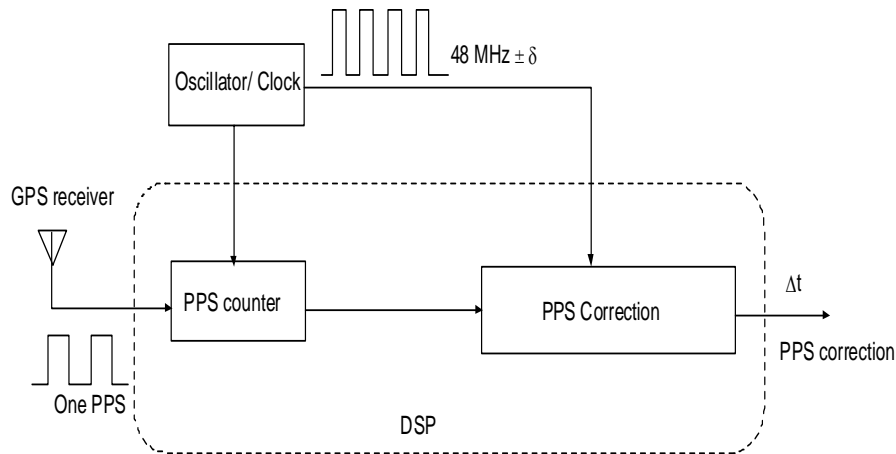


Figure 13: Schematic of 1 PPS Sampling

Thus, we can determine the absolute frequency of the local oscillator, and thus the absolute sample rate of the sensor. This can be used to compute an accurate time from the prior 1 PPS (Pulse Per Second) for any given sample. This number will still be influenced by the GPS jitter.

We can also use multiple FRC readings to counter the GPS jitter. By using a Least Means Square (LMS) algorithm, we can predict the correct location, or a correction factor, for any single GPS pulse. The length of data used in the LMS prediction is dependent on the stability/drift of the local oscillator. We have used ~20 seconds of data as a compromise between oscillator stability and correction accuracy.

Real-Time Processing

Maintenance of real-time behavior in all aspects of the signal processing chain is critical in this system. The figure below shows how real-time is maintained.

The data acquisition rate process is almost entirely a hardware activity. The FPGA on the data acquisition module generates the 2 MHz A/D clock via a division by 24 of the 48 MHz clock. Results from the A/D's enter directly into a 32K FIFO memory. This FIFO decouples the A/D from the data transfer process, such that no sample is ever lost. The Data Acquisition FPGA controls a high speed interface to the DSP FPGA. Time-stamped data is sent to the DSP FPGA, which is transferred directly into DSP memory via a DMA processor on the DSP.

In the same way, the GPS 1 PPS data is transferred, using another high speed interface.

The software process is shown in Figure 14. After an initialization process, a pair of real-time processes concurrently:

1. Use the DMA engine on the DSP to transfer data into memory. On completion of a 64K block, an interrupt occurs, where,
2. The signal energy is computed and PPS data is logged. When the energy threshold is triggered the data buffer is filled. After filling the buffer, the real-time mode is finished.

TOA estimation occurs as a non-real-time task. After computation, results are reported via the RF modem and the system restarts to the INIT mode.

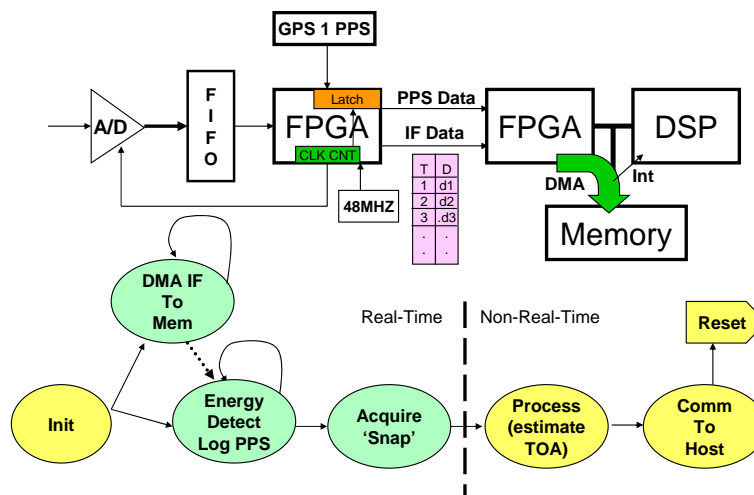


Figure 14: Structure of Node-based Real-time Processing

User Interface and Communications Infrastructure

A user interface was created to monitor the system and visualize the results. Figure 15 shows the WASP user interface.

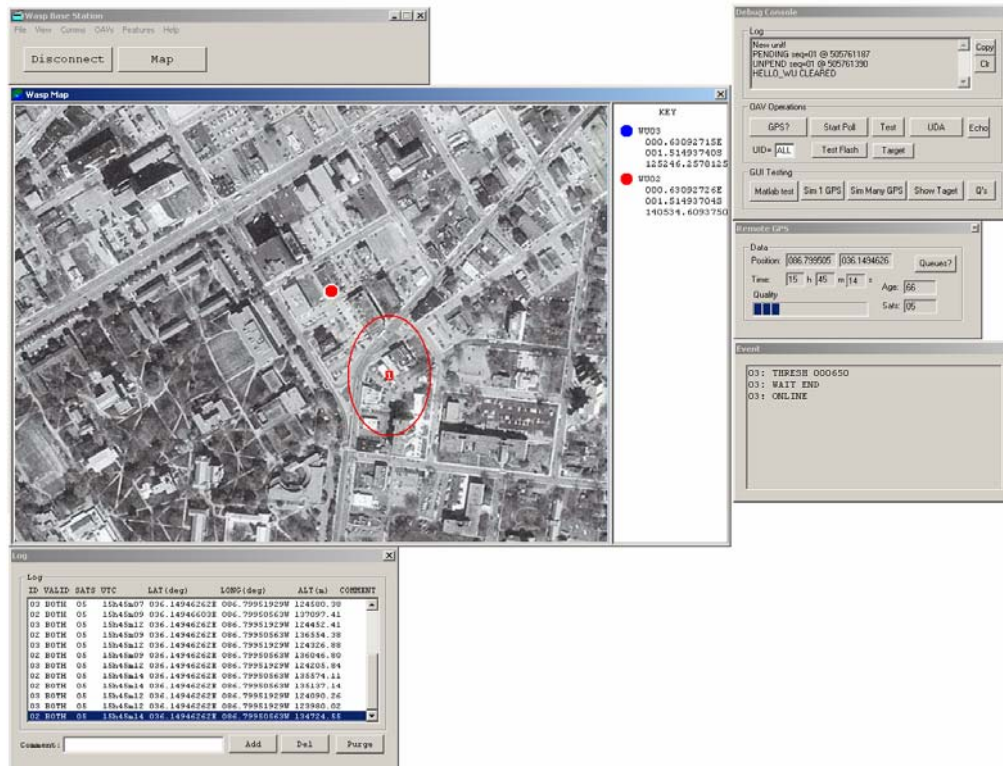


Figure 15: WASP User Interface

The large window shows an overhead satellite picture, registered with GPS coordinates. A color-coded circle annotates the position on the picture. A legend bar on the right shows the exact coordinates.

A control dialog box provide high-level control, connecting/disconnecting the base station console from the real-time data stream. Various debugging windows show and control the contents of messages traveling across the network. A status window shows RF modem/GPS signal strength. Another dialog shows node status. Finally, the log at the bottom shows the individual TOA messages from the nodes.

Results and Discussion

Basic Radio Performance

The TDOA algorithms and data were tested extensively in several laboratory and field configurations. The initial test evaluated the raw capability of the FRS radios as a TDOA radio front-end. The configuration of the test is shown in Figure 16.

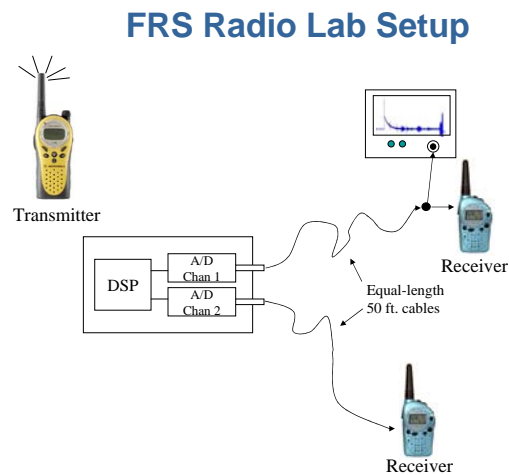


Figure 16: FRS Laboratory Experiment

Two FRS Radios were attached to a single sensor node, using the dual A/D capability of the analog input module. Varying lengths of cable were used to simulate a physical separation between the nodes. For each reading, the baseband data was cross correlated to determine lag. Matlab was used as a flexible experimentation platform. The results indicated the non-linear phase distortion within the radio's analog section. From these results, the project was redirected to include a software-based radio approach using the IF data.

FRS Radio Attenuation Performance

Signal time of arrival estimation accuracy is highly dependent on the signal-to-noise ratio (SNR) of the signals. To evaluate the effect of distance and obstructions, and to prepare for a planned demonstration at Honeywell. Figure 17 shows the measurement locations and measured SNR.

The results show a dramatic effect of obstructions on signal strength. The topography of the area included forested regions and small gulleys and hills. The results indicate the large gains expected in SNR with the ability to move antennas to avoid obstructions and to get closer to the signal.




Figure 17: SNR Measurements at Honeywell

The results of the SNR translate directly to accuracies that can be obtained in TOA estimates. Figure 18 shows the impact.




Figure 18: Expected Accuracy vs SNR vs Average Time

Signal Processing Measurements

The following figures show the progression of the signal through the onboard signal processing.

The initial data, as sampled is shown in Figure 19. The signal quality for a nominal strength signal is on the order of 30 dB.

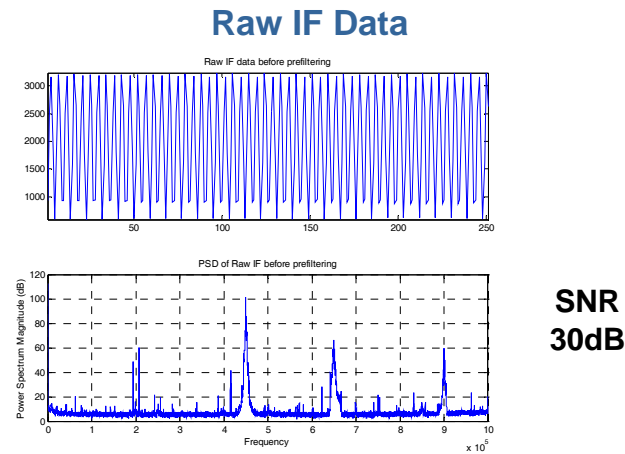


Figure 19: Raw Input from Radio

Figure 20 shows the processing gain after filtering and demodulation. In the example shown below, the SNR is approximately 85 dB. Results from other measurements show SNR up to 110 dB. The gain in SNR due to the digital demodulation is pivotal in achieving accurate TOA, and thus geolocation accuracy.

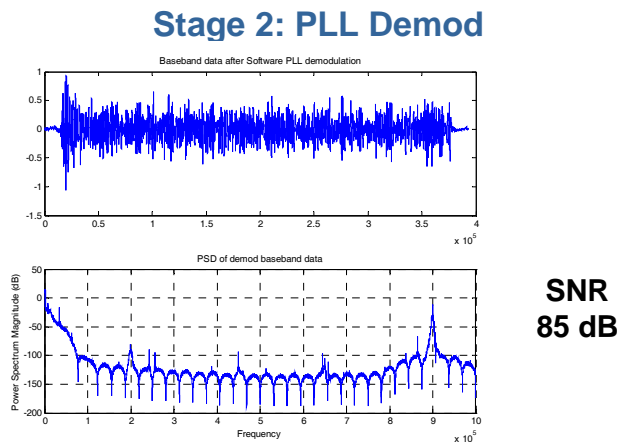


Figure 20: Results of Demodulation

The final GCC (Generalized Cross Correlation) processing results in a highly interpolated cross correlation array. The peak of the cross correlation indicates the delay between the template and measured signals. Figure 21 shows an example of this cross correlation function. The quality of the cross correlation w.r.t. TOA estimation is related to the sharpness and smoothness of the peak. The results show adequate quality to achieve desired accuracies.

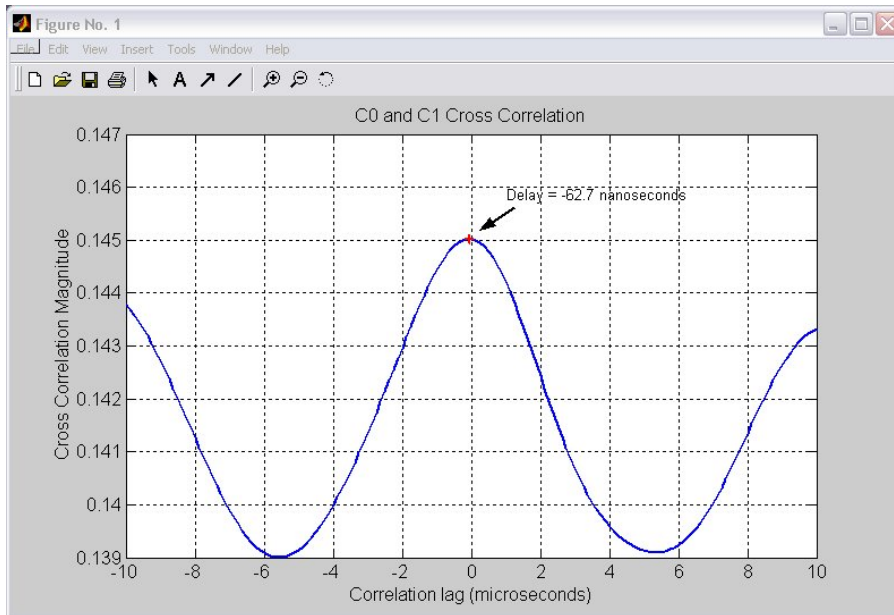


Figure 21: Cross Correlation Plot

Results from experiments to assess the quality of radio and digital demodulation on a pair of FRS radios are shown below. The results in Figure 22 show a wide variation during warm-up of the system, followed by relative stability. General trends in timing variation indicate a dependence on temperature. This temperature sensitivity can be overcome with frequent calibrations.

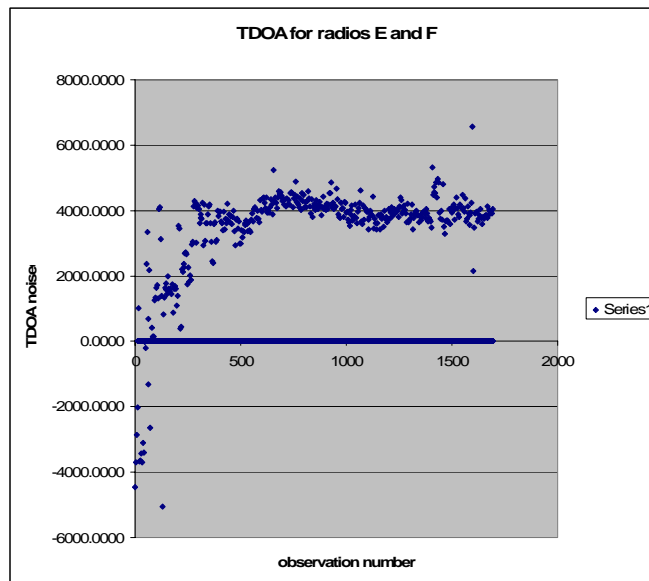


Figure 22: Long Term TOA Results from Radio Pair

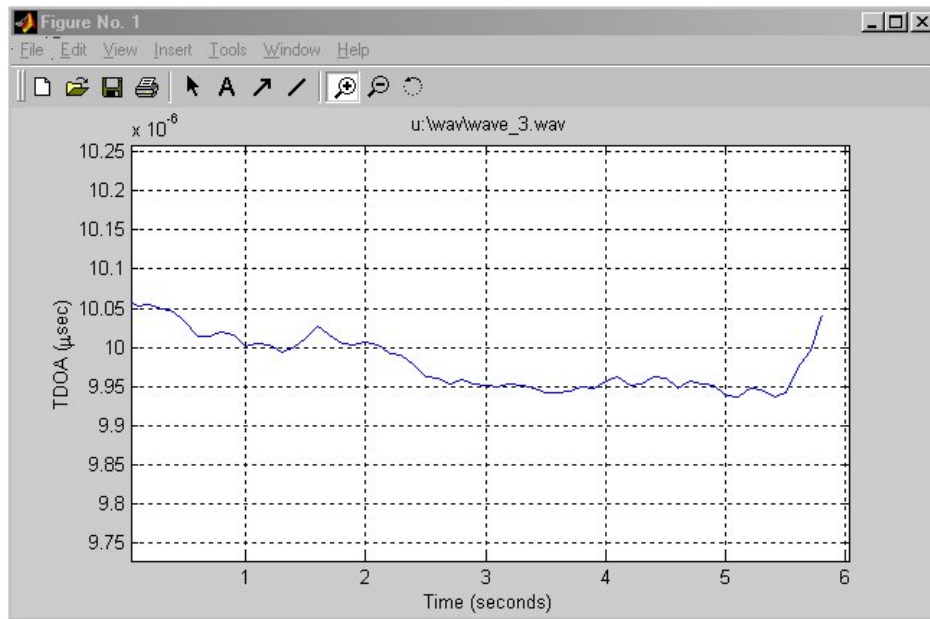


Figure 23: Incremental TOA Measurements

The prior results show TOA measurements spanning over several hours. While drift can be compensated for over this time scale, short term drift would invalidate the results of a TOA. The short term TOA tests are shown below. The data indicates that there is approximately 100 ns drift over the 6 second measurement period. Over some sections (e.g. the last 3 seconds), the results are very stable (approximately 20 ns) as seen in Figure 23.

Overall, the results indicate that the approach, algorithms, and the initial hardware prototype can form a successful TDOA platform. Data indicates that the limitations of the inexpensive components and other hardware can be overcome via post-processing.

Flight Package:

The hardware was assembled into a flyable package. Initial prototypes were designed to fit the Allied eMAV. The packaging was designed to fit in the upper payload bay. Figure 24 shows the planned physical mounting.




Figure 24: Mounting on eMAV




Figure 25: Planned Flight Package

After the eMAV was ruled out, we designed a generic rectangular flight package. This package is shown in Figure 25.

UAV-based Antenna Arrays

Using specifications for the UAV and position sensing capabilities of the UAV, simulations were performed for antenna gains achievable using multiple UAV's. The diagram below demonstrates the array antenna concept.

The figure below shows the end-fire array concept. UAV's are flown into a linear array, to the best ability of the controller. The relative position of the UAV's are measured (via the onboard

GPS/INS) and a deviation from the optimal position is calculated. These deviations are used to compute corrections for the antenna phase shifts. The corrected phase shifts are then used to transmit RF energy.

Figure 26 shows several antenna radiation patterns. The ideal pattern shows performance if the antennas are positioned exactly as desired. The raw pattern shows the degradation in on-axis energy and energy leakage due to position errors. The phase-corrected pattern shows the impact of the correction.

Simulations show the benefit of the approach for low power, stealth communications. The main problem with the use of this technique is that the position measurement accuracy required is not available on the current UAV's.




Figure 26: Antenna Phase Correction

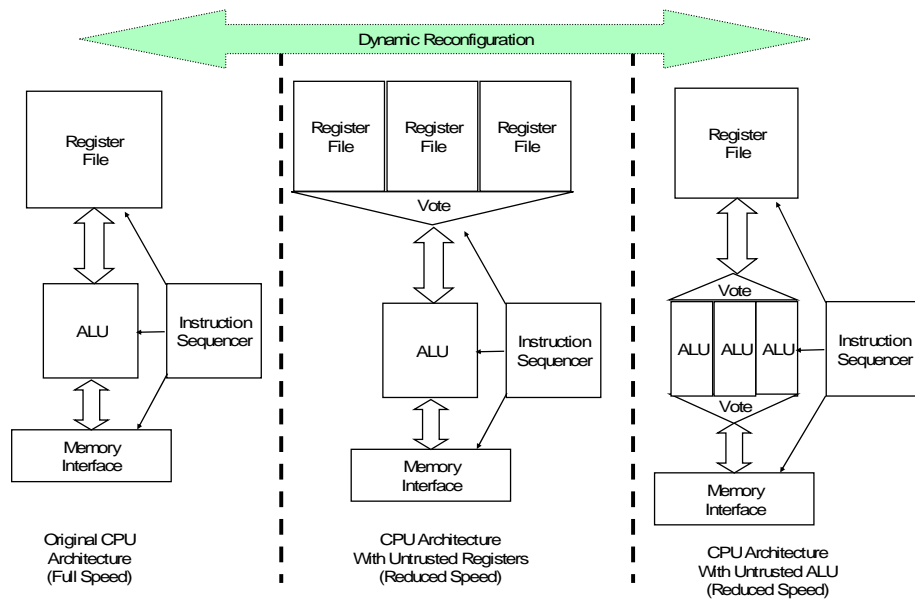


Figure 27: Hardware Fault Tolerance

Hardware Fault Tolerance

Fault tolerance was explored at the hardware level. Several variations on a CPU architecture were implemented as Soft Cores on a Xilinx Spartan 3 Field Programmable Gate Array (FPGA) Figure 27 shows the 3 variations implemented.

- Variation 1 is the standard, low resource, low tolerance version of the CPU. It uses a standard architecture (8051) with a single version of each of the components (Arithmetic Logic Unit, Instruction Sequencer, Register File, and Memory Interface).
- Variation 2 implements replication and voting (Triple mode redundancy, TMR) on the register file, assuming high SEI rate on the memory components.
- Variation 3 implements TMR on the ALU, assuming potential logic faults.

Since the processor is implemented on a dynamically reprogrammable core, dynamic swapping and fault resiliency can be used. All of the hardware modifications are completely transparent to the software.

The results of the processors were as expected. Higher levels of fault tolerance produced a processor with a lower maximum frequency. Mode 1 ran at 66 MHz. Mode 2 ran at 42 MHz. Mode 3 ran at 35 MHz.

CONOPS

A set of CONOPS were developed, with deployment strategies, costs, and expected performance. The WASP CONOPS were compared with traditional methods for area surveillance. Those CONOPS are described in detail in the attached presentation.

The basic operation (Figure 28) is to use sparse sets of UAV's to perform a wide area surveillance. This will keep costs of the system low. When a transmission is detected, UAV's relocate based upon a gross estimate. At the next transmission, a tighter position fix is computed. This sequence repeats, until the error probability is low enough for direct action.

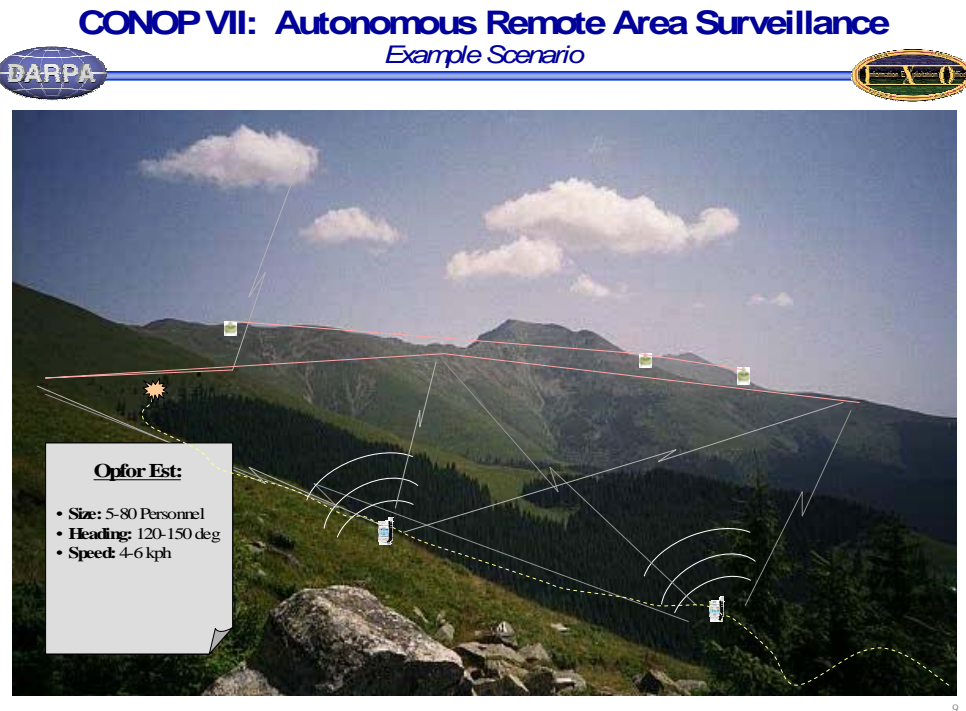


Figure 28: Concepts of Operations for Mountainous Surveillance

Conclusions

Experiments using the WASP hardware and software produced results that show the approach is highly feasible. The system was produced with low cost, commodity hardware, using standard radios. Results are competitive with high-end TDOA systems. The specific points that allow this high accuracy results stem from the following observations:

- Analog circuit quality is extremely important. The stability of the measurements is directly affected by the phase stability of the measurements.
- Software radio is mandatory for phase linearity and temperature immunity.
- Mobile sensor units contribute significantly to the accuracy of the system.

- The inexpensive, mobile TDOA system is flexible and can contribute to force situational awareness in ways not possible otherwise.

Recommendations

Work on the TDOA package has continued under the PAM seedling effort. Further development of software and hardware has proven the system performance in full-scale testing. To prove the system out for the WASP CONOPS, several optimizations are needed:

1. Power: The initial prototype was done cheaply, leveraging existing computer and acquisition platforms. Power was not a consideration in these designs. A power optimized design should lower consumption by 10x or more. This will allow a 10x improvement in runtime.
2. Cost: The current TDOA package costs approximately \$2500 to build. Design optimization can reduce this to 1/5th the cost.
3. Sensor Placement: The seedling work shows the expected dependency of SNR on sensor location. To enable effective deployment, a sensor placement is critical. These placement decisions are multi-dimensional optimizations, planning for maximum accuracy while minimizing energy expenditures, threat to assets, speed of position fix, etc.
4. Accuracy enhancement: Initial correction algorithms show large potential for accuracy enhancement. This, combined with a dynamic calibration process will result in consistent, high quality results.
5. Sensor fusion: Radiolocation is just one part of the wide area surveillance CONOPS. The sensor node must be extended or combined with sensors for video, dismount detection, vehicle detection, etc. Optimum platform efficiency will be obtained if all sensor modalities are considered in the design of the sensor platform.
6. Multiple CONOPS are possible. Due to the software radio approach, the system is amenable to other radio waveforms and units. This would enable a TDOA system to be rapidly deployed using captured radio units. A lengthy design and characterization process could be avoided by treating radio electronics as a 'black box'.