Cautious Greedy Strategy for Bearing-based Active Localization: Experiments and Theoretical Analysis

Joshua Vander Hook, Pratap Tokekar, Volkan Isler

Abstract— We study the problem of minimizing the time to accurately localize a target using radio-based telemetry. The directional nature of the antenna allows us to obtain bearing-to-target sensor measurements. There are two critical attributes that seperate our setup from the majority of bearing-only tracking literature: ambiguity and long measurement time. We provide a sensing strategy which midgates the effect of ambiguity, and prove that the time required to localize a target is less than a constant times any bearing-based localization strategy.

I. INTRODUCTION

In active localization, the objective 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 future measurement locations. 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. In this setting, several fish are caught and surgically implanted with low-power radio transmitters. The tagged fish are released, and eventually associate with larger groups. During the winter months, the majority of a lake's population tends to aggregate and remain stationary. By tracking the radio-tagged fish, which eventually join these aggregations, the population can be found and targeted for removal [1]. Since manual monitoring of tagged fish is time-consuming and potentially dangerous in harsh conditions, we are developing a robotic system to continuously monitor tagged fish. Our ultimate goal is to monitor many 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 servo motor. 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 radio tag. 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. Since the radio tags transmit with a low duty cycle (less than 1 Hz), it can take up to one minute to construct

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 Gervais, MN.

a bearing measurement. During this time the robot must remain stationary to avoid changing the range or bearing to the target. Our objective is to find a series of discrete measurements which provide a good final estimate of the target location. We seek to minimize the time taken to localize the target for two reasons: (i) We would like to minimize the probability that the target moves during the measurement process, and (ii) to possibly track multiple targets. Similar objectives have been studied in the context of bearings-only localization, but the long measurement time and ambiguity of this sensing modality set this problem apart from the majority of the literature on bearings-only tracking.

The infinite-line sensor was considered by [2], in the context of pursuit-evasion games. A finite-time capture strategy was provided by using pairs of measurements to resolve the ambiguity. But with sensor noise, this strategy cannot guarantee capture, because large measurement noise makes it increasingly difficult to resolve the true bearing (See [3] for more detail).

This problem has been considered in the wireless sensor network community, since direction of arrival, and therefore a bearing measurement, can be estimated from received radio messages (see [4] for a high-level survey). Ultimately, these works solve a fundamentally different problem and cannot

provide sensing locations to minimize the resulting error. Similarly, Derenick et al. [5] considered radio-based bearing measurements in the context of cooperative localization for teams of mobile robots. However, the robots were considered to be arbitrarily moving, and no motion strategy was provided.

Disregarding ambiguity, the problem of tracking a maneuvering target using bearing-only measurements has been well studied (e.g., [6]–[8]). The main focus of this thread of literature is on optimizing an estimator to be robust when provided arbitrary input. Fewer results exist for structuring the input itself, i.e., designing an optimal sequence of discrete measurements. Some relevant results are [9], who provided an action-space search which used the determinant of the posterior covariance as a utility function and [10], who numerically calculated an optimal approach using the Fischer Information Matrix (FIM). These results are useful for approximating an optimal continuous trajectory, but are designed for cameras or sonar, which have negligible measurement time. Bishop, [11] directly evaluated the determinant of the FIM to find sufficient conditions for an optimal measurement sequence. But because the FIM is evaluated around the true target location, which is unknown in practice, this result cannot provide time-to-localize guarantees.

In our previous work, [12], we provided an algorithm based on state-space enumerative search, and compared 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. Finally, there was no systematic approach to the ambiguity of measurements. Instead, we implicitly assumed the most likely measurement was correct. We address each of these issues in this work: We derive a new greedy step which provides the next measurement location in closed-form. This removes the need for a state-space search. We show how to select measurement locations to limit the effect of ambiguity, and bound the time spent by the robot while providing a guarantee of posterior covariance.

The rest of the paper is organized as follows. In the next section we provide a brief review of the necessary background and notation. In Section III we show how we can structure the measurement sequence to deal with ambiguous measurements, and introduce our algorithm. In Section IV, we provide a theoretical performance bound on the algorithm. We evaluate the resulting algorithm in simulation in Section V and we report on field trials of the algorithm which demonstrate the feasibility for the intended application. Due to space limitations, we refer the reader to the technical report [13] for detailed proofs.

II. PROBLEM FORMULATION

Our problem setup is shown in Figure 2. The goal is to localize a target, whose true position with respect to a fixed global frame $\{G\}$ is $x^\star$. The robot is initially located at s_0. The robot moves to a location denoted by s_i to take
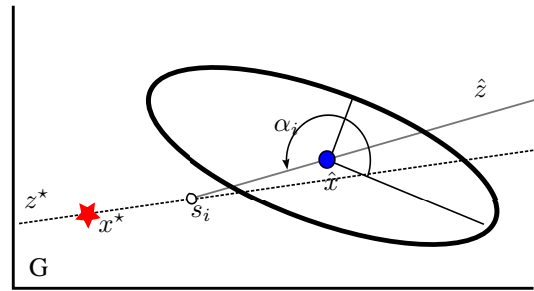

Fig. 2. The sensor location, s_i is described by $r_i = ||\hat{x} - s_i||$ and α_i as shown. $\hat{x}_i$ is the current target hypothesis. The predicted measurements are shown as $\hat{z}_i$ (gray line), and the actual noiseless measurement is given by $z^\star$ (dashed line).

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. We refer to the eigenvalues of Σ_i as $\sigma_i^2(1)$ and $\sigma_i^2(2)$ for the larger and smaller at time-step i, respectively. For convenience of notation and analysis, we specify sensing locations with respect to the frame defined by the major and minor axis of the covariance ellipse. Let r_i and α_i be the distance and relative angle between s_i and the major axis of the covariance ellipse Σ_i, i.e., the eigenvector corresponding to $\sigma_i^2(1)$.

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. We use an EKF because of its closed-form representation and low demand on processing power. The EKF update equations for i^{th} measurement are given as [14]:

$$\hat{x}_i = \hat{x}_{i-1} + K_i y_i$$
$$\Sigma_i = \Sigma_{i-1} - \Sigma_{i-1} \left(H_i^T S_i^{-1} H_i \right) \Sigma_{i-1}.$$

with:

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

The EKF assumes $p(z_i|\hat{x}_i) \sim \mathcal{N}(h(\hat{x}_i, s_i), S_i)$. 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 a stationary target.

An inverse covariance form of EKF called Extended Information Filter (EIF) is defined for the same problem, but in a form more useful for our analysis: $\Sigma_i^{-1} = \Sigma_{i-1}^{-1} + \frac{1}{\sigma_s^2} H_i^T H_i$.

We also use the fact that both the covariance matrix Σ^{-1}

and $H^T H$ admit a unitary decomposition, which results in:

$$\Sigma_i^{-1} = R(\theta_{i-1}) \cdot \delta \left[\frac{1}{\sigma_{i-1}^2(1)}, \frac{1}{\sigma_{i-1}^2(2)} \right] \cdot R^T(\theta_{i-1})$$
$$+ R(\alpha_i + \theta_{i-1}) \cdot \delta \left[0, \frac{1}{r_i^2 \sigma_s^2} \right] \cdot R^T(\alpha_i + \theta_{i-1}). \quad (1)$$

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. Here, θ is the orientation of the covariance with respect to G.

Let t_m be the time it takes to obtain a single measurement. Our objective is to find a set of measurement locations $S = \{s_1, s_2, \ldots, s_N\}$ which reduce both the eigenvalues of the prior covariance $(\sigma_0^2(1), \sigma_0^2(2))$ by a given factor $c_d^2 < 1$ in minimal 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, $\sigma_N(1) \leq c_d \cdot \sigma_0(1)$, $\sigma_N(2) \leq c_d \cdot \sigma_0(2)$, and $N \geq 1$, where $D(S)$ is the total distance traveled by the robot for the given set of sensing locations.

III. ACTIVE LOCALIZATION ALGORITHM

We define a greedy algorithm which incorporates all prior information and outputs the next measurement location. At each step, the input to our algorithm is the prior target estimate $\mathcal{N}(\hat{x}_{i-1}, \Sigma_{i-1})$. The output of our algorithm is the 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.

To show the intuition for a cautious strategy, consider the scenario shown in Figure 2. 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 as shown, 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. In the update step of the EKF, we are forced to chose from the two possible measurements ($\hat{z}$ which agrees with $\hat{x}$, or π+this measurement) or maintain two hypotheses (one for each). According to the EKF, the lower probability measurement ($\pi + \hat{z}$) 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. Or, if $y_i > \frac{\pi}{2}$ or $y_i < -\frac{\pi}{2}$, it is likely that the target was behind the sensor. Using this intuition, 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,

$$p_b(s_i, x^\star) = p\left(y_i | \hat{x}_i > \frac{\pi}{2}\right) + p\left(y_i | \hat{x}_i < -\frac{\pi}{2}\right) \quad (2)$$

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. In [13] we prove that the caution constraint provides an upper bound on the acceptable variance of y_i, given by σ_β, which constrains only the range at which the sensor can take measurements. The range constraint is given by:

$$r_i \geq \sqrt{\frac{\sigma_i^2(1) \cdot \sin^2 \alpha_i + \sigma_i^2(2) \cdot \cos^2 \alpha_i}{\sigma_\beta^2 - \sigma_s^2}} \quad (3)$$

where $\sigma_\beta = \dfrac{\pi}{2 \cdot \phi^{-1}(1 - \frac{\beta}{2})}$ and $\phi(\cdot)$ is the cumulative Gaussian distribution. 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 V 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 little effect on the hypothesis. However, the number of measurements follows in closed form (see Section IV), so the effect of β can be evaluated off-line.
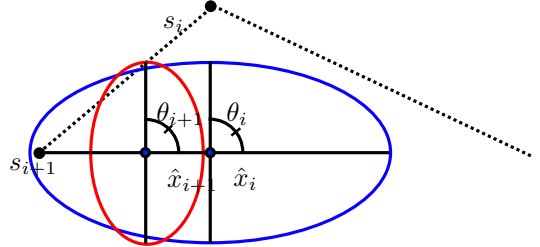

Fig. 3. One measurement step of the cautious strategy presented in Algorithm 1.

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

1: $\sigma_\beta^2 \leftarrow \dfrac{\pi}{2 \cdot \Phi^{-1}(1 - \frac{\beta}{2})}$
2: $\sigma_0^2(1), \sigma_0^2(2) \leftarrow \text{eigenvalues}(\Sigma_0)$
3: $i \leftarrow 1$
4: **while** $\sigma_i(1) \geq c_d \cdot \sigma_0(1)$ or $\sigma_i(2) \geq c_d \cdot \sigma_0(2)$ **do**
5: $\quad$ Polar frame at $\hat{x}_{i-1}$ aligned with $\sigma_{i-1}(1)$
6: $\quad r_i \leftarrow \sqrt{\dfrac{\sigma_{i-1}^2(1)}{\sigma_\beta^2 - \sigma_s^2}}$
7: $\quad$ Let s_i: (r_i, α_i) be the closer of $(r_i, \frac{\pi}{2})$ or $(-r_i, \frac{\pi}{2})$.
8: $\quad$ Collect measurement z_i from s_i
9: $\quad \hat{x}_i, \Sigma_i \leftarrow \text{ekf\_update}(z_i, \sigma_s, \hat{x}_{i-1}, \Sigma_{i-1})$
10: $\quad \sigma_i^2(1), \sigma_i^2(2) \leftarrow \text{eigenvalues}(\Sigma_i)$
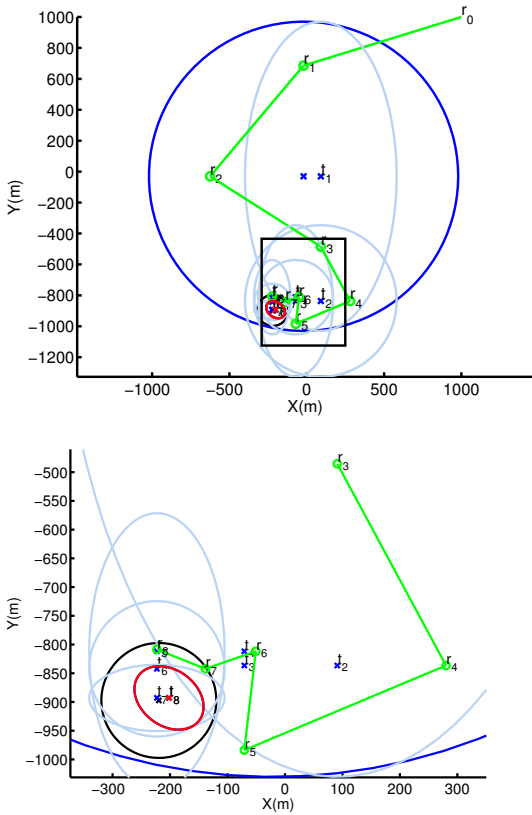11: $\quad i \leftarrow i + 1$
12: **end while**

Fig. 4. A simulation execution of the cautious strategy. Top: the full run. Bottom: Detail of the boxed region in the first figure. Red denotes the final target, black is the true target, and green are the sensor locations. In this example, the robot successfully keeps the target hypothesis on the same side as the true target location, with no knowledge of the true target.

Algorithm 1 presents the details of our strategy and Figure 3 shows a measurement step of Algorithm 1 in action. $\hat{x}_{i-1}$ and $\Sigma_{i-1} = \delta(\sigma_{i-1}^2(1), \sigma_{i-1}^2(2))$ is the estimate of the target before the i^{th} 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 $\sigma_{i-1}^2(1)$. Algorithm 1 chooses the measurement location $s_i : \left(\sqrt{\frac{\sigma_1^2(i-1)}{\sigma_\beta^2 - \sigma_s^2}}, \frac{\pi}{2}\right)$.

After updating the target estimate using the EKF update equation with measured bearing z_i, the target estimate shifts to a new location $\hat{x}_i$, $\sigma_{i-1}^2(1)$ decreases, and $\sigma_{i-1}^2(2)$ remains unchanged. The new measurement location is chosen perpendicular to the eigenvector corresponding to $\sigma_{i-1}^2(2)$. We show in Lemma 1 the covariance matrix does not rotate. The robot then travels to this new location s_{i+1} (shown as a dashed line in Figure 3).

This process continues until both the eigenvalues are reduced by the desired factor. The algorithm is guaranteed to terminate since for every measurement steps, the larger eigenvalue is guaranteed to decrease. In fact, we bound the number of measurements taken by our strategy in Lemma 4.

In practice, there are two candidate points which satisfy 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.

Figure 4 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.

Because the sensing range, r_i monotonically decreases, we expect that the time taken to travel to new sensing locations would also decrease. We also expect the target estimate to shift less with each successive measurement as the certainty of the estimate increases. In the next section we use these intuitions to prove an upper bound on the time required to localize the target.

IV. ANALYSIS

Our analysis proceeds in two parts: We first bound the total time taken by the cautious algorithm, for a given input instance. We call this T_{ALG}. We then derive a lower bound on the time for any bearing-based localization strategy, not-necessarily cautious, which we call T_{LB}. We then take the expectation of these times with respect to the possible true target locations, and show that the ratio is bounded above by a constant. The main structural lemmas are stated here to provide intuition, but the proofs are deferred to [13].

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

Lemma 2: 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.

Lemma 3: The maximum shift in target hypothesis is,

$$||\hat{x}_i - \hat{x}_{i-1}|| \le \pi \frac{\sigma_i^2}{\frac{\sigma_i^2}{r_i} + r_i \sigma_s^2}$$

where the measurement location was chosen according to Lemma 2.

A. Upper Bound on Algorithm Time

To derive T_{ALG} We begin by proving a bound on the number of measurements required for a given c_d, then show the sum of the distances between these measurement locations is bounded.

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

Lemma 5: If Σ_0 is circular, the total distance traveled by the robot during a complete execution of Algorithm 1 is bounded i.e.,

$$D_A \le d(s_0, s_1) + 2\sigma_0(1) \frac{1 - c_d}{\sqrt{\gamma} - 1} \left(\frac{\sqrt{2} + \pi \frac{\gamma - 1}{\gamma}}{\sigma_s \sqrt{\gamma - 1}}\right) \quad (4)$$

A natural extension shows that this is an upper-bound for the non-circular case. By combining Lemmas 4 and 5, 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,

$$T_{ALG} \leq N_z \cdot t_m + D_A \tag{5}$$

with constants N_z and D_A from Lemmas 4 and 5, respectively.

B. Lower Bound on Bearing-based Localization

We seek to bound the effect of ambiguity on the time to localize a target. To do this, we derive T_{LB}, a lower bound on the time required to satisfy the same objectives for any strategy, regardless of ambiguity. To eliminate the effect of the prior estimate, we allow the LB strategy to have access to the true target location. We analyze a simplified objective: Reduce only one eigenvalue by the required constant.

Lemma 6: Given an optimal measurement sequence S_O 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.

Lemma 7: Let S_O 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.

We combine these observations to produce a lower bound on any not-necessarily cautious, optimal measurement sequence. We denote this sequence as $S^\star$, and its time cost as T_{LB}.

Theorem 2: Let $S^\star$ be the measurement sequence described in Lemma 7. The minimum time taken by $S^\star$, which is a lower bound for any strategy, is given by

$$T_{LB} \geq T(S^\star) \geq \max(r_0^\star - \frac{\sigma_{max}^2}{2t_m(\frac{1}{c^2}-1)\sigma_s^2} + t_m, t_m) \tag{6}$$

where $\sigma_{max} = \max(\sigma_{1,0}, \sigma_{2,0})$, $r_0^\star$ is the range to the true target location at time 0, and $r_0^\star \geq r_k^\star$.

Let the sequence of sensor locations given by a β-cautious strategy be described by S_β with time cost $T_{ALG}(\mathcal{X})$ for a specific problem setting $\mathcal{X}$. We can state the following.

Theorem 3: The cautious strategy is C-competitive.

Proof: We show:

$$\mathbf{E}[T_{ALG}(\mathcal{X}))] - C \times \mathbf{E}[T_{LB}(\mathcal{X}))] \leq b$$

Where b does not depend on the input instance and $\mathbf{E}[\cdot]$ is the expectation over $x^\star$ [13]. ∎

V. Evaluation

In the previous section we reported the derivation of the expected performance of the algorithm. In this section we evaluate this performance in a variety of simulations and real-world experiments. We first derive the expected performance of our sensing platform as an example.

To detect tagged fish, we use a loop antenna and attached radio receiver. A bearing measurement is constructed by
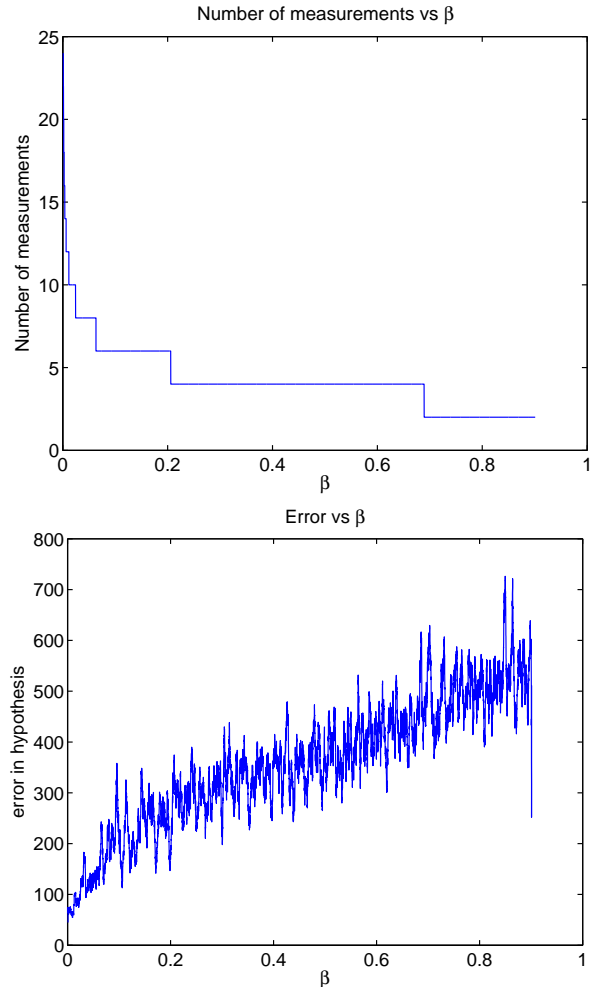

Fig. 5.   Simulation study of the effect of β on number of measurements, (top), and final accuracy (bottom). Note that a smaller value of β enforces more caution in the measurements.

sampling the RSSI as the antenna is rotated. This takes approximately 1 minute due to the low duty cycle of the radio tags. Since our robot travels at approximately 1 meter per second, $t_m \approx 60$. We have empirically estimated σ_s as 15 degrees [12]. Supposing our initial hypothesis satisfies $\sigma_0(1) = \sigma_0(2) = 50$ meters, and we begin our tracking at $r_0 \approx 60$ meters, we can calculate the expected performance using (5). In this case, we find that $N = 4$ measurements will be sufficient to reduce the one sigma bound to less than 6 meters. During this time we expect the robot to travel less than 220 meters, resulting in a competitive ratio of less than about 4.

A. Simulations

To gain further insight into the expected performance of a β-cautious algorithm for various starting parameters, we conducted simulation studies. We simulated a real-world localization scenario using the noise and estimate uncertainty we have encountered.

In Table I we show the results of approximating the

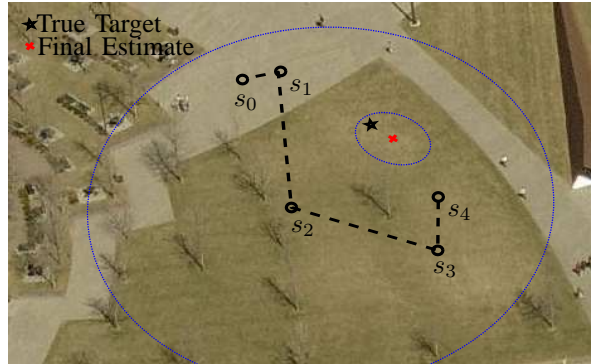| $\hat{r}_0$	c_d	t_m	β	Simulated Ratio	Theoretical Upper Bound
{20, 50, 200}	.1	60	.10	{3.6366, 2.8853, 1.6651}	{7.2853, 5.4477, 2.5477}
100	{.25 .1 .01}	60	.10	{1.3862, 2.1418, 2.8044}	{2.2164, 3.5869, 4.9432}
100	.1	{20, 60, 110}	.10	{1.6046, 2.1398, 2.5568}	{2.3268, 3.5869, 4.5461}
100	.1	60	{.10, .2, .5}	{2.6450, 2.5699, 1.5946}	{3.9104, 3.8353, 2.2264}

Fig. 6. 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 1 σ uncertainty ellipses are shown for the prior and final estimates. The final error was less than 5 meters after only 4 measurements. The total experiment area was approximately 70 meters by 64 meters.

ratio $C = \dfrac{\mathbf{E}[T_{ALG}]}{\mathbf{E}[T_{LB}]}$. The initial hypothesis was fixed with $\hat{x}_0 = (0,0)$ and $\sigma_0(1) = \sigma_0(2) = 60$. We varied the starting parameters initial range $\hat{r}_0$, desired reduction c_d, caution parameter β, and measurement time t_m. True target locations were sampled with replacement from the target hypothesis, the cautious algorithm run, and the mean of distance traveled and number of measurements were used to estimate the real-world execution time. This was compared directly to the predicted execution given by T_{LB} from Theorem 2.

In Figure 5 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 significantly increase the number of measurements required.

B. Field Experiments

We also implemented the algorithm and tested it using a mobile chassis. The chassis is pictured in Figure 1. A transmitting radio tag was placed in a field measuring approximately 64 by 70 meters. Previously, we determined this to be approximately the maximum detection distance of our receiver. The robot was given a starting hypothesis which encompassed the field. The starting parameters and performance closely match the parameters derived at the beginning of this section. Specifically, the final covariance was less than 6 meters, one sigma bound, and four measurements

were required.

VI. CONCLUSION

We presented a novel variant of a greedy bearing-only 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 shown that the expected performance times for real world applications is often better than the theoretical result.

ACKNOWLEDGMENT

REFERENCES

[1] P. G. Bajer, C. J. Chizinski, and P. W. Sorensen, "Using the Judas technique to locate and remove wintertime aggregations of invasive common carp," *Fisheries Management and Ecology*, Oct. 2011.

[2] S. D. Bopardikar, "Pursuit Strategies for Autonomous Vehicles," Ph.D. dissertation, University of California, Santa Barbara, 2010.

[3] 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.

[4] N. Patwari, J. Ash, S. Kyperountas, A. Hero III, R. Moses, and N. Correal, "Locating the nodes: cooperative localization in wireless sensor networks," *Signal Processing Magazine*, vol. 22, no. 4, pp. 54–69, 2005.

[5] J. Derenick, J. Fink, and V. Kumar, "Localization Using Ambiguous Bearings from Radio Signal Strength," in *Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2011.

[6] X. R. Li and V. P. Jilkov, "Survey of maneuvering target tracking: III. Measurement models," *Proceedings of SPIE*, vol. 4473, no. August, pp. 423–446, 2001.

[7] ——, "Survey of maneuvering target tracking: I: dynamic models," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 39, no. 4, pp. 1333–1364, Oct. 2003.

[8] ——, "Survey of maneuvering target tracking: V: multiple-model methods," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 41, no. 4, pp. 1255–1321, Oct. 2005.

[9] E. W. Frew, "Observer Trajectory Generation for Target-Motion Estimation Using Monocular Vision," Ph.D. dissertation, Stanford University, 2003.

[10] 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.

[11] A. Bishop and P. Pathirana, "Optimal trajectories for homing navigation with bearing measurements," in *Proceedings of the 2008 International Federation of Automatic Control Congress*, 2008.

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

[13] J. Vander Hook, P. Tokekar, and V. Isler, "Cautious greedy strategy for bearing-based active localization: Technical report," University of Minnesota, MN, USA, Tech. Rep., 2011.

[14] Y. Bar-Shalom, X. Li, and T. Kirubarajan, *Estimation with applications to tracking and navigation*. Wiley Inter Science, 2001.