

Localization with Sparse Acoustic Sensor Network Using UAVs as Information-Seeking Data Mules

DANIEL J. KLEIN, SRIRAM VENKATESWARAN, and JASON T. ISAACS, University of California, Santa Barbara

JERRY BURMAN, Teledyne Scientific Company

TIEN PHAM, US Army Research Lab

JOÃO HESPANHA and UPAMANYU MADHOW, University of California, Santa Barbara

30

We propose and demonstrate a novel architecture for on-the-fly inference while collecting data from sparse sensor networks. In particular, we consider source localization using acoustic sensors dispersed over a large area, with the individual sensors located too far apart for direct connectivity. An Unmanned Aerial Vehicle (UAV) is employed for collecting sensor data, with the UAV route adaptively adjusted based on data from sensors already visited, in order to minimize the time to localize events of interest. The UAV therefore acts as an information-seeking data mule, not only providing connectivity, but also making Bayesian inferences from the data gathered in order to guide its future actions. The system we demonstrate has a modular architecture, comprising efficient algorithms for acoustic signal processing, routing the UAV to the sensors, and source localization. We report on extensive field tests which not only demonstrate the effectiveness of our general approach, but also yield specific practical insights into GPS time synchronization and localization accuracy, acoustic signal and channel characteristics, and the effects of environmental phenomena.

Categories and Subject Descriptors: C.2 [Computer-Communication Networks]: Network Architecture and Design; C.3 [Special-Purpose and Application-Based Systems]: Signal Processing Systems; C.4 [Performance of Systems]: Measurement Techniques, Performance Attributes; G.3 [Probability and Statistics]: Correlation and Regression Analysis

General Terms: Algorithms, Performance

Additional Key Words and Phrases: Routing, acoustic source localization, large scale, heterogeneous, time of arrival, angle of arrival, UAV routing

ACM Reference Format:

Klein, D. J., Venkateswaran, S., Isaacs, J. T., Burman, J., Pham, T., Hespahanha, J., and Madhow, U. 2013. Localization with sparse acoustic sensor networks using UAVs as information-seeking data mules. *ACM Trans. Sensor Netw.* 9, 3, Article 30 (May 2013), 29 pages.
DOI: <http://dx.doi.org/10.1145/2480730.2480733>

This work adds significantly to our previous results [Burman et al. 2009, 2010; Klein et al. 2010]. The work was supported by the Institute for Collaborative Biotechnologies through grant W911NF-09-0001 from the U.S. Army Research Office. The content of the information does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred.

Authors' addresses: D. J. Klein (corresponding author), S. Venkateswaran, and J. T. Isaacs, ECE Department, University of California, Santa Barbara, CA 93106-9560; email: djklein@gmail.com; J. Burman, Teledyne Scientific Company, Thousand Oaks, CA 91360; T. Pham, U.S. Army Research Laboratory, Adelphi, MD 20783; J. Hespahanha and U. Madhow, ECE Department, University of California, Santa Barbara, CA 93106-9560.

©2013 Association for Computing Machinery. ACM acknowledges that this contribution was authored or co-authored by a contractor or affiliate of the U.S. Government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only. Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2013 ACM 1550-4859/2013/05-ART30 \$15.00

DOI: <http://dx.doi.org/10.1145/2480730.2480733>

1. INTRODUCTION

We propose and demonstrate an architecture for data collection and on-the-fly inference in sparse sensor networks where the sensor nodes do not have direct connectivity with each other. A data collector traverses the network, adapting its route as it collects data from each sensor node in order to draw reliable inferences as quickly as possible about the events of interest that the sensors are reporting on. We illustrate our approach for the problem of acoustic source localization. We are interested in homeland security and defense applications in which we wish to monitor large areas (of the orders of 100 square kilometers), with the purpose of detecting and localizing acoustic events, such as explosions or artillery, as quickly as possible. Such events can be detected at large ranges, so that a density of $\sim 5\text{--}6$ sensors/ 1 km^2 suffices for event detection and localization, provided that we can pool data from multiple sensors. For typical sensor radio range, however, sensors in such a sparse deployment cannot communicate directly with each other. The data mule in our experimental setting is a UAV, which reoptimizes its route as it receives data from each sensor node that it visits, with the goal of localizing the source as quickly as possible.

We demonstrated our system at the Center for Interdisciplinary Remotely-Piloted Aircraft Studies (CIRPAS) facility at Camp Roberts using a Fury UAV from Aeromech Engineering, Inc. and a heterogeneous collection of sensor nodes consisting of several single-microphone sensors, providing Time of Arrival (ToA) measurements, as well as a microphone array from the U.S. Army Research Lab, providing Angle of Arrival (AoA) measurements. The UAV route is recalculated after each sensor measurement is processed within a Bayesian framework, and a downward-looking camera on the UAV is used to image the source (a Zon Mark IV propane cannon) once successful localization is deemed to have taken place.

We develop a modular system architecture, shown in Figure 1, which enables a smooth transition from simulation to deployment: algorithms that do not work effectively can be replaced quickly, and a flight simulator can be used in place of the real UAV during development and testing. Each sensor is fitted with one or more sensitive microphones and processes the live audio stream to achieve two goals: (1) detect an event in the presence of ambient noise and (2) produce a Time-of-Arrival (ToA) or an Angle-of-Arrival (AoA) measurement corresponding to each event. Local signal processing captures “useful” data succinctly: large raw audio files that would require enormous bandwidths to transmit are reduced to ToA/AoA information that is easily transmitted over relatively slow links. Each sensor forwards its ToA or AoA information to the UAV when it comes within radio range. Instead of employing a fixed route for data collection, routing is coupled with localization to make smart decisions about where the UAV is to go next. Specifically, freshly acquired information is combined with prior data in Bayesian fashion in order to decide which sensor to visit next to locate the source as quickly as possible, or if the uncertainty regarding the source location is small enough, to go to the estimated source location. A significant reduction in time to localize comes in scenarios in which data from all sensors is not required to localize the source up to the desired accuracy. From our field deployment, we found that this situation was the norm rather than the exception.

We now summarize the key challenges addressed in this article and the resulting contributions.

Challenges. Three key issues must be addressed before we can effectively build a system based on the preceding architecture.

- (1) *Data-Driven UAV Routing.* The value of the information held by a sensor depends on its location relative to the source, but the location of the source is unknown.

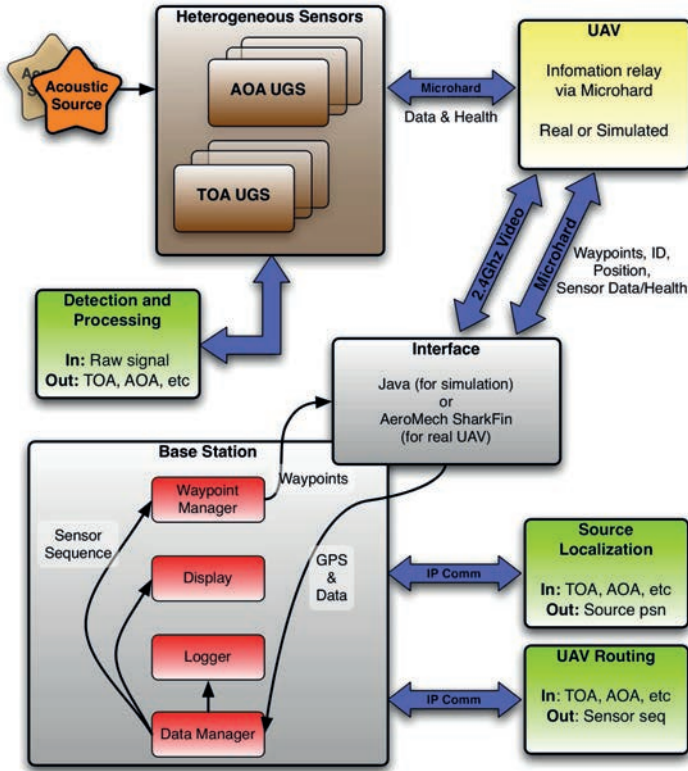


Fig. 1. The modular software architecture used for simulation and field testing.

Furthermore, the value of the information held by a sensor must be weighed against the time required for the UAV to reach the sensor. A fundamental challenge, therefore, is to devise an online data-driven routing algorithm that accounts for these trade-offs.

- (2) *Long-Range Characteristics of Acoustic Sources.* At short range (e.g., <400 m), the acoustic disturbance from field artillery is impulsive, but does this continue to hold at longer ranges? Is the acoustic channel consisting of the atmosphere, the earth, foliage, and low hills coherent over large distances? Is it reasonable to ignore atmospheric disturbances such as wind and the variations in the speed of sound (primarily with temperature, but also with pressure and relative humidity)?
- (3) *Sensor Localization and Synchronization.* Since the sensors do not communicate with each other, we cannot synchronize them with standard synchronization and self-localization techniques, and hence rely on GPS. What are the hardware and software choices for sufficiently accurate localization and synchronization at acceptable cost and complexity?

Our development and field deployment were slowed by uncertainties stemming from the preceding issues. We hope that the answers provided in this article will aid others in developing related systems.

Contributions. Our overall contribution is the introduction and demonstration of a novel architecture for rapid localization using information-seeking data mules in a

sparsely deployed sensor network. Core contributions that are critical to successful realization of the proposed architecture are as follows.

- (1) A Bayesian algorithm couples source localization with UAV routing to choose a future route that minimizes the expected source localization time. The algorithm is heterogeneous in that it explicitly incorporates both time-of-arrival and angle-of-arrival measurements. A novel “minimal sensor subsets” approach results in a dramatic reduction in computation overhead compared to exhaustively searching for the best route.
- (2) We demonstrate that the acoustic channel maintains coherence over large distances using a matched-filter-style signal-processing algorithm. Additionally, we show that ToA estimates from such an algorithm can be fused to localize a source efficiently and robustly.
- (3) Insights are provided into acoustic signal and channel characteristics for large-scale deployments, and the effect of environmental variations (e.g., temperature) on localization accuracy.
- (4) Demonstration is given of the integrated system using the CIRPAS facility within Camp Roberts, a Fury UAV from Aeromech Engineering, Inc., radios from Microhard Systems, Inc., a Zon propane cannon from Sutton Agricultural Enterprises, Inc., and a heterogeneous collection of nodes consisting of one AoA sensor from the Army Research Lab, and six custom ToA sensors.

Related Work. Over the past decade, sensor networks have been deployed to address a wide variety of problems such as habitat monitoring [Mainwaring et al. 2002], source localization [Ali et al. 2007; Wang et al. 2005], sniper detection [Simon et al. 2004], classification and tracking [Arora et al. 2004], indoor location services [Priyantha et al. 2000], structural monitoring [Xu et al. 2004], and volcanic studies [Werner-Allen et al. 2006]. An acoustic sensor network was used in Ali et al. [2007] and Wang et al. [2005] to localize marmots and woodpeckers from their calls. This deployment, which used AoA sensors (in contrast to our hybrid AoA-ToA system) was on a much smaller scale than ours, with maximal sensor separations on the order of 10m and 150m, respectively. Furthermore, unlike our sparse deployment, the arrays were close enough to exchange recorded waveforms and estimate an Approximate Maximum Likelihood (AML) source direction in multipath environments. The large-scale deployment in Arora et al. [2004] also investigates a heterogeneous sensor network to detect, classify, localize, and track targets. However, it employs magnetic and radar sensors (in contrast to acoustic sensors), and the deployment is dense enough for the sensors to form a connected network. The deployment closest in spirit to ours is the sniper detection network in Simon et al. [2004], where a network of ToA sensors detects and localizes a sniper based on the muzzle blast and shock wave. However, the nodes in Simon et al. [2004] are deployed densely on a smaller spatial scale, so that they can form a connected network to exchange data, and they use ToA sensors alone, unlike our hybrid AoA-ToA approach. A hybrid AoA-ToA approach has been considered in Volgyesi et al. [2007], but again on a small scale.

A data mule architecture to transport data in a disconnected sensor network was proposed, and scaling laws for this setting were investigated in Shah et al. [2003] and in several other papers on delay-tolerant networking, including Henkel and Brown [2005] and Henkel et al. [2006]. However, to the best of our knowledge, this is the first article in which the mules process the data they collect in order to guide their decisions, with a view to optimizing a specific application.

A number of algorithms have been proposed for ToA-based source localization [Abel and Smith 1987; Beck et al. 2008; Chan and Ho 1994] as well as hybrid ToA-AoA-based

localization [Bishop et al. 2008; Venkatraman and Caffery 2004]. However, most of these algorithms assume that the ToAs/AoAs are available at a single location and we believe that this is the first attempt to couple a UAV routing protocol that acquires ToAs/AoAs sequentially with the problem of source localization.

The UAV routing problem itself is a combinatorial optimization in that it seeks the “best” sensor sequence from the set of all possible sequences. Classic examples of such combinatorial optimization problems include the traveling salesman problem [Vazirani 2001] and the vehicle routing problem [Dantzig and Ramser 1959]. However, our problem differs from static routing problems in that the value of the information held at each *unvisited* sensor depends on the sensors already visited. This leads to a dynamic vehicle routing problem which has been considered in a different context in Bertsimas and Van Ryzin [1993] and Pavone et al. [2009].

Our routing algorithm is aimed at minimizing the expected volume of the Cramer-Rao ellipse. The Cramer-Rao matrix and associated error ellipse has been used extensively as an optimization criteria in the literature. Notable examples include work in the area of sensor selection and optimal observer steering. Information-driven algorithms have been studied in the context of sensor networks [Zhao et al. 2002; Chu et al. 2002] and active sensing [Ryan and Hedrick 2010; Hoffmann and Tomlin 2010], but the objective differed significantly and a data mule was not present. In Oshman and Davidson [1999], Doğançay [2007], and Frew et al. [2005] the trajectories of mobile sensors are optimized using information-based criteria. In these works the mobile agents were the sensors, thus the combinatorial issues of picking sensors were not addressed.

This article expands significantly on our previous conference publications in this general area. A broad overview of the project focusing on bio-inspired event classification and discovery algorithms can be found in Burman et al. [2009]. This work also provides a broad overview of an early version of the UAV routing algorithm, but does not provide technical details. A version of the UAV routing algorithm that did not include the minimal sensor subsets approach and was therefore much less efficient can be found in Klein et al. [2010]. Both of these publications preceded the field demonstration, and therefore lacked all the results and discussion related to hardware, field test, and the insights that are of primary concern here. The conference paper [Burman et al. 2010] provided an overview of the field tests, but did not include any of the technical details presented here and was focused on multisource classification and helicopter-based detection, which are topics not included in the present article.

Organization. In Section 2, we describe the algorithms for processing the acoustic information at each sensor node and for routing the UAV. In Section 3, we present each component in detail, focusing on the hardware. Results from field tests and a Monte Carlo simulation study are presented in Section 4. Section 5 provides a detailed analysis of the acoustic signal and channel characteristics and is followed by concluding remarks in Section 6.

2. ALGORITHMS

We now describe the algorithms needed to accomplish acoustic source localization using a sparse sensor network serviced by an unmanned aerial vehicle. We begin by explaining the acoustic signal processing performed to detect the event and estimate the ToAs/AoAs. We then provide the key ideas behind the UAV routing algorithm that uses Bayesian inference to choose the sequence of sensors to visit based on the data gathered from sensors already visited. The routing algorithm is computationally expensive, so we introduce a “minimal sensor subsets” approach to efficiently compute the optimal UAV route.

2.1. Detection and ToA Estimation

ToA Model. For the purposes of localization and routing, the output of the k^{th} ToA sensor is modeled as

$$z_k = T_0 + \frac{d_k(\mathbf{p})}{v} + w_z, \quad (1)$$

where T_0 is the unknown event time, $d_k(\mathbf{p})$ is the separation distance between the source located at unknown position \mathbf{p} and the k^{th} sensor, v is the speed of sound, and w_z is noise drawn from a zero-mean Gaussian distribution with standard deviation σ_z .

The first ToA measurement of an event is not immediately useful in inferring anything about \mathbf{p} due to the uncertainty in the event time T_0 . With two ToA measurements, the uncertainty in the event time can be eliminated by considering the Time-Difference-of-Arrival (TDoA) between two sensors, for example, with respect to sensor N_{ToA} ,

$$y_k := z_k - z_{N_{ToA}} = \frac{d_k - d_{N_{ToA}}}{v} + w_y, \quad k = 1, \dots, N_{ToA} - 1. \quad (2)$$

The noise w_y is distributed as a zero-mean Gaussian with standard deviation $\sigma_y = \sqrt{2}\sigma_z$.

ToA Signal Processing. The ToA estimation algorithm is a modified version of the traditional matched filter [Poor 1994] that adapts to an unknown background noise level. While the algorithm is straightforward, its value lies in the experimental conclusion that acoustic channels maintain “sufficient” coherence for the purpose of detection, even over large distances. Consider a window of length W seconds of the recorded signal $s(t)$ that begins, without loss of generality, at $t = 0$. We sample this continuous time signal at $f_s = 44100$ Hz to obtain a discrete-time sequence $s[k] = s(\frac{k}{f_s})$, $k \in \{0, 1, \dots, f_s W - 1\}$. We have a prerecorded template¹ of the propane cannon shot, of duration B seconds, denoted by $b[k]$, $k \in \{0, 1, \dots, f_s B - 1\}$. We hypothesize that the recorded signal at any sensor resembles the template and use a matched-filtering-style algorithm for detection.

We begin by forming the crosscorrelation between the recorded signal and the template $R_{sb}[k]$ as

$$R_{sb}[k] = \sum_{n=k}^{k+f_s B-1} s[n]b[n-k], \quad k \in \{0, 1, \dots, f_s W - 1\}, \quad (3)$$

with zero padding where necessary. We report a detection if the crosscorrelation is much more than the “typical variation” from the “average.” The median of the crosscorrelation samples, denoted by μ_R , is a metric that captures the average without being influenced by transient, loud signals.

$$\mu_R = \text{median}(R_{sb}[k]), \quad k \in \{0, 1, \dots, f_s W - 1\} \quad (4)$$

The only assumption we make here is that the duration of the transient acoustic event is much less than the size of the window, which is easily satisfied since the typical choices of window length W and the template duration B are 30 s and 0.04 s respectively. For a similar reason, we quantify the “typical” variation around the mean by the median absolute deviation (rather than the more traditional standard deviation) of the crosscorrelation samples, denoted by σ_R .

$$\sigma_R = \text{median}(|R_{sb}[k] - \mu_R|), \quad k \in \{0, 1, \dots, f_s W - 1\} \quad (5)$$

¹A single template was obtained by recording a cannon shot in an LoS environment before the flight tests and used at all the sensors.

Finally, the decision statistic for detection is the dimensionless quantity Z given by

$$Z = \max_k \left| \frac{R_{sb}[k] - \mu_R}{\sigma_R} \right|. \quad (6)$$

We report an event if Z exceeds a threshold κ , which is typically chosen to be 20. The sources of noise that determine the value of σ_R include the UAV flying overhead, the inverter used to power the laptop, chirping birds, passing cars, and wind, in addition to the sensor noise. In addition to noise, the success of the algorithm depends on the frequency selectivity of the channel between the acoustic source and the sensors. In Section 5, we describe experimental results which illustrate that the channel is indeed frequency selective, but is coherent enough for matched-filtering-style detection to work. The ToA of the event is the index \hat{k} where Z is achieved

$$\hat{k} = \arg \max_k |R_{sb}[k] - \mu_R|/\sigma_R, \quad (7)$$

normalized by the sampling frequency. For a window that begins at τ_0 , the ToA is given by $\tau_0 + \hat{k}/f_s$.

We note that the preceding algorithm requires the acoustic signature of the source in the form of a template. The performance of the algorithm depends on the normalized correlation ρ ($0 \leq \rho \leq 1$) between the source waveform and the template. We can calculate the degradation in performance with a “nonideal” template using standard techniques from detection theory. For a tolerable false alarm probability of ϵ_{fa} and an operating Signal-to-Noise ratio of SNR, the probability of miss is given by $p_{miss} = 1 - Q(Q^{-1}(\epsilon_{fa}) - \rho\sqrt{\text{SNR}})$, where $Q(\cdot)$ is the standard Q-function associated with Gaussian random variables and $Q^{-1}(\cdot)$ is its inverse. For example, consider a scenario where the probability of false alarm $\epsilon_{fa} = 5\%$, the SNR = 10 dB and the correlation $\rho = 0.8$. In this case, the probability of miss worsens from 6.45% to 18.8%, only a 12% increase for a fairly significant variation in the template. Finally, we would like to remark here that we consider only a single source and thus need only one template. However, we have related work [Burman et al. 2010] on acoustic source classification using a bio-inspired particle swarm optimization technique to differentiate among several sources.

2.2. AoA Estimation

AoA Models. The k^{th} of N_{AOA} AoA sensors produces an angular measurement modeled as

$$s_k = \theta_k + w_s, \quad (8)$$

where θ_k is the angle made by the line joining the center of the k^{th} array and the source to magnetic north and w_s is noise drawn from a zero-mean Gaussian² distribution with standard deviation σ_s . We now describe the array geometry and explain the processing techniques used to estimate the AoA of the source.

Array Geometry. The AoA sensor consists of four microphones located at the corners of a triangular pyramid. Although the array has a three-dimensional structure (see Figure 5(b)), we neglect the “vertical” dimension and approximate the array to be planar. This approximation is reasonable when the source is in the far field. With this approximation, the microphones may be assumed to be located at $\mathbf{r}_0 = (0, 0)$, $\mathbf{r}_1 = (1, 0)$, $\mathbf{r}_2 = (\cos(2\pi/3), \sin(2\pi/3))$, and $\mathbf{r}_3 = (\cos(4\pi/3), \sin(4\pi/3))$ as shown in Figure 2. While we have assumed here that the array center coincides with the origin

²A von Mises distribution is more appropriate here because s_k is an angle, but σ_s is sufficiently small so that the Gaussian distribution is a very good approximation.




Fig. 2. A top view of the Angle of Arrival (AoA) sensor. The explosion is assumed to occur in the far-field of the array.

for ease of exposition, the generalization is straightforward. Using a compass, the array is oriented so that the array arm M_0M_1 is aligned with magnetic north.

AoA Signal Processing. The AoA estimation technique is similar in spirit to the one used in Yli-Hietanen et al. [1996] and is described here for completeness. Events are detected using the matched filtering algorithm described in the previous section. We now describe the algorithm used to estimate the source direction once an event has been detected. Consider a source S located at $\mathbf{r}_s = (R_s \cos \theta, R_s \sin \theta)$ where R_s is the distance to the array center and θ is the angle $\angle SM_0M_1$ shown in Figure 2. Assuming a Line-of-Sight (LoS) channel between the source and the array, the difference in the propagation times from S to M_i and from S to M_0 is given by $\Delta\tau_{i0} \triangleq (||\mathbf{r}_s - \mathbf{r}_i|| - ||\mathbf{r}_s||)/v$. When the source is in the far field of the array ($R_s \gg 1$ m), we can show that $\Delta\tau_{i0} \approx \mathbf{e}_\theta^T \mathbf{r}_i/v$, where $\mathbf{e}_\theta = (\cos \theta, \sin \theta)$ is a unit vector in the direction of the source. We use this fact, in conjunction with an estimate of $\Delta\tau_{i0}$, denoted by $\widehat{\Delta\tau}_{i0}$, from the received signals to estimate the source direction.

Denoting the source waveform by $s(t)$ and assuming a LoS channel between the source and the array, the waveform recorded at microphone M_i is given by $s_i(t) = s(t - ||\mathbf{r}_s - \mathbf{r}_i||/v) (+ \text{noise})$, where v is the speed of sound. $\Delta\tau_{i0}$ is obtained by crosscorrelating $s_i(t)$ with the reference $s_0(t)$ in a fashion exactly analogous to the ToA estimation algorithm. We estimate the source to be in a direction $\hat{\theta}$ that best explains the propagation time differences between M_0 and each of the other three microphones in the least squares sense.

$$\hat{\theta} = \arg \min_{\phi} \sum_{i=1}^3 (\widehat{\Delta\tau}_{i0} - \mathbf{e}_\phi^T \mathbf{r}_i)^2 \quad (9)$$

The minimization was done by gridding the one-dimensional angular space and finding the optimum over the discrete set of points.

2.3. UAV Routing

When monitoring a large area in which not all sensors detect the event, the routing algorithm could be initialized by having the UAV follow a fixed route until it encounters a sensor with an event to report. For example, the UAV could be instructed to fly a minimum time circuit, computed using a traveling salesperson solver, to minimize the maximum time between detection of the event and the UAV visit. For large areas, the elapsed time before the source is localized could be dominated by the first detection time if a single UAV is employed. However, timely detection could be ensured by partitioning such areas into smaller subareas, each assigned a separate data mule which could follow the algorithm proposed here. Investigation of such scenarios, including algorithms for coordinating multiple UAVs, is beyond the scope of this article.

Once the initial detection occurs, the key question to answer is: *in which sequence should the remaining sensors be visited so as to localize the source as quickly as possible?* The routing optimization we propose uses the latest available information about the source in a Bayesian manner. The objective of the UAV routing is to minimize the time required to determine the location of the source (up to some confidence); additional time required to fly over and image the source is not considered. We say that the source has been localized when the area of the region containing the source with a specified confidence level falls below a threshold value. Due to the threshold in the area of the confidence region, we refer to this problem as Threshold Time Minimization (TTM). This problem is challenging due to the nonlinear nature of acoustic source localization. This nonlinearity means that some sensors may carry more information about the source than others, a discrepancy that must be balanced against the time required to service each sensor. Another ramification of this nonlinearity is that sensors close to the source do not generally provide more informative data than faraway sensors. A further complicating factor is that optimal route is data dependent and thus must be computed on-the-fly.

To explain the routing protocol design problem in more detail, we begin by describing how to compute the confidence region using the Cramer-Rao bound. We then use this confidence region to formally write the routing optimization problem. However, directly solving this optimization problem is intractable, so we then introduce minimal sensor subsets and evaluate the overall computational complexity.

Confidence Ellipse from Cramer-Rao Bound. To avoid coming up with a routing protocol that depends explicitly on the particular estimation algorithm used to compute the source location estimate $\hat{\mathbf{p}}$ from the available data, we make use of the Cramer-Rao matrix, denoted \mathcal{C} , and its determinant in particular. In short, the \mathcal{C} matrix is important because it is the smallest possible covariance any unbiased estimator could give using the available measurements. The Cramer-Rao matrix is defined as the inverse of the Fisher information matrix evaluated at the true value of the parameters [Van Trees 1968].

$$\mathbf{F}(\mathbf{p}, \mathcal{I}, \mathcal{J}) \triangleq E_{\mathbf{Y}, \mathbf{S}|\mathbf{p}} \left[\left(\frac{\partial \log p(\mathbf{y}_{[\mathcal{I}]}, \mathbf{s}_{[\mathcal{J}]|\mathbf{p}})}{\partial \mathbf{p}} \right) \left(\frac{\partial \log p(\mathbf{y}_{[\mathcal{I}]}, \mathbf{s}_{[\mathcal{J}]|\mathbf{p}})}{\partial \mathbf{p}} \right)^T \right] \quad (10)$$

$$\mathcal{C}(\mathbf{p}, \mathcal{I}, \mathcal{J}) \triangleq \mathbf{F}^{-1}(\mathbf{p}, \mathcal{I}, \mathcal{J}) \quad (11)$$

Here, $p(\mathbf{y}_{[\mathcal{I}]}, \mathbf{s}_{[\mathcal{J}]|\mathbf{p}})$ is the posterior probability of the available ToA ($\mathbf{y}_{[\mathcal{I}]}$) and AoA ($\mathbf{s}_{[\mathcal{J}]}$) measurements, respectively, given the source position, \mathbf{p} . The sets \mathcal{I} and \mathcal{J} contain indices of ToA and AoA measurements that have been collected by the UAV.

For problems in which the likelihood of the parameters given the data is a multivariate Gaussian, as is assumed and later verified here, the Fisher information matrix has a special form. Denoting by \mathbf{q}_k the position of the k^{th} sensor, the Fisher information matrix for TDoA localization can be written as [Chan and Ho 1994]

$$\mathbf{F}(\mathbf{p}, \mathcal{I}) = \frac{1}{v^2} \mathbf{G}_{[\mathcal{I}]}^T \mathbf{Q}^{-1} \mathbf{G}_{[\mathcal{I}]}, \quad (12)$$

where $\mathbf{Q} = \sigma_z^2(\mathbb{I} + \mathbf{1}\mathbf{1}^T)$ of appropriate dimension and

$$\mathbf{G} = \begin{bmatrix} \mathbf{g}_1^T - \mathbf{g}_{N_{ToA}}^T \\ \vdots \\ \mathbf{g}_{N_{ToA}-1}^T - \mathbf{g}_{N_{ToA}}^T \end{bmatrix} \quad \text{with} \quad \mathbf{g}_k = \frac{\mathbf{p} - \mathbf{q}_k}{d_k}, \quad (13)$$

and $d_k = \|\mathbf{p} - \mathbf{q}_k\|$ is the distance between the source and the k^{th} sensor. The notation $\mathbf{G}_{[\mathcal{I}]}$ selects rows of \mathbf{G} corresponding to available measurements. Each vector \mathbf{g}_k is a

unit vector pointing from the k^{th} sensor towards the source. It is important to note that the Fisher information matrix is not full rank until three noncollinear ToA sensors have been visited.

For the angle-of-arrival sensors, the Fisher information matrix can be written as [Dogancay and Hmam 2008]

$$F(\mathbf{p}, \mathcal{J}) = \frac{1}{v^2} L_{[\mathcal{J}]}^T R^{-1} L_{[\mathcal{J}]}, \quad (14)$$

where R is an appropriately sized identity matrix scaled by σ_s^2 and

$$L = \begin{bmatrix} \mathbf{h}_1^T / d_1 \\ \vdots \\ \mathbf{h}_{N_{AOA}}^T / d_{N_{AOA}} \end{bmatrix} \quad \text{with} \quad \mathbf{h}_k = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} \frac{\mathbf{p} - \mathbf{q}_k}{d_k}. \quad (15)$$

The combined Fisher information matrix from all visited AoA and ToA sensors is the sum of the individual 2×2 Fisher information matrices

$$F(\mathbf{p}, \mathcal{I}, \mathcal{J}) = F(\mathbf{p}, \mathcal{I}) + F(\mathbf{p}, \mathcal{J}), \quad (16)$$

from which the Cramer-Rao matrix, \mathcal{C} , can be computed using Eq. (11).

The physical intuition behind the \mathcal{C} matrix can be understood through the notion of a *confidence ellipse* [Van Trees 1968]. A confidence ellipse is an ellipsoidal region of the parameter space containing the true parameter value with a specified certainty (or confidence), much like a confidence interval for a one-dimensional variable. The overall shape of the ellipse is determined by the correlation matrix, as given by the estimation error covariance. The utility of the Cramer-Rao matrix is that it produces a confidence ellipse that fits within the ellipse produced by any unbiased estimator. The volume of the confidence ellipse, denoted V , for a particular confidence level is proportional to the square root of the determinant of the correlation matrix.

Online minimization of the confidence ellipse volume is complicated by the fact that the \mathcal{C} matrix depends on the true source location \mathbf{p} , which is of course unknown. One can instead compute the *expected value* of the volume of the uncertainty ellipse with respect to the posterior probability of the parameters [Uciński 2004]

$$\bar{V}(\mathbf{p}, \mathcal{I}, \mathcal{J}) := E_{P|Y,S} [V(\mathbf{p}, \mathcal{I}, \mathcal{J})]. \quad (17)$$

Threshold Time Minimization UAV Routing Protocol. The Threshold Time Minimization (TTM) protocol aims to minimize the time at which the volume of the expected uncertainty ellipse (17) falls below a threshold value, V_{th} . This threshold value represents the largest uncertainty which is determined by the field-of-view of the camera onboard the UAV, and gives the area of the region in which the source is likely to be found with a specified probability.

Sensor data (TDoA and AoA values) available at the current time t are collected in vectors $\mathbf{y}_{[\mathcal{I}(t)]}$ and $\mathbf{s}_{[\mathcal{J}(t)]}$, and let $\mathcal{I}(\tau, \mathbf{r})$ and $\mathcal{J}(\tau, \mathbf{r})$ be vectors of indices of sensors whose data will be available at time $\tau \geq t$ along route \mathbf{r} . Here, a route is a sequence of sensors to be visited by each UAV. The resulting optimization problem to accomplish this objective can be written as

$$\begin{aligned} & \min_{\mathbf{r} \in \mathcal{R}(t)} \quad \tau \\ & \text{subject to} \quad \bar{V}(\mathbf{p}, \mathcal{I}(\tau, \mathbf{r}), \mathcal{J}(\tau, \mathbf{r})) - V_{th} \leq 0, \end{aligned} \quad (18)$$

where $\mathcal{R}(t)$ is the set of all routes to unvisited sensors, and the expected uncertainty volume is computed using measurements that will be available at time τ along route \mathbf{r} .

The numerical computation needed to fully carry out the routing optimization can be burdensome due to the fact that the general nature of the posterior probability makes closed-form computation of the expected value in (17) intractable. The expected value we are interested in evaluating is with respect to the posterior probability of the parameters given the data

$$\mathcal{E} = \int p(\mathbf{p}|\mathbf{y}, \mathbf{s}) f(\mathbf{p}) d\mathbf{p}, \quad (19)$$

for some function f . A standard approximation technique is to consider K samples drawn from the posterior $p(\mathbf{p}|\mathbf{y}, \mathbf{s})$, denoting the k^{th} sample by \mathbf{m}_k . Then, the expected value is well-approximated by

$$\mathcal{E} \approx \frac{1}{K} \sum_{k=1}^K f(\mathbf{m}_k). \quad (20)$$

This approximation approaches the true value as K becomes large by the weak law of large numbers.

Drawing samples from a posterior distribution is nontrivial, but a number of techniques like Markov Chain Monte Carlo (MCMC) are well-documented in the literature [Andrieu et al. 2003; Hastings 1970]. The approach we employ is a bio-inspired MCMC technique based on a model of the motion of *E. coli* bacteria known as Optimotaxis [Mesquita et al. 2008]. Viewing the posterior probability density as a food source, the bacteria wander using a stochastic tumble-and-run mechanism that allows their positions to be seen as samples drawn from the posterior. Increasing the number of bacteria results in a better approximation of the expected value at the cost of computational resources. It is important to note that the posterior probability density changes after each measurement is collected, and thus it is necessary to resample the posterior frequently.

2.4. Minimal Sensor Subsets for Efficient UAV Routing

Due to the combinatorial nature of the routing problem, the optimization over all possible routes in Eq. (18) grows as $N!$. Exhaustively searching each and every route quickly becomes intractable. Here, we introduce the concept of “minimal sensor subsets” to efficiently find the best route. The key insight here is that the expected value of the volume of the uncertainty ellipse in (17) depends only on which sensors have been visited, not on the particular order in which the sensors were visited. For example, the uncertainty volume after completing route [6, 1, 4] will be the same as that for route [4, 1, 6], although the time required to perform these two routes could differ significantly. This special structure allows the problem of selecting sets of sensors to visit to be decoupled from the problem of selecting the order in which to visit the sensors.

Accordingly, the minimal sensor subsets approach to UAV routing proceeds in two main steps.

Step 1. A small number of unordered sets of sensors (i.e., minimal sensor subsets), one of which necessarily contains the sensors visited by optimal route, are identified using a method to be described shortly.

Step 2. The optimal route is computed for each minimal sensor subset, using exact or approximate techniques. Among these optimal routes, the UAV selects the one that can be serviced in the least amount of time.

We proceed by formally introducing minimal sensor subsets, explaining the two steps of the minimal sensor subsets approach to UAV routing, and then analyzing the complexity of the problem.




Fig. 3. An example of the hierarchy of unvisited sensor subsets for the case of $n = 4$ unvisited sensors. The algorithm to find all minimal sensor subsets would begin by evaluating the uncertainty volume of the column containing $\{1, 2\}$, making inferences after each and every evaluation.

Minimal Sensor Subsets. A single minimal sensor subset is a collection of *unvisited* sensors which, when combined with information provided by sensors already visited, reduces the expected volume of the uncertainty ellipse below the user-specified threshold, V_{th} . Further, this set is minimal in the sense that the exclusion of any one unvisited sensor would increase the uncertainty volume above the threshold. More formally, we provide the following definition.

Definition 2.1 (Minimal Sensor Subset). Let $\mathcal{V} \subset \{1, 2, \dots, N\}$ be the set of sensors already visited, $\mathcal{U} = \{1, 2, \dots, N\} \setminus \mathcal{V}$ be the set of unvisited sensors, $n = |\mathcal{U}|$ be the number of unvisited sensors, and let u be a nonempty subset of the unvisited sensors, $u \subseteq \mathcal{U}$. Denote by u^{-k} the set created by removing the k^{th} element from set u . Then, the set u is said to be *minimal* if the following two conditions are met.

- (1) The volume of the uncertainty ellipse after visiting all of the sensors in the set $\mathcal{V} \cup u$ is below the volume threshold, V_{th} .
- (2) The volume of the uncertainty ellipses after visiting all but any one of the unvisited sensors, for example, $\mathcal{V} \cup u^{-k}$ for all $k = 1, \dots, |u|$, are all above the volume threshold, V_{th} .

Note that the optimal UAV route is necessarily a permutation of one of the minimal sensor subsets. Visiting a subset of a minimal sensor subset would not provide enough information whereas visiting a superset would be unnecessary. We now explain an algorithm to find all minimal sensor subsets (step 1), and then select the TTM-optimal route (step 2).

Algorithm Step 1. Sensor subsets have a natural hierarchy (see Figure 3), that allows us to make inferences that reduce the number of uncertainty volume computations that need to be made compared to an exhaustive search of the power set of unvisited sensors, \mathcal{U} . Specifically, every superset of a set of sensors having an expected uncertainty ellipse volume below the threshold will have an expected uncertainty ellipse volume below the threshold (adding sensors can only reduce uncertainty volume). Similarly, all subsets of a set of sensors having an expected uncertainty ellipse volume above the threshold will have an expected uncertainty ellipse volume above the threshold (removing sensors can only increase uncertainty volume).

The algorithm begins by evaluating the uncertainty volume of all sets of unvisited sensors of size $\lfloor n/2 \rfloor$. Sets for which the expected uncertainty volume, when combined with information from already visited sensors, is above (below) the threshold are denoted with a plus (minus). Inferences are made after each evaluation by marking all supersets (subsets) with a plus (minus), accordingly. The few sets remaining without

any plus or minus after evaluating all sets of size $\lfloor n/2 \rfloor$ are evaluated exhaustively, but again inferences are made after each uncertainty volume evaluation. When the algorithm terminates, we will know if the expected uncertainty volume is above or below the threshold for every sensor subset, and thus can easily identify which are minimal.

Algorithm Step 2. Once the minimal subsets have been identified, the second step is to compute the optimal (minimum time) route for each minimal subset (note that routes do not need to be computed for nonminimum sets of sensors). To do so, we use a standard Traveling Salesperson Problem (TSP) solver. For each minimal set, the solver takes as input the distances between the sensors in the minimal subset and returns the optimal order in which to visit the sensors and the route completion time. The fastest route through any one of the minimal subsets is the optimal route in the TTM optimization (18), with the possible exception that the TSP solver may produce a suboptimal route. In practice, we have found the approximations provided by TSP sufficient, and have not performed exhaustive search.

Complexity Analysis. The computational savings of using the minimal sensor subsets approach to route determination are dramatic. There are at most 2^n unique subsets of $n = |U|$ unvisited sensors, and most ${}_nC_{\lfloor n/2 \rfloor}$ of these subsets can be minimal³. The worst-case performance of the algorithm to find minimal sensor subsets, described in step 1 before, computes the uncertainty volume for approximately half of the sets (and makes inferences for the other half). In particular, we have the following result.

THEOREM 2.2. *The number of uncertainty value calculations used in finding all minimal sensor subsets will not be greater than*

$$\begin{cases} 2^{n-1} + \frac{\prod_{k=n/2+1}^n k}{2 \prod_{k=1}^{n/2} k} & \text{if } n \text{ even} \\ 2^{n-1} + {}_nC_{\lfloor n/2 \rfloor} & \text{if } n \text{ odd.} \end{cases} \quad (21)$$

PROOF. The main idea is that the worst case for the procedure described in step 2 of the previous algorithm occurs when there is just one minimal sensor subset. There are at most 2^n unique subsets of $n = |U|$ unvisited sensors. Let $f(n)$ represent number of uncertainty value calculations used in finding all minimal sensor subsets.

n Even. For n even, there are ${}_nC_{n/2}$ sets of size $n/2$. The algorithm begins by searching these ${}_nC_{n/2}$ sets of size $n/2$, from which inferences can be made about all sets containing fewer (more) than $n/2$ elements, provided the lone minimal sensor subset contains more (fewer) than $n/2$ elements. At the completion of this step $2^{n-1} - 1/2 ({}_nC_{n/2})$ sets remain to be searched. The upper bound comes from assuming that the algorithm will have to exhaustively search all remaining sets (i.e., that no further inferences are made), which is clearly an upper bound on the number of evaluations that would actually be performed.

$$f(n) \leq ({}_nC_{n/2}) + 2^{n-1} - \frac{1}{2} ({}_nC_{n/2}) \quad (22)$$

$$= 2^{n-1} + \frac{1}{2} ({}_nC_{n/2}) \quad (23)$$

$$= 2^{n-1} + \frac{\prod_{k=n/2+1}^n k}{2 \prod_{k=1}^{n/2} k} \quad (24)$$

³Here ${}_nC_k$ denotes the number of combinations of n elements, selected k at a time.



Fig. 4. The propane cannon used to create acoustic disturbances.

n Odd. For n odd, there are ${}_nC_{\lfloor n/2 \rfloor}$ sets of size $\lfloor n/2 \rfloor$ and ${}_nC_{\lceil n/2 \rceil}$ sets of size $\lceil n/2 \rceil$. The algorithm begins by searching the ${}_nC_{\lfloor n/2 \rfloor}$ sets of size $\lfloor n/2 \rfloor$, from which inferences can be made about all sets containing fewer (more) than $\lfloor n/2 \rfloor$ elements, provided the lone minimal sensor subset contains more (fewer) than $\lfloor n/2 \rfloor$ elements. At the completion of this step 2^{n-1} sets remain to be searched. The upper bound comes from assuming that the algorithm will have to exhaustively search all remaining sets (i.e., that no further inferences are made).

$$f(n) \leq \binom{n}{\lfloor n/2 \rfloor} + 2^{n-1} \quad \square \quad (25)$$

As compared to exhaustively evaluating the cost of $n!$ routes, we compute at most ${}_nC_{\lfloor n/2 \rfloor}$ TSP solutions, after finding minimal sensor subsets. Using Christofides' TSP approximation algorithm [Christofides 1976], the running time for each TSP is $O(\eta_k^3)$, where $\eta_k \leq n$ is the number of sensors in the k^{th} minimal subset. It should be noted that the minimal sets approach efficiently provides results corresponding to an infinite planning horizon, thereby avoiding the need to consider a limited planning horizon for computational tractability.

3. SYSTEM COMPONENTS FOR FIELD DEMONSTRATION

In this section, we describe the system components, focusing primarily on the hardware choices made and the rationale behind these choices. We begin with the acoustic source and then describe the components that go into each sensor: GPS, laptop, microphone, microphone array, and the radio used for communication. We finally describe the UAV used in the field deployment and the base station.

3.1. Acoustic Source

We used a Zon Mark IV propane cannon to create acoustic events with characteristics similar to live artillery. The sound level at the muzzle of the cannon is approximately 120 dB. A ToA sensor node (to be described shortly) was placed close to the cannon to record the true event location and time. Figure 4 shows the propane cannon mounted in the bed of a pickup truck.

3.2. Time Synchronization and Localization

Localizing an acoustic source based on ToAs requires the sensors to have a common notion of space and be accurately synchronized in time. Given the rapidly dropping cost of GPS receivers, both synchronization and sensor localization can be achieved economically by equipping each sensor with a GPS unit. However, the choice of the

GPS unit is critical, especially for accurate timing. Furthermore, the GPS receivers that provide tight timing synchronization are not “plug-and-play” units. Consequently, in addition to a good choice of the GPS unit, hardware modifications and software choices (e.g., FreeBSD operating system) are needed to achieve μs -level timing accuracy across nodes.

Time Synchronization. To achieve μs -level timing synchronization, it is critical to choose a GPS unit that has a Pulse Per Second (PPS) output. Consequently, we chose the Garmin 18-LVC units. The PPS output is a logical signal which indicates the start of a second very precisely. This is done by making the rising edge of the logical signal coincide with the second, whose value is reported in a separate string. When used in conjunction with the Network Time Protocol (NTP) to perform filtering and set the system clock, the PPS signal helps us achieve μs -level timing synchronization. Note that NTP for our purpose is simply an application running separately at each node, since nodes are unable to communicate directly. More details on achieving timing synchronization on the order of μs with GPS units can be found in CVE-2009-1368 [2009].

GPS-Based Localization. We convert the latitude-longitude output of the GPS unit to a local two-dimensional Cartesian coordinate system that is then used by the localization and routing algorithms. We use the geodetic-to-East North Up (ENU) conversion rules described in CVE-2008-1368 [2008] for this purpose. The location of each sensor is then filtered with a moving average filter to account for GPS measurement noise.

3.3. Time-of-Arrival Sensor

To expedite the field demonstration for proof of concept, we built the ToA sensor with off-the-shelf components. The resulting sensors are extremely heavyweight and not energy efficient, but the sensors could easily be downsized to a small form factor using standard mote hardware.

Hardware. We interfaced a condenser microphone to a laptop to acquire and process the raw audio signals. *Microphone:* We chose the multipattern Samson C03U microphone for its sensitivity and flat frequency response over a wide range of frequencies (50 Hz–5000 Hz). A foam windscreen prevented unwanted noise. *Laptop and Software:* We chose a Dell Latitude E5500 laptop running FreeBSD and custom Java applications to process and store the raw audio signals. *Power Source:* The laptop was powered for an entire day by a marine deep cycle car battery through an inverter. *Radio:* The radio was a n920 unit from Microhard Systems, Inc., described in Section 3.5. We had six such ToA sensors, with one of them positioned very close to the propane cannon to obtain ground truth. Therefore, we had $N_{TOA} = 5$ ToA sensors to make measurements. One of these sensors is shown in Figure 5(a).

3.4. Angle-of-Arrival Sensor

The AoA sensor is an Acoustic Transient Detector System (ATDS) provided by the Army Research Lab (ARL). It consists of four microphones located at the corners of a triangular pyramid as shown in Figure 5(b). A compass is used to orient one of the “arms” of the array along magnetic north. The Earth’s magnetic declination needs to be taken into account while converting the AoA back to the “global” East North Up coordinate system; on the days of our tests, the magnetic north was oriented $13^{\circ}32'$ east of geographic north at Camp Roberts.




Fig. 5. (a) The time-of-arrival sensor consisting of a Dell laptop, Samson microphone, Garmin GPS, Microhard radio, battery, and an inverter; (b) the angle-of-arrival sensor provided by the Army Research Lab.

3.5. Communication

Communication Topology and Protocols. Each sensor transmits data such as event ToAs/AoAs and sensor location as well as sensor status information like battery life and temperature over a radio link to the UAV. Due to UAV payload and control restrictions, we were unable to perform computations onboard the UAV. Instead, we performed the signal processing envisioned for the UAV using a base station on the ground. Thus, Microhard radios were used to establish a network for communication between the sensors deployed in the field, the airborne UAV, and a base station. The radio onboard the UAV acted as a relay, ferrying messages between the sensors and the base station. Let us now look at the typical flow of information in this network.

The sensors first transmit information to the UAV, which relays this information to the base station. The base station provides an acknowledgment upon receiving the data. The radio onboard the UAV forwards the acknowledgment to the sensor, concluding the information flow for the event under consideration. In case the sensor does not receive an acknowledgment, the data is retransmitted a maximum of five times.

Emulation of Short-Range Radio Links. Our system emulates short-range radio links typical of long-life sensor nodes, low-profile antennas, rough terrain, and low-flying UAVs. However, the emulation was simplified by taking advantage of high-power radios, high-altitude UAV, and relatively benign terrain of operation for our demonstration, which meant that most of the sensors had radio connectivity with the UAV. Thus, in our experiments the sensor can transmit ToAs/AoAs to the UAV right after detection, which simplifies the message exchange protocol. This connectivity to the base station has the additional benefit of enabling the debugging of the system as physical access to the sensors is a very laborious process. In order to emulate short radio ranges, we “unlock” sensor data only when the UAV comes within an “effective” communication footprint of the sensor. This footprint is roughly a circular region within 200 m of the sensor, but we account for NLoS effects using a terrain map to further reduce the footprints; see Figure 6(a). To facilitate simultaneous transmissions, the radios were set in a point-to-multipoint TDMA configuration, and built-in multiple access schemes resolved potential conflicts.

The Microhard Nano n920 radios that we employed operate in the 902–928 MHz band and employ frequency hopped spread spectrum for communication. While these “high power” radios have an advertised range of 100 km in LoS scenarios, in Nonline of Sight (NLoS) settings, which are typical of field deployments, we found that the communication radius of these radios is less than 500 m. And of course, as described before, we restrict the effective communication range to 200 m or less.

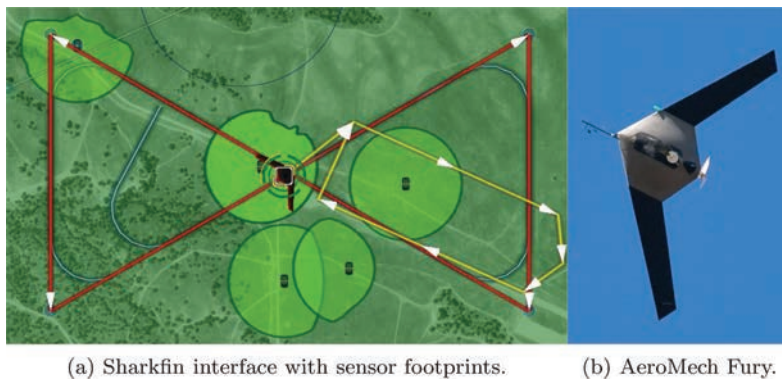


Fig. 6. This image shows a portion of the SharkFin interface (a) and the Fury UAV (b). The bright green blobs are communication regions for their respective sensors and the red hourglass is a figure-eight UAV flight pattern, used to image the source once localized.

3.6. Unmanned Aerial Vehicle

For the flight test, we used a Fury UAV (shown in Figure 6(b)) from AeroMech Engineering, Inc. This is a high-performance UAV capable of 18 hours of sustained flight and speeds up to 40 m/s. The payload on the UAV consists of a Microhard radio and a downward-looking (nadir) camera. The camera is used to image the acoustic source once it is located. While our eventual plan is to process the ToAs/AoAs onboard the UAV, we did not have control over the payload and could not locate the base station functionality on the UAV.

The interface to the Fury UAV is achieved by communicating with SharkFin, the UAV's ground control software and visualization package. For Monte Carlo simulation studies, we interfaced the base station to FlightGear, an open-source flight simulator, which provides realistic vehicle dynamics and also models the effects of wind. More details on the simulations are provided in Section 4.3.

3.7. Base Station

The base station performs the critical tasks of source localization and UAV routing from the ToA/AoA measurements that are unlocked by the UAV. In addition, it provides a number of services such as interfacing to the UAV (real or simulated), data logging, display, waypoint management, and debugging output such as sensor temperature, remaining battery life, GPS positioning, and NTP timing statuses.

AeroMech modified SharkFin to enable communication of UAV routes, fly-over commands, and status communications via a TCP connection. However, range safety and liability concerns prevented us from routing the UAV in a completely automated fashion. A human operator observed the routes (i.e., sequence of sensor waypoints to visit) output by our UAV routing protocol and manually specified a “safe” flight path for the UAV. The human operator essentially always adhered the output of our UAV routing protocol.

4. EXPERIMENTS AND RESULTS

We have conducted three types of demonstrations/experiments. The first demonstration consists of a pair of flight tests where we deploy 6 sensors over a 1 km² region and localize acoustic sources within this area. This emulates a portion of a large sensor network responsible for localizing an acoustic source and illustrates the efficacy of the overall system. In the second set of tests, we fired a propane cannon repeatedly

to characterize the statistical performance of the detection algorithms in a realistic environment. Finally, we conducted a series of Monte Carlo simulation trials to quantify the UAV routing performance.

The flight and statistical tests used the facilities of the Center for Interdisciplinary Remotely-Piloted Aircraft Studies (CIRPAS). McMillan Airfield, Camp Roberts, CA provides CIRPAS dedicated airspace for UAV testing that is remote from populated areas and free of interference from commercial or military air traffic. The McMillan Airfield is located near the southern boundary of Camp Roberts at an elevation of 280m and is surrounded by lightly wooded rolling hills and open grasslands. Figure 5(b) provides a rough idea of the terrain near the airfield.

4.1. Flight Tests

We conducted flight tests on November 3rd and 6th 2009 at Camp Roberts with five ToA sensors and one AoA sensor. The sensor locations on the two days are shown in Figure 7. The distance between the sensors and the propane cannon varied between 96 m and 880 m on November 3rd and 31 m and 603 m on November 6th.

Flight Test Protocol. In its “default” mode, the UAV was commanded to loiter about a sensor in the center of the surveillance region. We chose the “loiter sensor” to be the AoA sensor since it identifies (with high probability) a large portion of the field where the source cannot lie; on the other hand, a single ToA measurement with no prior information is of no immediate use in localizing the source because of the uncertainty about the event time. When an event is detected at the AoA sensor, the TTM algorithm is used to determine the sequence of sensors the UAV should visit. Event ToAs obtained by visiting other sensors are combined with any available measurements to arrive at a better estimate of the source location; if the volume of the uncertainty ellipse drops below the predefined threshold V_{th} , the UAV is instructed to abandon its route and fly over the estimated source location. Otherwise, a new route is computed using the TTM routing algorithm until sufficient confidence is obtained on the source location.

The time consumed in locating the source was dominated by the time taken by the human operator to specify the course (see Section 3.7) and therefore, does not convey the time savings that would be obtained by a fully automated implementation. We therefore report only on the accuracy of source localization (and the number of sensors visited before the event was localized) from the field experiments. The time savings from our data-driven routing algorithm are quantified using results from a flight simulator, described in Section 4.3.

Results. We conducted nine flight tests spread over two days. The results are shown in Table I. We find that the estimated source location is always within 20 m of the true source and on a few occasions, the error is only a couple of meters. Upon review, we discovered that the two instances where the source location error exceeded 15 m were due to large errors in the AoA estimates. Subsequent tests showed that such errors in the AoA estimates could be caused by the signal arriving along multiple directions (after reflections off near-by objects). These outlier estimates can be handled in a Bayesian framework by giving a low confidence to measurements from the AoA sensor. However, this was not part of the original measurement model and we could not repeat the hardware experiments with a more appropriate model. Instead, we omitted the erroneous data and found that this reduced the localization error to a few meters. On eight of the nine trials the source is “localized” without needing to visit all six sensors. Here, we deemed the source to be localized when the volume of the 5% uncertainty ellipse fell below a predetermined threshold of 420m² (the threshold was chosen so that a UAV could, with high probability, localize the source with a single fly-over). Note

(a) Configuration for Nov. 3rd, 2009.(b) Configuration for Nov. 6th, 2009.

Fig. 7. In both cases the position of the acoustic source is indicated by the megaphone symbol and the positions of the sensor nodes are indicated by the microphone symbols.

that localization time is our primary routing objective, and that choosing to skip some of the sensors is simply a byproduct.

4.2. Statistical Tests

We performed further field tests to address three issues: (a) quantify the performance of the detection algorithm statistically, (b) validate the ToA model (1), and (c) obtain

Table I. Localization Results

Date			Nov. 3, 2009		
Trial	Source Location Error (m)	# of Sensors Visited			
1	20.3	5			
2	4.3	6			
3	4.5	4			
Date			Nov. 6, 2009		
Trial	Source Location Error (m)	# of Sensors Visited			
1	5.6	4			
2	2.0	5			
3	1.8	5			
4	2.6	4			
5	19.6	5			
6	10.4	4			

statistics of localization accuracy. To this end, we conducted two sets of tests with ToA sensors on January 28, 2010; the picture in Figure 8 shows the sensor locations in the morning with S_0 providing the ground truth. In the afternoon, the propane cannon and S_0 exchanged their location with S_4 . The duration of each test was roughly three hours with the propane cannon being fired twice a minute. We will now describe the results of postprocessing the recorded data which provides answers to the questions raised.

Detection Performance. We characterize the performance of the detection algorithm by the traditional false alarm and missed detection metrics: a cannon shot that was not detected at a sensor is said to be a missed detection, while a positive detection made when there was no shot in reality is called a false alarm. Ground truth was obtained from a sensor placed right next to the propane cannon. To separate the performance of the detection algorithm from timing issues, we deemed a detection to be correct if it fell within ± 0.045 seconds (corresponding to approximately “ 3σ ”) of the expected ToA. If there were no detections in this window, then the cannon shot was deemed missed. Any detections outside this window were declared false alarms. The results are shown in Table II and a few points are worth noting: (1) During the morning tests, the cannon shot was successfully detected at all four sensors on 270 occasions out of 384 shots. In the afternoon, we detected 257 out of 288 shots at all four sensors. (2) In the morning, sensor 2 contributes to a large fraction of the misses, primarily because of significant foliage that lies between the sensor and the propane cannon (seen in Figure 8). Surprisingly, a small change in the propane cannon location in the afternoon improves the probability of detection at sensor 2 dramatically: 92% of the shots are detected at sensor 2, in spite of the foliage that is still present between the sensor and the cannon. (3) At the other sensors, the performance of the matched-filtering-style algorithm is very good with more than 95% of the shots being detected.

ToA and Localization Statistics. In the model, the expected ToA at sensor k for an event occurring at T_0 is given by $T_0 + d_k(\mathbf{p})/\nu$ where $d_k(\mathbf{p})$ is the distance between the source at \mathbf{p} and sensor k and ν is the speed of sound taken to be 340.29 m/s. Furthermore, the variation around the mean is assumed Gaussian. To verify this model, we placed a sensor close to the propane cannon to provide the event time T_0 and the distance was estimated from GPS data. We found a significant bias on the order of 20 ms in the measurements. Curiously, the bias also changed sign over the course of the day: the measured ToAs were larger than the expected ToAs in the morning and became smaller than the expected ToAs as the day went on. This led us to speculate that the change in speed of sound with temperature was the reason behind the bias. In Bohn



Fig. 8. Sensor configuration for January 28, 2010 morning tests. In the afternoon tests, the location of sensor 4 was switched with the position of the acoustic source.

Table II. Detection Performance

Sensor:	#1	#2	#3	#4	Sensor:	#1	#2	#3	#4
False	24	49	10	6	False	64	28	17	52
Missed	2	109	14	4	Missed	4	23	9	2
Correct	382	275	370	380	Correct	284	265	279	286

(a) Jan. 28, 2010 Morning

(b) Jan. 28, 2010 Afternoon

[1988], it is shown that speed of sound depends on the temperature $T(^{\circ}\text{C})$ as

$$v(T) = 331.45 \sqrt{1 + \frac{T}{273.15}} \text{ m/s.} \tag{26}$$

While we did not directly measure the temperature during the test, we obtained hourly temperature measurements from a near-by weather station [History for CA62 2010]. The speeds of sound corresponding to these temperature readings based on Eq. (26) can be found in Table III. We see that the speed of sound varied by as much as 6 m/s over the day, leading to ToA errors that grow with distance: at a 1 km range,

Table III. Speed of Sound Estimates

Date	Jan. 28, 2010	
Time	Temperature ($^{\circ}\text{C}$)	v (m/s)
10:20	7.004	335.6725
10:53	7.782	336.1384
11:53	11.729	338.4912
1:53	15.620	340.7950
3:53	16.731	341.4504
5:53	12.785	339.1181
7:53	10.617	337.8301

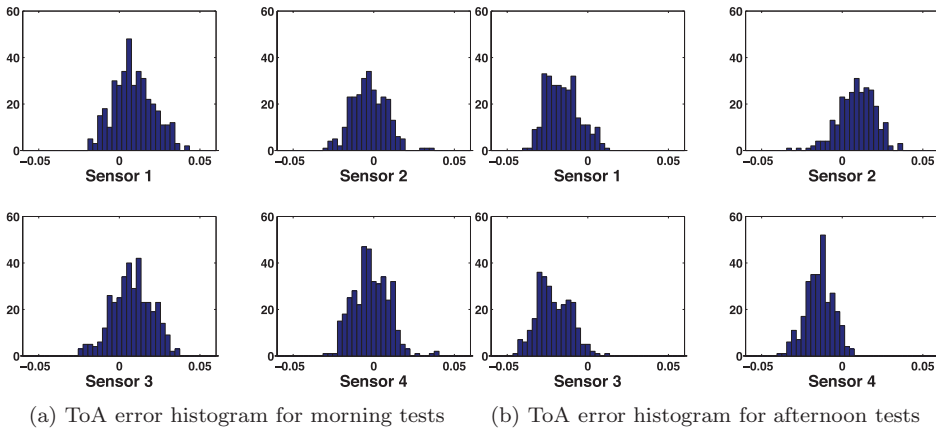


Fig. 9. Distribution of ToA error (in seconds) for each of four sensors.

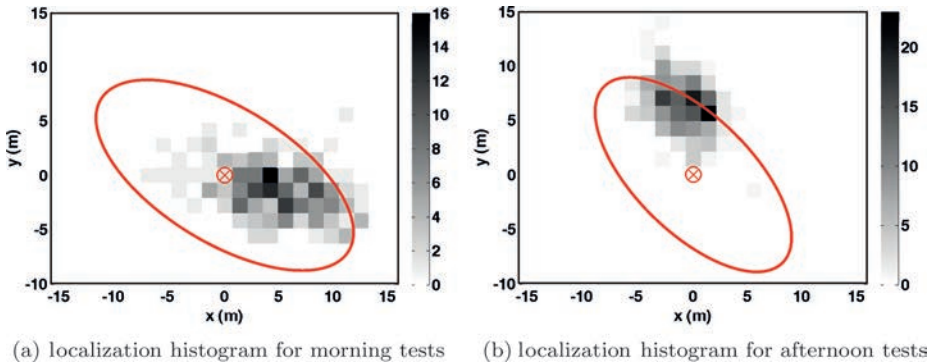


Fig. 10. In both plots the true source location is indicated by the red x, and the 95% confidence ellipse centered at the true source location is shown in red.

assuming a nominal value for the speed of sound leads to a ToA error of 50 ms which roughly translates to 17 m of location error.

We recomputed the ToA statistics taking the temperature variations into account and plot the results for the morning and afternoon tests in Figures 9(a) and 9(b), respectively. The localization error histograms along with the 95% confidence ellipses can be seen in Figures 10(a) and 10(b) respectively. Localization errors were computed only using shots in which all four sensors detected the event.

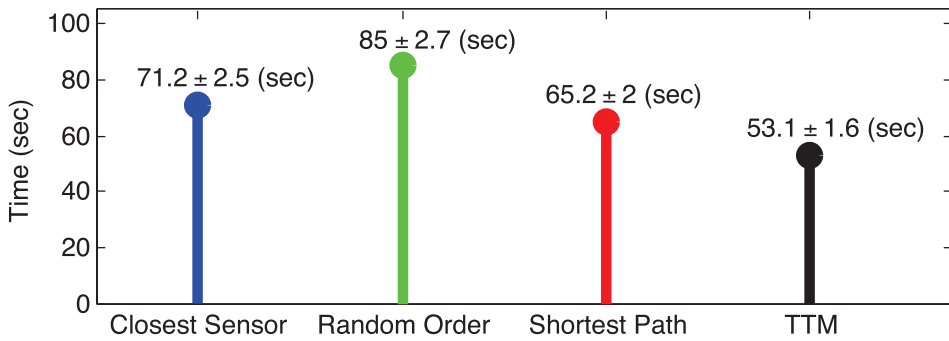


Fig. 11. Monte Carlo simulation results comparing the TTM algorithm against three other routing protocols.

For the tests in the morning, the bias is significantly reduced from 20 ms to less than 10 ms when temperature variations are taken into account. For the tests in the afternoon, there is a significant bias in the measurements even after correcting for temperature variations. The measured ToAs are biased from the expected ToAs by as much as 22 ms. The only explanation we can provide for the bias in the measurements is that the registered ToAs are caused by sound waves reflected off objects, rather than the direct path between the source and the sensor. However, we cannot prove this assertion rigorously.

The experiments illustrate that temperature variations and the propagation environment can have a significant effect on the bias of the ToA measurements. This bias causes slightly more than 5% of the localization errors to fall outside the 95% confidence ellipse. In spite of these inaccuracies in modeling the sound propagation, the localization is robust in that the maximum error during the entire day of testing was less than 14 m.

4.3. UAV Routing Performance in Simulation

To gain a better understanding of the performance of the UAV routing protocol, we exchanged the UAV module for FlightGear, a high-fidelity simulation platform. In simulation, the UAV has a waypoint controller that eliminates the need for human intervention. FlightGear offers realistic vehicle dynamics, including effects from wind. The particular vehicle model we simulated was a Sig Rascal 110, capable of an airspeed of 22 m/s.

The results of 500 simulation trials are shown in Figure 11 for four routing protocols that differ in how the next sensor is selected: (1) *Closest Sensor* greedily chooses the nearest sensor, (2) *Random* chooses a random sensor, (3) *Shortest Path* is a traveling salesperson tour, and (4) *Threshold Time Minimization* is the algorithm presented in this work. The closest sensor protocol visits sensors that are close together, and thus have less information compared to sensors that are widely spread. The fact that the localization time is large is not surprising, however, as the optimization is myopic and does not consider value of the sensor information. The random order protocol takes the most amount of time to localize the source, but visits relatively few sensors. This is attributed to the fact that the randomly selected sensors tend to be wide spread, and thus have high information content although they take a long flight time to reach. The shortest path protocol localizes the source fairly quickly, but tends to visit many sensors. As with the myopic closest sensor protocol, the information value of the data contained at each sensor is not considered. The threshold time minimization protocol results in a localization time that is on average 12 s faster than the shortest path protocol. This is a direct result of the optimization algorithm seeking to balance the




Fig. 12. The peak value of the normalized derivative is much larger at the near-by sensor (a) than the faraway sensor (b), indicating the cannon shot loses its impulsive characteristic with distance. The recorded signal at the faraway sensor has a high Signal-to-Noise Ratio (SNR) (c), indicating that the drop in normalized derivative is mainly because of the loss in impulsive nature with distance.

utility of the information of faraway sensors with the cost to reach those sensors. For each simulation trial, one AoA sensor and seven ToA sensors were randomly placed in a 1 km by 1 km area and the source was placed randomly in a 700 m by 700 m area with each area centered at the origin. These results complement the simulation results from our previous work [Klein et al. 2010] over a larger area.

5. SIGNAL AND CHANNEL CHARACTERISTICS

In this section, we provide detailed experimental results that show the loss in “impulsiveness” of a cannon shot with increasing distance from the source. We then analyze the spatiotemporal characteristics of the effective acoustic channels seen by sensors in such large-scale deployments.

Loss of Impulsiveness with Distance. Consider two sensors S_{near} and S_{far} placed on the runway, at distances of 400 m and 600 m from the cannon. This setting represents a “best-case scenario”: (1) the cannon is pointed directly at the sensors, (2) there is a LoS path between the cannon and the sensors, and (3) the environment is quiet, with no UAV flying overhead. We denote the recorded samples at a generic sensor over a 2 s window with frequency $f_s = 44100$ Hz by $s[k], k \in \{0, 1, 2, \dots, 88199\}$. If the cannon shot is much more impulsive than the typical background noise, we would expect the derivative of the signal when the shot occurs to be significantly larger than its typical value. We define a quantity called the normalized derivative to capture this. Let the sample difference $\Delta s[k] = (s[k] - s[k - 1])f_s$ be an approximation to the derivative of the continuous time signal. The typical magnitude of $\Delta s[k]$ is given by $\sigma = \text{median}(|\Delta s[k]|)f_s$. We now define the normalized derivative to be $\Delta_{nds}[k] = \Delta s[k]/\sigma$, which is expected to be large when there is a cannon shot.

For the nearby sensor S_{near} , we see from Figures 12(a) and 12(b) that the normalized derivative reaches a peak value of about 150 when there is a cannon shot, and this is substantially greater than its typical values. On the other hand, for the sensor that is only 200 m further away, the normalized derivative attains a peak value of only about 20 when there is a cannon shot. Worse still, this is not significantly larger than its typical values, implying that impulsiveness cannot be used as a reliable decision statistic at distances greater than about 400 m even in quiet, LoS environments. This is not a result of simple attenuation of the signal: we see from Figure 12(c) that the signal is well above the noise level at S_{far} . Therefore, this phenomenon is solely because the acoustic channel, consisting primarily of the atmosphere and the Earth, filters out the high-frequency components at larger distances from the source.

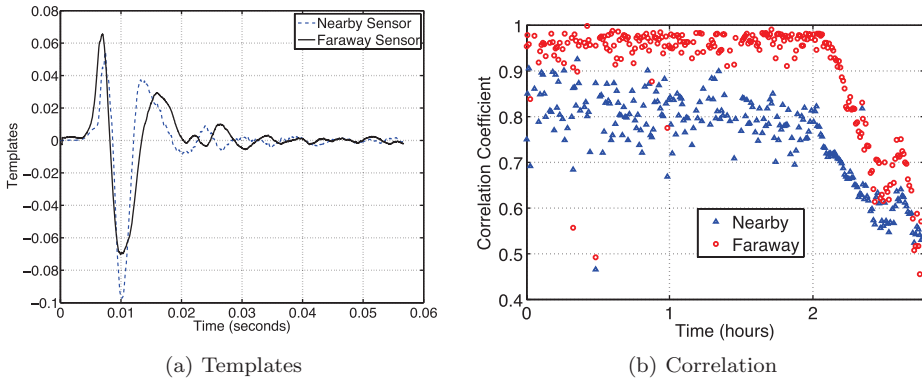


Fig. 13. (a) The templates used to correlate the readings at S_3 . We can see that the templates look fairly similar indicating significant coherence in the acoustic channel over large distances. (b) Correlation between the sensor readings at S_3 and the templates. The template recorded at S_3 correlates better with further readings at S_3 , but both correlations drop off drastically towards the end.

Channel Characteristics. Consider the static deployment in Figure 8 with the propane cannon located at S_4 (this corresponds to the afternoon round of tests). We analyze the signals recorded at two sensors S_1 and S_3 to understand the spatial and temporal variations in the recorded signals. The sensor S_1 has LoS to the cannon (and is also called “near-by sensor”) while low hills and trees block the LoS path between the cannon and S_3 (also referred to as “faraway sensor”). We pick a cannon shot recorded at S_1 and S_3 around $t = 1500$ s to be candidate templates and call them \mathcal{T}_1 and \mathcal{T}_3 respectively. The templates are shown in Figure 13(a). We correlate shots recorded at S_3 from $t = 0$ to about $t = 2$ hours 45 minutes with \mathcal{T}_1 and \mathcal{T}_3 and plot the correlation coefficient in Figure 13(b). Note that different templates are only used in postprocessing to understand the channel characteristics; during the deployment, we simply used one template recorded in an LoS environment at all sensors. We will now summarize our observations.

- (1) \mathcal{T}_3 correlates nearly perfectly with recordings at S_3 from $t = 0$ until about $t = 7500$ seconds. In this time interval, the channel from the source to S_3 can be modeled to be static.
- (2) We now take a closer look at the waveforms recorded after $t = 7500$ seconds to understand the curious phenomenon of the rapid and persistent fall in the correlation in this time window. From Figure 14, we see that early recordings of the cannon shot (for example, at $t = 25$ minutes) have a distinct “N” shape, which lends them the name N-wave. However, later recordings begin to develop a pronounced “hump” in the lower part of the N-wave, which grows with time. The progressively increasing hump, which eventually leads to a sign flip in parts of the N-wave, causes the rapid fall in correlation.
- (3) The waveforms with a hump can be approximated by passing \mathcal{T}_1 through a linear channel with a relatively small number of taps; from Figure 15, we see that 5–6 channel taps suffice to explain the recorded data. While we do not understand the physical phenomena and changes in the environment that caused the hump, the implications are clear: it is necessary to handle time-varying multipath channels, with significant temporal correlations, even between a static source and a sensor to extend the range of detection algorithms.
- (4) From Figure 13(b), we see that the correlations obtained with \mathcal{T}_1 are always smaller than those obtained with \mathcal{T}_3 , indicating that S_3 does not see an exact replica of the

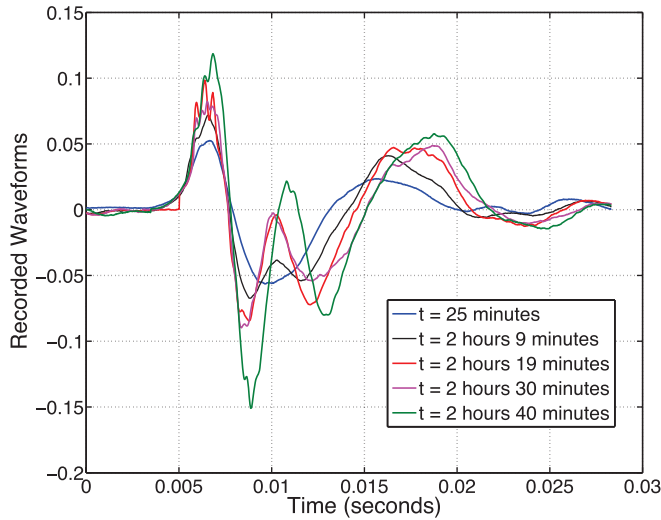


Fig. 14. We see the lower part of the “N-wave” developing a hump whose size increases with time. This causes the correlation to drop rapidly.

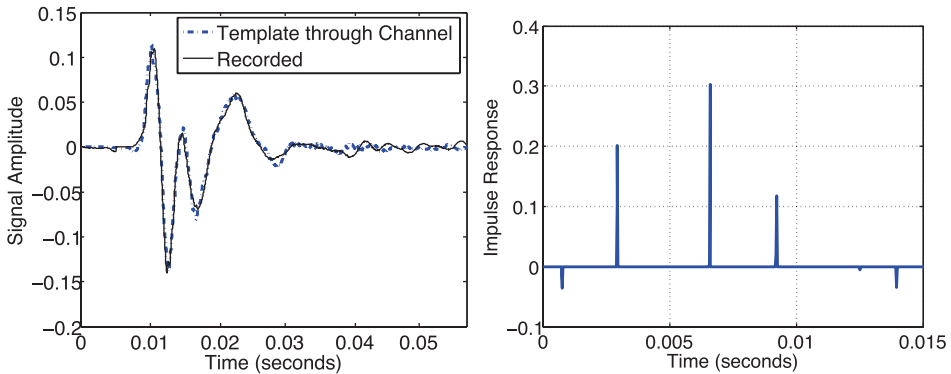


Fig. 15. We can see that the recorded data is approximated nearly perfectly by the signal propagating through a multipath channel with only six taps.

signal recorded at S_1 . However, the correlations obtained with \mathcal{T}_1 are also good (mean correlation coefficient of about 0.8 until $t = 7500$ seconds), indicating that the distortion is not too severe and that the channel is fairly coherent even over large separations.

6. CONCLUSIONS

We have introduced the novel concept of data mules which adapt their future actions based on-the-fly inference from data that they collect. We have demonstrated, through both UAV field tests and flight simulation, that this approach is effective for rapid acoustic source localization using sparsely deployed sensors. Key contributions of this work include a minimal sensor subsets approach to Bayesian UAV routing that couples source localization with path planning, a demonstration of channel coherence using a matched filtering technique, and detailed investigation of the effects of the environment on the recorded signals. We hope that the detailed description of our design choices and

findings provided in this article will prove helpful in future design and deployment efforts.

A natural next step is to develop architectures for coordinating multiple UAVs for covering larger deployment areas, under reasonable assumptions for their capabilities for communicating with each other. For larger deployment areas, it may also be useful to have a two-tier approach to the data provided by the sensors, such as a long-range, very low bit rate (even one bit) “alert” signal that draws in UAVs when the sensor detects an event of interest, together with shorter-range, higher bit rate communication that provides detailed measurements based on which the UAVs can make inferences. An important open problem is deriving bounds on the performance of the TTM algorithm. The TTM algorithm can potentially be improved by including UAV kinematic constraints in the shortest path portion of the optimization. This would ensure that the costs of paths planned for the UAV would be closer to the true cost of the UAV to fly these paths. It is also of interest to develop a better understanding of the acoustic channel (e.g., incorporating the effect of multipath) as well as of incorporating various sensing modalities into our architecture. Finally, it is of theoretical interest to reformulate the UAV routing problem to consider the impact of future measurements on future actions, and to find provably good approximations to such a problem.

REFERENCES

- ABEL, J. AND SMITH, J. 1987. The spherical interpolation method for closed-form passive source localization using range difference measurements. In *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP'87)*. Vol. 12, 471–474.
- ALI, A., YAO, K., COLLIER, T., TAYLOR, C., BLUMSTEIN, D., AND GIROD, L. 2007. An empirical study of collaborative acoustic source localization. In *Proceedings of the 6th International Symposium on Information Processing in Sensor Networks (IPSN'07)*. 41–50.
- ANDRIEU, C., DE FREITAS, N., DOUCET, A., AND JORDAN, M. 2003. An introduction to mcmc for machine learning. *Mach. Learn.* 50, 1–2, 5–43.
- ARORA, A., DUTTA, P., BAPAT, S., KULATHUMANI, V., ZHANG, H., NAIK, V., MITTAL, V., CAO, H., DEMIRBAS, M., GOUDA, M., ET AL. 2004. A line in the sand: A wireless sensor network for target detection, classification, and tracking. *Comput. Netw.* 46, 5, 605–634.
- BECK, A., STOICA, P., AND LI, J. 2008. Exact and approximate solutions of source localization problems. *IEEE Trans. Signal Process.* 56, 5, 1770–1778.
- BERTSIMAS, D. AND VAN RYZIN, G. 1993. Stochastic and dynamic vehicle routing in the Euclidean plane with multiple capacitated vehicles. *Oper. Res.* 41, 60–76.
- BISHOP, A., FIDAN, B., DOGANCAI, K., ANDERSON, B., AND PATHIRANA, P. 2008. Exploiting geometry for improved hybrid AOA/TDOA-based localization. *Signal Process.* 88, 7, 1775–1791.
- BOHN, D. A. 1988. Environmental effects on the speed of sound. <http://www.rane.com/pdf/eespeed.pdf>.
- BURMAN, J., HESPANHA, J. P., MADHOW, U., KLEIN, D., ISAACS, J., VENKATESWARAN, S., AND PHAM, T. 2010. Heterogeneous battlefield sensor networks: A bio-inspired overlay architecture. In *Proceedings of the MSS Battlespace Acoustic and Magnetic Sensors (BAMS) and Advanced Signatures Technology Symposium (ASTS)*.
- BURMAN, J., HESPANHA, J. P., MADHOW, U., KLEIN, D. J., PHAM, T., AND SWAMI, A. 2009. Heterogeneous sensor networks: A bio-inspired overlay architecture. In *Proceedings of the Unattended Ground, Sea, and Air Sensor Technologies and Applications Conference, and the SPIE Defense and Security Symposium (DSS)*.
- CHAN, Y. AND HO, K. 1994. A simple and efficient estimator for hyperbolic location. *IEEE Trans. Signal Process.* 42, 8, 1905–1915.
- CHRISTOFIDES, N. 1976. Worst-case analysis of a new heuristic for the travelling salesman problem. Tech. rep. 388, Graduate School of Industrial Administration, Carnegie Mellon University, Pittsburgh, PA.
- CHU, M., HAUSSECKER, H., AND ZHAO, F. 2002. Scalable information-driven sensor querying and routing for ad hoc heterogeneous sensor networks. *Int. J. High Performance Comput. Appl.* 16, 3, 293.
- CVE-2008-1368. 2008. Geodetic system. http://en.wikipedia.org/wiki/Geodetic_system.
- CVE-2009-1368. 2009. Using a garmin gps 18 lvc as ntp stratum-0 on Linux 2.6. <http://time.qnan.org/>.
- DANTZIG, G. B. AND RAMSER, J. H. 1959. The truck dispatching problem. *Manag. Sci.* 6, 1, 80–91.

- DOGANÇAY, K. AND HMAM, H. 2008. Optimal angular sensor separation for AoA localization. *Signal Process.* 88, 5, 1248–1260.
- DOGANÇAY, K. 2007. Optimal receiver trajectories for scan-based radar localization. In *Proceedings of the Conference on Information, Decision and Control (IDC'07)*. 88–93.
- FREW, E., DIXON, C., ARGROW, B., AND BROWN, T. 2005. Radio source localization by a cooperating uav team. In *AIAA Infotech@Aerospace*, Vol. 1, AIAA, 10–20.
- HASTINGS, W. 1970. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika* 57, 1, 97–109.
- HENKEL, D. AND BROWN, T. 2005. On controlled node mobility in delay-tolerant networks of unmanned aerial vehicles. In *Proceedings of the International Symposium on Advance Radio Technologies (ISART'05)*. 7–9.
- HENKEL, D., DIXON, C., ELSTON, J., AND BROWN, T. 2006. A reliable sensor data collection network using unmanned aircraft. In *Proceedings of the 2nd International Workshop on Multihop Ad Hoc Networks: From Theory to Reality*. ACM Press, New York, 125–127.
- HISTORY for CA62, California. 2010. Dwunderground.com. History for CA62, California, Hourly Observations for January 28, 2010.
- HOFFMANN, G. AND TOMLIN, C. 2010. Mobile sensor network control using mutual information methods and particle filters. *IEEE Trans. Automatic Control* 55, 1, 32–47.
- KLEIN, D. J., SCHWEIKL, J., ISAACS, J. T., AND HESPANHA, J. P. 2010. On uav routing protocols for sparse sensor data exfiltration. In *Proceedings of the American Control Conference (ACC'10)*. 6494–6500.
- MAINWARING, A., CULLER, D., POLASTRE, J., SZEWCZYK, R., AND ANDERSON, J. 2002. Wireless sensor networks for habitat monitoring. In *Proceedings of the 1st ACM International Workshop on Wireless Sensor Networks and Applications (WSNA'02)*. 88–97.
- MESQUITA, A. R., HESPANHA, J. P., AND ASTROM, K. 2008. Optimotaxis: A stochastic multi-agent on site optimization procedure. In *Proceedings of the 11th International Workshop on Hybrid Systems: Computation and Control*. M. Egerstedt and B. Mishra, Eds., Lecture Notes in Computer Science, vol. 4981, Springer, 358–371.
- OSHMAN, Y. AND DAVIDSON, P. 1999. Optimization of observer trajectories for bearings-only target localization. *IEEE Trans. Aerospace Electron. Syst.* 35, 3, 892–902.
- PAVONE, M., BISNIK, N., FRAZZOLI, E., AND ISLER, V. 2009. A stochastic and dynamic vehicle routing problem with time windows and customer impatience. *Mobile Netw. Appl.* 14, 3, 350–364.
- POOR, H. 1994. *An Introduction to Signal Detection and Estimation*. Springer.
- PRIYANTHA, N., CHAKRABORTY, A., AND BALAKRISHNAN, H. 2000. The cricket location-support system. In *Proceedings of the 6th Annual International Conference on Mobile Computing and Networking (MOBICOM'00)*. 32–43.
- RYAN, A. AND HEDRICK, J. K. 2010. Particle filter based information-theoretic active sensing. *Robotics Auton. Syst.* 58, 5, 574–584.
- SHAH, R., ROY, S., JAIN, S., BRUNETTE, W., RES, I., AND SEATTLE, W. 2003. Data mules: Modeling a three-tier architecture for sparse sensor networks. In *Proceedings of the 1st IEEE International Workshop on Sensor Network Protocols and Applications*. 30–41.
- SIMON, G., MAROTI, M., LEDECZI, A., BALOGH, G., KUSY, B., NADAS, A., PAP, G., SALLAI, J., AND FRAMPTON, K. 2004. Sensor network-based countersniper system. In *Proceedings of the 2nd International Conference on Embedded Networked Sensor System (SenSys'04)*. 1–12.
- UCINSKI, D. 2004. *Optimal Measurement Methods for Distributed Parameter System Identification*. CRC Press, Boca Raton, FL.
- VAN TREES, H. 1968. *Detection, Estimation, and Modulation Theory, Part I*. Wiley Interscience.
- VAZIRANI, V. V. 2001. *Approximation Algorithms*. Springer.
- VENKATRAMAN, S. AND CAFFERY, J. 2004. Hybrid TOA/AOA techniques for mobile location in non-line-of-sight environments. In *Proceedings of the Wireless Communications and Networking Conference (WCNC'04)*. Vol. 1, IEEE, 274–278.
- VOLGYESI, P., BALOGH, G., NADAS, A., NASH, C., AND LEDECZI, A. 2007. Shooter localization and weapon classification with soldier-wearable networked sensors. In *Proceedings of the 5th International Conference on Mobile Systems, Applications and Services*. ACM Press, New York, 126.
- WANG, H. B., CHEN, C. E., ALI, A. M., ASGARI, S., HUDSON, R. E., YAO, K., ESTRIN, D., AND TAYLOR, C. 2005. Acoustic sensor networks for woodpecker localization. In *Proceedings of the SPIE Conference on Advanced Signal Processing Algorithms, Architectures, and Implementations XV (SPIE'05)*. Vol. 5910.
- WERNER-ALLEN, G., LORINCZ, K., RUIZ, M., MARCILLO, O., JOHNSON, J., LEES, J., AND WELSH, M. 2006. Deploying a wireless sensor network on an active volcano. *IEEE Internet Comput.* 10, 2, 18–25.

- XU, N., RANGWALA, S., CHINTALAPUDI, K., GANESAN, D., BROAD, A., GOVINDAN, R., AND ESTRIN, D. 2004. A wireless sensor network for structural monitoring. In *Proceedings of the 2nd International Conference on Embedded Networked Sensor Systems (Sensys'04)*. 13–24.
- YLI-HIETANEN, J., KALLIOJARVI, K., AND ASTOLA, J. 1996. Robust time delay based angle of arrival estimation. In *Proceedings of the IEEE Nordic Signal Processing Symposium*. 219–222.
- ZHAO, F., SHIN, J., AND REICH, J. 2002. Information-driven dynamic sensor collaboration for tracking applications. *IEEE Signal Process. Mag.* 19, 2, 61–72.

Received August 2010; revised October 2011; accepted February 2012