Technical Report

Department of Computer Science
and Engineering
University of Minnesota
4-192 Keller Hall
200 Union Street SE
Minneapolis, MN 55455-0159 USA

TR 11-018

Cautious Greedy Strategy for Bearing-based Active Localization: Technical Report

Joshua Vander Hook, Pratap Tokekar, and Volkan Isler

September 16, 2011

Cautious Greedy Strategy for Bearing-based Active Localization: Technical Report

Joshua Vander Hook, Pratap Tokekar, Volkan Isler

***Abstract*— We study a novel sensing model which occurs in wildlife monitoring applications. The sensor can be approximated as a bearing-to-target sensor, but with several critical attributes that set it apart from those addressed in bearing-only tracking literature. We propose a measurement strategy that can guarantee reduction in uncertainty in competitive time, regardless of the problem instance. We provide theoretical analysis in the form of a competitive proof, numerical studies, and results from field trials which confirm the applicability of the algorithm.**

I. Introduction

In active localization, the goal is to plan measurement locations for a mobile robot so as to accurately localize a target. In most applications, the true location of the target is unknown. Instead, the robot must estimate the target's location from its measurements. The primary source of difficulty in active localization problems is their online nature: With each measurement, the robot obtains more information about the target's location which should be incorporated into the choice of next measurement location. The online nature of the decision making process makes it difficult to provide theoretical guarantees about the performance of an active localization algorithm.

The motivating application for our study is localizing radio-tagged invasive fish in lakes. These fish tend to aggregate and remain stationary in large groups during the winter months (when the lakes are frozen). If these aggregations can be found, they can be effectively removed. Since manual monitoring of tagged fish is difficult in harsh conditions, we are developing a robotic system to continuously monitor tagged fish. Our ultimate goal is to monitor hundreds of fish in a lake to detect aggregations.

Our platform, shown in Figure 1, is composed of a mobile robot and a directional antenna mounted on a pan-tilt unit. By rotating the antenna and finding the direction with the strongest signal strength, one can compute a line that passes through the robot and the fish. (The exact bearing is unknown due to the symmetric nature of the antenna.) For this reason, we refer to this type of sensor as an infinite-line sensor.

This setting differs from previously studied bearing-only localization problems in a number of ways. In contrast to traditional bearing-only sensors (e.g. cameras, sonars) which can take measurements nearly continuously, measurements in our application take significant time. The radio-tags transmit a single pulse at less than 1Hz, and a full bearing measurement can take up to two minutes depending on the sampling

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 mobile robot, loop antenna, and sensing equipment. The picture was taken during field trials on Lake Casey, MN.

interval (rotation angle). During this time the robot must remain stationary. In addition to the measurement time, the measurement uncertainty in our application is significantly higher than traditional sensors (A typical sampling interval is fifteen degrees.) Finally, due to the symmetric nature of the antenna, we can not distinguish between target bearings α and $\pi + \alpha$. This ambiguity becomes especially significant in the presence of large uncertainties (See Figure 2).

In this paper, we study the active localization problem for this novel type of sensing platform. Specifically, we study the problem of scheduling measurement locations to minimize the time spent, and present an active localization strategy which provides high quality measurements in a short amount of time. To provide bounds on these quantities, we compare to an optimal algorithm which has access to the true target location. We show that we can achieve an arbitrary reduction in uncertainty in an amount of time bounded by a constant factor times the optimal.

The rest of the paper is organized as follows. In the next section we review some relevant literature. We then provide a brief review of the necessary background and notation. In Section IV we show how we can structure the measurement sequence to deal with ambiguous measurements, and introduce our algorithm. In Section V, we provide a theoretical performance bound and evaluate the resulting algorithm in simulation in Section VI. We also report on a field trial of the algorithm which demonstrates the feasibility for the intended

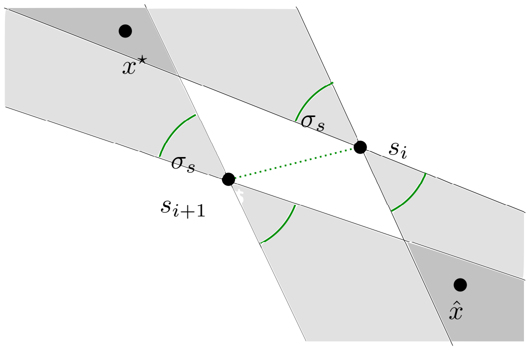

Fig. 2. Two measurements from locations s_i and s_{i+1}. The hypothesis is at $\hat{x}$, and the true target is at $x^\star$. The lower-right cones are predicted most likely, whereas the true cones are less likely, according to the EKF. Given the high noise in measurements, it is desirable to displace large distances between measurements to avoid creating two overlapping regions (shown as dark gray areas).

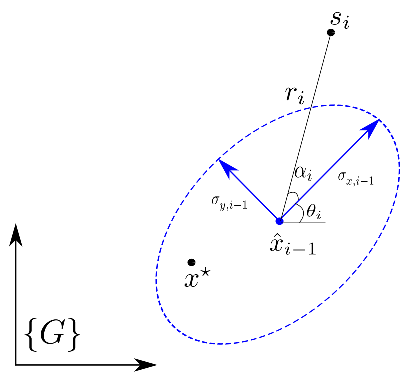

Fig. 3. Coordinate frames and notation: The 1-σ uncertainty ellipse of target's covariance Σ_{i-1} is shown in dashed-blue.

application.

II. Related Work

The literature on continuous bearing-only sensors is quite mature. Much of the work in this domain focuses on optimizing an estimator for tracking a maneuvering target (e.g. [1] and references). Estimators for bearing-only localization and tracking are mainly based on the Extended Kalman Filter. Our contribution is not in presenting a new estimation framework for bearing-only tracking. Instead we focus on structuring the measurement sequence for use with this common estimation technique.

Existing results in this direction include Hammel [2] who numerically approximated an optimal approach for the case of stationary target and continuous bearing-only sensor. This result was generalizable to various velocity and time constraints. Logethesis [3] used mutual information, and found a spiraling approach to the target. These are considered offline strategies, because the trajectory is precomputed, and cannot incorporate new information during execution.

Much recent work has focused on online algorithms. For example, Frew [4] provided an action-space search which used the determinant of the posterior covariance as a utility function. The result is useful for approximating an optimal continuous trajectory, but measurement time was assumed to be small and hence ignored.

Similar to this work was Kronhamn [5], who proposed a greedy trajectory choice, but used an ad-hoc measure which captured the uncertainty in range across a number of target hypotheses. Over time the proposed range information function decreased, driving the sensor in a new direction. This produces a zig-zag motion pattern for the sensor. This work does not easily extend to the case when the platform must stop to measure.

The infinite line sensor is rarely addressed with the exception of Bopardikar et al. [6] in the context of pursuit evasion problems. They provided a strategy which guaranteed capture of an evading target. The resulting zig-zag path is similar to the trajectories we consider. However, they neglect sensor noise and use a geometric approach, which is difficult to generalize, especially given the levels of noise which we deal with in this application.

In our previous work, ([7], [8]) we provide an algorithm based on state-space search, and compare it directly to a one-step greedy algorithm. Both algorithms used the determinant of the posterior covariance but neither could provide a guarantee about final uncertainty. Furthermore, discretization and search over state space can be costly, and both were limited to a fixed-sized displacement between measurements. We address each of these issues in this work by deriving a closed form solution for the one-step greedy algorithm, bounding the time spent and providing a guarantee of posterior covariance.

Before proceeding we review some of the concepts and notation necessary to understand our analysis.

III. Notation and Technical Preliminaries

Our goal is to localize a target, whose true position with respect to a fixed global frame $\{G\}$ is $x^\star$ (see Figure 3). The robot is initially located at s_0. The robot moves to a location denoted by s_i to take the i^{th} measurement. The prior target estimate is a 2-D Gaussian distribution given by $\mathcal{N}(\hat{x}_0, \Sigma_0)$. The mean and the covariance of the estimate after the i^{th} measurement are denoted as $\hat{x}_i$ and Σ_i respectively. Let r_i and α_i be the distance and relative angle between s_i and $\hat{x}_{i-1}$ as shown in Figure 3 (we sometimes drop the subscript for convenience; r and α can be assumed to be with respect to the current target estimate in such a case).

We use an Extended Kalman Filter (EKF) for incorporating the bearing measurements to localize the target. Throughout the rest of the paper, we repeatedly make use of the standard EKF update equations and so present them here.

The i^{th} bearing measurement is given by $z_i = h(x^\star, s_i) + n$, where $n \sim \mathcal{N}(0, \sigma_s^2)$ and $h(a,b) = \tan^{-1}(\frac{b_y - a_y}{b_x - a_x})$. An Extended Kalman Filter (EKF) is used to update the target estimate. The update equations for i^{th} measurement are given as [9]:

$$
\begin{aligned}
\hat{x}_i &= \hat{x}_{i-1} + K_i(z_i - \hat{z}_i), \\
\Sigma_i &= \Sigma_{i-1} - \Sigma_{i-1}\left[H_i^T S^{-1} H_i\right]\Sigma_{i-1}.
\end{aligned} \tag{1}
$$

where,

$$H_i = \nabla_{x^\star} h(\hat{x}_{i-1}, s_i) = \frac{1}{r_i}\begin{bmatrix}-\sin\alpha_i & \cos\alpha_i\end{bmatrix}$$
$$S_i = H_i \Sigma_{i-1} H_i^T + \sigma_s^2$$
$$K_i = \Sigma_{i-1} H_i^T S_i^{-1}.$$

The term $(z_i - \hat{z}_i)$ is the difference between the actual measurement z_i and the expected measurement $\hat{z}_i$ and called as *innovation*. The EKF assumes that $p(z_i|\hat{x}_i) \sim \mathcal{N}(\hat{z}_i, S_i)$. In some applications, particularly mobile target tracking, the equations for the EKF update include an additive noise term. It is not included in our algorithm, because we assume stationary target.

An inverse covariance form of EKF called Extended Information Filter (EIF) can be obtained from the EKF by applying the matrix inversion lemma. The resulting update equations are given as:

$$\Sigma_i^{-1} = \Sigma_{i-1}^{-1} + \frac{H_i^T H_i}{\sigma_s^2}.$$

The two eigenvalues of Σ_i are denoted as $\sigma_{x,i}^2$ and $\sigma_{y,i}^2$. Let θ_i be the angle made by the polar axis with the X-axis of $\{G\}$. Hence Σ_{i-1}^{-1} and $\frac{H_i^T H_i}{\sigma_s^2}$ can be diagonalized to yield,

$$\Sigma_i^{-1} = R(\theta_{i-1})\delta(\frac{1}{\sigma_{x,i-1}^2}, \frac{1}{\sigma_{y,i-1}^2})R^T(\theta_{i-1}) + R(\alpha_i)\delta(0, \frac{1}{r_i^2\sigma_s^2})R^T(\alpha_i). \quad (2)$$

where $R(\theta)$ represents a 2-D rotation matrix with angle θ, and $\delta(x,y) = \begin{bmatrix} x & 0 \\ 0 & y \end{bmatrix}$ is a diagonal matrix.

With this review, we formally state our objective and problem statement next.

IV. Active Localization Algorithm

In this section, we present an algorithm that achieves any desired reduction of uncertainty, despite the constraints imposed by the sensing model. We bound the total execution time to no worse than a constant factor times that of an optimal opponent who knows the true target location.

A. Objective

Let t_m be the time it takes to obtain a single measurement and t_t be the total time spent by the robot moving between measurement locations. Our objective is to find a set of measurement locations $S = \{s_1, s_2, \ldots, s_N\}$ to reduce both the eigenvalues of the prior covariance $(\sigma_{x,0}^2, \sigma_{y,0}^2)$ by a given factor $c_d < 1$, to minimize the total time $N \cdot t_m + t_t$. We assume that the robot moves with unit velocity and can turn in negligible time.

Hence our objective is:

$$\underset{S=\{s_1,\ldots,s_N\}}{\text{minimize}} \quad N \cdot t_m + D(S)$$

subject to,

$$\sigma_{x,N}^2 \le c_d \cdot \sigma_{x,0}^2$$
$$\sigma_{y,N}^2 \le c_d \cdot \sigma_{y,0}^2$$
$$N \ge 1$$

and $D(S)$ is the total distance traveled by the robot for the given set of sensing locations. To quantify the performance of this algorithm, we compare to an optimal algorithm which has access to the true target location.

The input to our algorithm is the prior target estimate $\mathcal{N}(\hat{x}_{i-1}, \Sigma_{i-1})$. An EKF is used to incorporate the prior measurements. The output of our algorithm is a location s_i from which the robot should obtain the next measurement. Measurement locations are chosen based on two factors. We seek measurement locations which can minimize the largest eigenvector at each time step. This is subject to a *caution* constraint, which represents the possibility of the target being behind the sensor.

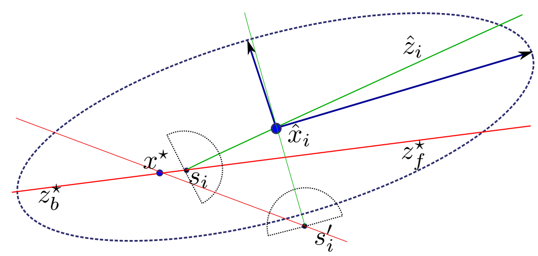

Fig. 4. Two possible locations for a measurement: s_i and s_i'. $\hat{x}_i$ is the current target hypothesis. The predicted measurements are shown as $\hat{z}_i$ (green line), and the actual noiseless measurement is given by $z^\star$ (red line) can be divided into two measurements, $z_f^\star$ and $z_b^\star$ for front and behind. The EKF predicts that $z_f^\star$ is more likely, for any sensor location. Then s_i' has less ambiguity, because the innovation provided by $\hat{z}$ can be applied directly to the hypothesis.

To show the intuition for a cautious strategy, consider the scenario shown in Figure 4. The EKF predicts measurements that correspond to the target hypothesis, $\hat{x}_i$, while the measurements are actually distributed around $x^\star$. If $x^\star$ is nearly collinear with the sensor and hypothesis, then the infinite line sensor is very likely to produce measurements which show a low innovation. This is in direct conflict with the true (but unknown) target location. We are forced to either chose one of the measurements to be used for the update step, or maintain two hypotheses. Worse, according to the EKF, the lower probability measurement is actually the correct one.

Intuitively, this situation occurs when the angle between $\hat{z}$ and $z^\star$ is greater than $\frac{\pi}{2}$ which implies the true target is behind the sensor. However, we can plan the measurement sequence to guard against this possibility, which allows us to use the most likely measurement while limiting the possibility that this is the incorrect choice.

Let the probability that the true target is behind the sensor, with respect to the target hypothesis be,

$$p_b(s, x^\star) \approx p(z|\hat{x} > \frac{\pi}{2}) + p(z|\hat{x} < -\frac{\pi}{2}) \quad (3)$$

Due to the Gaussian prior, this probability is non-zero for any candidate sensor location. Therefore we define a risk term β and seek measurement locations such that $p_b(s_i, \hat{x}) < \beta$. If an algorithm can assure this for each measurement location, we call it β-cautious.

Instead of numerically approximating this probability for each sensor location, we can use the following lemma to exploit some structure provided by the EKF.

Lemma 1: A measurement strategy is β-cautious when the innovation at every time step is less than σ_β^2, given by

$$S_i \leq \sigma_\beta^2 = \frac{\Delta\theta}{\Phi^{-1}(1-\frac{\beta}{2})} \tag{4}$$

Proof: We desire a value for the variance of the measurements such that Equation 3 is less than β. Or, $p(z|\hat{x} \leq \Delta\theta) \geq 1 - \frac{\beta}{2}$. (The other direction is symmetrical). Let the innovation be $\hat{y} = z_i - h(\hat{x})$ as described in Equation 1. Note that $p(z|\hat{x})$ is $\mathcal{N}(\hat{y}, S_i)$. Then $\frac{\hat{y}}{S_i}$ is $\mathcal{N}(0,1)$.

Then $p(\frac{\hat{y}}{S_i} \leq \frac{\Delta\theta}{S_i}) = 1 - \frac{\beta}{2} \implies \frac{\Delta\theta}{S_i} = \Phi^{-1}(1-\frac{\beta}{2})$, where Φ^{-1} is the inverse Gaussian CDF with mean 0 and variance 1. As an inequality,

$$p(\frac{\hat{z}}{S_i} \leq \frac{\Delta\Theta}{S_i}) \geq 1 - \frac{\beta}{2} \implies \frac{\Delta\theta}{S_i} \geq \Phi^{-1}(1-\frac{\beta}{2})$$

Note that everything is constant except S_i which depends directly on the candidate sensor location and distribution of the target hypothesis. By solving for S_i the result follows. ∎

Lemma 2: A measurement location is β-cautious if the range at time i is given by,

$$r_i \geq \sqrt{\frac{\sin^2\alpha_i\sigma_x^2 + \cos^2\alpha_i\sigma_y^2}{\sigma_\beta^2 - \sigma_s^2}} \tag{5}$$

Proof: By Equation 4, a location is cautious if

$$\begin{aligned} S_i &\leq \sigma_\beta^2 \\ H_i\Sigma_{i-1}H_i^T &\leq \sigma_\beta^2 - \sigma_s^2 \\ \frac{1}{r_i^2}(\sin^2\alpha_i\sigma_x^2 + \cos^2\alpha_i\sigma_y^2) &\leq \sigma_\beta^2 - \sigma_s^2 \end{aligned}$$

The result follows by inverting the inequality and gathering the terms. ∎

This allows a single offline calculation of a parameter σ_β. Most mathematical sofwares, including Matlab, have a simple routine to calculate $\Phi^{-1}(x)$. We assume $\sigma_\beta > \sigma_s$. If this is not the case, intuitively, the sensor is too noisy to satisfy caution given the value of β assigned. In Section VI we show that reducing β arbitrarily drives the time spent by the algorithm to infinity. This is because the required range r_i becomes so great that the measurements have litte effect on the hypothesis. However, the number of measurements follows in closed form (see Section V), so the effect of β can be evaluated offline.

Algorithm 1 presents the details of our strategy.

At the i^{th} measurement step, we define a local polar coordinate frame with pole at $\hat{x}_{i-1}$ and the polar axis aligned with one of the eigenvectors.

Figure 5 shows a measurement step of Algorithm 1 in action. $\hat{x}_{i-1}$ and $\Sigma_{i-1} = \delta(\sigma_{x,i-1}^2, \sigma_{y,i-1}^2)$ is the estimate of the target before the i^{th} measurement. Suppose that i is odd; hence we reduce $\sigma_{x,i-1}^2$ with this measurement. Recall that (r_i, α_i) are the polar coordinates in the frame of $\hat{x}_{i-1}$ with the axis aligned with the eigenvector corresponding to $\sigma_{x,i-1}$. Algorithm 1 chooses the measurement location $s_i : (\sqrt{\frac{\sigma_{x,i-1}^2}{\sigma_\beta^2 - \sigma_s^2}}, \frac{\pi}{2})$. After updating the target estimate using the EKF update equation with measured bearing z_i, the target estimate shifts to a new location $\hat{x}_i$, $\sigma_{x,i-1}^2$ decreases, and $\sigma_{y,i-1}^2$ remains unchanged. Since $i+1$ is even, the new measurement location is chosen perpendicular to the eigenvector corresponding to $\sigma_{y,i}^2$ (We show in Lemma 3 the covariance matrix does not rotate). The robot then travels to this new location s_{i+1} (shown as dashed in Figure 5).

This process continues until both the eigenvalues are reduced by a desired factor. The algorithm is guaranteed to terminate since for every two measurement steps, both eigenvalues are guaranteed to decrease. In fact, we bound the number of measurements taken by our strategy in Lemma 6.

In practice, there are two candidate points which satisfy

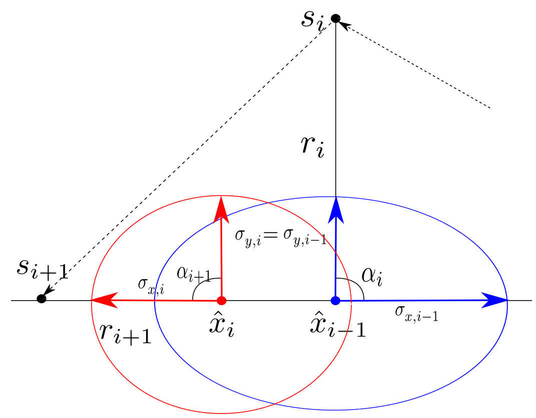

Fig. 5. One measurement step of the cautious strategy presented in Algorithm 1.

Algorithm 1 β-Cautious Strategy$(s_0, \hat{x}_0, \Sigma_0, \Delta\theta, \beta, \sigma_s^2)$

$\sigma_\beta^2 \leftarrow \dfrac{\Delta\theta}{\Phi^{-1}(1-\beta)}$
$(\sigma_{x,0}^2, \sigma_{y,0}^2) \leftarrow$ eigenvalues(Σ_0)
$i \leftarrow 1$
while $\sigma_{x,i-1}^2 \geq c_d \cdot \sigma_{x,0}^2$ or $\sigma_{y,i-1}^2 \geq c_d \cdot \sigma_{y,0}^2$ **do**
  if i is odd **then**
    Polar frame at $\hat{x}_{i-1}$ aligned with $\sigma_{x,i-1}$
    Move to s_i: $(r_i, \alpha_i) \leftarrow (\sqrt{\frac{\sigma_{x,i}^2}{\sigma_\beta^2 - \sigma_s^2}}, \frac{\pi}{2})$.
  end if
  if i is even **then**
    Polar frame at $\hat{x}_{i-1}$ aligned with $\sigma_{y,i-1}$
    Move to s_i: $(r_i, \alpha_i) \leftarrow (\sqrt{\frac{\sigma_{y,i-1}^2}{\sigma_\beta^2 - \sigma_s^2}}, \frac{\pi}{2})$.
  end if
  $(\hat{x}_i, \Sigma_i) \leftarrow$ ekf_update$(\hat{x}_{i-1}, \Sigma_{i-1})$
  $(\sigma_{x,i}^2, \sigma_{y,i}^2) \leftarrow$ eigenvalues(Σ_i)
  $i \leftarrow i + 1$
end while

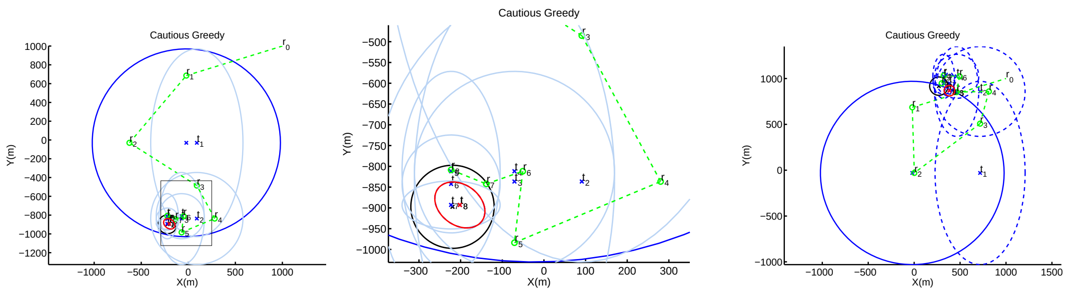

Fig. 6. A simulation execution of the cautious strategy. Left: the full run. Middle: Detail of the boxed region in the first figure. The robot successfully keeps the target hypothesis on the same side as the true target location, with no knowledge of the true target. Right: Another run from the same setup. The true target begins behind the first measurement location, but subsequent measurements correct because the caution constraint prevents the target covariance from becoming too small too quickly. In all figures, red is the final target, black is the true target, and green are the sensor locations.

the cautious requirement: at $(r_i, \frac{\pi}{2})$ and reflected across the target eigenvector, at $(r_i, -\frac{\pi}{2})$. We can chose the closer of the two to save travel time. However, we will proceed assuming we always chose the more distant option, because we seek an upper bound.

Figure 6 shows a complete simulated run of our algorithm. The robot travels in a zig-zag fashion towards target, and after each measurement update, the target estimate shifts (again in a rectilinear sense) towards the true target location. Instead of the alternating between the two eigenvalues for consecutive measurements, we could have proceeded with reducing one at a time, eliminating the need for the robot to zig-zag. However, alternating between the two eigenvalues has the desirable characteristic that we never get too close to the target hypothesis. Additionally, we bound the total distance traveled by the robot using our strategy (Lemma 8) and thus the total time as well (Theorem 1), presented in the next section.

V. Analysis

Our analysis proceeds in two parts: We first bound the total time taken by the cautious strategy, for a given input instance. We then derive a lower bound on the time for optimal strategy and show that the competitive ratio is bounded.

We would like to point out some of the geometry provided by the EKF update equations, and of bearing measurements in general. These derivations provide the framework for our analysis.

Suppose that the robot always takes a measurement of the form $\alpha_i = 0$ or $\alpha_i = \frac{\pi}{2}$.

Lemma 3: For the i^{th} measurement if $\alpha_i = 0$ or $\alpha_i = \frac{\pi}{2}$, then only one eigenvalue of the posterior covariance ${}^{i-1}\Sigma_i$ decreases, and the rotation of the posterior is the same as the rotation of the prior.

Proof: In the local frame of $\hat{x}_{i-1}$, Σ_i is a diagonal matrix, i.e., ${}^{i-1}\Sigma_{i-1} = \delta(\dfrac{1}{\sigma^2_{x,i-1}}, \dfrac{1}{\sigma^2_{y,i-1}})$. The EIF update equation (Equation 2), when written in the local frame of $\hat{x}_{i-1}$ becomes:

$$ {}^{i-1}\Sigma_i^{-1} = \delta(\frac{1}{\sigma^2_{x,i-1}}, \frac{1}{\sigma^2_{y,i-1}}) + R(\alpha_i)\delta(0, \frac{1}{r_i^2\sigma_s^2})R^T(\alpha_i) \tag{6} $$

The proof follows directly from Equation 6. When $\alpha_i = 0$, $R(\alpha_i) = \begin{bmatrix}1 & 0\\ 0 & 1\end{bmatrix}$. By substituting in Equation 6 we get,

$$ {}^{i-1}\Sigma_i^{-1} = \begin{bmatrix}\frac{1}{\sigma^2_{x,i-1}} & 0\\ 0 & \frac{1}{\sigma^2_{y,i-1}}\end{bmatrix} + \begin{bmatrix}1 & 0\\ 0 & 1\end{bmatrix}\begin{bmatrix}0 & 0\\ 0 & \frac{1}{r_i^2\sigma_s^2}\end{bmatrix}\begin{bmatrix}1 & 0\\ 0 & 1\end{bmatrix} $$

$$ \begin{bmatrix}\frac{1}{\sigma^2_{x,i}} & 0\\ 0 & \frac{1}{\sigma^2_{y,i}}\end{bmatrix} = \begin{bmatrix}\frac{1}{\sigma^2_{x,i-1}} & 0\\ 0 & \frac{1}{\sigma^2_{y,i-1}} + \frac{1}{r_i^2\sigma_s^2}\end{bmatrix} $$

We can observe that ${}^{i-1}\Sigma_i$ is a diagonal matrix. Also we have: $\sigma^2_{x,i} = \sigma^2_{x,i-1}$ and $\sigma^2_{y,i} = \frac{\sigma^2_{y,i-1}\sigma_s^2 r_i^2}{\sigma^2_{y,i-1}+\sigma_s^2 r_i^2}$. The case when $\alpha_i = \frac{\pi}{2}$ is symmetric yielding: $\sigma^2_{x,i} = \frac{\sigma^2_{x,i-1}\sigma_s^2 r_i^2}{\sigma^2_{x,i-1}+\sigma_s^2 r_i^2}$ and $\sigma^2_{y,i} = \sigma^2_{y,i-1}$. ■

Lemma 4: For a given range from target, r_i, and an eigenvalue we wish to reduce, a maximal reduction occurs when we set $\alpha_i = \frac{\pi}{2}$ with respect to the corresponding eigenvector.

Proof: The posterior, prior, and measurement jacobian are related by Equation 2.

$$ \Sigma_i^{-1} = \Sigma_{i-1}^{-1} + \frac{H_i^T H_i}{r_i^2\sigma_s^2} $$

The trace of a matrix equals the sum of the eigenvectors. Taking the trace of the above equation, we have,

$$ \begin{aligned} tr(\Sigma_i^{-1}) &= tr(\Sigma_{i-1}^{-1}) + tr(\frac{H_i^T H_i}{\sigma_s^2 r_i^2}) \\ \sigma_{x,i}^{-2} + \sigma_{y,i}^{-2} &= \sigma_{x,i-1}^{-2} + \sigma_{y,i-1}^{-2} + \frac{1}{\sigma_s^2 r_i^2} \\ \Delta\sigma_x^{-2} + \Delta\sigma_y^{-2} &= \frac{1}{\sigma_s^2 r_i^2} \end{aligned} $$

Without loss of generality assume we wish to decrease σ_x. Then,

$$ \Delta\sigma_x^{-2} \leq \frac{1}{\sigma_s^2 r_i^2} \qquad \because \Delta\sigma_y^{-2} \geq 0 \tag{7} $$

which shows that the decrease for any single eigenvalue is bounded above by $\frac{1}{r_i^2\sigma_s^2}$. By examining Equation 6, we can change an eigenvalue by this maximum amount when $\alpha_i =$

$\frac{\pi}{2} + k\pi$ or $0 + k\pi$, (for some integer k) depending on the choice of σ_x^2 or σ_y^2. ∎

The following lemma bounds the maximum distance the target estimate shifts after an update.

Lemma 5: After the i^{th} measurement, the maximum distance between the posterior hypothesis and the prior hypothesis is given by,

$$||\hat{x}_i - \hat{x}_{i-1}|| \leq \begin{cases} \pi \frac{\sigma_{y,i-1}^2}{\frac{\sigma_{y,i-1}^2}{r_i} + r_i\sigma_s^2}, \ \alpha_i = 0 \\ \pi \frac{\sigma_{x,i-1}^2}{\frac{\sigma_{x,i-1}^2}{r_i} + r_i\sigma_s^2}, \alpha_i = \frac{\pi}{2} \end{cases}$$

Proof: From Equation 1 we can write

$$||\hat{x}_i - \hat{x}_{i-1}|| = (z_i - \hat{z}_i)||K_i|| \leq \pi||K_i||,$$

since the innovation $(z_i - \hat{z}_i)$ can be at most π.

We express all quantities in the local coordinate frame of $\hat{x}_{i-1}$ (note that $||\hat{x}_i - \hat{x}_{i-1}||$ remains the same). By substituting the diagonal form of Σ_{i-1} into the EKF update equations and simplifying we get

$$S_i = \frac{\sigma_{x,i-1}^2 \sin^2\alpha_i + \sigma_{y,i-1}^2 \cos^2\alpha_i}{r_i^2} + \sigma_s^2$$

$$K_i = \frac{1}{\frac{\sigma_{x,i-1}^2 \sin^2\alpha_i + \sigma_{y,i-1}^2\cos^2\alpha_i}{r_i} + r_i\sigma_s^2} \begin{bmatrix} -\sigma_{x,i-1}^2 \sin\alpha_i \\ \sigma_{y,i-1}^2 \cos\alpha_i \end{bmatrix}$$

Hence we have,

$$||K_i|| = \begin{cases} \frac{\sigma_{y,i-1}^2}{\frac{\sigma_{y,i-1}^2}{r_i} + r_i\sigma_s^2}, \ \alpha_i = 0 \\ \frac{\sigma_{x,i-1}^2}{\frac{\sigma_{x,i-1}^2}{r_i} + r_i\sigma_s^2}, \alpha_i = \frac{\pi}{2} \end{cases}$$

$$\therefore \ ||\hat{x}_i - \hat{x}_{i-1}|| \leq \begin{cases} \pi \frac{\sigma_{y,i-1}^2}{\frac{\sigma_{y,i-1}^2}{r_i} + r_i\sigma_s^2}, \ \alpha_i = 0 \\ \pi \frac{\sigma_{x,i-1}^2}{\frac{\sigma_{x,i-1}^2}{r_i} + r_i\sigma_s^2}, \alpha_i = \frac{\pi}{2} \end{cases}$$

∎

A. Upper Bound on T_{ALG}

To analyze the performance of our algorithm, we first show that for a given $c_d < 1$, the number of measurements taken by the cautious strategy is bounded. We then show that the distance between consecutive sensing locations is also bounded, regardless of the actual measurements obtained. We prove these results for a circular-shaped prior, i.e. $\sigma_x = \sigma_y$, then generalize to the case where they can be different. Since our strategy assumes no knowledge of the true target location, we show the performance bounds hold for all starting conditions.

We begin by proving a bound on the number of meausurements required for a given c_d.

Lemma 6: For a given acceptable risk parameter β, $N = \lceil 4\log_\gamma\left(\frac{1}{c_d}\right)\rceil$ measurements are sufficient for the cautious strategy to achieve the given desired reduction c_d, where $\gamma = \dfrac{\sigma_\beta^2}{\sigma_s^2} > 1$.

Proof: The cautious strategy presented in Algorithm 1 at each iteration chooses a measurement location which is perpendicular to one of the eigenvectors of the covariance matrix at a distance $r_i = \sqrt{\frac{\sigma_i^2}{\sigma_\beta^2 - \sigma_s^2}}$ away from the target estimate, where $\sigma_i = \sigma_{x,i}$ or $\sigma_{y,i}$ depending on whether i is even or odd. In the following, consider all even ($j = 2i$) iterations, such that σ_j refers to $\sigma_{x,j}$.

By Lemma 3, we know that when $\alpha = \frac{\pi}{2}$, EKF updates only one of the eigenvalues and the other one remains unchanged. In this case we have,

$$\frac{1}{\sigma_{j+1}^2} = \frac{1}{\sigma_j^2} + \frac{1}{\sigma_s^2 r_j^2}.$$

Substituting the value of r_j we get,

$$\frac{1}{\sigma_{j+1}^2} = \frac{1}{\sigma_j^2}\left[1 + \frac{(\sigma_\beta^2 - \sigma_s^2)}{\sigma_s^2}\right]$$

Let $\gamma = \frac{\sigma_\beta^2}{\sigma_s^2}$. Hence back-substitution shows,

$$\begin{aligned} \frac{1}{\sigma_{j+1}^2} &= \frac{1}{\sigma_j^2}\gamma = \frac{1}{\sigma_{j-1}^2}\gamma^2 = \frac{1}{\sigma_{j-2}^2}\gamma^3 = \dots \\ &= \frac{1}{\sigma_0^2}\gamma^j. \end{aligned}$$

Let N be the number of measurements required by the cautious strategy to achieve some desired reduction. Then,

$$\frac{\sigma_0^2}{\sigma_n^2} = \gamma^N \geq \frac{1}{c_d^2}$$

$$\therefore \ N \geq 2\log_\gamma\left(\frac{1}{c_d}\right)$$

Hence, $N = \lceil 2\log_\gamma\left(\frac{1}{c_d}\right)\rceil$ measurements are sufficient for the cautious strategy to reduce one of the eigenvalues by c_d. With an additional $\lceil 2\log_\gamma\left(\frac{1}{c_d}\right)\rceil$ measurements, the cautious strategy can achieve the desired reduction for the other eigenvalue. ∎

Next we show that the total distance the robot travels for these N measurements is bound. The next lemma computes a bound on the total distance traveled for the case when Σ_0 is circular. Lemma 8 then extends this bound for the case when Σ_0 is not necessarily circular. Finally, Theorem 1 gives the bound on the total distance traveled by the cautious strategy during its entire run.

Lemma 7: If Σ_0 is circular, the distance traveled by the robot performing Algorithm 1 moving between all N sensor locations i.e., between s_1 to s_N is bounded i.e.,

$$D_{CIRC} \leq d(s_0, s_1) + 2\sigma_0\frac{1 - c_d}{\sqrt{\gamma} - 1}\left(\frac{\sqrt{2} + \pi\frac{\gamma - 1}{\gamma}}{\sigma_s\sqrt{\gamma - 1}}\right) \quad (8)$$

where $\gamma = \dfrac{\sigma_\beta^2}{\sigma_s^2}$.

Proof: Consider the situation shown in Figure 5 for the i^{th} measurement of Algorithm 1. Without loss of generality, assume that $i \geq 1$ is odd, hence the robot takes a measurement perpendicular to the eigenvector corresponding to $\sigma^2_{x,i-1}$.

For taking this i^{th} measurement the robot moves to s_i given by $r_i = \sqrt{\frac{\sigma_i^2}{\sigma_\beta^2 - \sigma_s^2}}$ and $\alpha_i = \frac{\pi}{2}$ in local frame of $\hat{x}_i$.

The robot takes a measurement from this location, and performs an EKF update. By Lemma 5 we know that the updated target estimate $\hat{x}_i$ shifts a bounded distance along the direction of the eigenvector corresponding to $\sigma^2_{x,i}$. By Lemma 3 we know that $\sigma_{y,i}$ does not change and is the same as σ_i. Hence, $r_{i+1} = r_i$.

By triangle law of inequality we know that the total distance traveled by the robot between s_i and s_{i+1} is bounded by,

$$\begin{aligned}
d_O(i) &\leq \sqrt{2} r_i + ||\hat{x}_i - \hat{x}_{i-1}|| \\
&\leq \sqrt{2} r_i + \pi \frac{\sigma_i^2}{\frac{\sigma_i^2}{r_i} + r_i \sigma_s^2} \quad \text{(By Lemma 5)} \\
&= r_i \left(\sqrt{2} + \pi \frac{\sigma_\beta^2 - \sigma_s^2}{\sigma_\beta^2} \right) \quad \left(\because \sigma_i^2 = r_i^2 (\sigma_\beta^2 - \sigma_s^2) \right).
\end{aligned}$$

Now, for the $(i+1)^{th}$ measurement the distance traveled is strictly less than $d_O(i)$. This yields $d_O(i) + d_E(i+1) \leq 2 d_O(i)$. Hence the total distance traveled for N measurements is bounded by,

$$\begin{aligned}
D_{CIRC} &\leq d(s_0, s_1) + 2 \sum_{i=1}^{N/2} d_O(2i-1) \\
&\leq d(s_0, s_1) + 2 \sum_{i=1}^{N} d_O(2i-1) \\
&= d(s_0, s_1) + 2 \left(\sqrt{2} + \pi \frac{\sigma_\beta^2 - \sigma_s^2}{\sigma_\beta^2} \right) \sum_{i=1}^{N/2} r_{2i-1} \\
&= d(s_0, s_1) + 2 \left(\sqrt{2} + \pi \frac{\sigma_\beta^2 - \sigma_s^2}{\sigma_\beta^2} \right) \sum_{i=1}^{N/2} \frac{\sigma_{2i-1}}{\sqrt{\sigma_\beta^2 - \sigma_s^2}} \\
&= d(s_0, s_1) + 2 \left(\frac{\sqrt{2} + \pi \frac{\sigma_\beta^2 - \sigma_s^2}{\sigma_\beta^2}}{\sqrt{\sigma_\beta^2 - \sigma_s^2}} \right) \sum_{i=1}^{N/2} \sigma_{2i-1}
\end{aligned}$$

Consider,

$$\begin{aligned}
\sum_{i=1}^{N/2} \sigma_i &= \sigma_0 \sum_{i=1}^{N/2} \frac{1}{\gamma^{i/2}} \\
&= \sigma_0 \frac{1 - c_d}{\sqrt{\gamma} - 1}
\end{aligned}$$

Substituting we get,

$$D_{CIRC} \leq d(s_0, s_1) + 2 \sigma_0 \frac{1 - c_d}{\sqrt{\gamma} - 1} \left(\frac{\sqrt{2} + \pi \frac{\sigma_\beta^2 - \sigma_s^2}{\sigma_\beta^2}}{\sqrt{\sigma_\beta^2 - \sigma_s^2}} \right)$$

from which Equation 8 follows. ■

We now extend the above lemma for the case when Σ_0 is not necessarily circular.

Lemma 8: The total distance traveled by the robot executing Algorithm 1 with prior covariance $\Sigma_0 = \delta(\sigma^2_{x,0}, \sigma^2_{y,0})$ is bounded by Equation 8 where $\sigma_0 = \max\{\sigma_{x,0}, \sigma_{y,0}\}$.

Proof: Without loss of generality assume that $\sigma_{x,o} \geq \sigma_{y,0}$. Using Lemma 3 we can show that by following the cautious strategy $\sigma_{x,i} \geq \sigma_{y,i}$ for any odd value of i. Hence, we can show $d_O(i) + d_E(i+1) \leq 2 d_O(i)$. The total distance can then be shown to be bounded as,

$$D_{ALG} \leq D_{CIRC}$$

■

By combining Lemmas 6 and 8, we can bound the total time spent by a cautious strategy, which holds for any problem instance.

Theorem 1: The total time spent by a β-cautious strategy is bounded such that,

$$\begin{aligned}
T_\beta &\leq N t_m + C_2 \sigma_0 + d(s_0, s_1) \\
N &= \lceil 4 \log_\gamma \left(\frac{1}{c_d} \right) \rceil \\
C_2 &= 2 \frac{1 - c_d}{\sqrt{\gamma} - 1} \left(\frac{\sqrt{2} + \pi \frac{\gamma - 1}{\gamma}}{\sigma_s \sqrt{\gamma - 1}} \right) \\
\gamma &= \frac{\sigma_\beta^2}{\sigma_s^2} > 1.
\end{aligned}$$

The previous analysis allows us to state a worst-case bounds on the performance of this algorithm. We can extend the result by showing that this time is below some constant factor times the optimal time for any problem instance.

B. Lower Bound on Optimal Strategy

To derive this we assume there is some player, OPT who knows the true target location.

Assume OPT has only one goal: to reduce the larger eigenvalue by the desired constant. This setup is pictured in Figure 7, and we refer to it as the 1D version of the problem.

Let the initial position of both algorithms be $s_0 = (\alpha_0, r_0)$, and the measurement sequence chosen by OPT be $S_{OPT} = \{s_1, s_2, \ldots, s_k\}$. Assume there are k measurements taken by the optimal strategy. Depending on the interrelation of measurement time and velocity cost, k is any positive integer.

Lemma 9: Given an optimal measurement sequence S for the one dimensional eigenvalue reduction problem there exists a sequence S', with measurements taken on a line perpendicular to the eigenvector which takes no longer to perform, and produces at least the same reduction in uncertainty.

Proof: Assume that the eigenvalue we wish to reduce has value σ_0 and corresponds to an eigenvector with orientation θ_0, both at time 0. We construct a new measurement sequence S' from $S = \{(\alpha_i, r_i)\}$ such that, $S' = \{(\frac{\pi}{2}, r_i)\}$ (refer Figure 7).

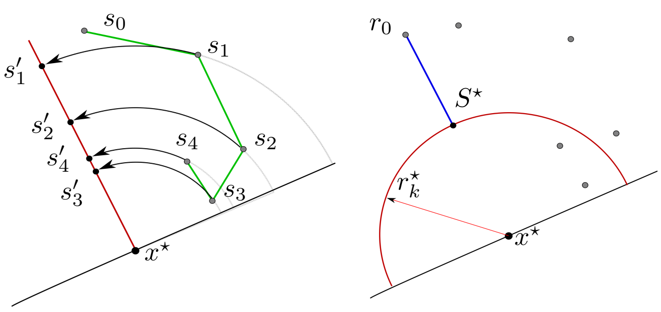

Fig. 7. Restructuring a sequence S (green line) into the sequence $S^\star$. S is first transformed to S' (red line), then all measurements are moved to $r_k^\star$ (red circle). The final sequence (blue line) consists of measurements from only the point $S^\star$ at range $r_k^\star$.

Note that the first sensor locations are then $s'_0 = (\frac{\pi}{2}, r_0)$ and $s_0 = (\alpha_0, r_0)$ and all measurements take place on concentric circles of radius r_i, centered on the true target location.

By induction on i, we can show how the transformation from S to S' produces the same reduction in the eigenvalue, without increasing the time spent.

Consider the situation after first measurement, s_1 and s'_1. We have,

$$\Delta\sigma'_1 = \frac{1}{\sigma_s r_1^2} \geq \Delta\sigma_1 \qquad \because \text{Lemma 4}$$
$$\theta'_0 = \theta'_1 \qquad \because \text{Lemma 3}$$

and the travel time is given by

$$d(s_0, s_1) = ||s_0 - s_1|| \geq d(s'_0, s'_1) = r_0 - r_1$$

which follows from the triangle inequality.

Assume for all $i < k$, $d(S_i) \geq d(S'_i)$ and $\Delta\sigma(S_i) \leq \Delta\sigma(S'_i)$. Then at the k^{th} measurement of the sequence S', the orientation of the eigenvector is the same as the original orientation. The effect of the sequences is given by,

$$\Delta\sigma(S_k) = \frac{H_k^T H_k}{r_k^2 \sigma_s^2} + \Delta\sigma(S_{k-1}) \leq$$
$$\Delta\sigma(S'_k) = \frac{1}{r_k^2 \sigma_s^2} + \Delta\sigma(S'_{k-1})$$

Which follows from Lemma 4, which states that the reduction in the target eigenvalue by $\dfrac{H_k^T H_k}{r_k^2 \sigma_s^2}$ is upper bounded by $\dfrac{1}{r_k^2 \sigma_s^2}$. The travel time taken by these sequences is,

$$D(S_k) = ||s_k - s_{k-1}|| + \sum_{i=0}^{k-2} ||s_i - s_{i+1}|| \geq$$
$$D(S'_k) = r_k - r_{k-1} + \sum_{i=0}^{k-2} r_i - r_{i+1}$$

Where we invoke the triangle inequality to show that $||s_k - s_{k-1}|| \leq r_k - r_{k-1}$ and the summation inequality follows from the induction hypothesis.

∎

Lemma 10: Let S be an optimal measurement sequence of size k measurements. Then a sequence $S^\star$ in which all k measurements occur at the same point produces the same change in eigenvalue in no worse time.

Proof: First we define a point at range $r_k^\star$ perpendicular to the eigenvector, such that k measurements taken from this point satisfies the objective. Then we transform the sequence into S' by the procedure given in the previous lemma.

If all of S' measurements take place at this point, then we are done. If S' contains a measurement with range less than $r_k^\star$, then it must travel farther, and this contradicts the assumption that S' is optimal. If S' contains a measurement with range r' greater than $r_k^\star$ then the reduction in uncertainty produced by the sequence S' is given by

$$\Delta\sigma(S') = \sum_{i=1}^{k-n} \frac{1}{\sigma_s^2 {r'}_i^2} + \sum_{j=n}^{k} \frac{1}{\sigma_s^2 r_k^{\star 2}}$$
$$\leq \sum_{i=1}^{k} \frac{1}{\sigma_s^2 r_k^{\star 2}}$$

because $r' > r_k^\star$. This again contradicts the assumption that S' is optimal.

∎

Theorem 2: Let $S^\star$ be the measurement sequence described in Lemma 10. The minimum time taken by $S^\star$, which is a lower bound for S_{OPT} is given by

$$t(S_{OPT}) \geq t(S^\star) \geq \max(r_0^\star - \frac{\sigma_0^2}{2t_m(\frac{1}{c^2} - 1)\sigma_s^2} + t_m, t_m)$$

where $\sigma_0 = \max(\sigma_{x,0}, \sigma_{y,0})$, $r_0^\star$ is the range to the true target location at time 0, and $r_0^\star \geq r_k^\star$.

Proof: Let the time for the optimal solution be given as,

$$T(S_{OPT}) = k \cdot t_m + \sum_{i=1}^{k} ||d_i^\star - d_{i-1}^\star||$$

where k is the (unknown) number of measurements taken by the optimal strategy. After applying the results of Lemmas 10 and 9 and assuming $r_0^\star \geq r_k^\star$,

$$t(S_{OPT}) \geq t(S^\star) = k \cdot t_m + r_0^\star - r_k^\star$$

We can find $r_k^\star$ as a function of the k measurements by looking at the final reduction in uncertainty, given by,

$$\Delta\sigma(S^\star) = \sum_{i=1}^{k} \frac{1}{r_k^{\star 2} \sigma_s^2} = \frac{k}{r_k^{\star 2} \sigma_s^2}$$

By solving the above for $r_k^\star$, we have

$$r_k^\star = \frac{\sqrt{\Delta\sigma(S^\star) k}}{\sigma_s}$$

Note that the function $\Delta\sigma$ is the change in eigenvalue for the sequence of measurements. Because $S^\star$ completes the objective (it is a transformation of S_{OPT}), the final change is given by

$$\begin{aligned}\Delta\sigma(S_{OPT}) &= \Delta\sigma(S^\star)\\ &= \sigma_k^{-2} - \sigma_0^{-2}\\ &= c_d^{-2}\sigma_0^{-2} - \sigma_0^{-2}\\ &= \sigma_0^{-2}(\frac{1}{c_d^2} - 1)\\ \implies r_k^\star &= \frac{\sqrt{k}\sigma_0}{\sigma_s(\frac{1}{c_d^2} - 1)}\end{aligned}$$

We have the change by both algorithms as an equality. This is because, if a measurement sequence produces a final eigenvalue smaller than $c_d^2\sigma_0^2$, we can move $r_k^\star$ slightly closer to $r_0^\star$, reducing the total travel time, until σ_n^2 exactly equals $c_d^2\sigma_0^2$

The total time taken is now

$$T(S_{OPT}) \geq k \cdot t_m + r_0^\star - \frac{\sqrt{k}\sigma_0}{\sigma_s(\frac{1}{c_d^2} - 1)}$$

We can minimize this time, and solve for k. Minimizing with respect to k produces

$$k_{min} = \frac{\sigma_0^2}{4t_m^2\sigma_s^2(\frac{1}{c_d^2} - 1)}$$

Substituting this back into $t(S^\star)$ and simplifying, we have the result,

$$\begin{aligned}t(S_{OPT}) &\geq r_0^\star - C_3\sigma_0^2 + t_m\\ C_3 &= \frac{1}{2\sigma_s^2 \cdot t_m(\frac{1}{c_d^2} - 1)}\\ &\because k \geq 1, t_m > 0\end{aligned}$$

Notice that, in general, $r_0^\star$ could be less than $C_3\sigma_0^2$. So the time bound should be expressed as $t(S_{OPT}) \geq \max(r_0^\star - C_3\sigma_0^2, 0) + t_m$. ■

Let the sequence of sensor locations given by a β-cautious strategy be described by S_β. We can state the following.

Theorem 3: The cautious strategy is C-competitive.

Proof: We show:

$$\mathbf{E}[t(S_\beta(\mathcal{X}))] - C \times \mathbf{E}[t(S^\star(\mathcal{X}))] \leq b$$

Where b does not depend on the input instance and $\mathbf{E}[\cdot]$ is the expectation over $x^\star$.

It suffices to show that

$$\frac{\mathbf{E}[t(S_\beta(\mathcal{X}))]}{\mathbf{E}[t(S^\star(\mathcal{X}))]} \leq C$$

for some constant C, which does not depend on the problem instance.

Expanding the results from Theorems 1 and 2, then grouping constant terms,

$$\frac{\mathbf{E}[t(S_\beta)]}{\mathbf{E}[t(S^\star)]} = \frac{\mathbf{E}[N \cdot t_m + C_2\sigma_0 + d(s_0, s_1)]}{\mathbf{E}[\max(r_0^\star - C_3\sigma_0^2, 0) + t_m]}$$

where N, C_2 are given by Theorem 1, and C_3 is given by Theorem 2.

Substitute in $\hat{r}_0$ for $d(s_0, s_1)$ because the first measurement location is always closer than the hypothesis, by definition. This makes $T(S_\beta)$ bounded, regardless of true target location. i.e., a constant with respect to true target. Then it remains to find the OPT expected bounds.

First, observe that the denominator is the expectation of the maximum of two convex functions, which is then a convex function. This allows us to apply Jensen's inequality, which states that $\mathbf{E}[f(x)] \geq f(\mathbf{E}[x])$ if the function $f(x)$ is convex. This provides a lower bound on the expectation over $T(S^\star)$ as

$$\frac{\mathbf{E}[t(S_\beta)]}{\mathbf{E}[t(S^\star)]} \leq \frac{N \cdot t_m + C_2\sigma_0 + \hat{r}_0}{\mathbf{E}[r_0^\star - C_3\sigma_0^2] + \mathbf{E}[t_m]}$$

$$\frac{\mathbf{E}[t(S_\beta)]}{\mathbf{E}[t(S^\star)]} \leq \frac{N \cdot t_m + C_2\sigma_0 + \hat{r}_0}{t_m - C_3\sigma_0^2 + \hat{r}_0}$$

Suppose $\hat{r}_0 \geq A_1\hat{r}_0 + A_2\sigma_0 + C_3\sigma_0^2$ for constants $0 < A_1 < 1$, $A_2 > 0$, and $C_3 > 0$. Then we have,

$$\begin{aligned}\hat{r}_0 - C_3\sigma_0^2 &\geq A_1\hat{r}_0 + A_2\sigma_0\\ \implies \frac{\mathbf{E}[t(S_\beta)]}{\mathbf{E}[t(S^\star)]} &\leq \frac{N \cdot t_m + C_2\sigma_0 + \hat{r}_0}{t_m + A_2\sigma_0 + A_1\hat{r}_0}\\ &\leq \frac{N \cdot t_m + C_2\sigma_0 + \hat{r}_0}{A_1(t_m + A_2\sigma_0 + \hat{r}_0)} \because A_1 < 1\end{aligned}$$

We define A such that,

$$A = \begin{cases} A_1 & \text{if } A_2 > 1\\ A_1 \cdot A_2 & \text{if } A_2 < 1\end{cases}$$

Then,

$$\begin{aligned}\frac{N \cdot t_m + C_2\sigma_0 + \hat{r}_0}{A(t_m + \sigma_0 + \hat{r}_0)} \leq&\\ &\frac{Nt_m + N\sigma_0 + N\hat{r}_0}{A(t_m + \sigma_0 + \hat{r}_0)}\\ &+ \frac{C_2t_m + C_2\sigma_0 + C_2\hat{r}_0}{A(t_m + \sigma_0 + \hat{r}_0)}\\ &+ \frac{+t_m + \sigma_0 + \hat{r}_0}{A(t_m + \sigma_0 + \hat{r}_0)}\\ &\leq \frac{N + C_2 + 1}{A}\end{aligned}$$

The last step follows by grouping like terms, factoring out the constants, and canceling, since all constants are greater than one. ■

VI. Evaluation

To gain insight into the expected performance of a β-cautious algorithm, we conducted simulation studies. We simulated a real-world localization scenario using the noise and estimate uncertainty we have encountered. We present both, the aggregate results, as well as some examples highlighting the major differences in cautious and non-cautious strategies.

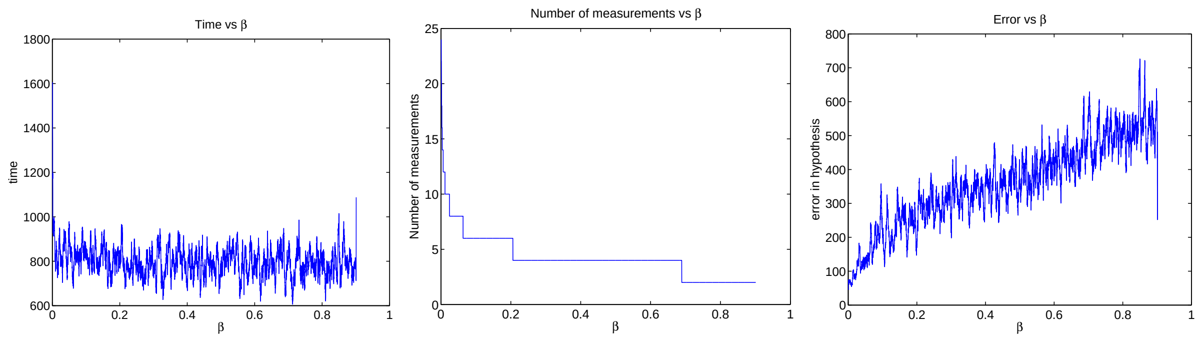

Fig. 8. Simulation study of the effect of β on time (left), number of measurements, (middle), and final accuracy (right). Note that a smaller value of β enforces more caution in the measurements.

In Figure VI we show the effect of introducing β-caution to a greedy measurement strategy. Measurements were chosen to minimize the posterior uncertainty for a range of values of β. Each execution was terminated when it produced an uncertainty with eigenvalues below some constant. From these we can see that a cautious strategy tends to reduce final error (measured as $||\hat{x} - x^\star||$), but excessively small values of β can increase the time spent by the algorithm.

We then varied the starting conditions, but kept the sensing model equal to the parameters for our radio receiver. Starting range to hypothesis was varied from tens of meters ($\mathcal{N}(10, 10)$) to a few thousand ($\mathcal{N}(1000, 500)$). This approximates a small sub-region of a lake, such as shallows where fish might aggregate, to searching an entire lake for a single tag. True target locations were sampled from the target hypothesis, the cautious algorithm run, and the distance traveled and number of measurements were used to estimate the real-world execution time. This was compared directly to the predicted optimal execution given by $S^\star$ from Theorem 2.

TABLE I

E[C] FOR VARIOUS ALGORITHM PARAMETERS.

$r_0^\star$	C_d	t_m	$\mathbf{E}[T(S_\beta)]/T(S^\star)$
100	.1	60	3.4
100	.1	1	4.9
100	.1	120	3.2
1000	.1	120	1.2
100	.1	120	6
10	.1	120	12

Table I shows the results of computing the expected performance. The last column is the expected competitive ratio computed by taking the expectation over all possible $x^\star$ in some finite state space. Notice that the ratio of times increases as the initial range, $r^\star$ decreases. This is because $t(S^\star)$ approaches its minimum value.

A. Field Experiments

In preparation for winter-time tracking of carp aggregations, we implemented the algorithm for use on a robot designed for the task. The chassis is pictured in Figure 1. We then conducted field trials consisting of localizing tags which were distributed in a field.

We show the result of an execution, where the robot was able to locate a reference tag from a distant starting location.

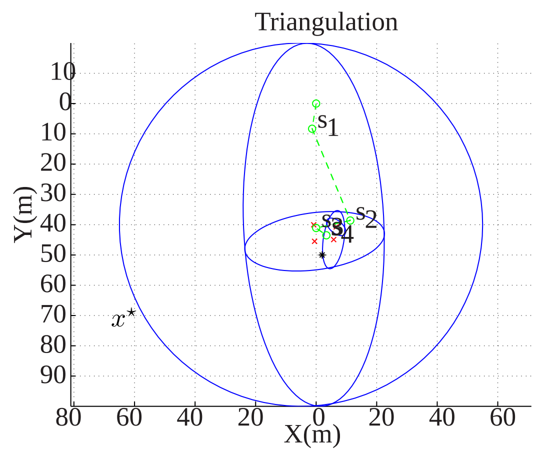

Fig. 9. A successful field trial of the cautious algorithm. Parameters for this trial were $\sigma_\beta^2 = 1$ and $\sigma_s^2 = \left(\frac{\pi}{8}\right)^2$. The 3-$\sigma$ uncertainty ellipses at each iteration are shown.

VII. Conclusion

We presented a novel variant of the greedy strategy, which we call cautious greedy. We have shown that a cautious strategy can minimize the effect of ambiguity with an infinite line sensor. We then bounded the performance of a cautious strategy, providing a worst-case guarantee on the time taken by the algorithm. We have compared the cautious strategy to other greedy strategies, and shown that the performance times for real world applications is often better than the expected result.

The strategy presented here could generalize to other bearing sensors. For instance, the strategy might be used to plan views which maximize the probability of the subject being in some field of view.

In this work we have bounded the time taken to track a single hypothesis. We also minimize the effect of ambiguity, which in turn can reduce the need for multiple hypotheses. Our next steps will focus on analyzing a strategy which admits multiple hypotheses.

Acknowledgment

This work is supported by NSF Awards #1111638, #0916209, #0917676, #0936710 and a fellowship from the Institute on the Environment at the University of Minnesota.

References

[1] X. Rong Li and V. Jilkov, "Survey of maneuvering targettracking . part I: dynamic models," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 39, no. 4, pp. 1333–1364, Oct. 2003. [Online]. Available: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=1261132

[2] S. E. Hammel, P. T. Liu, E. J. Hilliard, and K. F. Gong, "Optimal observer motion for localization with bearing measurements," *Computers and Mathematics with Applications*, vol. 18, no. 1-3, pp. 171 – 180, 1989.

[3] A. Logothetis, A. Isaksson, and R. 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.

[4] E. Frew, "Observer trajectory generation for target-motion estimation using monocular vision," Ph.D. dissertation, Stanford University, 2003.

[5] T. Kronhamn, "Bearings-only target motion analysis based on a multihypothesis Kalman filter and adaptive ownship motion control," *IEE Proceedings - Radar, Sonar and Navigation*, vol. 145, no. 4, p. 247, 1998. [Online]. Available: http://link.aip.org/link/IRSNE2/v145/i4/p247/s1\&Agg=doi

[6] S. Bopardikar, "Pursuit strategies for autonomous vehicles," Ph.D. dissertation, University of California, Santa Barbara, 2010.

[7] P. Tokekar, D. Bhadauria, 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.

[8] P. Tokekar, J. V. Hook, and V. Isler, "Active target localization for bearing based robotic telemetry," in *Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2011.

[9] Y. Bar-Shalom, X. Li, and T. Kirubarajan, *Estimation with applications to tracking and navigation*. Wiley Online Library, 2001.