

Position and Orientation for Distributed Sensors: The PODIS Network

August 24, 2004

Jason Wilden and Jim Agniel
Nova Engineering
9928 Windisch Road
West Chester, Ohio 45069

Randy Moses and Josh Ash
Department of Electrical Engineering, The Ohio State University
2015 Neil Avenue, Columbus, OH 43210

ABSTRACT

In order for the information provided by networks of unattended ground sensors (UGS) to be of use to the tactical planner, the location of each sensor must be obtained. Sensor localization is typically achieved by careful hand-emplacement, or facilitated by anchor nodes whose positions are precisely known. This paper describes a new sensor node implementation designed to meet the growing need for UGS networks that can self-localize without the use of anchor nodes in a GPS-denied environment. The Position and Orientation for Distributed Sensors (PODIS) application facilitates traditional threat identification and tracking by self-localizing a network of sensor nodes. The PODIS node hardware consists of a COTS acoustic signal processor, a four element microphone array and a low power, networked radio for remote sensing with multi-hop routing capability. PODIS utilizes a messaging protocol that facilitates system calibration. While the PODIS application is independent of the ranging medium, the prototype network is implemented using time of arrival estimates of acoustic ranging signals. In this paper we describe the applicability of electronically generated pseudonoise calibration signals. We describe a two-step localization procedure that partitions the network calibration into relative and absolute solutions. We describe the typical anchor nodes as a special case of *aimpoints*, prior location information about a subset of nodes with associated uncertainty. We first compute relative node locations and orientations using time and direction of arrival estimates, then use aimpoint information provided by the user at deployment to rotate and translate the network onto an absolute frame of reference. We provide a simple channel model to exercise the system in the simulation domain, with results for self-localizing networks of various geometries and levels of connectivity. Finally, simulations of typical network deployments predict the PODIS localization accuracy.

1. Introduction

The reliance on GPS, anchor nodes, and hand emplacement for locating nodes is a shortcoming of the typical unattended ground sensor (UGS) CONOPS. This knowledge is unavailable at times due to the deployment mode or hostile operating environment. PODIS is aimed at satisfying the military's immediate need for a self-localizing sensor network. This paper represents the first implementation based on the self-localization research reported in [2-5]. The work is significant in that it addresses the critical and immediate need of forces on the ground to detect and localize threats such as sniper fire, armored vehicles and enemy combatants. Self-calibrating sensor networks also show promise in many commercial applications such as ocean search and rescue, disaster recovery, distributed sensing of large public areas for homeland security, and smart buildings.

The PODIS system is comprised of a network of UGS nodes, each containing a digital signal processor (DSP), networked radio, four microphones, and a loudspeaker. After the UGS have been deployed, with little or no prior information about their position, the central information processor (CIP) that controls the network initiates a self-localization calibration. Each node on the network then broadcasts an acoustic pseudonoise (PN) sequence while the other nodes receive the signal and compute a time of arrival (TOA) and possibly a direction of arrival (DOA) estimate from the received signal. The estimates are then sent back to the CIP which performs the self-localization algorithm computations. The result of the calibration is a relative and absolute position estimate. The relative estimate refers to the position relative to other nodes in the network, while absolute refers to the position estimate relative to a known reference.

The PODIS program applies a new body of theoretical development to a sensor radio platform that is currently entering production. The purpose of the program is to demonstrate a self-calibrating acoustic sensor network. Each node on the network is equipped with a microphone array, a small, low-cost signal processor and a networked sensor radio. In section two we describe deployment scenarios of PODIS sensors, the networking protocol, and generalize the concept of aimpoints vs. anchor nodes. In section three we present the self-localization incorporating aimpoints. Section four describes acoustic ranging methods using PN sequences and the results from our outdoor propagation experiments.

2. PODIS Concept of Operations

2.1. Deployment Scenarios

For the tactical planner, PODIS can be an invaluable tool operating under a wide set of deployment scenarios. The PODIS network enables source localization and tracking applications by automatically self-localizing sensors within a sensing network. In this section we introduce several nominal scenarios in simulation describing the operation and accuracy of the PODIS acoustic localization application. In all cases, ranging measurements are taken from a Gaussian process with a 0.3 m standard deviation.

Roadside deployment: In this scenario, 24 sensors are deployed in a 60 m square area around a road intersection. At deployment time, the user emits a calibration signal at nine locations along the road. At several locations, the user logs a GPS measurement in order to calibrate the system in absolute lat-long coordinates. We assume the GPS readings are accurate to within ± 3 m. Figure 1 illustrates that PODIS calibrates the network to an average error of 31 cm. We also find that reducing the number of GPS measurements to three or two does not significantly affect the localization in this scenario.

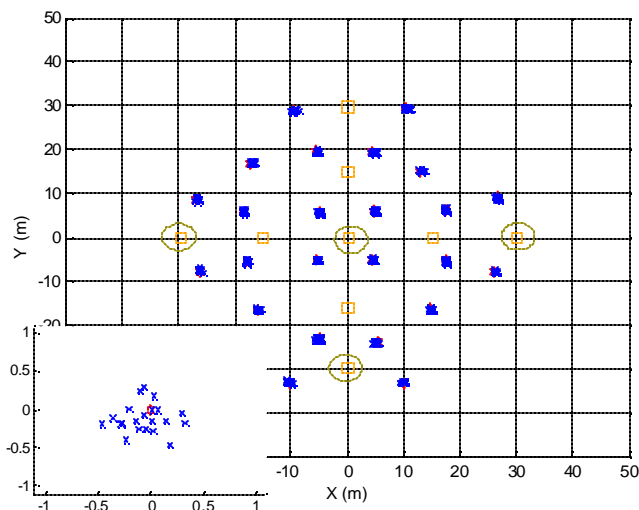


Figure 1: Roadside network in a 60 m x 60 m area with 24 nodes. Four nodes are equipped with GPS and used as aimpoints giving a localization accuracy of 31 cm.

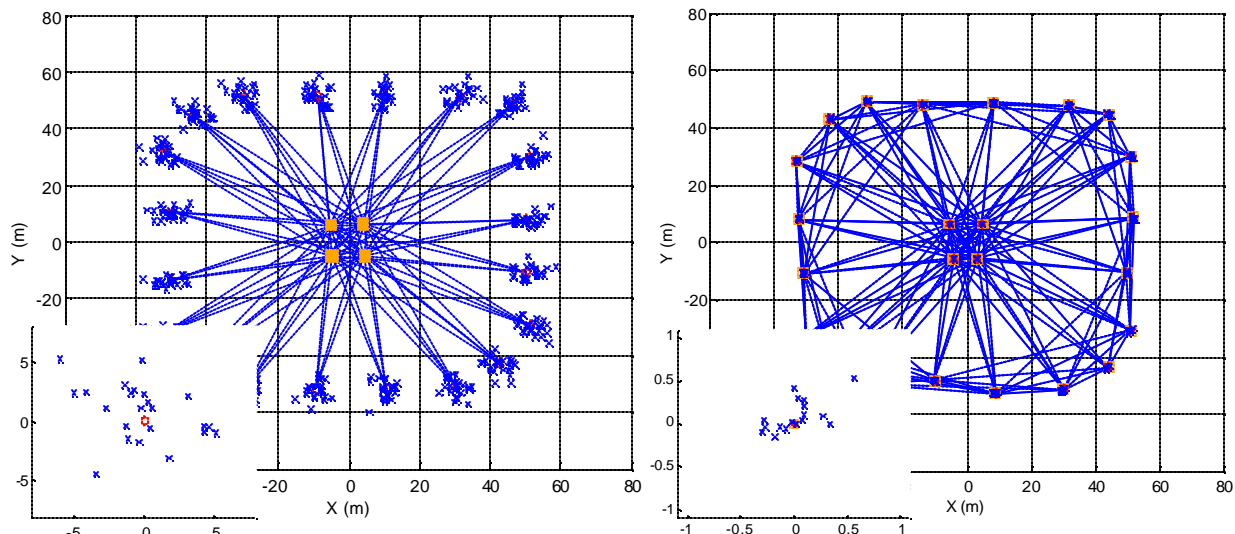


Figure 2: Perimeter network in a 100 m x 100 m area with 20 sensors and four internal anchor sources. **(Left)** With calibration sources solely on the network interior, dilution of precision limits the location accuracy to 3.3 m. **(Right)** When sensors are colocated with sources or ranging actuators, the localization accuracy improves to 36 cm.

Perimeter deployment: A second scenario shown in Figure 2 illustrates a perimeter sensing network with 20 sensors and four interior anchor nodes with known position. In the first example, all sensors calibrate based on measurements from sources on the interior of the perimeter. The narrow range of angles of arrival of the calibration signals at each node results in a high dilution of precision in the localization estimates, with an average error of 3.3 m. Supposing that each sensor in the network were equipped not only with sensing capability, but also with the ability to transmit its own calibration signal, we arrive at the network shown in Figure 3. Here we assume that the acoustic connectivity is limited to 70 m, so the network is not fully connected, but each sensor sees more calibration sources from a wider distribution of angles. Though the amount of prior information remains unchanged, the improved dilution of precision yields an average localization error of 36 cm.

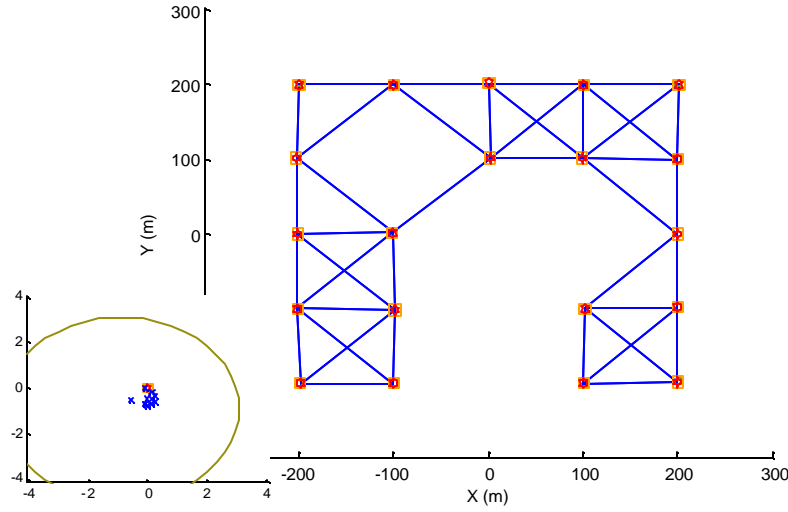


Figure 3: Rooftop network in a 400 m x 400 m area with 20 nodes sparsely deployed each with sensor and source. The network has three GPS nodes at three of the four corners that function as aimpoints giving a localization error of 86 cm.

Rooftop deployment: In this scenario we expand the scene and reduce the sensor connectivity to test the effects of a more sparsely connected network. We assume that sensor connectivity is limited to 160 m, and that three nodes with 4 m GPS capability occupy three of the four corners for an absolute calibration. In this case, the average absolute position error is 86 cm.

2.2. Sensor Networking Protocol

After the nodes have been deployed, the CIP uses networked radios for network discovery utilizing CSMA for channel access. The CIP creates a table of all the nodes in the network, and then asks each node to broadcast its calibration signal. Every node is equipped with a four element microphone array that is used to detect the presence of a calibration signal. A TOA estimate is then made preceding the DOA estimate. Other nodes listen and report to the CIP if the signal was received. If the signal was received then the TOA and DOA estimate is transmitted back to the CIP.

In the event that no node hears another node's PN calibration signal, the CIP will ask the node in question to retransmit the calibration signal three more times. If no other node receives the calibration signal after three transmissions the measurement will be dropped from the self-localization algorithm.

Erroneous measurements impact the self-localization solution of the entire network, not just the node in error. Identifying and discarding erroneous measurements improves the network self-localization solution. Erroneous measurements occur most often in low SNR environments, particularly as SNR approaches the threshold SNR given in [10].

Another method under investigation to add reliability to the TOA and DOA estimates takes advantage of the microphone array geometry. The distance from the center microphone to a microphone on the edge of the star configuration is known and imposes an upper bound on the difference between TOA estimates. If a differential TOA measurement exceeds this bound then an error flag is set for that microphone, with a maximum of three error flags per node possible. The number of error flags set for a particular node is returned to the CIP for analysis. As in the previous case, up to three transmissions occur in order to

eliminate any error flags present; if after three attempts error flags are present, a confidence metric may be assigned to the node indicating the reliability of the data to be used in the localization estimate.

2.3. Aimpoints vs. Anchor Nodes

Most sensor self-localization approaches assume that a subset of the nodes have known location (and possibly also known orientation). Such nodes are usually referred to as *anchor nodes* or *beacon nodes*. Anchor nodes are obtained by careful hand emplacement or by equipping a few special nodes with GPS systems. The anchor nodes are typically assumed to have *exactly* known locations. In practice, of course, it is rarely possible to provide exact locations; emplacement errors or GPS location errors are always present.

For the PODIS system, we take an alternate approach by assuming that a subset of nodes have uncertain prior location information associated with them. We refer to this prior information as node *aimpoints*. The aimpoints can arise from a prior measurement of location, as provided by a person who hand-emplaces the node or by a GPS location estimate. Alternately, the prior location information could be an aimpoint from an air drop deployment or from a munition-based emplacement system such as Volcano. In any case, the prior information is characterized by a nominal location of the sensor and by a location uncertainty. We model the prior uncertainty as a bivariate Gaussian quantity.

The aimpoint approach generalizes the idea of anchor nodes in a way that provides great flexibility in applications. If anchor nodes with known locations are available, they can be considered as special cases in the aimpoint approach as nodes with very low uncertainty in the aimpoint. For other cases in which there is uncertainty in the prior knowledge of certain node locations, the uncertainty is quantified and can be addressed in the localization procedure in a statistically rigorous manner. Finally, the aimpoint approach handles correlated errors; for example, aimpoint errors due to wind might be correlated across sensors, and this correlation can be readily encoded in the prior statistical information as described in Section 3 below.

3. Network Multilateration

The goal of network multilateration is to estimate the locations and orientations of sensor nodes, and to quantify the location and orientation uncertainty of the resulting estimates. Information available to us includes limited and uncertain prior information along with calibration measurements. The calibration measurements are estimates of the range and direction to neighboring sensors derived from TOA and DOA estimates of calibration signals.

Our solution to the multilateration problem involves defining a statistical model that incorporates both prior information and calibration measurements, then estimating sensor locations and orientations using a maximum *a posteriori* probability (MAP) estimator. A detailed derivation of the algorithm and its statistical properties is described in [3–4]; here we summarize the main points and specialize the more general algorithms there to the PODIS application. We encode aimpoint information as a prior probability density function (pdf) of a subset of the node locations. The calibration measurements provide additional information on the node locations and orientations, which is modeled as a posterior pdf. Our estimate of the node locations and orientations are at the maximum on this posterior pdf.

UNCLASSIFIED

A mathematical problem statement and solution is formulated as follows. Assume we have a set of A sensors in a plane, each with unknown location $\{r_i = (x_i, y_i)\}_{i=1}^A$ and unknown orientation angle q_i with respect to a reference direction (*e.g.*, North). Define the parameter vector

$$\mathbf{a} = [x_1, y_1, \mathbf{q}_1, \dots, x_A, y_A, \mathbf{q}_A]^T \quad (3A \times 1) \quad (1)$$

Prior Information: A subset of sensors and sources are assumed to have prior information about their locations or orientations, *e.g.*, from GPS-equipped sensors or (uncertain) sensor emplacement aimpoints. We quantify this uncertain prior information using a probability density function $f_0(\mathbf{a})$. We assume $f_0(\mathbf{a})$ is Gaussian distributed with known mean and covariance, so

$$f_0(\mathbf{a}) = N(\mathbf{a}_0, \Sigma_0) \quad (2)$$

where \mathbf{a}_0 and Σ_0 are given; however, other prior pdfs can also be used. In this case, \mathbf{a}_0 encodes the aimpoints and Σ_0 encodes the aimpoint uncertainty.

If there is no prior information on a parameter, the corresponding row and column of Σ_0^{-1} is zero; if a parameter is known with certainty, it is encoded with zero (or very small) variance in Σ_0 .

Calibration Measurements: Each sensor emits an acoustic calibration signal at time t_j . We assume that sensor i , if it detects source j , obtains an estimate of the (TOA) t_{ij} measured with respect to a time base established across the sensor network using the RF communication link, along with an estimate its DOA with respect to the sensor's local reference frame. Let H be the set of (i, j) pairs for which calibration measurements are available, and let N be the size of H . Let X be the $2N \times 1$ vector of TOA measurements T and DOA measurements Θ :

$$X = \begin{bmatrix} \text{vec}(T) \\ \text{vec}(\Theta) \end{bmatrix}^T \quad (2N \times 1) \quad (3)$$

The actual TOA and DOA of source signal j at sensor i can be computed from \mathbf{a} as

$$t_{ij}(\mathbf{a}) = t_j + \|r_i - \tilde{r}_j\|/c \quad (4)$$

$$\mathbf{f}_{ij}(\mathbf{a}) = \mathbf{q}_i + \angle(r_i, \tilde{r}_j) \quad (5)$$

respectively, $\|\cdot\|$ is the Euclidean norm, $\angle(\mathbf{x}, \mathbf{h})$ is the angle between the points $\mathbf{x}, \mathbf{h} \in R^2$, and c is the (known) signal propagation velocity.

Each element of X has measurement uncertainty modeled as

$$X = \mathbf{m}(\mathbf{a}) + E \quad (6)$$

where $\mathbf{m}(\mathbf{a})$ is the vector of actual TOA and DOA values given by equations (7) and (8), and where E is a zero mean random vector. If we adopt a Gaussian measurement error model then the pdf of X for a particular value of \mathbf{a} is

$$f_X(x|\mathbf{a}) = N(\mathbf{m}(\mathbf{a}), \Sigma_X) \quad (7)$$

where Σ_x is a known covariance matrix.

The self-calibration problem, then, is: Given the measurement vector X , measurement error pdf $f_x(x|\mathbf{a})$, and prior information pdf $f_0(\mathbf{a})$, estimate \mathbf{a} .

3.1. Maximum A Posteriori Calibration Algorithms

We consider \mathbf{a} as a random vector with prior probability density function $f_0(\mathbf{a})$. The measurement vector X informs us about \mathbf{a} as quantified by the posterior probability $f(\mathbf{a}|X)$. From Bayes' rule,

$$f(\mathbf{a}|x) = \frac{f(x|\mathbf{a})f_0(\mathbf{a})}{f(x)} \quad (8)$$

We choose as our estimate the maximum *a posteriori* estimate, which is the value of \mathbf{a} that maximizes the posterior probability density of \mathbf{a} :

$$\hat{\mathbf{a}} = \underset{\mathbf{a}}{\operatorname{argmax}} f(\mathbf{a}|X) = \underset{\mathbf{a}}{\operatorname{argmax}} f(X|\mathbf{a})f_0(\mathbf{a}) \quad (9)$$

For the case of Gaussian pdfs in equations (5)–(10), equation (12) becomes

$$\hat{\mathbf{a}} = \underset{\mathbf{a}}{\operatorname{argmin}} [X - \mathbf{m}(\mathbf{a})]^T \Sigma_x^{-1} [X - \mathbf{m}(\mathbf{a})] + [\mathbf{a} - \mathbf{a}_0]^T \Sigma_0^{-1} [\mathbf{a} - \mathbf{a}_0] \quad (10)$$

Equation (10) gives the MAP estimate of \mathbf{a} when both measurements and prior information are available. Its solution involves solving a nonlinear least squares (NLLS) problem, and a number of standard software and algorithms are available to compute the solution.

The solution to (10) requires obtaining an initial estimate of \mathbf{a} . Our approach is as follows: first, we compute an initial estimate of \mathbf{a} with respect to any one sensor. Assume Sensor 1 is located at the origin and oriented at $\mathbf{q}_1 = 0^\circ$. From the DOA and TOA of any neighbor node i whose signal was detected by Sensor 1, we can compute an initial location of Sensor i . The orientation of Sensor i can be computed from the DOA estimate \mathbf{q}_{i1} if it is available. We proceed similarly using neighbors of these initialized sensors, averaging location or orientation initial estimates in cases where nodes have two or more already-initialized neighbors. We continue until all sensor locations and orientations are initialized. This gives an initial localization with respect to Sensor 1. Finally, we rotate and translate all network estimates to best align with the prior aimpoint information, as illustrated in Figure 4.

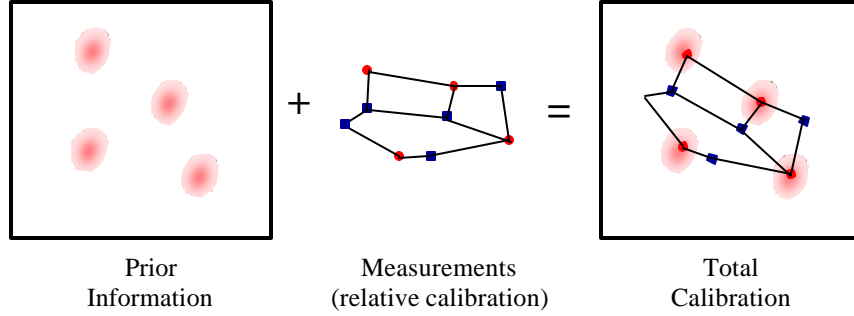


Figure 4: Relative calibration combined with prior aimpoint information provides an absolute calibration of the sensor nodes.

3.2. Location and Orientation Accuracy

For the case of Gaussian prior information, the covariance of the MAP estimate $\hat{\mathbf{a}}$ is given by (see [3] for a derivation)

$$E\{[\hat{\mathbf{a}} - \mathbf{a}][\hat{\mathbf{a}} - \mathbf{a}]^T\} \geq [\Sigma_0^{-1} + I_a]^{-1} \quad (11)$$

where I_a is the Fisher information matrix of \mathbf{a} due to calibration measurements, and is given by

$$I_a = E\left\{[\nabla_{\mathbf{a}} \ln f_X(X; \mathbf{a})][\nabla_{\mathbf{a}} \ln f_X(X; \mathbf{a})]^T\right\}. \quad (12)$$

An interpretation of the above result is as follows. Since the prior information and the measurement information are independent, the information matrix for \mathbf{a} is the sum of the individual information matrices. The Fisher information matrix for the Gaussian prior with covariance Σ_0 is Σ_0^{-1} . The Fisher information from the measurements is I_a . We note that equation (14) involves an expectation over both the random variable \mathbf{a} and the random measurement errors.

In many applications the uncertainty in relative calibration obtained from the calibration measurements is much lower than the absolute location uncertainty provided by the aimpoints. In this case, as is shown in [3], the MAP algorithm is approximated by one in which relative calibration is computed from X measurements only, followed by translation and rotation of this solution to a location most consistent with the prior information (see Figure 5). This results in a computational simplification in solving (13). In this case, the prior information plays a role only in the rotation-translation step. If some prior information uncertainty is on the order of calibration uncertainty, then the prior information plays a role in both the relative and absolute calibration components.

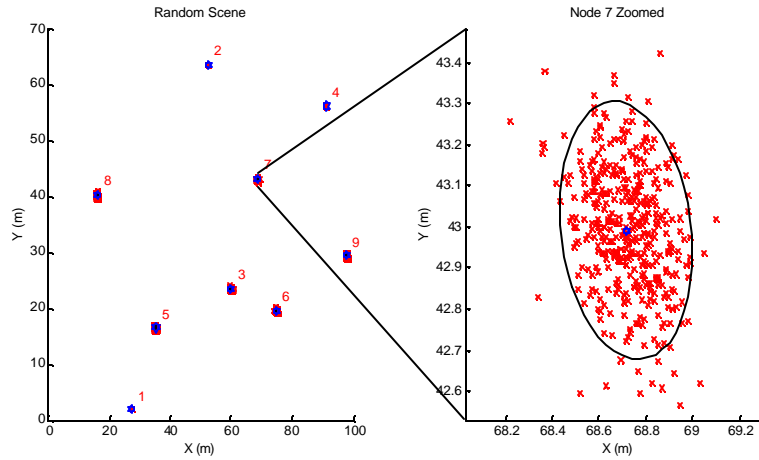


Figure 5: The PODIS position estimate includes a 2-sigma uncertainty ellipse, providing an upper bound on 86% of the localization estimates. This bound is a function of network size, connectivity and measurement variance. The PODIS calibration is independent of the measurement medium used for ranging. The points in the right plot are several position estimates over time.

3.3. Network Density and Geometry Considerations

In this section we summarize some of the important design considerations in network density and network geometry to ensure accurate self-localization.

Of primary importance is a network sensor density high enough so that a sufficient number of sensors can detect transmitted signals and accurately estimate the TOA and DOA of these signals. Each sensor must receive at least two signals from neighbors in order for calibration to succeed. Moreover, recent studies on network density effects [6–7] have shown that localization accuracy improves substantially as the network density increases to approximately six neighbors per node; after that point, localization accuracy increases modestly. The maximum distance between neighbors depends on transmit power, environment noise, and propagation considerations, some of which are discussed below and in [1].

In addition, results in [6–7] suggest a number of other geometric considerations. First, distance and angle measurements exhibit similar trends in localization error with varying network density. The error when angle-only measurements are used is much higher than when distance measurements are used, for uncertainties typical of acoustic array processing. Network localization appears to be scalable since error propagation is very slow and since localization accuracy appears to be relatively insensitive to increases in the number of anchor nodes or aimpoints. Finally, when the aimpoint uncertainty is high, most location error is an overall translation and rotation error of the entire network, meaning that nodes may be accurately located relative to each other if not to an absolute frame of reference.

4. Acoustic Propagation Results

4.1. A Simple Channel Model

We developed a first order channel model to predict the effects of atmospheric loss do to the propagation of noise. In order to simulate and evaluate candidate signal processing algorithms under a variety of environmental conditions, an acoustic propagation model was applied. We selected the model described below based on a survey of current acoustic propagation literature and our own outdoor measurements. The three main components of the model are spherical radiation, atmospheric absorption, and ground interactions.

Spherical radiation accounts for the simplest component of the model and refers to the $1/r$ loss in the sound pressure of a point source due to the radial dilution of sound intensity. The actual loss in sound pressure level (SPL), however, is greater than that predicted by spherical radiation alone.

Atmospheric absorption represents the dissipative effects of the atmosphere on the SPL and consists of two major sources: thermal conduction and viscosity of air, and relaxation losses of oxygen and nitrogen molecules in air. Unlike spherical radiation, atmospheric absorption is frequency dependent. For temperature T in Kelvins and pressure P in Pascals, the absorption coefficient in dB per meter is given by

$$a = 8.686 \cdot f^2 \sqrt{\frac{T}{293.15}} \cdot \left(1.84 \times 10^{-11} \frac{101325}{P} + \left(\frac{T}{293.15} \right)^{-3} [b_1 + b_2] \right), \quad (13)$$

where b_1 and b_2 are given by

$$b_1 = \frac{0.1068 \cdot \exp\left(\frac{-3352}{T}\right)}{f_{r,N} + \frac{f^2}{f_{r,N}}}$$

$$b_2 = \frac{0.01275 \cdot \exp\left(\frac{-2239.1}{T}\right)}{f_{r,O} + \frac{f^2}{f_{r,O}}}. \quad (14)$$

$f_{r,N}$ and $f_{r,O}$ are the relaxation frequencies of nitrogen and oxygen and can be computed from expressions in [11] with knowledge of the pressure, temperature, and molar concentration of water vapor – which can be derived from the relative humidity. The effects of atmospheric absorption have been codified into an international standard, ANSI Standard S1-26:1995.

Finally, because the sensor nodes are expected to be deployed near the ground, we must take into account ground interactions. These can be significant when either the source or receiver is near the ground surface. Ground interactions account for both absorption by the ground and reflections from the ground that can be received either constructively or destructively at the receiver. Ground interactions are a function of the locations of the source and receiver relative to the ground and each other, the speed of sound, the frequency of the sound, and the type of ground: dirt, grass, asphalt, etc.

Mathematically, ground interactions can be represented by the complex pressure amplitude at the receiver [8],

UNCLASSIFIED

$$p_c = \frac{S \exp(jkr_1)}{r_1} + Q \frac{S \exp(jkr_2)}{r_2}. \quad (15)$$

The first term in (18) represents the contribution of the direct ray from the source of distance r_1 away. The second term represents the component that has been reflected off of the ground, with total reflected path length r_2 . Here Q is the spherical-wave reflection coefficient, given by

$$Q = 1 - 2 \frac{k_1}{Z} \cdot \frac{r_2}{\exp(jk_1 r_2)} \int_0^{q_{\max}} \exp\left(\frac{-qk_1}{Z}\right) \frac{\exp(jk_1 \sqrt{r^2 + (z + z_s + jq)^2})}{\sqrt{r^2 + (z + z_s + jq)^2}} dq, \quad (16)$$

where

$$r_2 = \sqrt{r^2 + (z + z_s)^2} \quad (17)$$

is the length of the reflected signal path, $k_1 = \frac{2\pi f}{c_s}$ is the wave number of air c_s is the speed of sound, z and z_s are the height of the receiver and source, respectively, and r is the distance between source and receiver. Different ground surfaces are characterized by their flow resistivity \mathbf{s} , which is used in the calculation of the complex ground impedance

$$Z = 1 + 0.0511 \left(\frac{\mathbf{s}}{f}\right)^{0.75} + i0.0768 \left(\frac{\mathbf{s}}{f}\right)^{0.73}. \quad (18)$$

The upper limit on the integration in (19) is given by

$$q_{\max} = \frac{I(Z_r^2 + Z_i^2)}{Z_r}, \quad (19)$$

where Z_r and Z_i are the real and imaginary parts of Z , and I is the wavelength of the sound.

Based on our survey of outdoor acoustic propagation literature and comparison with our field measurements (see next section), we believe the three-component propagation model presented above will suit our modeling needs for Phase I. The limitation of this model is that it does not account for inhomogeneities in the atmosphere, such as temperature and wind gradients. These gradients cause a gradient in the speed of sound, which in turn results in refraction of propagating acoustic pressure waves due to Snell's law. Hence, this scenario is typically referred to as atmospheric refraction.

Although we have elected not to adopt the full sophistication of an atmospheric refraction model, we will be able to model the effects of a constant wind via an effective sound speed approximation. In [8], Salomons derives the following approximation to the speed of sound:

$$c_{eff} = c + u, \quad (20)$$

where c_{eff} is the effective speed of sound, c is the adiabatic speed of sound, and u is the component of the wind velocity in the direction of sound propagation. The approximation is valid in situations where the

sound wave travels with relatively low elevation angles. Because all PODIS nodes are expected to be placed at nearly the same height, this approximation is valid in the PODIS environment. In the PODIS system simulation, we model a uniform wind component according to [9].

4.2. Acoustic Ranging Experiments

The PODIS acoustic ranging function uses PN sequences to enable precise estimates of TOA. Their unique correlation properties are well suited for TOA estimates in the noisy environments where PODIS is intended to operate. We investigated several methods of performing the TOA correlation during PODIS development, and we performed outdoor experiments to test our hypotheses and the possibility of achieving good acoustic TOA estimates for PODIS localization.

The team developed a suite of outdoor propagation experiments to build upon data previously published in [1], validate the simple channel model presented above, assess the utility of low-cost audio components for generating and processing calibration signals, and characterize various signal sources for TOA estimation. An array of eight microphones was set up outside the Nova offices. Eleven different data sets were collected in different physical configurations with the same 230 PN sequences for each set. The PN sequences ranged in duration from 0.1 to ten seconds, the center frequencies varied from 100 Hz to 8 kHz, and the bandwidths varied from 25 Hz to 9600 Hz where applicable. Different configurations resulted from moving the loudspeaker in 12 m increments away from and around the array in both linear and circular configurations.

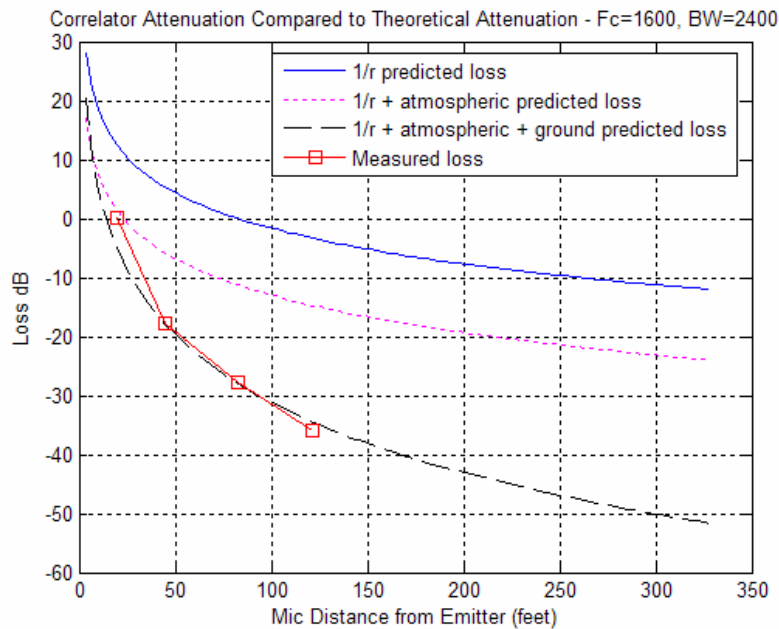


Figure 6: The experimental and model-predicted propagation losses for this calibration sequence show good agreement.

Figure 6 shows the experimental data compared to our channel model presented in the preceding section. The $1/r$ loss does not accurately predict the loss experienced in the channel nor does the $1/r$ loss with the atmospheric loss as seen in Figure 6. When all three components are included in the model it accurately predicts the loss experienced in the channel.

UNCLASSIFIED

One result we found was that cross-correlating the received PN sequence against a locally stored copy resulted in a poor quality output with high peak ambiguity, presumably because of channel effects not accounted for. We then stored as the local correlation base a reference signal that was recorded in an open-air, but benign, acoustic environment. The performance in this case was much improved, but the need for a ranging method that is robust to channel effects motivated the investigation of what we call "differential correlation," a technique inspired by an element of some synchronization schemes in wireless communications. With this technique the calibration sequence is transmitted twice, and the received signal is simply correlated against a delayed version of itself. Assuming stationarity over the correlation window, any channel effects would affect both halves equally. This method is effectively the zero-lag autocorrelation function as the incoming signal slides through the window; hence, in a noiseless channel, the correlator output is a triangle function. Results using all three correlation methods are shown in Figure 5. In each case, the transmitted signal is the same, and the correlation is performed in MATLAB using a different method.

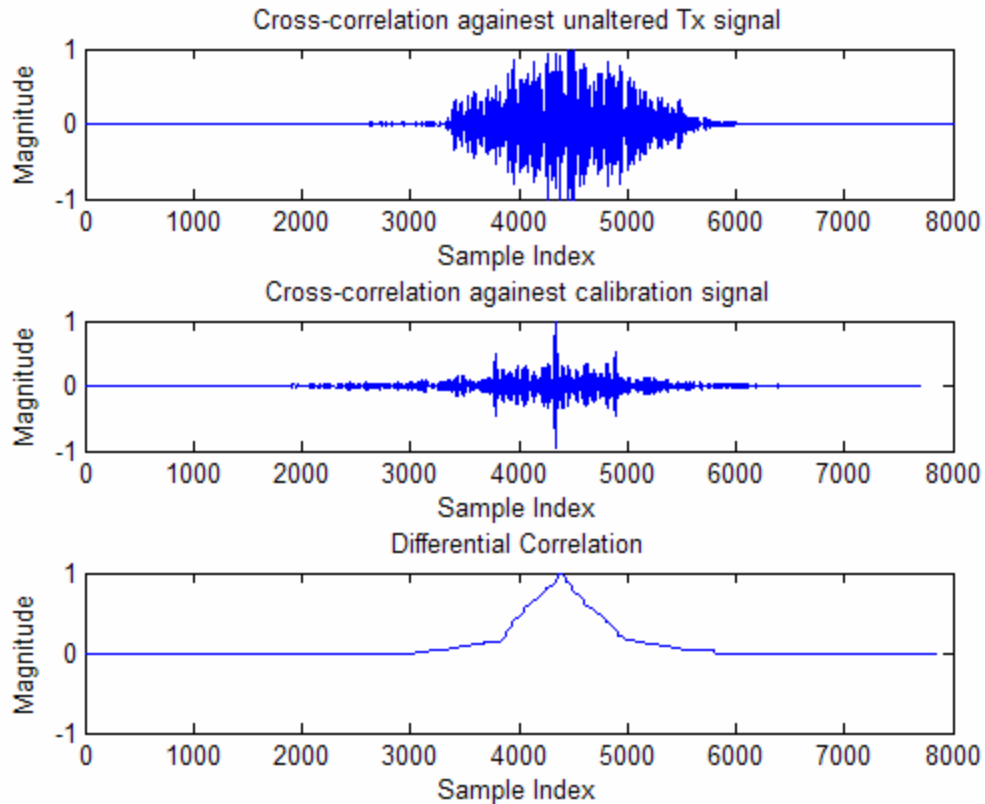


Figure 7: A comparison of the three different correlation methods with $F_c=800$ Hz. $BW=600$ Hz. and $T=0.2$ s. The cross-correlation of a received signal (**top**) with an ideal reference signal shows peak ambiguity, while correlating with a prerecorded reference signal (**middle**) provided improved results. The differential correlation (**bottom**) produces a triangular output, but in low SNR conditions becomes peak-ambiguous, too.

In field experiments we generated and detected signals using these three correlation methods. Clearly the top case in Figure 5 fails to provide a reliable TOA estimate, even in high SNR. In the second case, the

correlation shows three peaks because the received differential signal was cross-correlated against a reference differential signal. For the case illustrated, correlating against a prerecorded version of the calibration signal performed as well as the differential correlation. However, in general the performance of differential correlation was impaired due to its high off-peak correlation response.

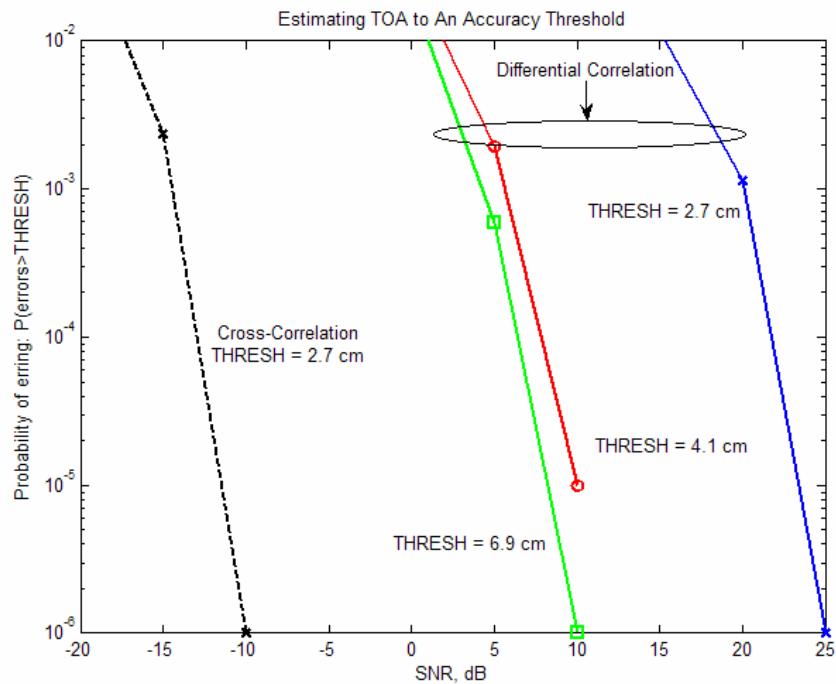


Figure 8: Monte Carlo simulation showing the performance of cross-correlation and differential correlation with a calibration signal center frequency of 1200Hz, 1200Hz bandwidth, and duration of 0.1s.

The differential correlation can accommodate networks with nodes comprised of different hardware making the algorithm easily transplantable. The differential method has a performance penalty associated with it when compared to cross-correlation. The variance of the estimate using the differential correlation peak is greater than the variance of the estimate using either cross-correlation method because the output of the differential correlation is proportional to the lag position. Figure 8 compares the performance of the differential correlation to the cross-correlation in terms of the probability of incorrectly picking the TOA at a given SNR.

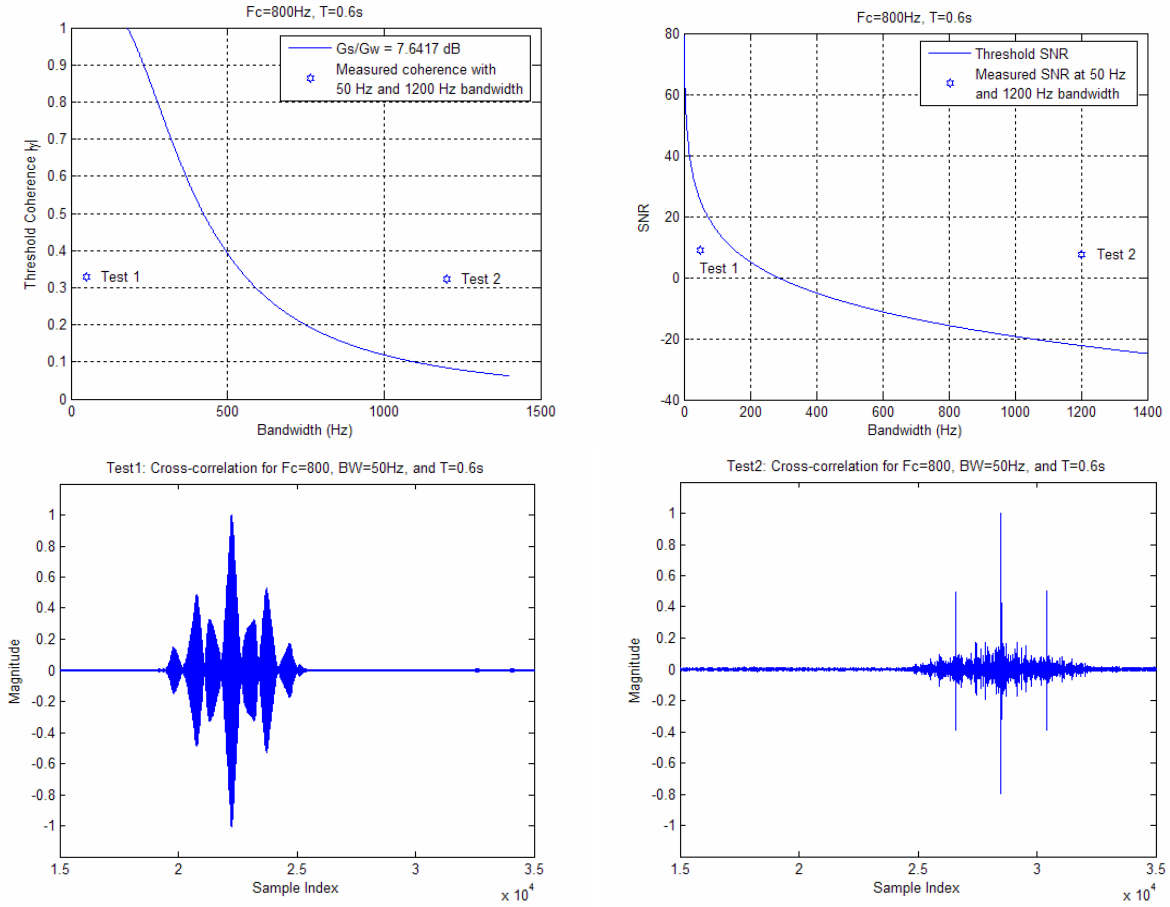


Figure 9: (Top) The plot on the left shows the threshold coherence as determined by equation (38) in [16]. The plot on the right shows the threshold SNR. For each metric we present two cases from the measured data, one that does not satisfy the threshold, and one that does. **(Bottom)** The correlation peaks corresponding to the two test cases.

Two conditions for time delay estimation (TDE) are the coherence threshold [16] and SNR threshold [15]. The lower right plot in Figure 9 shows the output of a cross correlation with both the coherence threshold and SNR threshold conditions are satisfied. If the threshold conditions are not satisfied a peak ambiguity problem occurs [1,15]. The lower left plot in Figure 9 shows the output of a cross correlation when both the coherence threshold and SNR threshold are not met and the peak ambiguity problem presents itself.

5. Conclusions

We have given a brief system overview of the PODIS network and the enabling algorithms. The aimpoint approach was presented as a way that generalizes the idea of anchor nodes that provides great flexibility in applications. We have also presented the results of our outdoor propagation experiments that built upon data previously published in [1], validated our simple channel model, and assed the utility of low-cost audio components for generating and processing calibration signals. We found that the best correlation method for the PODIS network is the cross-correlation method with a reference signal that was recorded

in an open-air, but benign, acoustic environment. Our data shows that the peak ambiguity problem does not present itself when the coherence threshold and SNR threshold are exceeded.

6. References

- [1] J.N. Ash and R.L. Moses, "Acoustic sensor network self-localization: experimental results," Military Sensing Symposia, 2003 MSS Specialty Group on Battlefield Acoustic and Seismic Sensing.
- [2] R. Moses, D. Krishnamurthy, and R. Patterson, "An Auto-Calibration Method for Unattended Ground Sensors," *Proceedings of the 2001 Symposium on Battlefield Acoustics and Seismic Sensing*, Laurel, MD, October 23-26, 2001.
- [3] R. Moses and R. Patterson, "Self calibration of sensor networks," Unattended Ground Sensor Technologies and Applications IV (Proceedings of SPIE Vol. 4743), Orlando, FL, 2-5 April 2002.
- [4] R. L. Moses, R. Patterson*, and W. Garber, "Self Localization of Acoustic Sensor Networks," *Proceedings of the 2002 Battlefield Acoustics Symposium*, Laurel, MD, September 24-26, 2002.
- [5] R. Moses, D. Krishnamurthy, and R. Patterson, "A Self-Localization Method for Wireless Sensor Networks," *Eurasip Journal on Applied Signal Processing, Special Issue on Sensor Networks*, March 15, 2003, pp. 148-158.
- [6] 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," in *Information Processing in Sensor Networks (IPSN '03)*, (Palo Alto, CA), April 2003.
- [7] A. Savvides, W. Garber, R. Moses, and M. B. Srivastava, "An Analysis of Error Inducing Parameters in Multihop Sensor Node Localization," to appear in *IEEE Transactions on Mobile Computing*.
- [8] E. M. Salomons, Computational atmospheric acoustics, Kluwer Academic Publishers, 2001.
- [9] M. West, K. Gilbert, and R. A. Stack, "A tutorial on the parabolic equation (PE) model used for long range sound propagation in the atmosphere," *Appl. Acoust.* vol. 37, pp. 31-49, 1992.
- [10] A.J. Weiss, and E. Weinstein, "Fundamental Limitations in Passive Time Delay Estimation, Part I: Narrow-band Systems," *IEEE Transactions on Acoustics, Speech, and Signal-Processing*, Vol. 31, No. 2, pp. 472-485, April 1983.
- [11] R.J. Kozick and B.M. Sadler, "Source Localization with Distributed Sensor Arrays and Partial Spatial Coherence," *IEEE Transactions on Signal-Processing*, Vol. 52, No. 3, pp. 601-616, March 2004.