

Local-Search Strategy for Active Localization of Multiple Invasive Fish

Joshua Vander Hook, Pratap Tokekar, Elliot Branson, Przemyslaw Bajer, Peter Sorensen and Volkan Isler

Abstract In this paper, we study a problem encountered during our ongoing efforts to locate radio-tagged fish aggregations with robots. The problem lies at the intersection of search-based methods whose objective is to detect a target, and active target localization methods whose objective is to precisely localize a target given its initial estimate. Real-world sensing constraints such as limited and unknown range, large measurement time, and ambiguity in bearing measurements make it imperative to have an intermediate initialization phase to transition from search to localization. We present a local search strategy aimed at reliably initializing an estimate for a single target based on observations from field experiments. We then extend this strategy to initialize multiple targets, exploiting the proximity of nearby aggregated tagged fish to decrease the cost of initialization per target. We evaluate the performance of our algorithm through simulations and demonstrate its utility through a field experiment where the robot successfully detects, initializes and then localizes nearby targets.

1 Introduction

We are developing a robotic system (Figure 1) and algorithms [10, 11] to enable a mobile sensor network to monitor the common carp (*Cyprinus carpio*), an invasive fish. The common carp is an ecologically damaging freshwater fish found in many regions around the world [14]. Biologists are interested in developing efficient methods for controlling carp populations. To this end, they catch a small sample of the population and implant each fish with radio transmitters (tags). These tagged

Joshua Vander Hook · Pratap Tokekar · Elliot Branson · Volkan Isler
Department of Computer Science, University of Minnesota, USA. e-mail: `jvander, tokekar, ebranson, isler@cs.umn.edu`

Przemyslaw Bajer · Peter Sorensen
Department of Fisheries, Wildlife and Conservation Biology, University of Minnesota, USA. e-mail: `bajer003, soren003@umn.edu`



Fig. 1 Our robotic system consists of radio tags, a radio antenna and receiver mounted on autonomous boat in summer and wheeled rover in winter (to operate on frozen lakes).

fish are reintroduced to the lakes and periodically tracked using radio receivers over the course of a year. When multiple tagged fish seem to aggregate, it is assumed a larger population is nearby. When these large aggregations of carp are found, typically during the winter, they can be removed by netting. This provides a safe and environmentally-friendly method for controlling the population of carp.

The radio tags (Figure 1) are small, low duty-cycle transmitters which are implanted under the skin of the fish. Each tag emits a pulsed signal on a dedicated frequency approximately once per second. A human operator carries a loop antenna and a receiver which converts the signal to a received signal strength indicator (RSSI). By monitoring the RSSI and rotating a directionally sensitive antenna, the operator can discern a bearing to the radio tag. Typically a human operator will take 2-3 bearing measurements to estimate the location of one tag. However, this manual tracking approach is tedious, time consuming and possibly inaccurate at times. Therefore, we believe that this problem is a good application for robotics.

Our overall objective in this application is as follows: Given a list of N frequencies (one per tagged fish), each of which can be detected by the robot at a unique range r_i , localize each target to a desired accuracy in bounded time. In Section 3.1, we discuss our previous work where we partition this problem in two separate phases: (i) *Search* phase where the objective is to find a location for the robot within the sensing range of each target, and (ii) *Localization* phase where the robot uses bearing measurements to reduce the uncertainty in the target's estimate.

During field tests of this system, we found that the localization routine was sensitive to the accuracy of the initial estimate. Constructing a consistent, reasonably certain prior estimate in limited time has proven to be a difficult task. The problem becomes further challenging because the sensing ranges of individual tags can vary based on the depth of the fish, the age of the tag, and other environmental factors. For example, Figure 2 shows a field trial where the robot could not complete the triangulation due to an incorrect initialization. The target was initialized with a 2D Gaussian distribution centered at the location where the robot first moved into the sensing range of the tag, with a variance based on empirical estimates of the sensing range. However, the variance was set too low and as such the initial estimate was not consistent. During triangulation, the robot moved to a location which fell outside the sensing range of the target, and the final estimate was wrong. The robot successfully triangulated the same tag in another run where the initial estimate (not shown for

clarity) was consistent. This indicates the importance of starting with a good initial estimate. Therefore, we present a local search strategy which, after detecting a target during the search phase,

1. Initializes a consistent estimate of the target location,
2. Maps a region from which bearing measurements are likely to succeed,
3. Exploits clustering behavior of the fish to locate nearby targets efficiently.

After presenting the details of the search strategy and its analysis in Section 4, we evaluate the strategy through simulations (Section 5), and present results from a field experiment (Section 6). The field trial demonstrates that our proposed initialization strategy is effective, and promising for large-scale future experiments. We believe our proposed approach of search, initialization, and localization should be applicable for other applications where one or more robots are tasked with accurately locating one or more targets in bounded time.

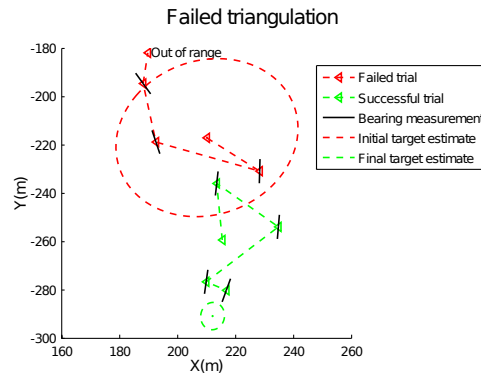


Fig. 2 Failed triangulation due to incorrect initialization for trials conducted on Lake Staring, MN. The initial estimate for the first trial was inconsistent and resulted in the localization to diverge and move the robot out of the tag’s sensing range. During a second trial, with a consistent initial estimate, the target was successfully localized.

2 Previous Work

Recently, there has been significant interest in developing algorithms for locating transmitting radio sources using mobile robots. Song et al. [8] considered the problem of localizing an unknown number of transient radio sources using a mobile robot. They used an occupancy grid in a Bayesian framework to update the probability of a radio source being located in a given grid cell. They further proposed a path-planning algorithm for the robot to improve the convergence time for locating all sources. In [6], Kim et al. presented a centralized multi-robot search algorithm

for the same problem setting, where the robots are controlled in pairs to allow detection of unknown transmission powers from the radio sources.

In [9], Tekdas et al. consider the problem of finding a point of high signal strength inside the sensing disc of transmitting sources. They assume a prior estimate of the source’s location is given but sensing range is unknown. Here, we consider the problem of finding a good point to begin triangulation, while estimating sensing range and target location simultaneously.

Fink and Kumar [3] presented methods to build a radio signal strength map in an unknown indoor environment and presented control laws for mobile robots to seek the transmitting radio source. Recently, Twigg et al. [12] addressed the problem of exploration while seeking a radio source. The algorithm builds a gradient of the RSSI by collecting samples locally. Their work involves indoor environments and areas with significant multi-path effects, and so is not directly applicable to our work. In addition, the directional sensitivity of our antenna makes it difficult to determine and follow a gradient.

The problem of simultaneously localizing a robot and multiple transmitting sources was considered in [4]. It was assumed that range could be explicitly recovered from the transmissions, and an arbitrary robot path was reconstructed while simultaneously estimating the position of each radio. An iterative, offline algorithm was proposed and evaluated. This problem is fundamentally different because we cannot recover range directly, and must solve the problem online, i.e., as measurements become available. Furthermore, we have direct control over the robot’s path. In fact, defining the robot’s path to aid the estimation problem is the what we address in the following sections.

3 Motivation


In this section, we present the details of our system and then discuss some intuitive methods for addressing the problem under consideration.

3.1 System


Our system consists of a wheeled rover to traverse frozen lakes and locate fish aggregations. We deploy a similar system during the summer using robotic boats. Our mobile rover is the A100 Husky by Clearpath Robotics, and our robotic boat is an OceanScience QBoat. Both chassis are fitted with a loop antenna, a servo motor to rotate the antenna, a radio receiver, and a laptop computer. The robots estimate their own pose and navigate using an Extended Kalman Filter (EKF) combining information from a Global Positioning System unit and a digital compass (on the boat) or encoders (on the wheeled rover).

The radio tags are shown in Figure 1 (Left). Each tag is programmed with a specific frequency in the 48-50MHz range and emits 30-50 signals per minute. We operate on lakes which have 10-20 tagged fish and the list of frequencies present in each lake is known a priori. Signals from the tags attenuate as a function of the humidity, salinity of the water, ice or snow thickness, and the depth of the fish it is attached to. These factors cause variations in the range at which tags can be detected. Therefore we do not use the absolute signal strength to estimate range, and instead use the directional nature of the antenna to estimate bearing.

The radio antenna and receiver are pictured in Figure 1 (middle and right), atop both robots. The sensitivity of the antenna varies with the relative angle with the tag. We rotate the antenna using a servo motor in 15 degree steps over 180 degrees. We sample the signal strength at each step and fit a smooth function to the data to estimate the direction with maximum RSSI. This direction is treated as the bearing towards the target. Because of the low signal rate, obtaining a bearing measurement takes about 1-2 minutes. Empirically we have found the bearing measurements to follow an approximately Gaussian distribution around the true target bearing ($\sigma \approx 15^\circ$). However, bearing measurements constructed in this way are ambiguous, or π -periodic. For any estimated bearing z , $z + \pi$ is also a valid bearing measurement (see Figure 3(b)).



(a) A typical search path.



(b) Ambiguous measurements

Fig. 3 Examples of search patterns (Figure 3(a)) and ambiguous bearing measurements (Figure 3(b)). High sensor noise, ambiguity and unknown sensing range makes it difficult to transition from search to localization.

In our previous trials we observed that the tags' radio signal is undetectable unless we are within 100-200 meters. This provides a natural task partitioning: *Search* and *Localization* [11]. The goal of the search phase is to cover the regions of the lake that are likely to contain tagged fish and move the robot to within sensing range of each tag. We then switch to *Localization* where the goal is to obtain multiple bearing measurements to localize the tag to a desired precision. Once a target is localized, the robot can resume its search for other tags. During the search phase, we simply

wait for a detection of a non-zero RSSI value, which takes significantly less time than obtaining a full bearing measurement.

Our current localization algorithm uses an EKF to estimate the position of the tag [13]. The localization subroutine takes time proportional to the area of initial uncertainty and the distance between the initial estimate and the robot. In simulation and experiments this method performs well, but only if the initial estimate of the target is consistent and not significantly uncertain. Obtaining an initial estimate of the target location with bounded uncertainty is challenging, as we discuss next.

3.2 The Initialization Problem

Before the localization algorithm can be deployed to precisely estimate tag locations, we must initialize a prior estimate as input. We briefly present some intuitive methods we have tried and discuss why they fail.

Measurement-based. As often recommended in bearing-only tracking literature, a small number of bearing measurements can be collected and processed in a batch. Given a set of k measurements $\mathbf{Z} = \{z_1, \dots, z_k\}$, we maximize the likelihood, $p(\mathbf{Z}|x)$ over target locations x . In practice, limited sensing range and long measurement time make this strategy infeasible. Also, consider Figure 3(b). The two dark regions show areas which are likely to contain the true target and we cannot easily determine which hypothesis is the origin of the measurements (\hat{x} or \hat{x}'). A third measurement, taken from a large baseline could disambiguate the two. However, a large baseline is likely to move the robot outside the sensing range of the target, producing no information while paying the full cost of a bearing measurement. Another solution could be to take a fixed number of measurements around the initial detection point. Again, the long bearing measurement time makes this an expensive strategy which must be repeated for each nearby tag. Further, it is not clear how these additional measurement locations should be chosen to guarantee a good estimate of the target.

Initial hypothesis. In contrast to the above, we can initialize a hypothesis by taking two measurements as shown in Figure 3(b). By drawing a wedge surrounding each measurement to represent its uncertainty, we can obtain an intersection representing the target hypothesis. We can fit a Gaussian distribution to this intersection area and use as an initial estimate. This is not robust in practice, since the intersection can be unbounded. Additionally, we have two intersection areas leading to two initial hypothesis. As such, this method provides no guarantees about initial estimate uncertainty or range.

Signal-strength based. We can attempt to use the signal strength to resolve the ambiguity of each measurement. The robot could travel toward one hypothesis and measure the signal strength. We expect the signal strength to increase if the robot travels towards the correct hypothesis. In practice, we found this strategy to be sensitive to sensor noise from the unknown and possibly complex spatial signal strength patterns. We found that for small movements near the edge of the sensing range this method was unreliable.




Fig. 4 (a) While on its coverage path (curved arrow), the robot, at O , detects a non-zero signal for some frequency X . (b) The initialization strategy determines the sensing circle for X by moving along search paths as shown until X is not detectable again. Shown is a case where three search paths fail to uniquely identify the sensing circle. (c) An example of a Four-path search.

Each of these initialization methods fails to provide a guarantee of time cost, uncertainty, or consistency of the estimate. In the next section, we describe our solution to this problem which relies on a local search strategy.

4 Local Search

The goal of the local search is two-fold: (1) determine whether an aggregation exists nearby and which targets are contained within the aggregation, and (2) form good initial estimates (mean and covariance) for each target in the aggregation. The initialization phase begins as soon as the robot first detects a non-zero RSSI from a radio tag while on the search path (Figure 4(a)). We assume that the detected tag X is at the center of a sensing circle C_X of radius r . Our objective is to establish an initial estimate of X and r . In this section, we first present our local search initialization strategy for a single target (i.e. X). We bound the worst-case and average-case time required for this strategy. We then extend this strategy for the case of an aggregation of multiple tagged fish.

4.1 Single-Target Local Search

Note that both X (the origin of C_X) and r are unknown. By finding three points on the perimeter of C_X we can solve for X and r . To find these points, the local search proceeds as follows:

1. From the point of first detection (O), the robot moves in a fixed direction with respect to the global frame (e.g., North or angle α).
2. When the robot can no longer detect the target X (position A in Figure 5) it reverses direction and returns to O .

The line segment traversed in these two steps is called as a *search path*. To analyze the time cost of this strategy, we establish the minimum number of search paths needed to find at least three points on the boundary of C_X . We can see that at least four equally spaced search paths are necessary and sufficient from Figure 4(b).




Fig. 5 The robot continues along an arbitrary but fixed direction until it cannot detect the signal from X (at position A). The robot then returns to O and repeats the same strategy along a perpendicular line (B). In general, the O can lie in the interior of the sensing circle, hence the robot also searches along C and D

We now establish the cost of using four search paths to find X and r . The analysis follows Figure 5. Let angle OAX be θ . By design, the angle AOB is $\frac{\pi}{2}$. The distance $|AB|$ is $2r$ and segment OA has length $2r \cos \theta$ while OB has length $2r \sin \theta$. Assume the robot moves with velocity v . Each of these lines must be traversed twice, for a total required time of,

$$T_{\text{single}} = \frac{4r}{v} \cos \theta + \frac{4r}{v} \sin \theta + 4 \cdot \varepsilon \quad (1)$$

where ε is the time taken to recognize the robot has left C_X , turn around, and re-enter C_X . Note that θ is unknown and can take any value between 0 and 2π , depending on the relative orientation of the target position with respect to the first search direction. To obtain the worst-case cost, we maximize the cost function with respect to θ . A straightforward derivation shows the cost is maximum when $\theta = 45$ degrees for a maximum cost of,

$$\max_{\theta} T_{\text{single}} = 2^{\frac{5}{2}} \frac{r}{v} + 4 \cdot \varepsilon \quad (2)$$

The expected search time, assuming θ is uniform in the range $[0, 2\pi]$ is $\mathbf{E}[T_{\text{single}}] = 2^{\frac{5}{2}} \frac{r}{v} + 4\varepsilon$.




Fig. 6 (a) The single target search fails to intersect all the sensing circles in the case of an aggregation. (b) An example of searching for an aggregation using two separate search steps. The robot first finds the boundary of C_X (dashed), centered at O' . Then, after moving to O' , searches along four paths to identify the boundaries of each sensing circle.

4.2 Multi-Target Local Search

To extend the local search strategy to multiple targets, we need a model for fish aggregations. While common carp are relatively broadly dispersed during summers, they tend to form tight aggregations under ice-covered lakes in winters [1, 2, 5, 7]. For example, while average distances between radio-tagged carp during summers are 300-500 meters, in winters, these distances decrease to 50-100 meters [1]. In some cases, entire populations of carp, usually several thousands of fish, have been shown to aggregate in areas that are only 100×100 meters in size [1]. We formalize the notion of an aggregation using the following definition.

Definition 1. Let $\mathcal{L} = \{X_1, \dots, X_i, \dots, X_N\}$ be a set of tagged fish, r_i be the sensing radius of X_i , and $r^* = \min_i r_i$. \mathcal{L} is called an aggregation if, $\forall i, j, \|X_i - X_j\|_2 < r^*$

Under this definition, we cannot directly use the local search strategy for a single target for multiple targets. Figure 6(a) illustrates an example case where the four search paths do not intersect the sensing circle of Y present in the aggregation.

We propose the following strategy: By Definition 1, for any target x , the distance to all other targets to x is less than r^* . Returning to the case of one target shown in Figure 5, we see that four search paths can provide an estimate of a target location as the center of the estimated sensing disk. In general, since we don't know which fish are contained in the aggregation, it might be necessary to search for all frequencies. As a practical step, we make the assumption that the true location of the first fish X is close to the center O' of the estimated sensing circle. This allows us to move to O' and determine which fish are nearby. We can then perform another multi-path




Fig. 7 In general, the starting location of local search can lie anywhere on the boundary or interior of the sensing circle. In each case, we obtain a different number of points as shown. For all cases, we can determine the sensing circle uniquely.

search to map the boundaries of all nearby frequencies (see Figure 6(b)). We call the resulting algorithm *Four-Path*.

Assuming we begin a search from the target location X , we can show that four paths are sufficient to detect the boundaries of each sensing circle in the aggregation. Consider Figure 7, which illustrates the possible configurations of the rest of the targets with respect to the first. We have three cases:

- The target Y is aligned with the search path starting at O' , and we detect two points of C_Y . This case has a unique solution: Y is at $\frac{1}{2}|XA'|$ along XA' .
- O' is on the boundary of C_Y . In this case we detect three points $O', A',$ and B' . We can solve C_Y directly.
- O' is inside the circle C_Y . We can detect four points at A', \dots, D' , and solve the sensing circle C_Y using least-squares fitting.

Each search path begins at $O' \approx X$. The robot moves until it cannot detect any nearby tags. By Definition 1, this can be a maximum of $2r$ in any direction (traveled twice) for a total cost of $16\frac{r}{v} + 4 \cdot \varepsilon$. A total of five targets are required to achieve the worst-case cost. Adding this to the worst-case cost of the initial search, plus the maximum displacement between the points O and O' gives,

$$T_{\text{multi}} = 17\frac{r}{v} + 2\frac{5}{2}\frac{r}{v} + 4 \cdot \varepsilon. \quad (3)$$

4.3 Discussion

The cost shown by Equation (3) may seem large. For example, given our system, v is approximately 2 meters per second and, for comparison, assume r is near 100 meters. Thus the total cost is approximately 19 minutes for the worst-case 5 targets. While we are not concerned with the aggregation displacing in this time, this may cause unnecessary drain on the limited operational life of the robot. To put this in context, compare this to the cost of taking two bearing measurements to initialize each target individually. Recall from Section 3.1 that a bearing measurement takes approximately 1-2 minutes. At least two measurements are required, resulting in 10-20 minutes for 5 targets, not counting the time to displace between measurement

locations. By amortizing the cost of a local search on a per-target basis, it is clear the search-based strategy will incur a lower cost to initialize larger aggregations.




Fig. 8 To extend the single target local search strategy, we need at least twelve search paths (separated by less than $\frac{\pi}{6}$) to intersect each sensing circle at least thrice.

To see the relative advantage of a two-phase search, consider the work required by a single-phase search. That is, upon detecting a non-zero signal strength, we could search along $K > 4$ search paths and attempt to intersect each nearby sensing circle. The necessary number of search paths can be found as follows. Refer to Figure 8. In this example two targets, X and Y are arranged along the x axis with respect to the starting location O . Assume the first search path moves along the x axis and the next search path is offset by an angle α . Then, to intersect the circle C_Y we require $2r \sin \alpha = r$. Solving, we get $\alpha = 30$ degrees, i.e., $K \geq 12$ search paths over 360 degrees. We call the resulting algorithm *Twelve-Path*. Note, unlike the *Four-Path* strategy, we must sample the entire list of frequencies in the lake over each of the twelve paths because we do not know until we are finished which tags belong to the aggregation. Hence the time taken to sample a frequency, and the total number of targets in the lake affect the cost of this strategy.

Because the distribution of the targets both in and between aggregations plays a large role in the *expected* search time, we compare these strategies in simulations.

5 Simulations

In the analysis presented in the previous section we assumed the time required to sample a frequency (t) was negligible. In practice, we may periodically stop the robot while sampling the frequencies to avoid radio interference from the electric drives, which takes some time. Second, we assumed the same sensing range r for all tags, when in practice it can be different for each tag. Finally, we evaluated the cost to initialize the targets in a single aggregation. In general, there can be more than one aggregation in the lake, each possibly containing different numbers of tagged fish. In this section, we investigate the role of the time spent in sampling the frequencies, the effect of multiple aggregations on total cost, and the effect of different sensing ranges on the time to initialize all targets.

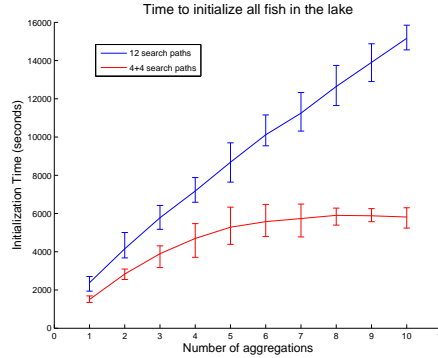


Fig. 9 Simulation comparing the time taken to initialize all 10 fish in the lake, as the number of aggregations varies. The Four-Path strategy performs better than the Twelve-Path. The bars indicate the minimum and maximum times, and the trend line plots the mean time of 50 instances.

We conducted simulations as follows. To evaluate a varying sensing range, r is drawn uniformly at random between $[50, 100]$ m for each tag. We vary the number of aggregations from 1 – 10 (with at least one fish each). The remaining fish are assigned randomly. The direction in which the robot enters the detection disk of the first target for each aggregation is also drawn uniformly at random between 0 and 2π radians. The velocity of the robot is given as v and is assumed fixed.

We compare Twelve-Path and Four-Path strategies presented in the previous section. Recall that the Twelve-Path (Figure 8) strategy moves along twelve search paths from the point of first detection, while sampling on the entire list of frequencies present in the lake. The Four-Path strategy (see Figure 6(b)) estimates the sensing circle for first tag detected, moves to the center of this estimated circle, samples all frequencies once to detect the list of frequencies present in the aggregation, and then moves along four search paths to estimate the sensing disks for only the subset of tags detected in the aggregation. Both produce an estimate of the sensing range and position of each nearby tag.

In Figure 9, we compare the mean, min and max time taken for executing both strategies for 50 iterations, as a function of the aggregation size M with total number of fish, $N = 10$. The sampling time per frequency is $t = 0.03$ sec (we obtain similar results for other choices of sampling time). We observe that the Four-Path strategy takes less time, as compared to the Twelve-Path strategy.

Figure 10 shows the time taken by the Four-Path strategy when size of one aggregation is increased (as opposed to the number of aggregations in Figure 9). For lower sampling time, we observe that the time to travel over the search paths dominates the time to sample for various frequencies. Since the distance traveled by the robot doesn't change significantly with increasing number of fish in the aggregation (by Definition 1), we see that the time taken scales well.

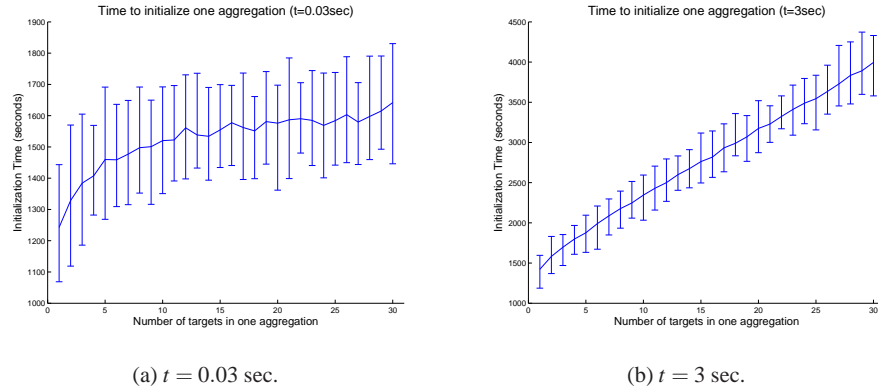



Fig. 10 Mean, minimum and maximum time taken as the number of fish increases in one aggregation for the Four-Path strategy. For lower sampling time, the time to travel dominates and thus scales well for larger aggregations.

6 Experiments


We implemented our initialization strategy on the mobile chassis shown in Figure 1. Three tags were deployed on Lake Gervais, MN, and their true locations were recorded for comparison (see Figure 11). The robot first detected the tag with frequency 48341 at the location marked *START* in Figure 11(a). The robot then executed the Four-Path strategy. After completing the first phase of the Four-Path strategy, we fit a circle to the points where we stopped detecting the signal for 48341 as shown. This circle was used as the $3\text{-}\sigma$ uncertainty ellipse of a 2D Gaussian distribution with the center of the circle used as the mean for initializing the estimate for this tag. The robot then traveled to the center of this circle and sampled the list of frequencies to detect nearby tags. The robot detected signals for frequencies 48931 and 48999 (48999 was due to radio interference and not an actual tag—the Localization strategy received no valid measurements and discarded this estimate).

The robot then executed the second phase of the Four-Path strategy, where it searched for frequencies detected at the center of the initial circle as shown in Figure 11(b). The corresponding hypothesis for all tags are shown relative to the true locations. Using this initial hypothesis, the robot then executed the active localization algorithm described in [13]. Figure 11(c) shows the execution of this localization algorithm, the measurement locations selected for each tag (triangles), and the bearing measured (black lines).

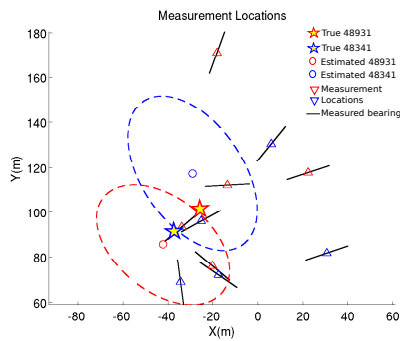
The final estimates for the two actual tags in the aggregation after five measurements (48341 and 48931) are shown using the $3\text{-}\sigma$ uncertainty ellipse. Figure 11(d) shows the GPS location of the tags along with the initial and final estimates. The final covariance for 48341 had eigenvalues 56m^2 and 168m^2 (corresponding to an error ellipse with radii 7m and 12m), starting from an initial covariance with eigenval-




(a) First Local Search (Section 4.1)



(b) Aggregation Search (Section 4.2)



(c) Localization Output (see [13])



(d) Final Estimates vs True

Fig. 11 A successful experiment demonstrating the local search strategy and localization steps.

ues 1380m^2 . The final covariance for 48931 had eigenvalues 49m^2 and 127m^2 (radii 7m and 11m), starting from an initial covariance with eigenvalues 1758m^2 . The final error for 48341 and 48931 were 27m and 23m respectively.

7 Conclusions

We are working toward the goal of localizing multiple targets in a known environment in bounded time. The complicated interplay of target distribution, sensing range, measurement noise, and ambiguous measurement model makes each phase independently interesting. Here we presented a strategy to initialize consistent hypotheses for multiple targets in an aggregation. In our future work, we plan to extend our system to multiple robots and incorporate fish mobility models. To extend

this algorithm to multiple robots, we must account for communication constraints between the robots and develop allocation algorithms which guarantee the work is distributed evenly. For mobile targets, we must both develop motion models for fish and develop new search and localization algorithms based on these models. One possible approach is to model the fish as adversarial—part of our ongoing work.

Acknowledgements The authors would like to thank Brett Miller and Mary Haedrick for their involvement in gathering telemetry data. This work is supported in part by NSF Awards #1111638, #0916209, #0917676 and #0936710.

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. P. G. Bajer and P. W. Sorensen. Recruitment and abundance of an invasive fish, the common carp, is driven by its propensity to invade and reproduce in basins that experience winter-time hypoxia in interconnected lakes. *Biological Invasions*, 12(5):1101–1112, Aug. 2009.
3. J. Fink and V. Kumar. Online methods for radio signal mapping with mobile robots. In *Proc. IEEE International Conference on Robotics and Automation*, pages 1940–1945, 2010.
4. G. Hollinger and J. Djughash. Target tracking without line of sight using range from radio. *Autonomous Robots*, (July 2011):1–14, 2011.
5. P. Johnsen and A. Hasler. Winter aggregations of carp (*Cyprinus carpio*) as revealed by ultrasonic tracking. *Transactions of the American Fisheries Society*, 106(6):556–559, 1977.
6. C. Kim, D. Song, Y. Xu, and J. Yi. Localization of multiple unknown transient radio sources using multiple paired mobile robots with limited sensing ranges. In *Proc. IEEE International Conference on Robotics and Automation*, pages 5167–5172, 2011.
7. C. Penne and C. Pierce. Seasonal Distribution, Aggregation, and Habitat Selection of Common Carp in Clear Lake, Iowa. *Transactions of the American Fisheries Society*, 137(4):1050–1062, July 2008.
8. D. Song, C.-Y. Kim, and J. Yi. Simultaneous localization of multiple unknown and transient radio sources using a mobile robot. *IEEE Transactions on Robotics*, 28(3):668–680, June 2012.
9. O. Tekdas, N. Karnad, and V. Isler. Efficient strategies for collecting data from wireless sensor network nodes using mobile robots. In *14th International Symposium on Robotics Research (ISRR)*, 2009.
10. P. Tokekar, D. Bhadauria, A. Studenski, and V. Isler. A Robotic System for Monitoring Carp in Minnesota Lakes. *Journal of Field Robotics*, 27(6):779–789, 2010.
11. P. Tokekar, E. Branson, J. Vander Hook, and V. Isler. Coverage and active localization for monitoring invasive fish with an autonomous boat. *IEEE Robotics and Automation Magazine*, 2012. Note: to appear.
12. J. Twigg, J. Fink, L. Paul, and B. Sadler. RSS Gradient-Assisted Frontier Exploration and Radio Source Localization. In *Proc. IEEE International Conference on Robotics and Automation*, 2012.
13. J. Vander Hook, P. Tokekar, and V. Isler. Cautious Greedy Strategy for Bearing-based Active Localization: Experiments and Theoretical Analysis. In *Proc. IEEE International Conference on Robotics and Automation*, 2012.
14. M. J. Weber and M. L. Brown. Effects of Common Carp on Aquatic Ecosystems 80 Years after Carp as a Dominant: Ecological Insights for Fisheries Management. *Reviews in Fisheries Science*, 17(4):524–537, Oct. 2009.