

Algorithms for Cooperative Active Localization of Static Targets with Mobile Bearing Sensors under Communication Constraints

Joshua Vander Hook, Pratap Tokekar, Volkan Isler

Abstract—We study the problem of actively locating a static target using mobile robots equipped with bearing sensors. The goal is to reduce the uncertainty in the target’s location to a value below a given threshold in minimum time. Our cost formulation explicitly models time spent in traveling as well as taking measurements. In addition, we consider distance-based communication constraints between the robots.

We provide the following theoretical results. First, we study the properties of an optimal offline strategy for one or more robots with access to the target’s true location. We derive the optimal offline algorithm and bound its cost when considering a single robot or an even number of robots. In other cases, we provide a close approximation.

Second, we provide a general method of converting the offline algorithm into an online, adaptive algorithm (that does not have access to the target’s true location) while preserving near optimality. Using these two results, we present an online strategy proven to locate the target up to a desired uncertainty level at near-optimal cost. In addition to theoretical analysis, we validate the algorithm in simulations and multiple field experiments performed using autonomous surface vehicles carrying radio antennas to localize radio tags.

I. INTRODUCTION

Systems of mobile networked sensors have already had a significant impact on environmental monitoring applications by automating tedious and potentially dangerous sensing tasks. One application area of interest monitoring radio-tagged invasive fish [1], [2]. We have been developing a system of Autonomous Surface Vehicles (ASVs) which carry radio antennas that detect radio tags attached to the fish (Figure 1). The system is intended to provide data on long-term motions of invasive fish and search large areas for stationary fish aggregations. The duration of such tasks forces us to consider the trade-off between system life and time spent to localize each radio tag. Our goal is to design algorithms that actively choose measurement locations which provide good information about each target, but do not produce an overly time-consuming trajectory.

Previously, we studied the problem of designing active localization strategies for a single robot with a sensor capable of measuring bearings of the radio tags [3]. Extending this to the case of multiple robots is not straightforward, since two or more robots must synchronize their estimates of the target location by communicating with each other. Furthermore, the communication range of the robots is limited in practice, and an optimal algorithm must include time spent while the robots meet to establish communication. Therefore, we study the problem of active localization for mobile robots subject to distance-based communication constraints.

In this paper, we provide three main results. First, we extend the existing body of work which analyzes the offline case: planning measurements with respect to a known target location. The optimal algorithm allows a direct comparison of system improvements (increased velocity, decreased measurement time, increased numbers of robots, or improved sensing), and the effect on the mission objectives such as time-to-localize. We present optimal offline algorithms and bounds on offline costs for any number of collaborating mobile bearing sensors in Section IV.

The authors are with the Department of Computer Science and Engineering, University of Minnesota, Minneapolis, MN, USA. {jvander,tokekar,isler}@cs.umn.edu



Fig. 1. The robotic systems used in field experiments. The 2-meter long boats were designed to track invasive fish autonomously. Each one is equipped with wireless communications, directional antenna used as bearing sensors, a navigation suite, and computing hardware. Our algorithm was implemented on this system and was shown to localize targets with small uncertainty.

Second, we extend the offline-optimal algorithm to include communication constraints (Section IV-C). While the optimal strategy must include at least one communication exchange to gather all the robots’ measurements, we show the optimal strategy might sometimes break communication to take better measurements. Thus, enforcing persistent communication in our setting is potentially suboptimal and our algorithm allows the robots to break communication when necessary. As a result, it also includes a rendezvous component.

Third, we address the more realistic case: when only a prior estimate of the true target location is known. To solve this *online* problem, we plan measurements to minimize the worst-case cost, even though the prior estimate may be uncertain or even misleading. We present a general method of adapting any offline measurement strategy for use in an online setting. In doing so, the cost is shown to be at most a logarithmic factor more than that of the offline optimal algorithm, as shown in Section V.

Our final contribution involves field experiments: We have implemented and tested the online algorithms on the robot system shown in Figure 1. We present field deployments in which two networked robots successfully locate a radio transmitter accurately and without requiring significant travel time. The field experiments are presented in Section VI-B.

II. PROBLEM FORMULATION

In this section we present the notation used and formalize the problem studied in this paper. Let the true target location be denoted by x^* . Each robot can take bearing measurements of the target location. Each measurement takes a fixed amount of time t_m , and is corrupted by zero-mean Gaussian noise with variance σ^2 . Sensor locations are given as $s_{u,i}$ for robot u and measurement i . When the time index is clear from context, we will simply specify s_u . Likewise s_i is used for time i when considering only a single robot’s trajectory. The full sequence of N_u measurements for robot u is $S_u = \{s_{u,1}, \dots, s_{u,N}\}$. S is the union of all measurements taken by all robots, and $N = |S|$ is the total number of measurements taken among all robots. We use the notation $d(a,b)$ to be the Euclidean distance between points a and b . Without loss of generality, we assume that traveling between locations a and b takes $d(a,b)$ units of time and is proportional to the Euclidean distance between the points.

For any robot u , the time required to travel to each point where the robot takes a measurement plus the measurement cost (for example,

in seconds) is

$$C(S_u) = t_m \cdot N_u + \text{len}(S_u) \quad (1)$$

$$\text{where } \text{len}(S_u) = \sum_{i=1}^{N_u} d(s_{u,i}, s_{u,i-1}).$$

The uncertainty in the target's estimate is measured by the covariance matrix of its Probability Density Function (PDF), Σ_i . The covariance matrix will depend on the type of estimator used. However, the Cramer Rao Lower Bound (c.f. Chapter 2.7 in [4]) establishes a lower bound on the covariance of any unbiased estimator as the inverse of the Fisher Information Matrix (FIM), i.e., the ordered eigenvalues of Σ are no smaller than the ordered eigenvalues of \mathbb{F}^{-1} . The design of estimators with output matching the CRLB is separate from this work.

The FIM is defined with respect to the true target location (x^*), and measurement locations given by S . Hence, the inverse of the FIM represents the uncertainty in the target's estimate when deriving the optimal offline trajectories in the following sections.

Our goal is to minimize the maximum diameter of the uncertainty ellipse, i.e., the maximum eigenvalue of \mathbb{F}^{-1} , denoted by $\lambda \mathbb{F}^{-1}$. Since \mathbb{F} is the inverse of the uncertainty, and the eigenvalues of a matrix inverse are the inverse of the matrix eigenvalues, our information constraint is given by the minimum eigenvalue,

$$\underline{\lambda} \mathbb{F}(S, x^*) \geq \lambda_d \quad (2)$$

For brevity, $\underline{\lambda} \mathbb{F}(S)$ is used whenever x^* does not change. Here λ_d defines the requested precision (information) in the final estimate. For example, in our application, it is desirable to obtain estimates accurate to 5 meter resolution ($\sqrt{\lambda \Sigma} = 5m$), so λ_d is equal to $1/25$ during field experiments. Our goal is to localize a static target as quickly as possible using mobile bearing sensors which may have communication constraints. The problem can be stated as follows.

Problem 1 (Active Target Localization). *Given n mobile robots, find a sequence (S_u) of measurement locations ($s_{u,i}$ for $i = 1$ to N_u) for each robot u , such that the maximum cost,*

$$\max_u C_u = N_u t_m + \sum_{i=1}^{N_u} d(s_{i,u}, s_{i-1,u}) \quad (3)$$

is minimized. Furthermore, the measurement locations must satisfy,

$$\underline{\lambda} \mathbb{F}(S, x^*) \geq \lambda_d \quad (4)$$

where $\mathbb{F}(S, x^)$ is the Fisher Information Matrix resulting from all measurement locations S , evaluated with respect to the true target location x^* , and $\underline{\lambda} A$ is the minimum eigenvalue of the matrix A .*

Beginning in Section IV-C, the distance-based communication constraint (which states that the robots may not communicate unless they are within a given distance r_c) will be introduced. If the robots do not communicate the results of their measurements, they cannot form a joint estimate of the target's position. Thus, the uncertainty in target estimate is a function of the measurements gathered up to the last time the robots met. For example, in the offline case (as in Section IV), the fact that the true target location (x^*) is known makes Equation (4) dependent only on the relative locations of the sensors with respect to x^* . Then, the robots do not need to meet more than once, since exchanging measurements will not affect x^* . However, in the online case (as in Section V), x^* is unknown. Therefore, the robots must either take enough measurements to ensure that for any x^* , Equation (4) is satisfied, or periodically meet, update their estimate of x^* , and adjust the measurement sequence to make use of the shared information.

III. RELATED WORK

There has been significant interest in designing estimators to achieve the Cramer Rao Lower Bound (CRLB). We focus on the complementary problem of choosing measurement points, thus our results apply to any estimator that approaches the CRLB in the limit.

Most active-tracking algorithms can be classified as locally optimal, gradient-ascent (e.g., work by Grocholsky et al. [5], [6] or Zhou and Roumeliotis [7]). Another suboptimal approach is track enumeration (e.g., Frew et al. [8], [9]). Zhou and Roumeliotis [7] considered constraints on the target motion to find a gradient ascent to minimize the trace of the target's covariance matrix. The work in [9] searches over the action space for a feasible sensor trajectory. These works do not bound the cost of the resulting trajectories.

A novel aspect of our formulation is that we optimize the trajectory of the robots with respect to the measurement cost and the distance traveled. Incorporating measurement time is relevant in a variety of real-world problem settings. For example, we previously used sensors to sample radio signal strength over one to two minutes to discern the bearing toward the target [2], [10], [11]. Another possible measurement cost is local maneuvers during a measurement. For example, Derenick et al. [12] used rotations in the robot chassis to construct bearing measurements to targets. Similarly, Forney et al. [13] use an S-shaped maneuver to resolve the direction to a target when using a hydrophone array. These maneuvers do not significantly change the robot-target configuration, but cost time and energy, a cost ignored by traditional active-localization literature.

A possible approach to the active-tracking problem is to formulate it as finding the optimal policy of a Markov Decision Process [14]. When the true target location is unknown and measurements are imprecise, the problem becomes finding the optimal policy for a Partially Observable Markov Decision Process (POMDP) [14], [15]. In our case, the state space would be given by the locations of their robots and their individual belief of the target space. The optimal strategy would choose measurement locations and a way for robots to exchange beliefs. Solving POMDPs in general is intractable, and we are not aware of dedicated methods to solve for the POMDPs which would arise from our setup.

The study of optimal, offline, active-localization algorithms using the Fisher Information Matrix dates to Hammel et al., [16], and has seen more recent results by Logothetis et al. [17], Bishop et al. [18], [19], and Martinez and Bullo [20]. Of these, only [16] considered time-constrained trajectories. However, the results were for a single robot with a continuous sensor, and are not directly applicable to the setup considered in the present work.

We also study the problem of including communication constraints among the robots. The problem of estimating the target state despite loss of connectivity has recently gained attention. Hollinger and Singh [21] considered the problem of re-establishing an estimation task after losing connectivity. In the same vein, Makarenko and Durrant-Whyte [22] studied estimation when connectivity was either enforced, or intermittent. This was similar to Leung et al. [23] who showed how to maintain a consistent estimate of a multi-robot system while relying on future reconnection. Spletzer and Taylor [24] studied the problem of assigning robots to targets while also enforcing network connectivity. In these works, the optimality of maintaining connectivity was assumed, but we provide an algorithm which may break connectivity between robots so they can reach better measurement locations, leading to quicker estimate convergence.

Our work is limited to bearing measurements. However, Bishop [18] proved that a solution to an optimal bearing sequence also applies (with minor modification) to range sensors, which are modelled to have range-dependent sensor noise.

IV. THE OPTIMAL OFFLINE ALGORITHM

In the offline problem, the true target location (x^*) is known. The goal is to design a minimum-cost measurement strategy S to satisfy the information requirements in Equation (4). We study the case with unbounded communication range before introducing communication constraints in Section IV-C. We start by discussing the structure of the matrix \mathbb{F} since the closed-form representation of \mathbb{F} is used to derive the optimal measurement sequences.




Fig. 2. The *target-local* coordinate frame. By expressing the measurement locations (black dots) with respect to the frame rotated by θ , \mathbb{F} is a diagonal matrix. The covariance ellipse's eigenvectors are aligned with the frame in which we express sensor locations.

Consider a measurement sequence S , and resulting Fisher Information Matrix, $\mathbb{F}(S)$. Define a coordinate frame, called the Target-Local (TL) frame, centered at x^* . Align the x axis of this frame with the eigenvector corresponding to the maximum eigenvalue of $\mathbb{F}^{-1}(S)$. In the TL frame, all sensor locations are specified in polar coordinates; the i^{th} measurement taken by the u^{th} robot is given as $s_{u,i} = (\alpha_{u,i}, r_{u,i})$. α is the angle formed with respect to the x axis, and r is the distance between the sensor location and x^* .

The TL coordinate frame is illustrated in Figure 2. In practice, a TL frame is obtained by applying a de-correlating transform, e.g., the Singular Value Decomposition or Eigen decomposition of \mathbb{F} [25]. The FIM has a convenient decomposition as the sum of all FIM from each individual observation as given below [3].

$$\begin{aligned} \mathbb{F}_{\text{TL}}(S) &= \sum_{u=1}^n \mathbb{F}_{\text{TL}}(S_u) = \sum_{u=1}^n \sum_{i=1}^{N_u} \mathbb{F}_{\text{TL}}(s_{u,i}) \\ &= R(\theta) \begin{bmatrix} \sum_{i=1}^N \frac{\sin^2(\alpha_i)}{r_i^2 \sigma^2} & 0 \\ 0 & \sum_{i=1}^N \frac{\cos^2(\alpha_i)}{r_i^2 \sigma^2} \end{bmatrix} R(\theta)^T \end{aligned} \quad (5)$$

The variable $N = |S|$ is the total number of measurements taken by all robots, and $R(\theta)$ is a transform that rotates coordinates to the TL frame from the world frame. In the TL coordinate frame, two useful properties of $\mathbb{F}(S)$ become evident.

First, the eigenvalues are simply the diagonal elements. Thus,

$$\lambda \mathbb{F}(S) = \sum_{i=1}^N \frac{\sin^2(\alpha_i)}{r_i^2 \sigma^2} \quad (6)$$

and

$$\bar{\lambda} \mathbb{F}(S) = \sum_{i=1}^N \frac{\cos^2(\alpha_i)}{r_i^2 \sigma^2}. \quad (7)$$

Note, if (6) is greater than (7), then the axes of the frame are switched so that (6) is less than (7). In general, the off-diagonal elements of $\mathbb{F}(S)$ are given by $\sum_{i=1}^N -\frac{\sin(2\alpha_i)}{r_i^2 \sigma^2}$. When $\mathbb{F}(S)$ is diagonalized this sum must equal 0, i.e.,

$$\sum_{i=1}^N -\frac{\sin(2\alpha_i)}{r_i^2 \sigma^2} = 0. \quad (8)$$

Second, the value of θ can be adjusted without affecting the eigenvalues, implying the following useful lemma.

Lemma 1. *All measurement locations can be rotated around the true target location without affecting the eigenvalues of \mathbb{F} .*

Proof. Changing the orientation of the world frame with respect to the covariance ellipse (θ in Equation 5) has no effect on eigenvalues since rotations are orthogonal transforms. \square

The remainder of the section is devoted to algorithms to find the optimal number of measurements and the correct assignment of robots to measurement locations.

A. Active Localization Using a Single Robot

In this section, we solve the special case of Problem 1 when $n = 1$. The derivation proceeds as follows. In Lemma 2, we show that the optimal one-robot trajectory has only two measurement locations. Lemma 3 establishes the optimal second location as a function of the first measurement location. The section ends by describing a method of searching for the optimal first measurement location.

Lemma 2 (Two Measurement Locations are Necessary and Sufficient). *There exists an optimal one-robot, offline, bearing-only measurement sequence consisting of exactly two measurement locations.*

Proof. First note that there must be at least two measurement locations to satisfy $\lambda \mathbb{F} > 0$. For contradiction, suppose not and consider Equation 8 with only one measurement location. Since $\sin(2\alpha) = 2 \sin(\alpha) \cos(\alpha)$, Equation 8 implies that one of Equation 6 or Equation 7 is equal to zero, contradicting the assumption that both eigenvalues are greater than zero.




Fig. 3. An illustration of Lemma 2. Three or more measurement locations are sub-optimal in the case of a single robot. Two measurement locations s_1 and s_2 can be moved closer together to produce a lesser-cost trajectory with the same information. This process can be repeated until the pair of measurement locations is collapsed to the same point.

To complete the proof, suppose there are three or more distinct measurement locations. Let $S = \{s_1, s_2, s_3\}$ be three consecutive measurements from the optimal trajectory. Consider the diagonalized \mathbb{F} resulting from the trajectory. Since $N \geq 3$, there is a pair of measurement locations (s_1, s_2) with either (i) $\alpha_1 \geq 0$ and $\alpha_2 \geq 0$ or (ii) $\alpha_1 < 0$ and $\alpha_2 < 0$. We will show that if three distinct measurement locations exist, the cost is not optimal. From Equation (6), we see

$$\lambda = \sum N_i \frac{\sin^2(\alpha_i)}{r_i^2 \sigma^2}, \quad (9)$$

where N_i is the number of measurements taken at location s_i . The locus of measurement locations yielding the same value for Equation (6) is defined by the circular contour of the form $r_i = C \sin \alpha_i$. See Figure 3.

If the pair is on the same contour, then we can reduce the sequence cost by taking $N_1 + N_2$ measurements from one location, with no affect on the information gains—a contradiction of the assumption of optimality of the original trajectory.

If not, then one measurement is “more informative” than another. Let $N_2 \frac{\sin^2(\alpha_2)}{r_2^2} \geq N_1 \frac{\sin^2(\alpha_1)}{r_1^2}$ (the proof for the opposite case is similar). This implies s_1 lies “inside” the circular contour of s_2 . Thus, s_2 could be moved closer to s_1 , increasing the information. By the triangle inequality, the cost of the path from previous measurement locations to s_1 then to s_2 is no longer than before. Since the sequence now has more information than required, s_1 can be moved closer to s_2 , producing a trajectory with lesser cost, but with the same information as the supposed optimal solution—a contradiction of the optimality of the original trajectory. \square

From both locations, there may be many measurements taken. For the moment, assume that the correct measurement counts N_1 and N_2 are known. We will later show how to search over possible values for N_1 and N_2 . The following lemma provides a relationship between the two measurement locations.

Lemma 3 (Structure of One-Robot Trajectory). *There exists an optimal solution with the first measurement location on the line s_0x^* . Furthermore, given the first measurement location, s_0 , and measurement counts for the two measurement locations, N_1 and N_2 , the second measurement location must satisfy,*

$$\sin^2 \alpha_2 = \lambda_d \frac{r_2^2 \sigma^2}{N_2} + \lambda_d \frac{r_1^2 \sigma^2}{N_1} - \lambda_d \frac{r_1^2 r_2^2 \sigma^4}{N_1 N_2} \quad (10)$$

and

$$r_2^2 = N_2 \frac{N_1 \sin^2 \alpha_2 - \lambda_d \sigma^2 r_1^2}{N_1 \lambda_d - \lambda_d^2 \sigma^4 r_1^2}. \quad (11)$$

Proof. Given a starting location s_0 , the goal is to find the optimal subsequent measurement locations, s_1 and s_2 . Note that there exists an optimal algorithm with the first measurement location, s_1 placed on the line between s_0 and x^* . If it was not, rotating both s_1 and s_2 (adjusting θ as stated in Lemma 1) would reduce the time to travel between s_0 and s_1 without affecting the eigenvalues or cost to visit s_2 from s_1 .

The eigenvalues of the resulting FIM can be found using the quadratic formula. For any 2×2 matrix A with trace $\text{tr}(A)$ and determinant $\det(A)$, the eigenvalues λ satisfy

$$\lambda(A) = \frac{1}{2} \text{tr}(A) \pm \frac{1}{2} \sqrt{\text{tr}(A)^2 - 4 \cdot \det(A)}. \quad (12)$$

When considering the optimal sequence, both $\text{tr}(\mathbb{F})$ and $\det(\mathbb{F})$ are positive. Fix the x axis of the coordinate frame to the line s_0x^* . Then $\alpha_1 = 0$, and it is possible to solve for α_2 and r_2 using the fact that

$$\det(\mathbb{F}) = \sum_{i=1}^{N_1 \cdot N_2} \frac{\sin^2 \alpha_2}{r_1^2 r_2^2 \sigma^4} \quad (\text{c.f. Equation. 6 [26]}). \quad (13)$$

Since $\underline{\lambda}$ is equal to λ_d , the desired information, solving the previous yields

$$\sin^2 \alpha_2 = \lambda_d \frac{r_2^2 \sigma^2}{N_2} + \lambda_d \frac{r_1^2 \sigma^2}{N_1} - \lambda_d \frac{r_1^2 r_2^2 \sigma^4}{N_1 N_2} \quad (14)$$

and

$$r_2^2 = N_2 \frac{N_1 \sin^2 \alpha_2 - \lambda_d \sigma^2 r_1^2}{N_1 \lambda_d - \lambda_d^2 \sigma^4 r_1^2}. \quad (15)$$

\square

The values of r_2 and α_2 from Lemma 3 describe a curve as shown in Figure 4 (for differing values of r_1 , and λ_d).

The optimal second measurement location is the closest point on the resulting curve described by Equations 14. Because the curve




Fig. 4. Examples of the contours described by Equation (14) and (15) for a target at $(0,0)$ (shown as a star), $r_1 \in \{1, 1.5, 2\}$ and λ_d set to .01 (black lines of lesser curvature) or .05 (blue). σ_s was 1 and N_1 was 1. For example, if the robot first travels to the position $s_1 = (2,0)$ and takes one measurement then it can take all remaining measurements from anywhere on the dashed curve (e.g., s_2 as labelled) to satisfy $\underline{\lambda}(\mathbb{F}) = \lambda_d$.




Fig. 5. These optimal one-robot offline trajectories corresponding to values of $\lambda_d \in \{1, 2, 4, 6, 12\}$. Note the robot begins at location $(5,0)$ and moves along the x -axis to the first marked location.

is convex (for a given r_1 , N_1 , and N_2) there exists a unique point on the curve closest to the first measurement location. Finding the optimal trajectory reduces to searching for the optimal range, r_1 , and measurement counts, N_1 and N_2 .

The cost of the sequence as a function of the two measurement locations is

$$C_{\text{one-robot}} = t_m(N_1 + N_2) + d(s_0, x^*) - r_1 + \sqrt{(r_2 \sin \alpha_2)^2 + (r_1 - r_2 \cos \alpha_2)^2}. \quad (16)$$

Minimizing this cost over N_1 and N_2 can be done by enumeration using a table of size $N \times N$, with $N = N_1 + N_2$. The table size is bounded since $N \leq \frac{d(s_0, x^*)}{t_m} + 2$ and N is a positive integer. If $N > \frac{d(s_0, x^*)}{t_m}$, the robot spends more time measuring than would be required to travel to the true target location. From close to x^* , any two measurements that are not collinear with the target are sufficient to achieve any information objective since $r_1 = r_2 \approx 0$ (Equation (6)). Each entry in the table corresponds to a minimization of Equation (16) over r_1 , which is accomplished using finite-difference methods. The one-robot trajectories shown in Figure 5 were calculated for values of $\lambda_d \in \{1, 2, 4, 6, 12\}$.

To summarize, the optimal one-robot trajectory can be found using Algorithm 1. For brevity, we have omitted boundary checking (e.g., $0 < r_1 \leq d(s_0, x^*)$) and initialization. As stated, the minimization on Line 7 is done using a finite difference approximation to gradient

cost to travel to the boundary of the circle of radius r_λ and take N^* measurements.

$$C = \sqrt{d(s_0, x^*)^2 + \frac{N^*}{2\lambda_d\sigma^2}} - \sqrt{\frac{N^*}{2\lambda_d\sigma^2}} + N^*t_m \quad (24)$$

To complete the optimization, a table of size $1 \times \lceil \frac{d(s_0, x^*)}{t_m} \rceil$ is used to search for the optimal value of N^* . The i^{th} cell of the table represents the evaluation of Equation (24) with $N^* = i$. The index of the cell containing the minimum value is the optimal number of measurements for a single robot. With the number of measurements solved, Equation (23) can be used to find the trajectories. The result is illustrated in Figure 8. The process described in this section is formalized in Algorithm 2.

Algorithm 2 Two Robot Offline Solution

Input: $\sigma, s_0, x^*, \lambda_d$

Output: s_u, s_v, N

- 1: $N_{max} \leftarrow \frac{d(s_0, x^*)}{t_m} + 2$
 - 2: $M \leftarrow N_{max} \times 1$ vector
 - 3: **for** each row i of M **do**
 - 4: $N \leftarrow i$
 - 5: $C_i \leftarrow$ Equation (24)
 - 6: **end for**
 - 7: $N \leftarrow \text{argmin } C$
 - 8: $r_\lambda = \frac{1}{2} \sqrt{\frac{2N}{\lambda_d\sigma^2}}$
 - 9: $s_u \leftarrow$ closest point on circle of radius r_λ centered at $(0, r_\lambda)$
 - 10: $s_v \leftarrow$ closest point on circle of radius r_λ centered at $(0, -r_\lambda)$
-

The proposed algorithm produces the optimal two-robot trajectory. For completeness, it can be shown that this strategy is optimal for k pairs of robots as well.

Theorem 1 (Optimality of Algorithm 2 for n Robots). *Let there be $n = 2k$ robots for some positive integer k . Computing the optimal two robot measurement strategy using sensor noise, $\sigma' = \sigma\sqrt{\frac{1}{k}}$ produces the optimal n robot measurement strategy.*

Proof. It must be shown that there exists an optimal symmetric n robot strategy to generalize the two-robot algorithm. Let S^* be an optimal set of n trajectories, one for each robot. Similar to Lemma 4, we can mirror the trajectories and choose the “most informative” of the n pairs of trajectories as follows. Recall for each pair, S_u and S'_u , $\alpha_u = -\alpha'_u$, $r_u = r'_u$, $N_u = N'_u$, and the FIM produced by the pair satisfies $\lambda(\mathbb{F}(S_u) + \mathbb{F}(S'_u)) = 2N_u \frac{\sin^2 \alpha_u}{r_u^2 \sigma^2}$. At least one of the n pairs satisfies $N \frac{\sin^2 \alpha}{r^2 \sigma^2} \geq \frac{\lambda_d}{n}$ since the summation of information from all pairs of trajectories satisfies $\sum_{i=1}^n 2N_i \frac{\sin^2 \alpha_i}{r_i^2 \sigma^2} = 2\lambda_d$. Thus, any optimal n robot trajectory has identical cost and information gains as a symmetric n robot trajectory when n is even.

Let S_u and S'_u be the pair of trajectories selected in the previous step. Now, repeat the steps of Corollary 1 to collapse the set of measurement locations for down to two: one for S_u and another for S'_u .

Since the optimal strategy consists of two symmetric paths with one measurement location, as before, we can solve for only one of them in closed form. In this case, the path derived will be travelled by $k = \frac{n}{2}$ robots, however.

To calculate the $\frac{n}{2}$ robot optimal trajectory, simply repeat the steps of the two-robot algorithm, but notice that each “measurement” is actually $\frac{n}{2}$ robots measuring simultaneously. Thus, each measurement produces a factor $\frac{n}{2}$ more information, which is equivalent to scaling down the variance of the sensor noise by the same factor. \square

We now move on to the case when communication among all the robots is required to form a final estimate of the target location.

C. With Distance-Constrained Communications

In this section, we describe an extension to the previous algorithm for the case when the robots have limited communication range. First consider the case of two robots that must be within distance r_c to communicate. A natural strategy is simply to execute the optimal unconstrained algorithm (e.g., Algorithm 2), then have all robots move towards the centroid of the robots’ positions until communication is possible among all robots. Figure 9 illustrates this strategy as the solid line. However, it may be more time-efficient to simply move to s'_u , which places the robots in communication range during measurements (as illustrated by the dotted line). However, the second option requires the robots to travel further before taking the same number of measurements. Algorithm 3 expands Algorithm 2 to incorporate this tradeoff.

While simple and easy to use in practice, it is not clear if Algorithm 3 is always optimal or if it extends to arbitrary numbers of robots. For example, when many robots are used, it may be more cost effective to form a long “chain” of robots, allowing the ends of the chain to spread out to informative locations, while the middle robots periodically establish links between the distant robots. However, in Theorem 2, we will show that no other strategy can do significantly better than the symmetric strategy given in Algorithm 3. The result of Theorem 2 will allow derivation of a near-optimal, online, and communication-adaptive algorithm in the next section. Finally, in light of Theorem 1, Algorithm 3 is useful for any number of pairs of robots.

Assuming α_u , N_u , and r_u can be found, expansion of Equation (24) produces the following equation.

$$\sqrt{d(s_0, x^*)^2 + r_u^2 - 2d(s_0, x^*)r_u \cos(\alpha_u)} + N_u t_m + \min(r_u \sin(\alpha_u) - \frac{1}{2}r_c, 0) \quad (25)$$

Using this new cost function, the previous two-robot algorithm changes to the following.

Algorithm 3 Two-Robot Communication-Constrained

Input: $\sigma, s_0, x^*, \lambda_d, r_c$

Output: s_u, s_v, N

- 1: $N_{max} \leftarrow \frac{d(s_0, x^*)}{t_m} + 2$
 - 2: **for** $i \in [1, N_{max}]$ **do**
 - 3: $N \leftarrow i$
 - 4: $A_i \leftarrow \text{argmin}_\alpha$ Equation (25)
 - 5: $C_i \leftarrow$ evaluate Equation (25) with A_i
 - 6: **end for**
 - 7: $i^* \leftarrow \text{argmin } C$
 - 8: $N \leftarrow i^*$
 - 9: $\alpha \leftarrow A_{i^*}$
 - 10: $r_\lambda = \frac{1}{2} \sqrt{\frac{2N}{\lambda_d\sigma^2}}$
 - 11: $s_u \leftarrow (r_\lambda \sin(\frac{\pi}{2} - \alpha), r_\lambda \cos(\frac{\pi}{2} - \alpha))$
 - 12: $s_v \leftarrow (r_\lambda \sin(\frac{\pi}{2} - \alpha), -r_\lambda \cos(\frac{\pi}{2} - \alpha))$
-

Note that as $r_c \rightarrow \infty$, the output matches the result from the previous section. The next result shows that Algorithm 3 is close to the optimal cost.

Theorem 2 (Algorithm 3 is a Two Approximation). *Algorithm 3 produces a measurement strategy of cost less than twice that of the optimal communication-constrained measurement strategy.*

Proof. Let C be the cost of Algorithm 3 when $r_c \rightarrow \infty$. By Theorem 1, C is the optimal cost for even numbers of robots. Let




Fig. 7. An illustration of the constraints on the measurement sequences. (a) by Equation 6, (b) by Equation 7, and (c), the intersection. Note the radius in (b) is never greater than the radius in (a). Assuming the other measurements are placed, the last measurement U , must fall in the regions specified, while traveling the least. The dotted line adjoining U and the starting location illustrates the shortest path.




Fig. 8. Optimal two-robot trajectories for various system parameters. In all figures the robots started at location $(3,0)$ and the true target was at $(0,0)$. Communication constraints were only considered in the top right figure and all other parameters were held fixed. Top-left: $\lambda_d \in \{.1, 1, 3, 8\}$. Top-right: $r_c \in \{2.1, 1.1, .1\}$. When $r_c = .1$, the robots rendezvous after measuring. When $r_c = 2.1$ the output is the same as the result from unbounded r_c . Bottom-right: $t_m \in \{.01, .1, 1\}$ Note that as t_m increases, the optimal algorithm travels to more informative locations so that fewer measurements are required. Bottom-left: $\sigma \in \{.1, 1, 4, 8\}$. Note that changing the sensor noise produces the same effect as requiring more information (compare left two figures).

C_r^* be the cost of the optimal strategy for any communication radius r and the same number of robots. Then $C \leq C_r^*$ since adding communication constraints can only increase the cost of the strategy. Let R be the cost for all robots to rendezvous after taking the measurements defined by Algorithm 3. Clearly $R < C$ since the cost to rendezvous is at most the cost to move back to the starting location. Then $C \leq C_r^* \leq C + R$ implies that $C_r^* < 2 \cdot C$. Thus, Algorithm 3 produces a measurement strategy of cost at most twice that of the optimal communication-constrained measurement strategy. \square

D. Discussion

The offline algorithms presented here require a repeated minimization of convex functions (See Line 7 in Algorithm 3). Interestingly, due to the symmetry of the offline, multi-robot optimal solution only one measurement location must be found regardless of the number of




Fig. 9. An illustration of two different choices for communication-constrained measurement locations. s_u and s_v represent the output of Algorithm 2. From these locations, the robots can move directly toward each other to communicate. Alternatively, they can move to s'_u and s'_v , and remain in communication during measurement. Algorithm 3 finds the optimal placement of symmetric measurements to minimize the cost while including the rendezvous cost.

robots. Thus, the computational complexity of optimizing N robot trajectories is the same regardless of N . However, up to $\frac{d(s_0, x^*)}{t_m}$ solutions must be evaluated (See Line 1 in Algorithm 2). For our application, $t_m \approx 100$ and the time to travel to the true target is only a few minutes. Thus, we typically evaluate less than 20 solutions before finding the optimal. Furthermore, Equation 24 is convex in N^* and thus newton-like methods could be applied for further speedup if necessary.

In an online setting, the true target location is not known, but the objective is the same: For any true target location, the measurement sequence must satisfy $\underline{\lambda F}(S) > \lambda_d$. Furthermore, the PDF representing the distribution of possible target locations changes with each measurement.

To remain competitive with the optimal case, an online measurement strategy should not allocate too much time moving towards a distant target estimate if it is highly likely to change given a few measurements. In the next section, we rely on the fact that Algorithm 3 is simple to implement and build an adaptive algorithm for real-world localization problems.

V. ONLINE ALGORITHM

In this section, we present a conversion of the offline strategy presented in the previous section into a near-optimal online strategy. In the online setting, algorithms do not have access to the true target location x^* . Instead, we start with an estimate in the form of a prior PDF. It is important to note that the goal is still to find the required amount of information with respect to the true target location irrespective of the initial estimate.

Given an estimate of x^* , one possible extension to the offline algorithm is to choose the most likely point in the PDF to be x^* , then execute the offline algorithm with respect to this point. For example, if the target estimate is a two-dimensional Gaussian, then the most likely point is the mean, \hat{x} . However, the true target location may be close to the robots' initial location, while \hat{x} may not, resulting in much more work than is necessary.

Instead, consider x , the *closest* point of the current PDF with "high enough" probability. For example, x could be the closest point lying within the 3-sigma bound of a two-dimensional Gaussian, which accounts for 99.7 percent of the probability mass. Given the high likelihood of the true target being inside the 3-sigma bound, it is "safe" to assume the optimal algorithm must travel at *least* to the closest point within this 3-sigma bound. The 3-sigma bound of a 2D Gaussian distribution is an ellipse, but in general, any convex shape containing the desired amount of probability mass can be used.

More formally, we adapt the previously discussed offline algorithms to an online problem setting as follows. Let the *offline* two-robot algorithm be described by the function $\mathcal{A}(s_0, x^*, r_c, n, \lambda_d, t_m)$ with cost $C(s_0, x^*, r_c, n, \lambda_d, t_m)$ for n robots starting at location s_0 with communication range r_c and measurement time t_m .

At each step, i , form a convex shape R_i containing the desired probability mass (e.g., the 3-sigma bounds of our Gaussian prior). Then, locate the point in the interior of the shape closest to the robot's starting location, label it x , execute $\mathcal{A}(s_0, x, r_c, n, \lambda_d, t_m)$, and pay cost $C(s_0, x, r_c, n, \lambda_d, t_m)$. The gathered measurements are used to update the hypothesis, a new x is selected from the posterior PDF, and the centroid of the robots' positions is assumed to be the robots' starting location. For example, Algorithm 3 satisfies all the requirements of the offline algorithm (\mathcal{A}) as discussed in the previous section.

We call such an algorithm MULTI-STEP, and illustrate the steps for a Gaussian target estimate in Algorithm 4. The remainder of this section shows that using repeated calls to \mathcal{A} produces near-optimal costs with high probability. Specifically, the cost is at most

a logarithmic factor worse than optimal assuming the true target is contained in all convex regions chosen at each time step.

In practice, we use a filtering algorithm to update the PDF after the robots take measurements of the bearing to x^* . We require some technical assumptions about the starting locations of the robots and the filtering method used. As mentioned, at each step of the algorithm, we use a region R_i to contain the possible locations for x^* . In Lemmas 5 and 6, as well as Theorem 3, we require that the robots begin outside the region and that the true target will fall within the region with probability $1 - \epsilon$ which is independent of the choice of region (i.e., $P(x^* \in R_i)$ is independent from $P(x^* \in R_j)$ for all $i \neq j$). The methods used in [11] ensure the robots begin outside a suitable region. The second condition is satisfied if the filter used is *consistent* and the last requirement is satisfied if each estimate of the target location is conditioned only on the measurements received and the measurements have independent noise [4]. In practice, we employ a batch-processed, maximum likelihood estimator [4], [28].

The rest of the section proceeds as follows. First, Lemma 5 shows the cost of each *individual* invocation is bounded. Then, Lemma 6 proves an upper bound on the number of calls to \mathcal{A} required to localize any target to required precision. This will produce the bound presented in Theorem 3.

Algorithm 4 MULTI-STEP($s_0, \hat{x}(0), \Sigma(0), r_c, \lambda_d, t_m, n$)

```

 $\Sigma(i) \leftarrow \Sigma(0)$ 
 $\hat{x}(i) \leftarrow \hat{x}(0)$ 
while  $\bar{\lambda}(\Sigma(i)) > \frac{1}{\lambda_d}$  do
   $R_i \leftarrow$  circle of radius  $3 \cdot \sqrt{\bar{\lambda}(\Sigma(i))}$  at point  $\hat{x}(i)$ 
   $x_d \leftarrow$  closest point on  $R_i$ 
   $s_{u,i}, s_{v,i} \leftarrow \mathcal{A}(s_i, x_d, r_c, n, \lambda_d, t_m)$ 
   $Z \leftarrow$  Collect measurements from  $s_u$  and  $s_v$ .
   $\hat{x}(i) \leftarrow$  Update target estimate using  $Z$ 
   $s_i \leftarrow$  centroid of  $s_{u,i}$  and  $s_{v,i}$ 
end while

```

Lemma 5 (Bounded Subroutine Cost). *Let \mathcal{R} be a convex region such that $P(x^* \in \mathcal{R} | Z) = 1 - \epsilon$, where Z is the measurements obtained by all previous steps. Let x be the closest point inside (or on the boundary of) \mathcal{R} to the robots' starting location, s_0 . Let $C(x)$ be the cost of the call to $\mathcal{A}(s_0, x)$. Let \mathcal{A} be a γ -approximation of optimal: the optimal offline cost C^* satisfies $C(s_0, x^*) \leq \gamma C^*$. Then the cost, $C(s_0, x) \leq \gamma C^*$ with probability equal to $1 - \epsilon$.*

Proof. By definition, x is the closest point in \mathcal{R} to the robots' starting location, and so $d(s_0, x^*) \geq d(s_0, x)$. We can prove that $C(x) \leq C(x^*)$ by contradiction. If $C(x^*) < C(x)$ we could take all the measurements taken with respect to x^* and place them in the same configuration around x . Now, both invocations spend the same time measuring. However, travelling to sensor locations with respect to x would take no more time than travelling to sensor locations near x^* , since $d(s_0, x) \leq d(s_0, x^*)$ by assumption. This contradicts the assumption that $C(x^*) < C(x)$.

Since $C(x) \leq C(x^*)$ and $C(x^*) \leq \gamma C^*$, it follows that $C(x) \leq \gamma C^*$. \square

Thus, the cost of each invocation of \mathcal{A} is bounded. However, it may take arbitrarily many calls to \mathcal{A} to reduce the uncertainty adequately, producing an unbounded cost compared to the optimal algorithm. The following result shows the number of calls is small compared to optimal, assuming the regions R_i contain the true target.

Intuitively, we will establish that all measurements are taken inside the K -sigma bound of the *true* covariance, with K possibly large, but bounded. The true covariance (CRLB) is proportional to $\sqrt{\frac{1}{\lambda F}}$ at




Fig. 10. An illustration of the two-robot MULTI-STEP algorithm (solid paths) and the optimal \mathcal{A} algorithm (dashed paths). The robots begin at location s_0 . The optimal choice is to move directly to x^* , but we are unsure of the location of x^* . Thus, we form a convex region R such that the probability x^* is in R is high. At each step of the MULTI-STEP algorithm, we do less work than $\mathcal{A}(s_0, x^*)$ by Lemma 5. Since the number of \mathcal{A} executions is bounded by Lemma 6, we do not do significantly more work than optimal by Theorem 3.

each measurement step. We will show the information gained at each step is inversely proportional to the range from the true target, and the range is inversely proportional to the *current* information. Thus, the information gained by each measurement is proportional to the sum of the information already obtained. That is, each measurement step produces a constant-factor increase in $\underline{\lambda}\mathbb{F}$ producing quick convergence to λ_d .

Lemma 6 (Number of Calls to \mathcal{A}). *The MULTI-STEP algorithm requires $\mathcal{O}(\log_b \lambda_d)$ calls to the optimal offline algorithm, where $b = \mathcal{O}(1 + \frac{1}{\sigma^2})$.*

Proof. In what follows, S is the set of all measurement locations visited during an execution of MULTI-STEP, and S_i is the measurement locations chosen by the i^{th} call to \mathcal{A} . The FIM of all previous measurement steps up to and including the i^{th} step is denoted \mathbb{F}_i . Let MULTI-STEP use T calls to \mathcal{A} . Then $\underline{\lambda}\mathbb{F}_T = \underline{\lambda}\mathbb{F}(S, x^*) \geq \lambda_d$ (we drop the x^* in further analysis).

Recall from Equation 5 that

$$\mathbb{F}_T = \mathbb{F}(S) = \sum_{i=1}^T \mathbb{F}(S_i). \quad (26)$$

By Weyl's theorem (Section 6.7 [27]),

$$\underline{\lambda}\mathbb{F}_T = \underline{\lambda}\mathbb{F}(S) \geq \sum_{i=1}^T \underline{\lambda}\mathbb{F}(S_i). \quad (27)$$

The right hand side of Equation 27 represents the information gains from each of the T measurement steps. Now consider the i^{th} measurement step and the resulting \mathbb{F} .

$$\underline{\lambda}\mathbb{F}_i \geq \underline{\lambda}\mathbb{F}(S_i) + \underline{\lambda}\mathbb{F}_{i-1} \quad (28)$$

$$\geq \sum_{j=1}^N \frac{\sin^2 \alpha_j}{r_j^2 \sigma^2} + \underline{\lambda}\mathbb{F}_{i-1} \quad (29)$$

where the second equation follows from Equation 5 expanded for all N measurements taken during the i^{th} step. From Lemma 2 we know that $N \geq 2$ measurements are taken.

Let R_i be the i^{th} region chosen by MULTI-STEP. We will assume that R_i contains the true target location and is convex. Since each region is convex, and the measurement locations chosen by \mathcal{A} are symmetric about the line $\overline{s_0 x}$ (Figure 10), there is no point in the intersection of all the regions which is collinear with the measurement

locations chosen. So, for all of the T sets of measurement locations chosen by MULTI-STEP, we have

$$\underline{\lambda}\mathbb{F}(S_i) = \sum_{j=1}^N \frac{\sin^2 \alpha_j}{r_j^2 \sigma^2} > 0 \quad (30)$$

which implies that each step provides non-zero progress of $\underline{\lambda}\mathbb{F}_i$ to λ_d .

To lower-bound the rate of convergence let

$$K = \max_i r_i \sqrt{\underline{\lambda}\mathbb{F}_i}. \quad (31)$$

By definition, for all steps i , $r_i \leq \frac{K}{\sqrt{\underline{\lambda}\mathbb{F}_i}}$. In light of Equation 30 and substituting Equation 31 into Equation 29 we see

$$\begin{aligned} \underline{\lambda}\mathbb{F}_i &\geq \underline{\lambda}\mathbb{F}(S_i) + \underline{\lambda}\mathbb{F}_{i-1} \\ &\geq \sum_{j=1}^N \frac{\sin^2 \alpha_j}{K^2 \sigma^2} \underline{\lambda}\mathbb{F}_i + \underline{\lambda}\mathbb{F}_{i-1} \\ &\geq \frac{\beta}{K^2 \sigma^2} \underline{\lambda}\mathbb{F}_{i-1} + \underline{\lambda}\mathbb{F}_{i-1} \\ &\geq \left(1 + \frac{\beta}{K^2 \sigma^2}\right) \underline{\lambda}\mathbb{F}_{i-1} \end{aligned}$$

where for brevity β was chosen to be $\min_i \sum \sin^2 \alpha_i$.

The previous implies at least a constant factor increase in information about the true target location (and corresponding decrease in the CRLB) at every time step. If MULTI-STEP makes $N \geq 1$ calls to \mathcal{A} , then

$$\begin{aligned} \underline{\lambda}\mathbb{F}_N &\geq \underline{\lambda}\mathbb{F}_0 \left(1 + \frac{\beta}{K^2 \sigma^2}\right)^N \\ N &\leq \log_b(\underline{\lambda}\mathbb{F}_N) - \log_b(\underline{\lambda}\mathbb{F}_0) \end{aligned} \quad (32)$$

with $b = \left(1 + \frac{\beta}{K^2 \sigma^2}\right) > 1$. \square

It is worth noting that the previous lemma establishes the cost of the MULTI-STEP algorithm as a function of the desired *uncertainty*, rather than the range to the true target or other uncontrollable variables. Our final result follows: We show that the cost of using MULTI-STEP(\hat{x}) is less than a constant factor worse than the optimal offline algorithm, $\mathcal{A}(x^*)$ with high probability.

Theorem 3 (Cost Bounds). *With probability, $(1 - \epsilon)^{\log_b \lambda_d - \log_b \underline{\lambda}\mathbb{F}_0}$, the ratio of the cost of the MULTI-STEP algorithm to the optimal offline algorithm satisfies $\frac{\text{MULTI-STEP}(\hat{x}(0))}{\mathcal{A}(x^*)} = \mathcal{O}(\log_b \lambda_d - \log_b \underline{\lambda}\mathbb{F}_0)$, where $b = \mathcal{O}(1 + \frac{1}{\sigma^2})$, $\underline{\lambda}\mathbb{F}_0$ is the "prior" information (if available), λ_d is the desired information, and $0 < \epsilon < 1$.*

Proof. By Lemma 6 MULTI-STEP makes $\mathcal{O}(\log \lambda_d - \log \underline{\lambda}\mathbb{F}(0))$ calls to \mathcal{A} , and by Lemma 5 each of these costs is less than a scalar multiple of the optimal cost. Thus, the first result follows.

When the regions are selected to independently contain the true target location with probability $1 - \epsilon$, the probability all regions contain the target is $(1 - \epsilon)^N$, for N regions. Given the value of N from Lemma 6, the probability follows as stated. \square

In this section we have shown that an optimal *offline* algorithm can be converted to an online algorithm by carefully selecting a conservative (nearby) point to serve as a proxy for the true target location. The method is general to any offline optimal algorithm or any filtering method, provided an independent convex region can be described which contains the true target with high probability. In the next section we verify the results of Theorem 3 in simulations before testing the MULTI-STEP algorithm in field experiments.

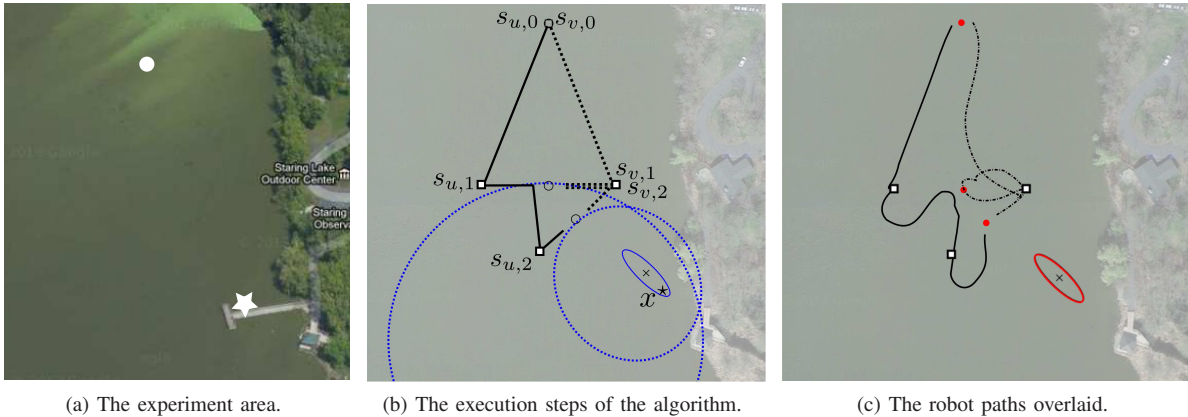


Fig. 11. Experiment results from Lake Staring, Minnesota, USA. **(a)**: The experiment area. The true target (and camera in [29]) were placed on the docks near the bottom right corner (labelled with a star). The robots began near the top-middle (circle). **(b)**: The two calls to Algorithm 3 produced the dark paths shown, and reduced the uncertainty (the blue circles). The final actual uncertainty was the solid ellipse. **(c)**: More execution details. The solid red circles are the points where the robots exchanged information. The figure covers an area approximately 200m vertically by 150m horizontally.

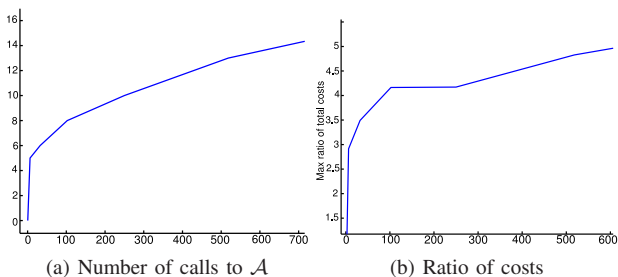


Fig. 12. The aggregate results of numerical studies. Left: the number of calls to \mathcal{A} as a function of $\frac{\lambda_d}{\lambda_{F_0}}$. Right: the ratio of costs of online algorithm to optimal offline algorithm, as a function of $\frac{\lambda_d}{\lambda_{F_0}}$. Shown is the maximum value encountered during simulations.

VI. IMPLEMENTATIONS AND EXPERIMENTS

We now explore the results of Theorem 3 through simulations and experiments. Our goal is to verify the logarithmic behavior of the bound presented in the previous section and test the effectiveness of the algorithm in locating radio-tags in real-world environments.

A. Simulations

We explore the effects of the ratio of prior information to desired precision through simulations. In each simulation, we provide a prior estimate of the target location with circular covariance and eigenvalues $\frac{1}{\lambda_{F_0}}$ and execute the MULTI-STEP algorithm until the uncertainty converges to the desired $\frac{1}{\lambda_d}$.

We repeatedly test the performance ratio by sampling a true target location from the prior PDF and executing Algorithm 3 using the true target location, and MULTI-STEP on the hypothesis. To give a real-world sense of scale to the simulations, note our choice represents a starting hypothesis (three sigma-bound) which grows to encompass a 254 square kilometer area, while requiring a final estimate which is as accurate as a commercial GPS fix (i.e., a few meter uncertainty).

The results are presented in Figure 12. The x -axis of the figures shows the ratio $\frac{\lambda_d}{\lambda_{F_0}}$ (desired gain in information). The ratio of prior uncertainty to final uncertainty is the inverse of this. According to Theorem 3, we expect the curve to roughly match $\log \lambda_d$.

First, we present the actual number of calls the MULTI-STEP algorithm makes to the \mathcal{A} subroutine. We notice a logarithmic trend to the ratios, as expected. These results are shown in Figure 12(a).

The number of calls is not necessarily reflective of the ratio of the costs between our online algorithm and the optimal offline algorithm, since the total distance traveled for each call will decrease.

To explore this, we present the ratio of the *actual* cost in Figure 12(b) from the same trials. The actual cost is given by the maximum distance traveled plus the maximum time spent measuring. We expect the cost ratio to significantly change, depending on the relative positions of the true target location, hypothesis location and uncertainty, and the starting robot positions. In Figure 12(b) we present the worst-case ratio of costs encountered during simulations for each prior hypothesis. Interestingly, the worst-case ratio of costs was less than 7 in these trials, which was less than the number of calls to the \mathcal{A} subroutine. This suggests in practice the cost of our online algorithm could be closer to the optimal offline algorithm than is suggested by the theoretical results since each subsequent invocation seems to cost less than the previous one.

B. Field Experiments

As described previously, we are building a working multi-robot system to search for invasive fish in lakes. To test the suitability of the algorithm to real-world conditions, we have implemented the algorithm for field trials on lakes in Minnesota, USA. We report five field experiments which show the feasibility of the algorithm in practice. The robots used were OceanScience QBoats, pictured in Figure 1. Although designed for remote operation, the boats were augmented with on-board laptops and motor control boards for autonomous navigation, and pan-tilt servos, antennas, and real-time spectral analyzers to produce bearing measurements. They are 2 meters in length and have an average speed of 1 meter per second. The robots were used to localize a low-power radio transmitter which is typical of models used by field biologists when tracking carp. The tags transmit an uncoded pulse approximately once per second and our system can detect a radio tag from 100 meters on average. Due to the low transmission rate, it can take up to one minute to construct a reliable bearing measurement by rotating the antenna while sampling the signal strength.

The experiment was run in Lake Staring, Minnesota, USA (shown in Figure 11) in 2012 and 2013. The transmitting tag was deployed at a known location in the environment, and the robots executed the MULTI-STEP algorithm. In our prior work on this system, we described a method for reliably constructing a consistent, bounded-uncertainty prior estimate of the target location [11], and so we assume a prior is available during the bearing-only localization phase.

The boats began 140 meters from the target and executed the algorithm given in Section V. To exchange measurements the robots used an ad-hoc wireless network. An example of the final result is shown in Figure 11(c), as a blue square. We show the algorithm steps in Figure 11(b), and the actual robot paths in Figure 11(c). After each measurement, the boats transmitted measurement values over the wireless network. We used a low power network which could not communicate more than 10-20 meters reliably. In each of the five experiments, the robots traveled a combined distance of one kilometer and localized the target to within 10 meters of its measured location. Note the position of the tag's location was accurate to within 5 meters due to GPS error. In all cases, the localization took the expected two steps to locate the target. The final expected error (distance of the final mean of the estimate from the true target mean) was 11.2, 7.1, 1.3, 10.1, and 23.9 meters across the five trials. We provide a video of the localization process at [29].

VII. CONCLUSIONS

In this paper, we make three contributions to the problem of active localization of static targets with mobile bearing sensors. We started by presenting the first derivation of the optimal deployment of mobile bearing sensors with respect to a known target location. We then extended this result to provide near-optimal deployment strategies when subjected to limited communication range. Next, we used the insights provided by these results to develop an online active localization strategy suitable for field deployments. We showed a theoretical upper bound of a logarithmic approximation of the optimal strategy. We verified the theoretical bounds through simulations studies and presented a working field implementation in our intended operating environment. In field trials, two communicating robots were able to repeatedly locate a radio transmitter.

A clear next step for this work is locating mobile targets. If the target loiters in a small region, we expect that the motion of the target will have negligible affect on the cost to localize.

Our future work will also focus on tightening the logarithmic approximation of our proposed online algorithm. We expect that a constant factor approximation is possible. Another avenue for future work is reducing the number of communication steps in the online algorithm. Design and analysis of active localization algorithms for multiple robots when multiple targets are nearby is also an important avenue for future work.

ACKNOWLEDGMENT

This work is supported by NSF Awards #1111638, #0916209, #0917676, #0936710. Josh gratefully acknowledges ongoing support from the Minnesota Chapter of the ARCS Foundation.

REFERENCES

- [1] P. Tokekar, D. Bhaduria, A. Studenski, and V. Isler, "A Robotic System for Monitoring Carp in Minnesota Lakes," *Journal of Field Robotics*, vol. 27, no. 6, pp. 779–789, 2010.
- [2] P. Tokekar, E. Branson, J. Vander Hook, and V. Isler, "Tracking Aquatic Invaders: Autonomous Robots for Monitoring Invasive Fish," *IEEE Robotics and Automation Magazine*, vol. 20, no. 3, pp. 33–41, 2013.
- [3] J. Vander Hook, P. Tokekar, and V. Isler, "Cautious Greedy Strategy for Bearing-Only Active Localization: Analysis and Field Experiments," *Journal of Field Robotics*, 2013, under review.
- [4] Y. Bar-Shalom, X.-R. Li, and T. Kirubarajan, *Estimation with Applications to Tracking and Navigation*. New York, USA: John Wiley & Sons, Inc., 2001.
- [5] B. Grocholsky, "Information-theoretic control of multiple sensor platforms," Ph.D. dissertation, University of Sydney. School of Aerospace, Mechanical and Mechatronic Engineering, 2006.
- [6] B. Grocholsky, A. Makarenko, and H. Durrant-Whyte, "Information-theoretic coordinated control of multiple sensor platforms," in *Robotics and Automation, 2003. Proceedings. ICRA'03. IEEE International Conference on*, vol. 1. IEEE, 2003, pp. 1521–1526.
- [7] K. Zhou and S. Roumeliotis, "Multirobot Active Target Tracking With Combinations of Relative Observations," *Robotics, IEEE Transactions on*, vol. 27, no. 4, p. 678, 2011.
- [8] E. Frew and S. Rock, "Exploratory motion generation for monocular vision-based target localization," vol. 7, 2002, pp. 3633–3643.
- [9] E. W. Frew, "Observer trajectory generation for target-motion estimation using monocular vision," Ph.D. dissertation, Stanford University, 2003.
- [10] P. Tokekar, J. Vander Hook, and V. Isler, "Active target localization for bearing based robotic telemetry," in *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on*. IEEE, 2011, pp. 488–493.
- [11] J. Vander Hook, P. Tokekar, E. Branson, P. G. Bajer, P. W. Sorensen, and V. Isler, "Local-search strategy for active localization of multiple invasive fish," in *Experimental Robotics*, B. Siciliano and O. Khatib, Eds., vol. 88. Springer Tracts in Advanced Robotics, 2013, pp. 859–873.
- [12] J. Derenick, J. Fink, and V. Kumar, "Localization Using Ambiguous Bearings from Radio Signal Strength," *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 3248–3253, Sept. 2011.
- [13] C. Forney, E. Manii, M. Farris, M. Moline, and C. Lowe, "Tracking of a Tagged Leopard Shark with an AUV: Sensor Calibration and State Estimation," in *IEEE International Conference on Robotics and Automation*, 2012, pp. 5315–5321.
- [14] R. Sutton and A. Barto, *Reinforcement Learning: An Introduction*, ser. Adaptive Computation and Machine Learning Series. MIT Press, 1998.
- [15] S. C. Ong, S. W. Png, D. Hsu, and W. S. Lee, "Planning under uncertainty for robotic tasks with mixed observability," *The International Journal of Robotics Research*, vol. 29, no. 8, pp. 1053–1068, 2010.
- [16] S. Hammel, P. Liu, E. Hilliard, and K. Gong, "Optimal observer motion for localization with bearing measurements," *Computers & Mathematics with Applications*, vol. 18, no. 1, pp. 171–180, 1989.
- [17] A. Logothetis, A. Isaksson, and R. J. Evans, "An information theoretic approach to observer path design for bearings-only tracking," in *Decision and Control, 1997., Proceedings of the 36th IEEE Conference on*, vol. 4. IEEE, 1997, pp. 3132–3137.
- [18] A. N. Bishop, B. Fidan, B. D. Anderson, K. Doançay, and P. N. Pathirana, "Optimality analysis of sensor-target localization geometries," *Automatica*, vol. 46, no. 3, pp. 479–492, Mar. 2010.
- [19] A. N. Bishop and P. N. Pathirana, "Optimal trajectories for homing navigation with bearing measurements," in *Proceedings of the 2008 International Federation of Automatic Control Congress*, 2008.
- [20] S. Martinez and F. Bullo, "Optimal sensor placement and motion coordination for target tracking," *Automatica*, vol. 42, no. 4, pp. 661–668, Apr. 2006.
- [21] G. Hollinger and S. Singh, "Multi-robot coordination with periodic connectivity," in *Robotics and Automation (ICRA), 2010 IEEE International Conference on*. IEEE, 2010, pp. 4457–4462.
- [22] A. Makarenko and H. Durrant-Whyte, "Decentralized data fusion and control in active sensor networks," in *Proceedings of the Seventh International Conference on Information Fusion*, 2004.
- [23] K. Y. K. Leung, S. Member, T. D. Barfoot, and H. H. T. Liu, "Decentralized Localization of Sparsely-Communicating Robot Networks: A Centralized-Equivalent Approach," vol. 26, no. 1, pp. 62–77, 2010.
- [24] J. R. Spletzer and C. J. Taylor, "Dynamic sensor planning and control for optimally tracking targets," *The International Journal of Robotics Research*, vol. 22, no. 1, pp. 7–20, 2003.
- [25] G. H. Golub and C. F. V. Loan, *Matrix computations (3rd ed.)*. Johns Hopkins University Press, 1996.
- [26] A. Bishop and B. Fidan, "Optimality analysis of sensor-target geometries in passive localization: Part I-Bearing-only localization," *ISSNIP 2007*, vol. 1, pp. 7–12, 2007.
- [27] J. N. Franklin, *Matrix Theory*, ser. Dover books on mathematics. Mineola, New York: Dover Publications, 1968.
- [28] T. Kirubarajan, Y. Bar-Shalom, and D. Lerro, "Bearings-only tracking of maneuvering targets using a batch-recursive estimator," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 37, no. 2, July 2001.
- [29] "Two-robot field experiment," 2013. [Online]. Available: <https://www.youtube.com/watch?v=jZWTUxqTskY>