

WASP: A Radio Geolocation System on Highly Resource Constrained Mobile Platforms

Kumar Chhokra, Ted Bapty, Jason Scott, Simon Winberg, Janos Sztipanovits
Vanderbilt University
{kumar, bapty, jscott, swinberg, sztipaj} @isis.vanderbilt.edu

Don van Rheeden, Ben Abbott
Southwest Research Institute
{donald.vanrheeden, babbott}@swri.edu

Abstract

In recent years, there has been an increased need for surveillance capabilities in both civilian and military arenas. Mobile unmanned sensor fleets have long been envisioned as a tool for acquiring such intelligence. This paper describes a geolocation system as a payload for unmanned aerial vehicles. In particular, we address the issues that arise in designing an inexpensive, distributed, mobile, RF sensor processing system: low payload cost; power-limited components; real-time constraints; absence of common clocks; and varying environmental conditions. We also target, qualitatively, the tradeoff decisions across the dimensions of accuracy (of time of arrival estimates and the final geolocation), computational effort, physical dimensions, and operational logistics. The final system architecture, algorithms, and accuracy figures are also described.

1. Introduction

Position information is essential in many applications. A hospital or health-care facility may wish to be informed of the locations of all its critical staff on its premises to deliver the best care to its patients at all times [1]; a military unit may wish to extend its reconnaissance capabilities beyond line of sight (BLOS) [2], [3]; a team of scientists studying animal movement patterns might want to log the positions of its subjects [4]; a search and rescue team can locate and provide help quicker if the location of the individual in distress may be accurately known in advance (the federal E911 program is such an effort) [5]. Having a system in place that can track and locate individuals is thus essential, particularly when the target is non-cooperative (i.e. does not transmit its location).

While time difference of arrival (TDOA) based radio-geolocation systems have existed for a long time, the current need for greater position accuracy in urban or

hostile environments demands that the traditional methods be augmented. Traditional radio-geolocation systems, such as those used by the military and cell-phone providers, are limited by large stand-off ranges, fixed measurement locations, and small numbers of sensors (due to high cost and large form-factors [6]). The city environment is cluttered with buildings and other large structures that restrict transmission / reception ranges, introduce back-scatter and multi-path interference. In a military scenario, the adversary may be sheltered by foliage or a complex terrain, necessitating close-range tracking and stealth. Targets with low-power transmitters further complicate long-standoff systems, due to low signal-to-noise ratios.

Unmanned aerial vehicles (UAVs) and organic aerial vehicles (OAVs) equipped with geolocation electronics can afford proximity, safety, security and stealth. However, the nature of these devices introduces several technical challenges: robustness and reliability under military use conditions demand that the devices be compact, rugged and mechanically stable. Stealth in reconnaissance implies a small, RF-silent design (i.e. limited communication bandwidth) while, endurance and usability demand an extensive operating life (i.e. low power). These UAVs are envisioned as a fleet of attritable mobile devices, implying a relatively strong constraint on the manufacturing and operating cost.

Widely Adaptive Signal Processing (WASP)[†] is an effort to produce a fleet of low-cost, low-power, mobile geolocation units. This paper describes the design of the TDOA based geolocation payload of a WASP unit (WU).

2. Design objectives and considerations

[†] This work has been supported by DARPA / IXO under the WASP contract F30602-02-2-0206

On highly resource and power constrained platforms, the TDOA and other signal processing algorithms must be adapted to have the following desirable features [7]:

2.1. Distributed processing

Low communication bandwidth and stealth behavior require that the signals sampled at each node be processed locally. This obviates transmission over the wireless network for common processing at the base station. The nodes must use *a priori* knowledge about the signal of interest (SOI) to extract the relevant discriminators. For TDOA algorithms, the most important statistic is the time of arrival (TOA) estimate and some measure of the confidence of the estimate. These statistics are much more compact than the received signals and may be communicated to a base-station (BS) for locus computation and tracking.

2.2. Goal-oriented on-demand processing

The nodes must expend power only when an interesting event occurs. The nodes communicate their location and TOA estimates only when requested by the base station. This aids in stealth because the nodes operate in a passive mode between transmission requests. Since each radio communication consumes micro-joules to milli-joules per transmitted bit (depending on the desired range and link geometry of the radio network), power is conserved by transmitting data only on demand [8].

2.3. Multi-resolution processing

The signal processing algorithms required to estimate the time of arrival typically are computation and power intensive. A higher temporal resolution demands the time series data be gathered and processed at a high rate. A higher bandwidth, however, implies larger power consumption. The nodes may thus save power by searching for an event of interest at a coarser resolution, and then switch to a more computation intensive higher resolution when a SOI has been detected.

3. WASP Unit Hardware

A prototype electronics payload architecture that meets the cost, form-factor and power constraints is shown in Figure 1. A commodity FRS-band radio is used as a receiver. An intermediate frequency (IF) signal is sampled by tapping into the analog circuits of the receiver. This signal is sampled at 2 MHz and brought into the memory of a DSP. A GPS module provides both the GMT time/position and a 1 pulse-per-second (PPS) signal accurate to ± 100 nano-seconds (ns). Using this architecture, several challenges arise:

3.1. Low quality radio units

Analog processing of the radio signal impedes high precision measurement. Circuit inconsistencies between different radio units, temperature drift, interference, varying stray capacitance and inductance etc. all affect phase performance.

3.2. Lack of common global clock

The TOA estimates from each unit must be referenced against a common time. The only available “global clock” is the GPS 1-PPS, specified accurate to ± 100 ns. Measurements observe the deviation to be ± 125 ns or more.

3.3. Drifting sample rate clocks

Each WU has an independent clock, accurate to ± 100 PPM [9]. The frequencies and/or phases of the clocks on different nodes will generally not align. The TOA algorithms must correct for frequency and phase drift.

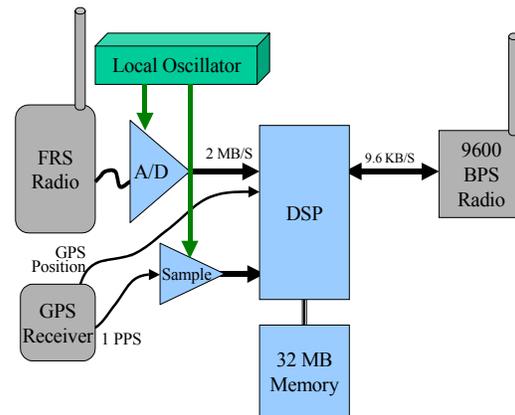


Figure 1. Wasp Unit Hardware Architecture

4. Signal Analysis and System Architecture

The following discussion shows how the overall signal analysis and software architecture, depicted in Figure 2, addresses the challenges recognized in sections 2 and 3.

4.1. Low quality radio units

Digital processing removes many of the variables that impact radio phase performance. Ideally, a fully digital radio could be employed, but the size, weight, power and

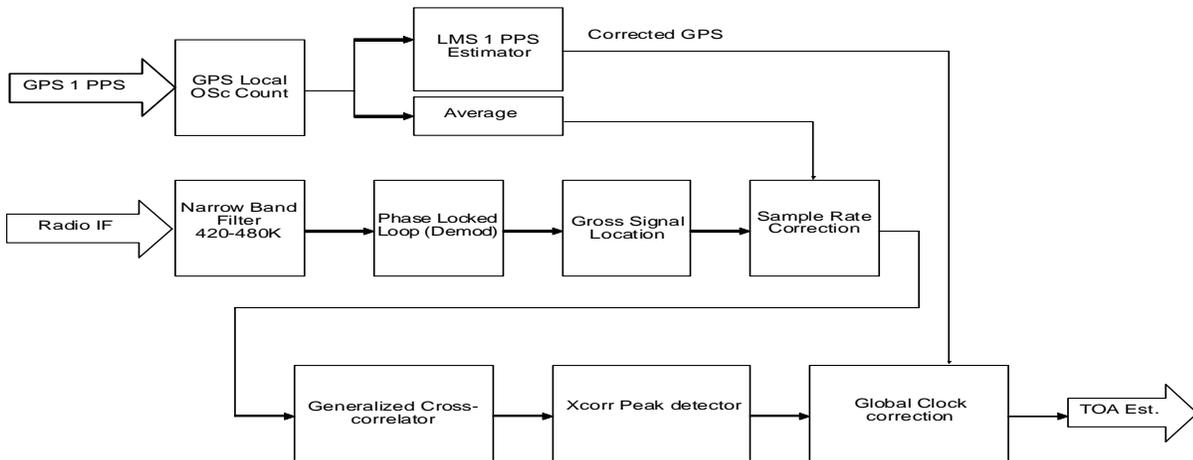


Figure 2. Overall signal processing architecture

cost (SWPC) parameters prohibit this approach. Instead, we compromise and sample the intermediate frequency of the FRS radios. The IF is a 450 KHz FM signal, with a 25 KHz channel bandwidth, sampled at 2 MHz. A narrowband filter is applied to the IF. Demodulation is performed using a software based phase-locked loop.

These software radio functions of filtering and demodulation are standard DSP operations and will not be discussed further in this paper. We concentrate instead on the processing and implementation of the PPS correction, Gross Signal Location (GSL), Generalized Cross Correlation (GCC) [10], and Sample Rate Corrections (SRC) blocks.

4.2. Distributed processing

The challenge of distributed processing manifests itself as two sub-problems: a) achieving global synchronization, and b) locally detecting the SOI and estimating its time of arrival. We address the issue of a common global clock first.

4.3. Global Clock Synchronization

The WUs are, by design, spatially separate: they cannot have a common physical clock. A high degree of external synchronization [11] may be obtained by using the GPS clocks. The 24 satellites of the GPS constellation transmit a synchronization pulse every second. Since each of these satellites carries within it Cesium atomic clocks, and are in external synchronization with each other, the pulse per second (PPS) signal provides a reliable mechanism for synchronization.

The GPS based synchronization achieves several aims concurrently: being a passive method of synchronization, the number of messages between participating nodes is reduced thus providing greater stealth and savings in transmission power. In most cases, it also improves the area coverage as the nodes needn't be in line of sight (LOS) and within each other's transmission radius. The spread spectrum encoding of the GPS signal also affords some level of security against malicious jamming.

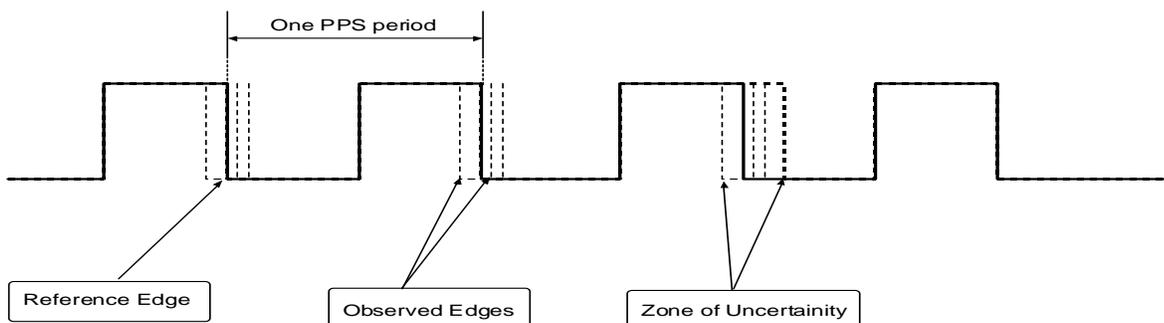


Figure 3. Schematic of clock jitter. The actual edge (represented by the dashed-lines) appears with a certain error around the “expected” or “reference” edge

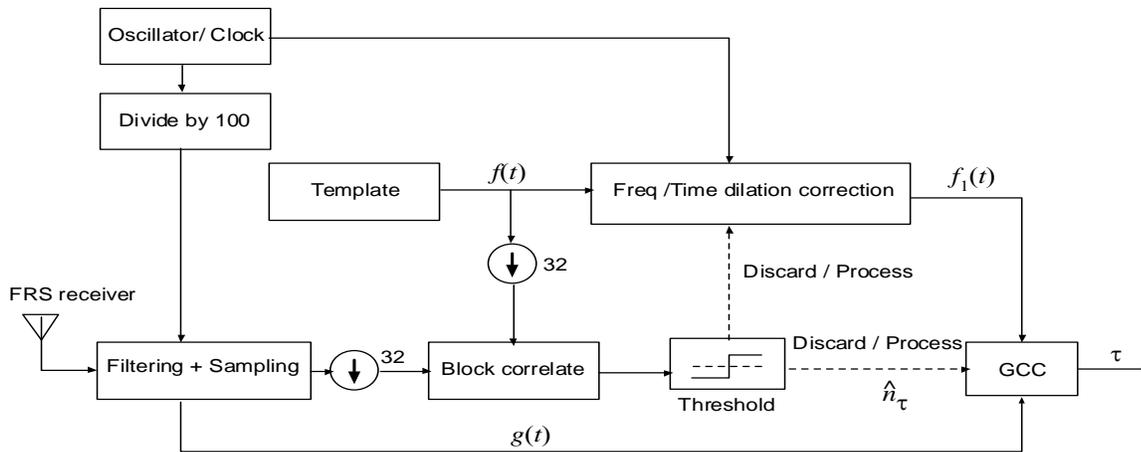


Figure 4. Coarse-grained search and Fine-grained estimation

However, the PPS signals received by the WUs in geographically distinct areas may contain jitter. These timing jitters may arise due to several factors including atmospheric effects (such as varying ionospheric charged particle and water-vapor concentrations) or echoes or obstructions in the GPS signal path. Moreover, the GPS devices themselves possess some tolerances that may introduce significant errors. Also, some GPS devices create time glitches when switching between satellite clusters, and lose time-lock with fewer than three satellites in view. Some GPS devices treat their serial output as a “low priority” task and can delay time code outputs by randomly varying amounts if they are busy computing.

High-end COTS GPS receivers have tolerances of a few tens of nano-seconds. However, due to SWPC constraints, their use is prohibitive. The embedded GPS is specified to be accurate to ± 100 ns. Since the WASP units use radio signals that travel at the speed of light ($\sim 3 \times 10^8$ m/s) every nano-second of error in the TOA estimate corresponds to roughly a foot of uncertainty in the source position estimate.

Figure 3 depicts clock jitter where the actual clock edge occurs not at the expected instant in time, but within some region of uncertainty around the expected time. Since the randomness of the edge is caused by a variety of reasons, the overall distribution of the clock jitter may be assumed to be Gaussian.

In addition to GPS jitter, the local clocks at different nodes generally start at different times and thus have an inherent offset. They also suffer from thermal drifts, leading to varying operating frequencies depending on environmental conditions. Figure 5 shows the variation of real TDOA estimates over time for two nodes. While

there are methods for reducing thermal drift (e.g. by using a temperature controlled environment for the crystal oscillator) and GPS jitter (better GPS receivers and antennas, differential GPS etc.) the SWPC constraints preclude the use of such existing solutions.

Each WU has a high resolution clock that operates at a nominal frequency of 48 MHz, which is used to derive a 2MHz sample clock. The same 48 MHz clock is used to sample the GPS 1 PPS signal. A high resolution (32-bit) counter measures the period of GPS synchronization pulses. The successive period values constitute a “history” of the duration of each GPS pulse interval. Since the GPS system is based on a high-accuracy Cesium clock, we can trust the PPS frequency to be accurate over the long term. We can also trust our local clock to have low jitter over the short-term. The long-term history is used to determine the current operating frequency of the WU local oscillator, which may in turn be used to compensate for the PPS jitters.

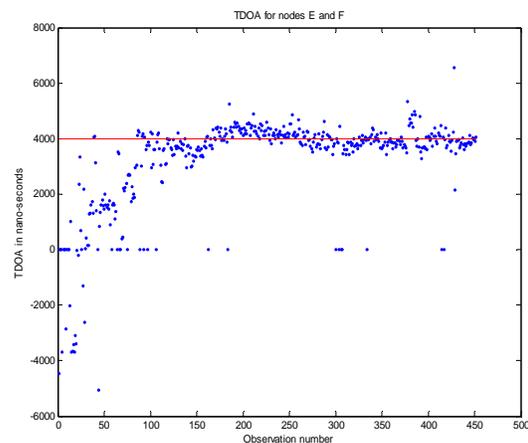


Figure 5. TDOA estimate variation over time

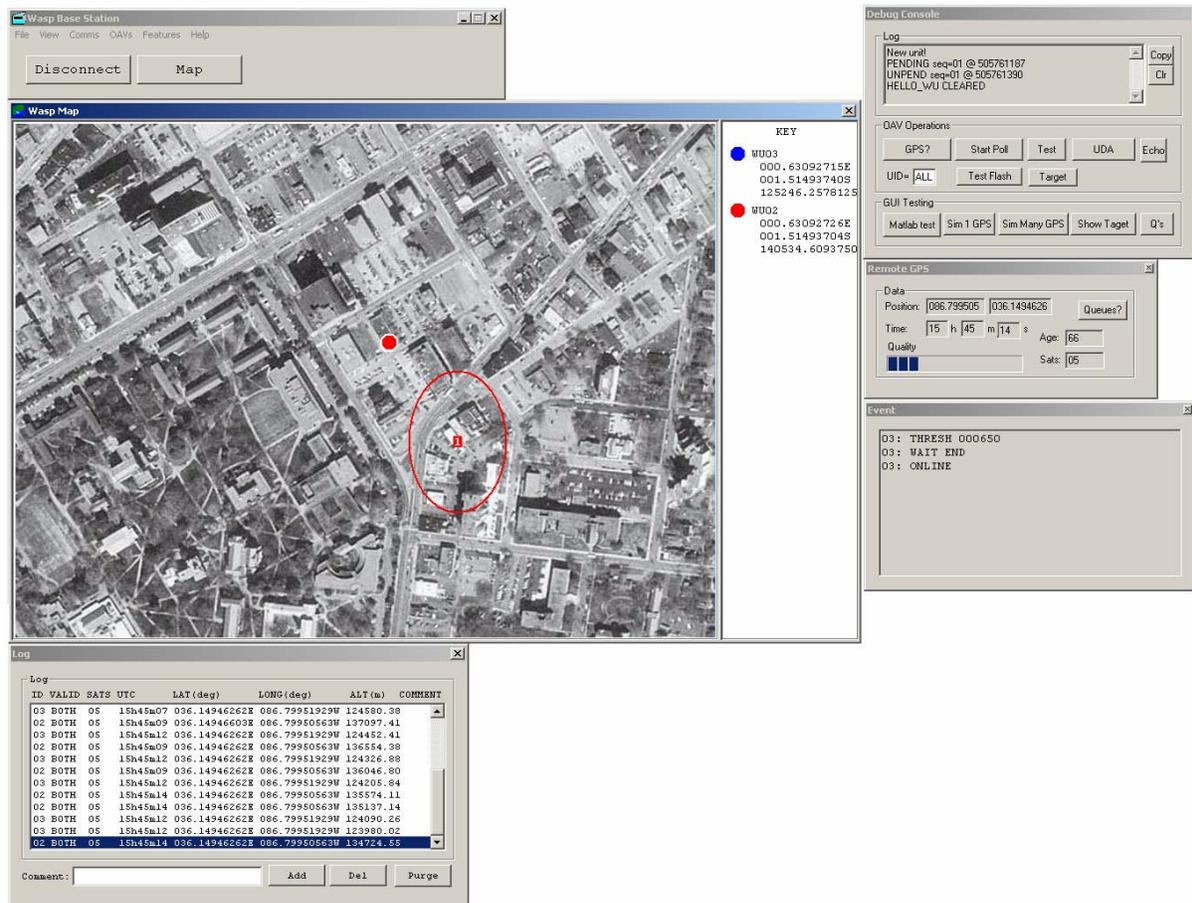


Figure 6. WASP Base-station GUI showing location of geolocation nodes

The PPS counts, when corrected for counter wrap-arounds, are a monotonically increasing series of values. We perform a LMS straight-line fit over a 22 second period to compute the current operating frequency as well as compensate for the GPS jitter. The parameters for the straight line fit, viz. its slope (operating frequency) and the time offset (beginning second of PPS history) may be computed from the closed form solutions of the LMS fit [12].

4.4. Multi-resolution SOI detection and TOA estimation

The WASP operating constraints warrant a technique that is heavily based on local processing. Typically known signal features, such as synchronization sequences, may be used as the SOI, or template. A good template should have the property of a sharp autocorrelation and low correlation with any other signal, allowing unique detection. A typical example could be the 26-bit pseudo-

random synchronization sequences [13] used in each TDMA packet burst. Such signal features are usually corrupted by additive noises (usually modeled as Gaussian disturbances) and may be reliably detected using a matched filter.

For a real signal $x_1(t)$ and its delayed and additive Gaussian noise corrupted version $x_2(t)$, the matched filter or correlation operation is given by [14]:

$$R_{x_1 x_2}(\tau) = \int_{-\infty}^{\infty} x_1(t) x_2(t + \tau) dt$$

The value of τ that maximizes the preceding equation gives an estimate of the delay between $x_1(t)$ and $x_2(t)$. This delay may also be interpreted as the "location" of the template, $x_1(t)$, in the received signal $x_2(t)$. The maximum value of R provides a measure of the "goodness" of the match: if the maximum is smaller than a heuristic threshold, we may deem the match as "invalid".

Matched-filtering and GCC are computationally expensive tasks. To maximize the accuracy of TOA estimates, the integration time for the correlation operation must be as large as possible. Typically, we collect up to a few seconds worth of data. Given a sampling frequency of 2 MHz, this represents a prohibitively large number of samples over which the cross-correlation must be performed. To circumvent this, we use a multi-resolution approach consisting of a coarse search of the SOI followed by a fine-grained estimate. Figure 4 illustrates the interaction between the two phases.

The first phase, gross signal location (GSL), attempts to locate the small template within the entire sampled signal. The gross signal location algorithm is described below:

1. Create low-resolution versions of the data. A low-pass filter is applied to the sampled data, followed by decimation. Low pass filtering prevents aliasing, while down-sampling reduces the number of data points that must be processed.
2. Perform a cross-correlation with a low-resolution version of the template. The down-sampled signal stream is split into 4096 point blocks with 50% overlap. A frequency domain cross-correlation is performed on each block.
3. Locate the peak in each block.
4. If the peaks per block exceed a detection threshold, select the maximum of the individual blocks. The “location” of maximum gives a low-resolution TOA estimate. If no value exceeds the threshold, then computation is stopped, due to the lack of signal or poor SNR, and the TOA estimate is discarded.

If the template is detected, then the second phase applies a generalized cross-correlation on the corresponding high-resolution versions of that data-block and the template.

This technique has the following advantages:

- Since the template is known in advance, its low and high resolution DFTs may be pre-computed, producing considerable savings in computational cost.
- The GSL phase acts as a discriminator: the more computationally expensive GCC is invoked only when a SOI has been detected.
- The reduced vector lengths which the GCC must process enable hardware and algorithm specific implementations, increasing the throughput.

5. Results

We present the results of the PPS jitter correction algorithm in Figure 7. Note that the variation is reduced nearly by a factor of 10. While intuition suggests that results improve as a greater history length is used, the long term instability of the local oscillator clocks makes the straight line fit assumption invalid. Figure 7 shows that even for a history size of 64 seconds a slight downward trend is visible in the corrected data.

The following setup was used to test the sample clock compensation techniques. A 1-sec, 5 kHz band-limited pseudo-random sequence served as the template. A signal source placed at a controlled distance from both receivers transmitted a sequence containing the template. The sensor nodes had on-board a pre-manufactured version of the template. The detection and TOA operations were performed locally on the nodes and the TOA estimates were transmitted to a PC (base-station) containing the TDOA location determining algorithms. The nominal sampling frequency at all the nodes was 2 MHz.

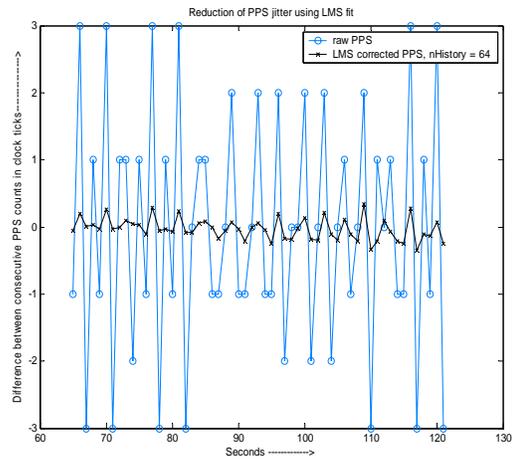


Figure 7. Effect of PPS correction. The LMS fit may be viewed as an FIR filter. The past 64 values are used to predict the next value. The black line with x’s show the result of the “filtering” operation. Note that the maximum deviation is less than one high clock tick

6. Summary

We have described the implementation of a geolocation system using highly constrained sensor nodes. The signal processing limitations from these physical constraints have been addressed by the system architecture design and a set of algorithms described in this paper. In particular,

1. The challenges of stealth and low communication bandwidth were met by designing a system that relies on local processing and uses *a priori* information

about the signals of interest, with minimal inputs to the remote base station.

2. The computational cost and memory foot-print issues were factored into the design of a multi-resolution detection and estimation algorithm. Accuracy lost during the initial phase of detection is compensated by a gain in computational savings. Once the SOI is detected, accuracy is recovered by processing a reduced data set at a higher resolution.
3. The inaccuracies introduced by analog processing on low-cost, low-quality radios are circumvented by implementing a software radio. FM demodulation and filtering are implemented in software on the DSP.
4. Global clock synchronization is achieved using passive GPS devices. The jitters in the GPS signals and local oscillator clocks are reduced using a LMS fit.

The combination of these algorithms and selected hardware yields a cost-effective system with an accuracy that approaches the ideal, unconstrained geolocation system.

7. References

- [1] Bahl, P.; Padmanabhan, V.N. "RADAR an in-building RF-based user location and tracking system", INFOCOM 2000. 19th Annual Joint Conference of the IEEE Comp. and Comm. Soc. Proc., Vol.2, Iss., 2000, pp 775-784
- [2] Capt. Lamm, L., et. al., "Networked unattended ground sensor fields: tradeoff study and configuration rules methodology", DSE-TR-02-10, http://www.dean.usma.edu/math/research/msce/10th_AUT_S/papers/Lamm_02.pdf
- [3] Kumar, S., Shepherd, D.; "SensIT: Sensor information technology for the war fighter"; Proc. Inter. Conf. on Information Fusion, Vol. 1, 2001
- [4] <http://www.npwrc.usgs.gov/resource/tools/telemetry/timediff.htm> "Time Difference of Arrival 4life Radiolocation System"
- [5] <http://www.fcc.gov/911/enhanced/>
- [6] Stotts, L.B.: "Unattended ground sensor related technologies; an army perspective", Proc. SPIE. Vol. 4040. (2000) 2-10
- [7] "Detection, Classification, and Tracking of Targets", Dan Li, Kerry D. Wong, Yu Hen Hiu, Akbar M. Sayeed IEEE Signal Processing Magazine, Mar 2002, pp 17-29
- [8] Boettcher, P., Shaw, G., "Energy-constrained collaborative processing for target detection, tracking, and geolocation", IPSN 2003, LNCS 2634, pp. 254-268, 2003
- [9] Epson crystal oscillator, SG-531C device specifications, [http://www.epsondevice.com/www/PDFS/epdoc.nsf/308956542197017449256a9c001d58b4/ee97b916bce00f0c49256a7e00399245/\\$FILE/SG-51531series.pdf](http://www.epsondevice.com/www/PDFS/epdoc.nsf/308956542197017449256a9c001d58b4/ee97b916bce00f0c49256a7e00399245/$FILE/SG-51531series.pdf)
- [10] Knapp, C., Carter, G.; "The generalized correlation method for estimation of time delay"; IEE. Trans. Acoustics Speech and Signal Processing Vol.24, Aug 1976 pp. 320- 327
- [11] Fetzer, C., Cristian, F., "Integrating external and internal clock synchronization", Real-Time Systems, vol. 12, num. 2, pp. 123-171, 1997
- [12] Chhokra, K., et. al., "Accuracy Enhancements for TDOA Estimation on Highly Resource Constrained Mobile Platforms", ISIS-Technical report, TR# :ISIS-03-402, <http://www.isis.vanderbilt.edu/>
- [13] Drane, C. et. al. "Positioning GSM telephones" IEE. Comm. Mag. Vol.36, Iss.4 Apr 1998 pp. 46-54, 59
- [14] Proakis J; Manolakis, D; "Digital Signal Processing, Principles Algorithms, and Applications", Prentice Hall, Third Ed., 1996.