

On UAV Routing Protocols for Sparse Sensor Data Exfiltration

Daniel J. Klein, Johann Schweikl, Jason T. Isaacs, and João P. Hespanha
Department of Electrical and Computer Engineering
University of California, Santa Barbara, CA 93106

Abstract—The problem addressed in this paper is data exfiltration from a collection of sensors that are unable to establish ad-hoc communication due to their widespread deployment, geographical constraints, and power considerations. Sensor data is exfiltrated by one or more uninhabited aerial vehicles (UAVs) that act as data mules by visiting each sensor in order to establish a communication link. In many applications, the sequence in which the UAVs visit the sensors can have large impact on the overall performance because some sensors have more informative data than others and because distant nodes take a long time to visit. One such application that we will focus on in this paper is the acoustic source localization problem in which the objective is to localize the source of a transient acoustic event as quickly as possible. We introduce two protocols, ACM and TTM, based on receding horizon optimization of the volume of the Cramer-Rao uncertainty ellipsoid and show significant performance benefits over several other routing protocols using a high-fidelity online simulation environment.

I. INTRODUCTION

In recent years, sensor networks have enabled persistent surveillance of large regions of interest for a wide variety of applications including environmental monitoring [1], agriculture [2], vehicular detection [3], and acoustic source localization [4]. In some sensor network scenarios, the individual sensors form a peer-to-peer ad-hoc communication network for information exchange. However, in low density sensor deployments the individual sensors may be too far apart to communicate effectively, thus demanding a different means of data exfiltration. We propose using one or more uninhabited aerial vehicles (UAVs) for this purpose, and present two protocols for determining the order in which the sensors should be visited so as to extract information optimally.

The UAV routing protocols developed in this paper are sufficiently general in nature so as to handle a large variety of applications. However, for concreteness we focus our discussion on the acoustic source localization problem. In this problem, an acoustic source creates a loud disturbance at an unspecified time and location, and the overall objective is to determine the position of the source as quickly as possible. To this end, we equip each sensor with a microphone, to listen for the event, and a GPS unit, for localization and time synchronization. The output of each sensor is a noisy measurement of the time-of-arrival (TOA) of the acoustic event.

Acoustic source localization problems have a long history in the literature. Many authors have addressed the problem

of estimating the source location from the noisy TOA measurements [5], [6], [7]. Others have addressed optimal sensor placements [8], [9], [10], [11]. To the best of our knowledge, this work is the first to explore protocols for routing one or more UAVs to collect TOA measurements.

The UAV routing problem is a combinatorial optimization in that we seek the “best” sensor sequence from the domain of all possible sequences. Classic examples of combinatorial optimization problems include the traveling salesperson [12] and general vehicle routing [13] problems, which are known to be NP-hard. In these problems, one or more agents are required to visit many customers in a time-optimal way. Each customer usually has the same priority, although heterogeneous priority classes have been considered recently [14]. Another generalization of the traveling salesperson problem is surveyed in [15], where a profit is associated with each customer, and the objective is to maximize profit while minimizing travel costs.

A main difference between the UAV routing problem in consideration here and other combinatorial optimizations is that the value of the information held at each unvisited sensor depends on which sensors have already been visited, due to correlations in the data. The key factor in the acoustic source localization problem for determining this correlation is the direction from which each sensor views the source. For example, two sensors viewing the source from orthogonal directions provide much more information than two sensors viewing the source from similar directions, regardless of the source-to-sensor distance. The problem is not as simple as choosing sensors that view the source from diverse directions because the location of the source is unknown and because the flight time between sensors must be factored into the selection process.

To quantify the informational value of visiting a particular sensor *after* visiting a subset of the other sensors, we use the Cramer-Rao (CR) matrix from estimation theory. The CR matrix quantifies the maximum extent to which a set of measurements can provide information about the unknown parameters in the form of a lower bound on the estimation error covariance provided by *any* unbiased estimator. Computing the CR matrix, however, requires knowledge of the true parameter values which are *never* available in practice. Instead, we take an expectation with respect to the posterior probability of the parameters given the measurements collected so far, which itself is approximated with a sampling approach for computational reasons.

The main contributions of this work include two routing

protocols, ACM and TTM, which seek to minimize the time-averaged expected volume of the CR uncertainty ellipsoid and the first time at which the volume of the CR uncertainty ellipsoid falls below a threshold, respectively. These protocols weigh the value of the information held at each sensor against the UAV flight time required to obtain the sensor data in a receding-horizon online optimization. Implementation shortcuts, including a *majority rules* algorithm, allow the protocols to run in real time. A high-fidelity online simulation environment is used to compare the proposed protocols to random, closest sensor, and shortest path protocols. The performance benefit is significant, the TTM protocol localizes the source 30% faster on average than the shortest path protocol.

The remainder of this paper is organized as follows. In the next section, the UAV routing problem is formalized and details are given regarding the acoustic source localization problem. In Section III, the lower-bound on the volume of the estimation error uncertainty ellipsoid is connected to the determinant of the Cramer-Rao matrix, and the CR matrix for the acoustic source localization problem is reviewed. The two routing protocols are described in Section IV, followed by several implementation shortcuts that are discussed in Section V. Simulation results are presented in Section VI, and concluding remarks follow.

II. UAV ROUTING PROBLEM FORMULATION

The UAV routing problem is set in a large region containing N sensors, each equipped with a means to localize and perform time synchronization, e.g. via GPS. The sensors do not communicate directly with each other due to large separation distances, terrain or urban signal blockages, and/or concerns for battery conservation. Instead, data from each sensor will be exfiltrated by one of M uninhabited aerial vehicles (UAV). The UAV will fly sufficiently close to the sensor to enable a brief communication channel over which the sensor data is communicated. The objective is to choose for each UAV a sequence of sensors to service so as to enable quick identification of a vector of unknown constant parameters θ .

Initial information about the unknown parameters is encoded in the prior probability density $p_0(\theta)$, which can be taken as a uniform distribution over the domain of possible values if no detailed knowledge is available a priori. The data vector \mathbf{y} recorded by the sensors is related to the unknown parameters through the conditional probability, $p(\mathbf{y}|\theta)$. This function can be used as a likelihood $L(\theta; \mathbf{y}) := p(\mathbf{y}|\theta)$ to reason about parameters θ from the available sensor data \mathbf{y} . The posterior probability of the parameters given the data is proportional to the product of the prior density and the likelihood,

$$p(\theta|\mathbf{y}) \propto L(\theta; \mathbf{y})p_0(\theta). \quad (1)$$

In the acoustic source localization problem that we will focus on in this paper, an acoustic source at an unspecified location generates a transient acoustic event at an unspecified time. The uncertain parameter vector θ contains the x and y

coordinates of the source. The data communicated between the k^{th} sensor and the UAV is the time-of-arrival (TOA) z_k of the acoustic disturbance. This TOA measurement is modeled as

$$z_k = T_0 + \frac{d_k}{\nu} + w_k, \quad (2)$$

where T_0 is the unknown event time, d_k is the separation distance between the source and the sensor, ν is the speed of sound, and w_k is noise due to timing errors modeled as a zero mean Gaussian distribution with standard deviation σ_k .

A single TOA measurement is not immediately useful in inferring anything about θ due to the uncertainty in the event time. With two TOA measurements, the event time uncertainty can be eliminated by taking the time-difference-of-arrival (TDOA) between two sensors, for example with respect to the N^{th} sensor,

$$y_k := z_k - z_N = \frac{d_k - d_N}{\nu} + w_{kN}, \quad k = 1, \dots, N-1. \quad (3)$$

The noise $w_{kN} := w_k - w_N$ is distributed as a zero-mean Gaussian with standard deviation $\sigma_{kN} = \sqrt{\sigma_k^2 + \sigma_N^2}$. The likelihood function for the source position θ given this TDOA data is

$$L(\theta; y_k) \propto \exp\left(-\frac{[y_k - (d_k(\theta) - d_N(\theta))/\nu]^2}{2\sigma_{kN}^2}\right), \quad (4)$$

which has a continuum of maxima along the hyperbola of points in the plane for which no noise must be added to the TOA data to explain the source, see Fig. 1

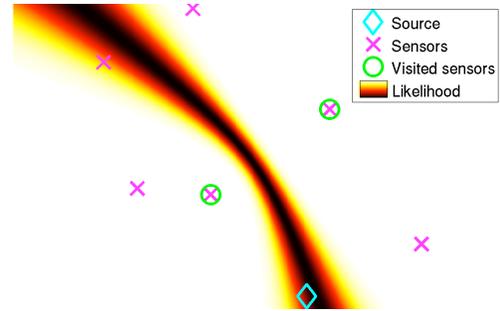


Fig. 1. The likelihood of the source given TOA measurements from two of six sensors. Dark regions are more likely to contain the source than light regions. Sensor noise was exaggerated to better show detail in this example.

Let the vector of all $N-1$ independent TDOA measurements, with differences taken with respect to the N^{th} sensor, be collected into a vector $\mathbf{y} \in \mathbb{R}^{N-1}$ and denote by $\mathcal{I} \subseteq \{1, 2, \dots, N-1\}$ the indices of sensors that have been visited in addition to sensor N . For a vector $v \in \mathbb{R}^{N-1}$, denote by $v_{[\mathcal{I}]} \in \mathbb{R}^{n-1}$ the projection of v onto the elementary subspace indexed by \mathcal{I} . A similar notation will be used for matrices using two indices. Then, after the UAVs have visited $n = 1 + |\mathcal{I}|$ of the N sensors, the available TDOA observations can be denoted $\mathbf{y}_{[\mathcal{I}]} \in \mathbb{R}^{n-1}$.

The likelihood of θ given $\mathbf{y}_{[\mathcal{I}]}$ is then

$$L(\theta; \mathbf{y}_{[\mathcal{I}]}) \propto \exp\left(-\frac{1}{2}(\mathbf{y} - \bar{\mathbf{y}})_{[\mathcal{I}]}^T Q_{[\mathcal{I}, \mathcal{I}]}^{-1} (\mathbf{y} - \bar{\mathbf{y}})_{[\mathcal{I}]}\right), \quad (5)$$

where $\bar{\mathbf{y}}$ is a vector containing expected TDOA measurements for a source located at $\boldsymbol{\theta}$, and assuming all sensors measurements have an identical error variance of σ^2 , the covariance matrix is

$$Q := E \left[(\mathbf{y} - \bar{\mathbf{y}}) (\mathbf{y} - \bar{\mathbf{y}})^T \right] = \sigma^2 (I + \mathbf{1}\mathbf{1}^T), \quad (6)$$

where $\mathbf{1}$ is the ones vector of length $N - 1$.

III. THE CRAMER-RAO UNCERTAINTY ELLIPSOID

The UAVs visit the sensors in order to gain information about the unknown parameters $\boldsymbol{\theta}$. This information will be used by an estimation algorithm to reduce the uncertainty in the parameter estimate. To avoid coming up with a routing protocol that depends explicitly on the particular estimation algorithm used to compute the parameter estimate $\hat{\boldsymbol{\theta}}$ from the available data $\mathbf{y}_{[\mathcal{I}]}$, we make use of the Cramer-Rao (CR) matrix, and its determinant in particular.

The CR matrix is computed by evaluating the inverse of the Fisher Information matrix at the true value of the parameters [16],

$$F(\boldsymbol{\theta}, \mathcal{I}) = E_{Y|\Theta} \left[\left(\frac{d \log p(\mathbf{y}_{[\mathcal{I}]|\boldsymbol{\theta}})}{d\boldsymbol{\theta}} \right) \left(\frac{d \log p(\mathbf{y}_{[\mathcal{I}]|\boldsymbol{\theta}})}{d\boldsymbol{\theta}} \right)^T \right] \quad (7)$$

$$\mathcal{C}(\boldsymbol{\theta}, \mathcal{I}) = F^{-1}(\boldsymbol{\theta}, \mathcal{I}). \quad (8)$$

The utility of the CR matrix comes from the fact that it provides a lower bound on the estimation error covariance matrix of any unbiased estimator in the sense that if

$$\Sigma(\boldsymbol{\theta}, \mathcal{I}) = E_{Y|\Theta} \left[\left(\hat{\boldsymbol{\theta}}(\mathbf{y}_{[\mathcal{I}]} - \boldsymbol{\theta}) \right) \left(\hat{\boldsymbol{\theta}}(\mathbf{y}_{[\mathcal{I}]} - \boldsymbol{\theta}) \right)^T \right] \quad (9)$$

is the estimation error covariance of an unbiased estimate, then $\Sigma(\boldsymbol{\theta}, \mathcal{I}) - \mathcal{C}(\boldsymbol{\theta}, \mathcal{I})$ is a positive semi-definite matrix. The Cramer-Rao lower bound is tight for *efficient* estimation algorithms, by definition.

The physical intuition behind the CR matrix can be understood through the notion of a *confidence ellipsoid* [16]. A confidence ellipsoid 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. To contain the true parameter value with a higher confidence, a larger ellipsoid is required. The overall shape of the ellipsoid is determined by the correlation matrix, as given by the estimation error covariance. An unbiased estimator providing a Gaussian estimate of the parameters with covariance Σ produces a confidence ellipsoid that is strictly *larger* than the ellipsoid produced using the CR matrix. In other words, the region containing the unknown parameters with, for example, 95% confidence is ellipsoidal, and the CR matrix gives a bound in the sense that the CR-ellipsoid fits strictly *within* the Σ -ellipsoid of the unbiased estimator. The uncertainty ellipsoid of an efficient estimator is equivalent to the ellipsoid given by the CR matrix.

The volume of the confidence ellipsoid for a particular confidence level is proportional to the square root of the

determinant of the correlation matrix. In particular, for a problem with p parameters, the volume of the $(1-\alpha) \times 100\%$ confidence ellipsoid is no smaller than

$$V_{\mathcal{C}}(\boldsymbol{\theta}, \mathcal{I}) = \gamma^{p/2} V_s \sqrt{\det \mathcal{C}(\boldsymbol{\theta}, \mathcal{I})}, \quad (10)$$

where γ is the value for which the p -dimensional chi-squared cumulative distribution reaches $(1 - \alpha)$,

$$V_s = \frac{\pi^{p/2}}{\Gamma(1 + p/2)} \quad (11)$$

is the volume of a unit sphere in p dimensions, and Γ is the complete gamma function. For a 95% confidence ellipse in two dimensions, $\gamma = 5.991$. Because the CR-ellipsoid is smaller than that of any unbiased estimator, the determinant of the CR matrix is necessarily smaller than that of the estimation error covariance of any unbiased estimator.

Online minimization of the confidence ellipsoid volume is not possible due to the fact that the CR matrix depends on the true parameters $\boldsymbol{\theta}$, which are of course unknown. However, one can instead compute the *expected value* of the volume of the uncertainty ellipsoid with respect to the posterior probability of the parameters,

$$\bar{V}_{\mathcal{C}}(\boldsymbol{\theta}, \mathcal{I}) = E_{\Theta|Y} [V(\boldsymbol{\theta}, \mathcal{I})]. \quad (12)$$

For problems in which the likelihood of the parameters given the data is a multivariate Gaussian, as is the case in the acoustic source localization problem, the Fisher Information matrix has a special form that permits closed form computation of the uncertainty ellipsoid. Denoting by \mathbf{q}_k the position of the k^{th} sensor, the Fisher Information matrix for the TDOA localization problem can be written as [17],

$$F(\boldsymbol{\theta}, \mathcal{I}) = \frac{1}{\nu^2} G_{[\mathcal{I}]\{1,2\}}^T Q_{[\mathcal{I},\mathcal{I}]}\^{-1} G_{[\mathcal{I}]\{1,2\}}, \quad (13)$$

where Q is from (6) and

$$G = \begin{bmatrix} \mathbf{g}_1^T - \mathbf{g}_N^T \\ \vdots \\ \mathbf{g}_{N-1}^T - \mathbf{g}_N^T \end{bmatrix}, \quad (14)$$

with

$$\mathbf{g}_k = \frac{\boldsymbol{\theta} - \mathbf{q}_k}{\|\boldsymbol{\theta} - \mathbf{q}_k\|}. \quad (15)$$

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 non-collinear sensor readings have been obtained, thus the Cramer-Rao matrix (8) is undefined until at least three sensors have been visited.

IV. ROUTING PROTOCOLS

In this section, we present two optimization-based routing protocols for the UAV routing problem. The protocols are related in that they are both minimizations of a function of the expected volume of the uncertainty ellipsoid, however they differ in overall objective. An aspect common to both problems is that the route through the remaining sensors

should be recomputed after each new piece of data arrives because the posterior probability changes.

Note that these protocols are based on the volume of the uncertainty ellipsoid, and thus only apply when enough sensors have been visited so as to render the Fisher Information matrix full rank. To get this matrix to be full rank as quickly as possible, each of the following two protocols requires prior information or should begin by taking the shortest path through the number of sensors needed to make the Fisher Information matrix invertible (three in the case of the acoustic source localization problem).

Good routes tend to collect a lot of information about the unknown parameters in a short amount of time. However, there is design choice to be made concerning when the information should be collected. For example, a protocol can optimize early information gain at the cost of average or long-term information gain. The following two protocols differ in how this tradeoff is balanced.

Protocol I: Average Cost Minimization

The first routing protocol (ACM) encodes this design freedom in a user-chosen monotonic cost function $C_v : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ which maps the volume of the uncertainty ellipsoid to a cost. The optimization is carried out from the current time t to a planning horizon of length T . Sensor data (TDOA values) available at the current time are collected in the vector $\mathbf{y}_{[I(t)]}$, and let $\mathcal{I}(\tau, \mathbf{r})$ be a vector 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 total cost for route \mathbf{r} is given by the average cost over the planning horizon,

$$J_1(\mathbf{r}, t) := \frac{1}{T} \int_t^{T+t} E_{\Theta|\mathbf{y}(\mathbf{z}(t))} [C_v(V_C(\boldsymbol{\theta}, \mathcal{I}(\tau, \mathbf{r})))] d\tau. \quad (16)$$

Because the volume of the uncertainty ellipsoid, $V_C(\boldsymbol{\theta}, \mathcal{I}(\tau, \mathbf{r}))$ does not depend on the *sequence* in which the sensors are visited, an infinite planning horizon can be accomplished by selecting T to be larger than the time required to service the slowest route through all remaining unvisited sensors,

$$T \geq T_\infty := \max_{\mathbf{r} \in \mathcal{R}(0)} S(\mathbf{r}), \quad (17)$$

where $S(\mathbf{r})$ is the service time for route \mathbf{r} and $\mathcal{R}(t)$ is the set of all possible routes through the unvisited sensors at time t . The optimal route is selected as the one yielding the lowest cost,

$$\mathbf{r}_1^*(t) = \arg \min_{\mathbf{r} \in \mathcal{R}(t)} J_1(\mathbf{r}, t). \quad (18)$$

While (18) is computed using the posterior probability of data currently available, the volume of the uncertainty ellipsoid is a function of time because it will change each time a UAV hypothetically visits a sensor along the planned route \mathbf{r} . Although the data from a sensor will not be available until a UAV visits, the information value of the sensor data is factored into the Fisher Information matrix (7), from which the uncertainty volume is updated using (10).

As an example cost function, one can take the identity function,

$$C_v(V_C(\boldsymbol{\theta}, \mathcal{I}(\tau, \mathbf{r}))) = V_C(\boldsymbol{\theta}, \mathcal{I}(\tau, \mathbf{r})). \quad (19)$$

This particular choice results in a minimization of the time-average expected volume of the uncertainty ellipsoid.

Protocol II: Threshold Time Minimization

The second protocol (TTM) aims to minimize the time at which the volume of the uncertainty ellipsoid falls below a threshold value. This threshold value represents the largest uncertainty which is “good enough.” Again, we use the lower bound on the uncertainty ellipsoid volume provided by the Cramer-Rao matrix to avoid having the optimization depend on a particular estimation algorithm.

The volume of the uncertainty ellipse in the acoustic source localization problem gives the area of the region in which the source is likely to be found with probability determined by γ in (10). Once this region is small enough, additional resources can be directed to the area for further investigation. The cost function and resulting optimization problem to accomplish this objective are as follows:

$$J_2(\mathbf{r}, t) = \int_t^{T+t} E_{\Theta|\mathbf{y}(\mathbf{z}(t))} [H(V_C(\boldsymbol{\theta}, \mathcal{I}(\tau, \mathbf{r})) - V_{th})] d\tau \quad (20)$$

$$\mathbf{r}_2^*(t) = \arg \min_{\mathbf{r} \in \mathcal{R}(t)} J_2(\mathbf{r}, t), \quad (21)$$

where $H : \mathbb{R} \rightarrow \{0, 1\}$ is the Heaviside step function. In other words, each route is penalized at unit rate until the first time at which the volume of the uncertainty ellipsoid is expected to fall below the volume threshold, V_{th} .

V. IMPLEMENTATION

The numerical computation needed to fully carry out the routing optimizations presented in the previous section can be burdensome due to the fact that the general nature of the posterior probability of the parameters given the measurements makes closed form computation of the expected value intractable in (16) and (20). Further, the optimization over all possible routes grows exponentially if a planning horizon is not incorporated properly. In this section, we discuss options for mitigating these two implementation challenges.

A. Expected Value Evaluation

The expected value we are interested in evaluating is with respect to the posterior probability of the parameters given the data,

$$\mathcal{E} = \int_{\boldsymbol{\theta}} p(\boldsymbol{\theta}|\mathbf{y}) f(\boldsymbol{\theta}) d\boldsymbol{\theta}, \quad (22)$$

for some function f . Now consider K samples drawn from $p(\boldsymbol{\theta}|\mathbf{y})$, and denote the k^{th} sample by \mathbf{s}_k . Then, the expected value is well approximated by

$$\mathcal{E} \approx \frac{1}{K} \sum_{k=1}^K f(\mathbf{s}_k). \quad (23)$$

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 [18], [19] are well documented in the literature. The approach we will employ in the simulation section is based on evaluating the posterior probability at every point on a fine regular grid of size w by h . The k^{th} sample \mathbf{s}_k is selected from the collection of grid points with a frequency ω_k equal to its probability density over the sum of the probability densities,

$$\omega_k = \frac{p(\mathbf{s}_k|\mathbf{y})}{\sum_{i=1}^{wh} p(\mathbf{s}_i|\mathbf{y})}. \quad (24)$$

A finer mesh and larger K yields a better approximation of the expected value at the cost of computational resources.

B. Planning Horizon Shortening

A planning horizon significantly shorter than the time required to reach all sensors can significantly reduce the computational burden of computing the optimization over all routes in the optimization problems (18) and (21). Further variations on the two protocols allow us to realize an even greater savings. Each of these variations can produce sub-optimal paths, but paths are recomputed online in a receding horizon manner, and thus we are primarily concerned with the first leg. These variations were designed with routing of a single UAV in mind, but could be extended to a multi-UAV setting.

The ACM protocol finds the route with the lowest average cost over the time-based planning horizon. However, determining which partial routes satisfy a time-based horizon is nontrivial, and the optimal route can change at any time when a sensor that was slightly more than T seconds away suddenly becomes reachable. Instead, we select the route with the lowest average cost over the next L sensors,

$$\tilde{J}_1(\mathbf{r}, \mathbf{s}_k) := \frac{1}{t_{\mathbf{r},L} - t_{\mathbf{r},0}} \sum_{\ell=1}^L \tau_{\mathbf{r},\ell} C_v(V_C(\mathbf{s}_k, \mathcal{I}(t_{\mathbf{r},\ell}, \mathbf{r}))) \quad (25)$$

$$\mathbf{r}^* = \arg \min_{\mathbf{r} \in \mathcal{R}_L} \frac{1}{K} \sum_{k=1}^K \tilde{J}_1(\mathbf{r}, \mathbf{s}_k), \quad (26)$$

where \mathcal{R}_L is the set of routes of length L through the remaining unvisited sensors, $t_{\mathbf{r},\ell}$ is the time at which the ℓ^{th} sensor will be reached along route \mathbf{r} , and $\tau_{\mathbf{r},\ell} := t_{\mathbf{r},\ell} - t_{\mathbf{r},\ell-1}$ is the travel time between sensors $\ell - 1$ and ℓ on route \mathbf{r} , assumed to be proportional to the distance separating the sensors. A number-of-sensors horizon is beneficial not only in determining which routes satisfy the horizon criteria, but also in the fact that the integrand becomes piecewise constant, changing only when a sensor is visited. This simplification is clearly suboptimal as the route will not change between two sensors even if a better route is available, yet produces good results at a significant computational savings.

The TTM protocol has in the integrand a Heaviside step, which means that routes need not be evaluated beyond the

point at which the volume of the uncertainty ellipsoid falls below the volume threshold. However, as with the time-based horizon in the ACM protocol, one would need to search the routes to determine where truncation can occur. The simplification we make comes from the fact that the volume of the confidence ellipsoid does not depend on the sequence in which the sensors are visited. Thus, we can find all minimal combinations¹ of sensors for which the volume is under the threshold, and then choose an order in which to visit the sensors in each combination using a shortest path protocol.

The route which achieves the threshold in the least amount of time for each sample \mathbf{s}_k is recorded on the k^{th} row of a routing matrix \mathcal{S} . Routes can vary in length because they terminate when the threshold is achieved, so entries of routing matrix after termination are filled with empty characters (e.g. \emptyset). This routing matrix is used in the following *majority rules* algorithm to choose a sequence of sensors for the UAV to visit.

Majority Rules Algorithm: After visiting n of N sensors, the majority rules algorithm selects a route \mathbf{r} through the remaining $N - n$ sensors in the following way.

- 1: **for** $i = 1$ to $N - n$ **do**
- 2: **if** the i^{th} column of \mathcal{S} contains only \emptyset entries **then**
- 3: Set r_i , the i^{th} sensor of route \mathbf{r} , equal to the closest unvisited sensor.
- 4: **else if** the mode of the i^{th} column is unique **then**
- 5: Set r_i equal to this mode.
- 6: **else**
- 7: Set r_i equal to a randomly selected sensor in the mode.
- 8: **end if**
- 9: Eliminate rows of \mathcal{S} that have an entry other than r_i in i^{th} column.
- 10: **end for**

As an example, consider the following routing matrix,

$$\mathcal{S} = \begin{bmatrix} 3 & 5 & \emptyset \\ 3 & \emptyset & \emptyset \\ 3 & 6 & \emptyset \\ 6 & 2 & 3 \end{bmatrix}, \quad (27)$$

which was generated from $K = 4$ samples although many more would be used in practice. The optimal route for the first sample, \mathbf{s}_1 , was $3 \rightarrow 5$, as entered on first row of \mathcal{S} . The majority rules algorithm begins by considering only the first column of \mathcal{S} , which is equal to $[3, 3, 3, 6]^T$. The mode of this column is unique, so $r_1 = 3$. The fourth row of \mathcal{S} is then eliminated because it visited sensor 6 instead of 3. The routing matrix thus becomes

$$\mathcal{S} = \begin{bmatrix} 3 & 5 & \emptyset \\ 3 & \emptyset & \emptyset \\ 3 & 6 & \emptyset \end{bmatrix}. \quad (28)$$

¹A minimal combination is one in which the removal of any one sensor brings the volume of the uncertainty ellipsoid above the threshold.

Then the second column of the reduced routing matrix, $[5, \emptyset, 6]^T$, is considered. Because the mode is not unique, r_2 is selected from 5 and 6 at random, and the other row is eliminated, e.g. $r_2 = 6$ and $\mathcal{S} = [3 \ 6 \ \emptyset]$. The third column is just \emptyset , so the nearest unvisited sensor is selected as r_3 .

VI. SIMULATION RESULTS

In this section, the effectiveness of the two proposed routing protocols will be compared against three other routing protocols in a real-time simulation. As before, the focus is on the acoustic source localization problem, but we will only consider routing a single UAV despite the fact the the routing protocols are more general in nature.

A. Routing Protocols in Comparison

Three routing protocols in addition to the two proposed in this paper will be compared. These other routing protocols are *source independent* in that they do not specifically factor in the probability distribution of the source location. The first of these is a *random* routing in which the sensors are visited in a random order. This protocol is not expected to perform very well, but requires the least amount of computation.

The second protocol chooses to always visit the *closest* sensor. In doing so, this protocol will quickly collect TOA data, but the data may be of low quality due to the sensors viewing the source from similar directions. While again only a small amount of memory and computation are required, this protocol is myopic in that it only considers where to go next.

The third routing protocol in consideration is the *shortest* path. This protocol chooses the route that minimizes the total time to visit all sensors from the initial location of the UAV. While significantly more memory and processing power are required to compute this route, it is advantageous in that the sensors are serviced as quickly as possible. In computing the route, we assume that UAV flight time from sensor to sensor is proportional to the distance between those sensors. The actual flight time used for evaluation could be significantly different, as it comes from a realistic UAV simulator.

B. Simulation Environment and Performance Metric

To ensure realistic UAV flight dynamics in a simulation environment, the UAV is simulated in FlightGear [20], an open source flight simulation software package. The aircraft model simulated by FlightGear is a Sig Rascal 110, which has an air speed of 80kph and performance characteristics closely resembling typical UAV. Realistic environmental conditions, such as wind, are present on each simulation trial.

The performance metric is the time required to “localize” the source, where we say that the source has been localized when the volume of the uncertainty ellipsoid, computed using the currently available sensor data, falls below a threshold. We note the TTM protocol aims to minimize the localization time directly, and is therefore expected to achieve the best performance despite the implementation shortcuts. In practice, the choice of TTM or ACM should be based on user preferences and mission parameters.

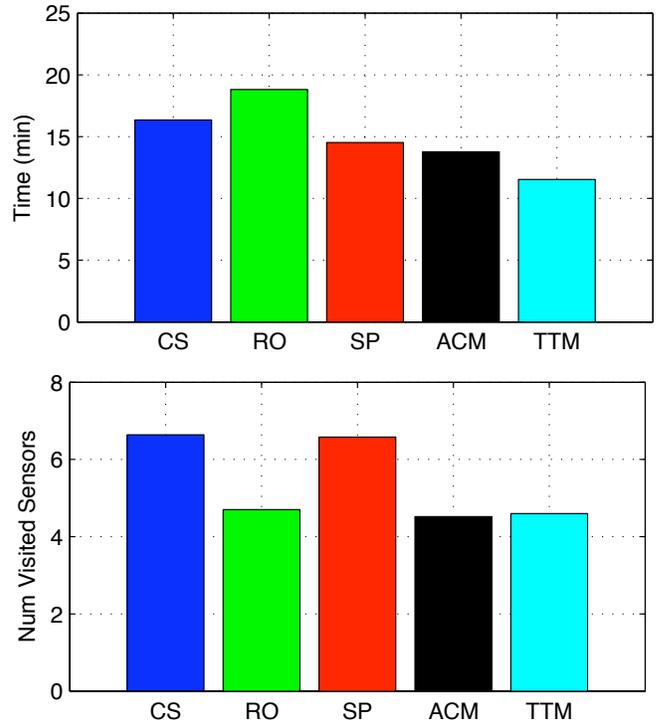


Fig. 2. Simulation results from Monte Carlo trials using FlightGear showing for each of the five protocols (top) the average time required to achieve the threshold, and (bot) the average number of sensors visited. The five protocols are: closest sensor (CS), random order (RO), shortest path (SP), and the proposed protocols for average cost minimization (ACM) and threshold time minimization (TTM).

Monte Carlo trials were conducted to measure the performance of the five routing protocols. On each trial, ten sensors and a single source are each placed randomly within a square region of size $10\text{km} \times 10\text{km}$. The time was recorded when the area of the 95% uncertainty ellipsoid fell below 420m^2 (about the size of a basketball court). These parameters are consistent with current U.S. Army interests [21]. To obtain consistency in initial conditions, the UAV flies through a sequence of fixed dummy waypoints before beginning the actual routing protocol. When one routing protocol terminates, the UAV is routed back through the dummy waypoints before beginning the next protocol.

C. Results

The result of extensive Monte Carlo simulations of all five routing protocols is shown in Fig. 2. The top chart shows the average time to required to localize the source, and the bottom chart shows the number of sensors that must be visited, on average, to obtain the threshold.

The closest sensor (CS) protocol visits sensors that are close together, and thus have less information compared to sensors that are widely-spread. The fact that both the localization time and the number of visited sensors are large is not surprising, however, as the optimization is myopic and does not consider value of the sensor information.

The random order (RO) 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 (SP) protocol localizes the source fairly quickly, but tends to visit many sensors. As with the myopic CS protocol, the information value of the data contained at each sensor is not considered.

The ACM protocol localizes the source quicker than any of the three other protocols, and visits fewer sensors in the process. The benefit here comes from balancing sensor information content with travel time, according to the cost function C_v . Notice here that we are evaluating this protocol with a cost functional that is significantly different from J_1 , yet the performance is already better than choosing the shortest path through the sensors.

The best performing protocol is TTM, which was designed specifically to minimize the localization time. Note that this protocol visits about as many sensors as the ACM protocol, but chooses a better subset of sensors so that the overall time-to-localization is smaller. Specifically, this protocol will fly past a nearby sensor if a slightly further sensor is expected to reduce the uncertainty below the threshold when the nearby sensor would not.

Both of the proposed protocols were implemented using $K = 30$ samples from the posterior probability, which was evaluated at points on a 1000×1000 uniform grid. The prior probability was taken as a uniform distribution and a 95% confidence was used for computing the area of the uncertainty ellipse.

VII. CONCLUSION

The work in this paper has introduced two routing protocols for data exfiltration from a group of sensors by one or more unmanned aerial vehicles. The two protocols are based on selecting the route that is expected to minimize a function of the lower bound on the volume of the uncertainty ellipsoid. While the protocols require significant computation as initially derived, implementation shortcuts were discussed to enable an on-line and real-time implementation. Simulation results with a realistic UAV model demonstrated the effectiveness of the routing protocols compared to three other protocols.

Future work will focus on extending this work in the following directions. First, we have in this work assumed that all sensors are capable of detecting the acoustic event. In an urban or mountainous setting, however, terrain and obstacles will attenuate the acoustic signature to the extent that distant sensors may fail to differentiate the event from the background noise. By adding a stochastic detection model to the likelihood function, optimal routes will factor in the effects of acoustic attenuation.

Second, the UAVs may be able to communicate simultaneously with several sensors from certain locations. Where in this paper we required a UAV to fly to one sensor at a time, routes could instead be planned using a detailed model of the communication *region* of each sensor. Also, a more

detailed vehicle dynamics model, like Dubins paths, could be used to approximate the flight time instead of using straight line distances.

Finally, the computational complexity of computing optimal routes increases drastically with the number of UAVs. To enable better scaling, the region over which the sensors are deployed could be partitioned, for example using a Voronoi decomposition. The multi-UAV problem then scales sub-linearly in the number of UAV because each UAV has only to plan routes through N/M sensors, on average.

REFERENCES

- [1] A. Mainwaring, D. Culler, J. Polastre, R. Szewczyk, and J. Anderson, "Wireless sensor networks for habitat monitoring," in *WSNA '02: Proceedings of the 1st ACM intl workshop on Wireless sensor networks and applications*. New York, NY, USA: ACM, 2002, pp. 88–97.
- [2] A. Baggio, "Wireless sensor networks in precision agriculture," in *ACM Workshop on Real-World Wireless Sensor Networks (REALWSN 2005)*, Stockholm, Sweden, 2005.
- [3] D. Li, K. Wong, Y. Hu, and A. Sayeed, "Detection, classification and tracking of targets in distributed sensor networks," *IEEE Signal Processing Magazine*, vol. 19, no. 2, pp. 17–29, 2002.
- [4] J. Smith and J. Abel, "Closed-form least-squares source location estimation from range-difference measurements," *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. 35, no. 12, pp. 1661–1669, 1987.
- [5] T. Ajdler, I. Kozintsev, R. Lienhart, and M. Vetterli, "Acoustic source localization in distributed sensor networks," *Signals, Systems and Computers, 2004. Conference Record of the Thirty-Eighth Asilomar Conference on*, vol. 2, pp. 1328–1332 Vol.2, Nov. 2004.
- [6] A. Beck, P. Stoica, and J. Li, "Exact and approximate solutions of source localization problems," *IEEE TRANSACTIONS ON SIGNAL PROCESSING*, vol. 56, no. 5, pp. 1770–1778, 2008.
- [7] A. Bishop, B. Fidan, B. Anderson, K. Dogancay, and P. Pathirana, "Optimal range-difference-based localization considering geometrical constraints," *Oceanic Engineering, IEEE Journal of*, vol. 33, no. 3, pp. 289–301, July 2008.
- [8] A. N. Bishop, B. Fidan, B. D. O. Anderson, P. N. Pathirana, and K. Dogancay, "Optimality analysis of sensor-target geometries in passive localization: Part 2 - time-of-arrival based localization," *Proceedings of the 3rd International Conference on Intelligent Sensors, Sensor Networks and Information*, pp. 13–18, Dec. 2007.
- [9] J. T. Isaacs, D. J. Klein, and J. P. Hespanha, "Optimal sensor placement for time difference of arrival localization," in *48th IEEE Conference on Decision and Control, 2009. CDC 2009.*, Dec. 2009, pp. 7878–7884.
- [10] S. Martínez and F. Bullo, "Optimal sensor placement and motion coordination for target tracking," *Automatica*, vol. 42, no. 4, pp. 661–668, 2006.
- [11] J. Neering, M. Bordier, and N. Maizi, "Optimal passive source localization," *Proceedings of the 2007 International Conference on Sensor Technologies and Applications*, pp. 295–300, Oct. 2007.
- [12] V. V. Vazirani, *Approximation Algorithms*. Springer, 2001.
- [13] G. Dantzig and J. Ramser, "The truck dispatching problem," *Management Science*, pp. 80–91, 1959.
- [14] S. Smith, M. Pavone, F. Bullo, and E. Frazzoli, "Dynamic vehicle routing with heterogeneous demands," in *47th IEEE Conference on Decision and Control, 2008. CDC 2008.*, Dec. 2008, pp. 1206–1211.
- [15] D. Feillet, P. Dejax, and M. Gendreau, "Traveling salesman problems with profits: An overview," *Transportation Science*, vol. 39, pp. 188–205, 2001.
- [16] H. Van Trees, "Detection, estimation, and modulation theory part I," 1968.
- [17] Y. T. Chan and K. C. Ho, "A simple and efficient estimator for hyperbolic location," *IEEE Transactions on Signal Processing*, vol. 42, no. 8, pp. 1905–1915, Aug. 1994.
- [18] W. Hastings, "Monte Carlo sampling methods using Markov chains and their applications," *Biometrika*, vol. 57, no. 1, pp. 97–109, 1970.
- [19] C. Andrieu, N. De Freitas, A. Doucet, and M. Jordan, "An introduction to MCMC for machine learning," *Machine learning*, vol. 50, no. 1-2, pp. 5–43, 2003.
- [20] A. Perry, "The FlightGear flight simulator."
- [21] Swami, A., Personal communication, 2009.