

Stochastic Source Seeking in Complex Environments

Nikolay Atanasov, Jerome Le Ny, Nathan Michael, and George J. Pappas

Abstract—The objective of source seeking problems is to determine the minimum of an unknown signal field, which represents a physical quantity of interest, such as heat, chemical concentration, or sound. This paper proposes a strategy for source seeking in a noisy signal field using a mobile robot and based on a stochastic gradient descent algorithm. Our scheme does *not* require a prior map of the environment or a model of the signal field and is simple enough to be implemented on platforms with limited computational power. We discuss the asymptotic convergence guarantees of algorithm and give specific guidelines for its application to mobile robots in unknown indoor environments with obstacles. Both simulations and real-world experiments were carried out to evaluate the performance of our approach. The results suggest that the algorithm has good finite time performance in complex environments.

I. INTRODUCTION

The ability to detect the source of a signal is a fundamental problem in nature. At a microscopic level, some bacteria are able to find chemical, light, and magnetic sources [1], [2]. At a macroscopic level similar behavior can be observed in predators who seek a food source using their sense of smell. Reproducing this behavior in mobile robots can be used to perform complex missions such as environmental monitoring [3], [4], intelligence, surveillance, and reconnaissance [5], and search and rescue operations [6].

This paper considers the following scenario. A mobile robot is placed in an unknown environment with the goal of locating the source of an unknown signal field, which could express the spatial distribution of magnetic force, heat, wireless signal, or chemical concentration. The robot has access to noisy measurements of the signal *only* at its current position. Since this source seeking problem arises in a variety of situations we make minimal assumptions about the robot, the environment, and the signal. Our approach consists of climbing the gradient of the signal field by using a stochastic approximation technique to deal with the underlying noise. The main advantage of our strategy is that it does not rely on a prior map of the environment or a model of the signal field and can be applied to robots with various types of dynamics, no localization capabilities, and limited computational power.

We do not claim that this approach performs better than model-based ones when an accurate model of the signal is available. However, when the environment is unknown, constructing a good model a priori is not possible. For example, the Navier-Stokes equations cannot be used to model the diffusion of a gas signal without a prior map.

This work was supported by the ARL MAST-CTA award W911NF-08-2-0004.

The authors are with the Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104, {atanasov, jeromel, michael, pappasg}@seas.upenn.edu

On-line modeling of the signal might not be feasible either because it requires time and computational resources, which are limited on small platforms and in time-critical missions.

Recent work developing source seeking and gradient climbing strategies includes [3], [7], which rely on a model and require several sensing robots moving in a formation, whereas we only have a single sensing robot at our disposal. Furthermore, the authors do not address the effect of model inaccuracies on the performance. Krstić and coworkers [8], [9] apply techniques from extremum-seeking control to find a local maximum of a field, assuming noiseless measurements. In our case, the stochastic variations in the measurements can be significant and cannot in general be neglected. Moreover, the extremum-seeking approach is applicable only to specific robot dynamics (e.g. unicycle, differential drive), while our method can be used with more general robots. Stanković and Stipanović [10] extended the extremum-seeking methodology to the case of noisy signals but considered obstacle-free environments only. Similarly, other recent approaches to source seeking [11], [12] consider only the obstacle-free case, while we apply our technique to environments with different obstacle configurations.

The simplicity of the stochastic approximation approach in the face of the complexity of the envisioned scenarios calls for an in-depth study of implementation issues. While existing stochastic gradient algorithms guarantee *asymptotic* convergence in *convex* environments, it is unclear if the theoretical guarantees translate into finite-time results in non-convex environments, which are meaningful in applications.

The contributions of this paper are as follows. First, we adapt an existing stochastic gradient algorithm to provide a trajectory in the robot's local coordinate frame. Second, we address the main challenge of applying this algorithm in practice, i.e. the choice of parameters, by providing only a few intuitive parameters to control the algorithm. Finally, we examine the *finite-time* performance of the algorithm and its applicability to a *real* wireless signal in a *non-convex* environment. We consider a wireless signal because it is very noisy and difficult to model and yet most approaches for wireless source seeking are model-based (e.g. [13]).

The rest of the paper is organized as follows. In section II we describe the considered scenario precisely. Section III describes the algorithm and states its theoretical convergence guarantees. We discuss practical modifications to the algorithm in section IV in order to incorporate the environment geometry and the robot characteristics. Finally, in section V we present results from simulations and real-world experiments and discuss the validity of our approach.

II. SOURCE SEEKING PROBLEM

Consider a robot with dynamics described in the world coordinate frame by $\dot{s} = f(s, u)$, where $u \in \mathbb{R}^m$ is the control input and $s(t) = (x(t), \theta(t))$ is the robot's state with translation $x(t) \in \mathbb{R}^n$ and orientation $\theta(t) \in \Theta \cong \mathbb{R}^l$. The robot evolves in a workspace $W \subset \mathbb{R}^n$, which is a priori unknown, and has a field of view $F(s) \subseteq W$. We assume that the workspace is the free space in the environment, excluding the obstacles. The task of the robot is to reach the minimum of a scalar field $L : W \rightarrow \mathbb{R}$, which represents a physical quantity of interest. Due to sensor and environment noise, the robot has access only to measurements of the form $\hat{L}(x) = L(x) + \nu(x)$ at its current position $x \in W$, where $\nu(x)$ is a random noise term.

Our approach is to divide the source seeking problem into two parts. First, we generate a trajectory converging to the unknown source and then steer the robot along it. This separation allows applying the stochastic source seeking strategy to robots which have various dynamics and employ different navigation schemes.

Problem 1 (Source Seeking): Generate a sequence of points $\{p_k\} \subseteq W$, which satisfies:

$$\lim_{k \rightarrow \infty} p_k = \bar{p} \in P^* \text{ almost surely (a.s.),}$$

where P^* is the set of points in the workspace, which minimize the signal field.

The second part of the task is to generate control inputs, which drive the robot along the trajectory specified by Problem 1 and we address it by making the following assumption:

Assumption 1: The source seeking robot is able to follow a sequence of reachable points $\{p_k\}$ in the free space W using a combination of planning and control.

The combination of planning and control is a widely used approach to robot navigation both in literature and practice and we adopt it to ensure that our source seeking strategy can be applied to robots with various dynamics, which already have a path-following scheme in place. As a result, the burden of generating points p_k , which are reachable and easy to follow is placed on the algorithm used for Problem 1.

III. STOCHASTIC APPROXIMATION ALGORITHM FOR SOURCE SEEKING

This section presents a solution to a relaxed version of Problem 1, in which P^* is a nonempty set of *local* minima of the signal field L and the workspace W is *convex*. We employ a modified version of the random direction stochastic approximation (RDSA) algorithm [14] in order to generate the required sequence $\{p_k\} \subseteq W$. The general form of a stochastic approximation algorithm is:

$$p_{k+1} = \Pi_W[p_k - a_k \hat{g}_k(p_k)], \quad k = 0, 1, \dots, \quad (1)$$

where $\hat{g}_k : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is an approximation to the true gradient of the signal field L at the point p_k , $a_k \in \mathbb{R}_{>0}$ is a sequence of gain coefficients, and Π_W denotes projection onto the

constraint set W . We propose the following gradient estimate:

$$\hat{g}_k(p_k) := \sqrt{n} \left[\frac{\hat{L}(p_{k+}) - \hat{L}(p_{k-})}{\|p_{k+} - p_{k-}\|_2} \right] T_k d_k, \quad (2)$$

where $d_k \in \mathbb{R}^n$ is a random *direction vector* with probability distribution $\mu_k : \mathbb{R}^n \rightarrow [0, 1]$, and T_k is a rotation matrix, whose role is explained below. Also, $p_{k+} := \Pi_W[p_k + c_{k+} T_k d_k]$ and $p_{k-} := \Pi_W[p_k - c_{k-} T_k d_k]$ are two sampling points determined by gain coefficients $c_{k-}, c_{k+} \in \mathbb{R}_{\geq 0}$.

At iteration k of the algorithm, the signal field is sampled at points p_{k-} and p_{k+} in order to estimate its directional gradient at p_k . Then, a step is taken in the negative direction of the gradient estimate to generate the next estimate of the source location p_{k+1} . Repeating this procedure generates a sequence of points $\{(p_k, p_{k-}, p_{k+})\}$, which represents a trajectory in the world coordinate frame. The role of the matrices T_k is to transform this trajectory to the body coordinate frame of the robot. At time t_k , which is the beginning of the k -th iteration, the sampling points can be written in the robot body frame:

$$\begin{aligned} z_{k-} &= \Pi_W[-c_{k-} R(-\theta(t_k)) T_k d_k] \\ z_{k+} &= \Pi_W[c_{k+} R(-\theta(t_k)) T_k d_k] \\ z_{k+1} &= \Pi_W[-a_k R(-\theta(t_k)) \hat{g}_k(0)], \end{aligned}$$

where $\theta(t_k)$ is the robot orientation and $R(-\theta(t_k))$ is a rotation matrix. Choosing $T_k = R(\theta(t_k))$, allows the robot to run the algorithm *without* the need for localization:

$$z_{k+1} = \Pi_W \left[-a_k \sqrt{n} \left[\frac{\hat{L}(z_{k+}) - \hat{L}(z_{k-})}{\|z_{k+} - z_{k-}\|} \right] d_k \right] \quad (3)$$

The *design parameters* of the algorithm are a_k , d_k (i.e. μ_k), c_{k-} , c_{k+} , and T_k . Successful application of this algorithm to a robot platform in a realistic environment requires a careful choice of the design parameters, which takes the robot capabilities and the geometric constraints of the environment into account in order to ensure that the resulting sequence $\{(p_k, p_{k-}, p_{k+})\}$ is easy to follow.

Proposition 1: If the assumptions listed below are satisfied, the algorithm in (1) and (2) is a solution to a relaxed version of Problem 1, in which P^* is a nonempty set of *local* minima of signal field L in a *convex* workspace W .

Assumptions on the problem parameters:

- (A1) L is bounded and three times cont. differentiable
- (A2) For all $k \in \mathbb{N}$, the random noise on L satisfies:

$$\begin{aligned} \mathbb{E}[\nu(p_{k+})^2] &< \infty, \text{ and } \mathbb{E}[\nu(p_{k-})^2] < \infty \\ \mathbb{E}[\nu(p_k) | p_0, \dots, p_k, d_k] &= 0 \text{ a.s.,} \end{aligned}$$

- (A3) The iterates p_k remain bounded: $\sup_k \|p_k\| < \infty$ a.s.

Assumptions on the design parameters:

- (Sample point gains) $c_k := \max\{c_{k-}, c_{k+}\} > 0$, $c_k \rightarrow 0$
- (Gradient gains) $a_k > 0$, $a_k \rightarrow 0$, $\sum_{k=0}^{\infty} a_k = \infty$, and $\sum_{k=0}^{\infty} a_k^2 / c_k^2 < \infty$
- (Direction vectors) d_k are i.i.d. with distribution μ_k , which is symmetric with respect to reflections about the coordinate axes, and satisfy $\mathbb{E}[\|d_k\|_2^2] = n$ and $\mathbb{E}[d_k d_k^T] = I$.

Proof: Let $d'_k := T_k d_k$ and notice that:

$$\begin{aligned} \mathbb{E} \left[\frac{\|p_{k_+} - p_{k_-}\|_2}{\sqrt{n}} \right] &= \frac{(c_{k_+} + c_{k_-}) \|T_k\|_2 \mathbb{E}(\|d_k\|_2)}{\sqrt{n}} \\ &= (c_{k_+} + c_{k_-}) \end{aligned}$$

Then, the result follows from a slight modification in the proof of the RDSA algorithm [14, Chapter 5.6 and 10.7], which takes into account that $c_{k_-} d'_k \neq c_{k_+} d'_k$. ■

Despite initial impressions, the assumptions on the problem parameters are quite weak. The smoothness condition on L (A1) guarantees that the estimate of the gradient \hat{g}_k is almost surely an unbiased estimate of g_k to within an $O(c_k^2)$ error, which is guaranteed to be small by the assumptions on the design parameters. In related work [3], [7], the assumptions on the signal field are stronger than ours because (A1) is satisfied whenever the real signal can be modeled reasonably by an almost everywhere three times differentiable and bounded function. For example, a wireless signal is typically modeled by a function, which depends on the inverse squared distance between the measurement point and the source and is smooth everywhere except at the source [15]. Assumption (A2) places no restrictions on the distribution of the signal noise and only requires that it has a finite second moment, which is valid for a variety of situations. For example, the noise associated with the free-space loss of a wireless signal is modeled as a Rayleigh distribution, which has finite variance. The second part of (A2) requires that the noise is a martingale difference but if necessary this can be weakened to the case of correlated noise (see [14, Chap 6]). Assumption (A3) is the most challenging to verify in the general theory of stochastic approximation but in our application it can be dealt with by assuming that the workspace W is bounded.

IV. INCORPORATING GEOMETRIC CONSTRAINTS AND ROBOT CHARACTERISTICS

This section provides specific guidelines for successful application of the source seeking algorithm to a mobile robot. The assumption that the workspace W is convex is relaxed since our goal is to apply the algorithm in an environment with obstacles. As a result, the theoretical almost sure convergence result from the previous section is lost but we would like to show that the algorithm still works in practice with the appropriate choice of parameters.

Of course, there are environments in which the obstacle configuration prevents the robot from finding the source even if it is reachable. Often, the estimate of the source location, which the robot makes based on its approximation of the directional gradient, is outside of the explored portion of the environment. The robot attempts to reach it for a certain time limit τ , after which the next iteration is started, i.e. the robot samples at its current position to get a new source estimate. The planned trajectory to the source estimate cannot be followed indefinitely because in the worst case, there might not be a feasible path. The stochastic approximation algorithm will not be successful in environments which require the robot to go in a direction opposite to the directional gradient

of the signal field for a period of time longer than τ because the robot might get stuck going back and forth.

The rest of this section concentrates on the choice of the design parameters a_k , μ_k , c_{k_-} , and, c_{k_+} with two goals in mind: first to take the geometric constraints of the environment and the robot characteristics into account so that the generated sampling points $\{(p_k, p_{k_-}, p_{k_+})\}$ are easy to follow and second to provide only a few high-level and intuitive parameters to the user of the algorithm. As a result, the choice of the design parameters is simplified to specifying only the following two constants:

- *Aggressiveness factor* $r > 0$: a constant determining the size of the steps of the algorithm. Intuitively, r is the amount of meters, by which the position of the robot is expected to change in the early iterations. The further away the source is expected to be, the larger the value of r should be.
- *Stability factor* $s \geq 0$: a constant which allows for large steps in the early iterations of the algorithm without causing instabilities. It should be set to 5 – 10 percent of the expected number of iterations of the algorithm.

A. Choosing the direction vectors d_k

Several choices for the distributions of d_k have been considered in literature [16], [14], [17] with a Bernoulli distribution in each coordinate being preferred in applications. We note that the Bernoulli distribution is an optimal choice only for signal fields, which are aligned with the coordinate axes in a way that their third cross-derivatives $\partial^3 L / \partial x_i \partial x_j \partial x_k$ are all zero [18]. Since the signal field will not be axis aligned in a non-convex W it is beneficial to choose d_k uniformly from all possible directions. In particular, we let μ_k be a *Shell distribution*, which is defined as follows: choose $d_i \sim N(0, 1)$ for $i = 1, \dots, n$ and then rescale the vector to guarantee that its magnitude is \sqrt{n} as required in *Proposition 1*.

B. Choosing the gradient gains a_k

The usual form used for the gradient gain coefficients in literature [16], [14] is:

$$a_k = \frac{a}{(k + 1 + s)^\alpha}, \quad k = 0, 1, \dots, \quad (4)$$

where $a > 0$ is a constant, s is the stability factor mentioned earlier, and $\alpha > 0$ governs the decay rate for the gains and can be set to $\alpha = 0.602$ as suggested in [16, chap. 6].

A modification to this choice is required for our application. If the numerator a is constant, the gain coefficients a_k are monotonically decreasing, which is not desirable because the robot will be taking decreasing steps along the gradient and it might get trapped in a location, where the magnitude of the gradient estimate is small. We replace a with a varying numerator $a_k^0 > 0$, which is inversely proportional to the magnitude of the gradient estimate. This is beneficial because when the magnitude of the gradient estimate is large, the robot takes small steps in a controlled manner towards the source but if the magnitude of the gradient estimate

decreases, the gain coefficients increase allowing the robot to follow the gradient even if the signal field is very flat. Based on these observations we propose:

$$a_k^0 = \frac{r(1+s)^\alpha}{\frac{1}{w} \sum_{j=k+1-w}^k \frac{1}{n} \|\hat{g}_j(p_j)\|_1}, \quad k = 0, 1, \dots, \quad (5)$$

where r is the aggressiveness factor, $w \in \mathbb{N}$ is a window over which the mean magnitude of the elements of \hat{g}_k is averaged, and n is the dimension of \hat{g}_k . The size of w determines the speed at which a_k reacts to changes in the magnitude of the gradient estimate. We used $w = 10$ in our experiments.

C. Choosing the sample point gains c_{k-} and c_{k+}

A typical schedule used in the stochastic approximation literature [16], [14] for the sample point gain coefficients is:

$$c_{k-} = \frac{c_-^0}{(k+1)^\gamma}, \quad c_{k+} = \frac{c_+^0}{(k+1)^\gamma}, \quad k = 0, 1, \dots, \quad (6)$$

where $c_-^0, c_+^0 \in \mathbb{R}_{\geq 0}$ are constants and $\gamma \in \mathbb{R}_{>0}$ is the gain decay rate. In practical applications, a slow decay leads to better finite sample performance and a good choice is $\gamma = 0.101$ (see e.g. [16]). The constants c_-^0 and c_+^0 are typically set to the standard deviation of the measurement noise in L at the current position of the robot by measuring the signal several times.

The gains c_{k-} , c_{k+} , and the direction vector d_k affect the position of the sampling points p_{k-} and p_{k+} as specified in (2). While the choice in (6) is applicable to the obstacle-free case, when dealing with a general environment it needs to be modified to accommodate for the constraints introduced by the obstacles. Let b_{k+} be the value for c_{k+} originally suggested by the form in (6). Suppose that the robot is traveling from its previous estimate p_{k-1} towards p_k . As soon as p_k is in the field of view $F(s)$, we choose p_{k+} in $F(s)$ to ensure that it is reachable. Algorithm 1 with $y \leftarrow p_k$ and $r \leftarrow b_{k+}$ shows how to sample d_k and simultaneously choose c_{k+} , with a magnitude as close to b_{k+} as possible. The shell distribution is sampled for a direction d_k , which

Algorithm 1 Rejection Sampling of d_k

- 1: **Input:** Position $y \in \mathbb{R}^2$ in the robot body frame, radius $r \in \mathbb{R}_{>0}$, and field of view $F(s) \subseteq \mathbb{W}$
 - 2: **Output:** Step size c_{k+} and direction vector d_k
 - 3: Let dS be a small area element on the surface of the hypersphere of radius \sqrt{n} centered at y
 - 4: $count \leftarrow 1$
 - 5: **repeat**
 - 6: Sample d_k from the shell distribution
 - 7: **if** $(1 - \mathbb{P}(\{d_k \in dS\}))^{count} < 0.05$ **then**
 - 8: decrease r ; $count \leftarrow 1$
 - 9: **else**
 - 10: $count \leftarrow count + 1$
 - 11: **until** $r = 0$ or there is a path between y and $(y + rd_k)$ in $F(s)$
 - 12: **return** d_k and $c_{k+} \leftarrow r$
-

selects a possible sample point $q = p_k + rd_k$. Line 11 checks if the path from p_k to q is within the robot's field of view and if so the chosen values for d_k and $c_{k+} \leftarrow r$ are returned. Otherwise, another sample for d_k is chosen. Thus, the allowable values for $p_{k+} = p_k + c_{k+}d_k$ lie on the



Fig. 1. A Scarab robot equipped with an XBee-PRO RF module.

intersection of the field of view $F(s)$ and the ball of radius $r\sqrt{n}$ centered at p_k . Due to the obstacles, there might not be a feasible choice for d_k with the specified radius r . When the probability of selecting a sampling point in any small region dS on the surface of the sphere of radius r is above 95% but a suitable direction has not been chosen yet, the radius is decreased (Line 6).

Once the direction d_k is known, the choice of c_{k-} is the maximum distance that can be traveled along $-d_k$ starting from p_k and up to b_{k+} or until an obstacle is reached.

V. PERFORMANCE EVALUATION

We carried out simulations and real-world experiments in an indoor environment using a differential drive robot to seek a wireless signal source in order to evaluate the validity of our approach to source seeking. We used an A^* planner to generate a feasible path along the sampling points supplied by the algorithm. The robot was controlled using a predictive controller, which simulates possible trajectories over a short time horizon and evaluates them based on a multi-objective cost function, which takes into account orientation correctness, proximity to the planned path, speed, obstacle avoidance, and distance to the goal. Details about the implementation of the controller are presented in [19].

A. Platform Description

1) *Robot:* The real-world experiments were carried out on a nonholonomic Scarab robot (Fig. 1), which is equipped with an on-board computer, power management system, wireless communication, and is actuated by stepper motors. It uses a Hokuyo URG laser range finder for obstacle detection. Its physical dimensions are 30 x 28 x 20 cm with a mass of 8kg. Two Scarabs were used in the source seeking scenario, one transmitting a wireless signal and the other acting as a seeker. The received signal strength (RSS) between the source and the seeker was measured using a XBee-PRO RF module.

2) *Simulation software:* Simulations were carried out in MATLAB using a simulated wireless signal field of the form: $L(z) = P_{tx}(z) + G_{tx}(z) - L_{tx}(z) - L_{fs}(z) - L_m(z) + G_{rx}(z) - L_{rx}(z)$, where P_{tx} is the transmitter output power, G_{tx} is the transmitter antenna gain, L_{tx} represents transmitter losses, G_{rx} is the receiver antenna gain, and L_{rx} represents receiver losses. Furthermore, L_{fs} is a free space loss calculated as follows:

$$L_{fs}(z) = -27.55 + 20 \log_{10} 2400 + 20 \log_{10} \|z - z_{src}\| - R$$

where z is the measurement position and $R = \sqrt{X^2 + Y^2}$ is the signal noise with parameters $X \sim N(\nu \cos \theta, \sigma^2)$ and




Fig. 2. A path followed by the robot after 20 iterations of the source seeking algorithm in an obstacle-free environment. The blue circles indicate the positions at which the robot measured the signal strength. The white dots indicate the starting and final positions of the robot and are 21.85 m and 1.47 m away, respectively, from the actual position of the source.




Fig. 3. Distance to the source calculated over 50 independent replications of the algorithm for several run lengths in an obstacle-free environment. The source was 21.85 m away. The bars show one standard deviation.

$Y \sim N(\nu \sin \theta, \sigma^2)$ for some real θ , ν , and σ^2 . Finally, a multi-wall path loss model L_m described in [20] was used.

B. Simulation Results

A set of simulations was run for an obstacle-free environment using the wireless field model described above. The positions of the source and the seeker as well as a sample trajectory after 20 iterations of the algorithm are shown in Fig. 2. The final distance to the source after running the algorithm for several run lengths is shown in Fig. 3. For each run length, the performance was averaged over 50 independent runs of the RDSA algorithm. These results confirm experimentally the convergence guarantees for a convex (obstacle-free) workspace. Starting 21.85 meters away and running the algorithm for only 30 iterations resulted in a mean error in the source estimate of less than 30 cm. For these simulations we used $r = 1.5$, $s = 4$, and $w = 10$.

Next, we applied the algorithm to non-convex environments, for which convergence to the source cannot be guaranteed theoretically. Fig. 4 shows an example trajectory followed by the source seeker during a simulated run on a map of our engineering department. We simulated 50 independent replications for seven different run lengths to investigate the finite-sample performance (Fig. 5). We can see that as the number of iterations increases, the mean distance from the estimated source position to the actual




Fig. 4. A path followed by the robot while running the source seeking algorithm for 20 iterations. The white dots indicate the starting and the final positions. The seeker and the source were 17.85 m apart initially and 0.75 m apart in the end.




Fig. 5. Distance to the source calculated over 50 independent replications of the algorithm for several run lengths in an environment with obstacles. The initial distance from the source was 17.85 m. The bars show one standard deviation from the mean.

one decreases. The variance in the position of the source estimate in this case is much larger than in the obstacle-free case because when obstacles are present the robot can get trapped as mentioned in the beginning of Section IV and then its estimate of the source does not improve regardless of how long the algorithm is run for. The parameters in this case were $r = 2.5$ and $s = 7$.

C. Experiment Results

The simulations show that the finite time performance of the algorithm is good even in non-convex environments but suffer from one drawback. It is difficult to develop a realistic RSS model for an indoor environment because temporal and spatial fluctuations in the signal due to fading and shadowing are very significant [15]. To evaluate the applicability of the stochastic gradient algorithm to a real wireless signal and to study the implementation issues in detail, we resorted to real-world experiments. Fig. 6 shows the map of the environment and the trajectory followed by the robot after 10 iterations of the algorithm in one of the experiments. The parameters were $r = 3.5$ and $s = 4$. Several experiments with various starting positions of the source and the seeker were run in the same environment. While we do not have enough data to reproduce the extensive convergence analysis shown in the simulations subsection, the experiments showed that the algorithm is applicable to real signals.




Fig. 6. Path followed by the robot in a real experiment on the fourth floor of the ESE Department, UPenn, for 10 iterations of the algorithm. The seeker was 17.9 m away from the source initially. The red circle shows the final estimate of the source location and is 2.2 m away from the actual one, denoted by the blue cross. See attached video or http://www.seas.upenn.edu/~atanasov/ICRA2012_StochasticSourceSeeking.mp4




Fig. 7. History of measurements taken by the robot while seeking for the source in a real experiment.

A video demonstrating the performance of the source seeking algorithm for the scenario shown in Fig. 6 is submitted as an attachment. The robot's movement is less erratic when further away from the source because the change in the gradient is significant and easy to discern. Closer to the source, however, the signal field is relatively flat and the noise in the measurements affects the gradient estimation significantly. Fig. 7 shows the wireless signal measurements taken by the robot in the experiment shown in the video.

VI. CONCLUSION

This paper presents a strategy for estimating the position of a signal source in a complex environment, using a stochastic gradient descent algorithm. Our approach is applicable to robots with various types of dynamics, signal fields corrupted with stochastic noise, and environments with general obstacle configurations. No prior knowledge of the environment or the signal field is required.

Among the drawbacks of the stochastic approximation algorithm is its sensitivity to the choice of parameters. We have provided specific guidelines which take the geometric constraints of the environment into account and reduce the choice to two intuitive high-level parameters. These new parameters are the main determinant of the algorithm's

performance and are easier to specify because they can be interpreted in physical terms.

Even though the convergence of the algorithm is theoretically guaranteed only asymptotically and for obstacle-free environments, our experiments suggest that the algorithm has good finite time performance in complex real environments. In future work we plan to carry out additional real-world experiments in order to reproduce the extensive convergence analysis which was performed for the simulations. Additionally, we would like to investigate how much leniency there is in the choice of the high-level parameters and to apply the algorithm to other signals, which are difficult to model.

REFERENCES

- [1] R. Lux and W. Shi, "Chemotaxis-guided movements in bacteria," *Crit Rev Oral Biol Med*, vol. 15, no. 4, pp. 270–220, July 2004.
- [2] R. Frankel, D. Bazyliniski, M. Johnson, and B. Taylor, "Magneto-aerotaxis in marine coccoid bacteria," *Biophys J*, vol. 73, no. 2, pp. 994–1000, 1997.
- [3] P. Ögren, E. Fiorelli, and N. Leonard, "Cooperative control of mobile sensor networks: Adaptive gradient climbing in a distributed environment," *IEEE Transactions on Automatic Control*, vol. 49, no. 8, pp. 1292–1302, August 2004.
- [4] G. S. Sukhatme, A. Dhariwal, B. Zhang, C. Oberg, B. Stauffer, and D. A. Caron, "Design and development of a wireless robotic networked aquatic microbial observing system," *Environmental Engineering Science*, vol. 24, no. 2, pp. 205–215, 2007.
- [5] P. E. Rybski, S. A. Stoeter, M. D. Erickson, M. Gini, D. F. Hougen, and N. Papanikolopoulos, "A team of robotic agents for surveillance," *Proceedings of the International Conference on Autonomous Agents*, pp. 9–16, 2000.
- [6] V. Kumar, D. Rus, and S. Singh, "Robot and sensor networks for first responders," *IEEE Pervasive Computing*, vol. 3, pp. 24–33, 2004.
- [7] W. Wu and F. Zhang, "Experimental Validation of Source Seeking with a Switching Strategy," *Proceedings of the IEEE Conference on Robotics and Automation*, pp. 3835–3840, 2011.
- [8] C. Zhang, D. Arnold, N. Ghods, A. Siranosian, and M. Krstić, "Source seeking with non-holonomic unicycle without position measurement and with tuning of forward velocity," *Systems and Control Letters*, vol. 56, pp. 245–252, 2007.
- [9] S. J. Liu and M. Krstić, "Stochastic source seeking for nonholonomic unicycle," *Automatica*, vol. 46, no. 9, pp. 1443–1453, 2010.
- [10] M. S. Stanković and D. M. Stipanović, "Extremum seeking under stochastic noise and applications to mobile sensors," *Automatica*, 2010.
- [11] S. Azuma, M. Sakar, and G. J. Pappas, "Nonholonomic source seeking in switching random fields," *IEEE Conference on Decision and Control*, December 2010.
- [12] N. Ghods and M. Krstić, "Source Seeking With Very Slow or Drifting Sensors," *Journal of Dynamic Systems, Measurement, and Control*, vol. 133, no. 4, July 2011.
- [13] J. Fink and V. Kumar, "Online Methods for Radio Signal Mapping with Mobile Robots," *IEEE International Conference on Robotics and Automation*, May 2010.
- [14] H. J. Kushner and G. G. Yin, *Stochastic Approximation and Recursive Algorithms and Applications*, 2nd ed. Springer, 2003.
- [15] A. Goldsmith, *Wireless Communications*. Cambridge University Press, 2005.
- [16] J. Spall, *Intro to Stochastic Search and Optimization*. Wiley, 2003.
- [17] J. Le Ny and G. J. Pappas, "Sensor-based robot deployment algorithms," in *Proceedings of the IEEE Conference on Decision and Control*, December 2010.
- [18] J. Theiler and J. Alper, "On the choice of random directions for stochastic approximation algorithms," *IEEE Transactions on Automatic Control*, vol. 51, no. 3, pp. 476–481, 2006.
- [19] N. Atanasov, "Control module for an unmanned ground vehicle," University of Pennsylvania, Tech. Rep., 2011. [Online]. Available: [http://www.seas.upenn.edu/~atanasov/\[ESE650-Project5\]Atanasov_report.pdf](http://www.seas.upenn.edu/~atanasov/[ESE650-Project5]Atanasov_report.pdf)
- [20] F. Capulli, C. Monti, M. Vari, and F. Mazzenga, "Path Loss Models for IEEE 802.11a Wireless Local Area Networks," *Int. Sym. Wireless Comm. Sys.*, 2006.