

Nikolay A. Atanasov¹

Department of Electrical and
Systems Engineering,
University of Pennsylvania,
Philadelphia, PA 19104
e-mail: atanasov@seas.upenn.edu

Jerome Le Ny

Department of Electrical Engineering and GERAD,
École Polytechnique de Montréal,
Montréal, QC H3T-1J4, Canada
e-mail: jerome.le-ny@polymtl.ca

George J. Pappas

Department of Electrical and
Systems Engineering,
University of Pennsylvania,
Philadelphia, PA 19104
e-mail: pappasg@seas.upenn.edu

Distributed Algorithms for Stochastic Source Seeking With Mobile Robot Networks

Autonomous robot networks are an effective tool for monitoring large-scale environmental fields. This paper proposes distributed control strategies for localizing the source of a noisy signal, which could represent a physical quantity of interest such as magnetic force, heat, radio signal, or chemical concentration. We develop algorithms specific to two scenarios: one in which the sensors have a precise model of the signal formation process and one in which a signal model is not available. In the model-free scenario, a team of sensors is used to follow a stochastic gradient of the signal field. Our approach is distributed, robust to deformations in the group geometry, does not necessitate global localization, and is guaranteed to lead the sensors to a neighborhood of a local maximum of the field. In the model-based scenario, the sensors follow a stochastic gradient of the mutual information (MI) between their expected measurements and the expected source location in a distributed manner. The performance is demonstrated in simulation using a robot sensor network to localize the source of a wireless radio signal.

[DOI: 10.1115/1.4027892]

1 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 discusses how to control a team of mobile robotic sensors to locate the source of a noisy signal, which represents a physical quantity such as magnetic force, heat, radio signal, or chemical concentration. We distinguish between two cases: *model-free* and *model-based*. The first scenario supposes that the sensors receive measurements without knowledge of the signal formation process. This is relevant when the signal is difficult to model or the environment is unknown a priori. Online modeling of the signal requires time and computational resources and might not be feasible, especially on small platforms and in time-critical missions. In contrast, the second scenario supposes the sensors have an accurate signal model which can be exploited to localize the source, potentially faster and with better accuracy.

1.1 Related Work. Our model-free source-seeking approach consists in climbing the gradient of the signal field by using a stochastic approximation (SA) technique to deal with the underlying noise. Our strategy is robust to deformations in the geometry of the sensor network and can be applied to sensors with limited computational resources and no global localization capabilities. Recent model-free source-seeking work which uses a sensor formation to ascend the gradient of the signal field includes Refs. [3] and [7–9]. Ögren et al. [3] use artificial potentials to decouple the formation stabilization from the gradient ascent. Centralized least squares are used to estimate the signal gradient. A distributed approach for exploring a scalar field using a cooperative Kalman filter is presented in Ref. [10]. The authors design control laws to

achieve a formation, which minimizes the estimation error. Similarly, in Ref. [9], a circular formation is used to estimate the signal gradient in a distributed manner based on a Newton–Raphson consensus method. A drawback of these works is the assumption that the sensor formation is maintained perfectly throughout the execution of the algorithm which is hardly possible in a real environment. In this paper, imperfect formations are explicitly handled by recomputing the correct weights necessary to combine the sensor observations at every measurement location. Choi et al. [11,12] present a general distributed learning and control approach for sensor networks and apply it to source seeking. The sensed signal is modeled by a network of radial basis functions (RBFs) and recursive least squares are used to obtain the model parameters. Instead of a sensor network, a *single* vehicle may travel to several sensing locations in order to collect the same measurements [13–17]. While costly maneuvers are required to climb the gradient effectively, in our previous work [18], we discussed the parameter choices which enable good performance.

In the model-based scenario, we choose the next configuration for the sensing team by maximizing the MI between the source location estimate and the expected measurements. Even if all pose and measurement information are available at a central location, evaluating the MI utility function is computationally demanding. Charrow et al. [19] focus on approximating MI when the sensed signal is Gaussian and the sensors use a particle filter to estimate the source location. Hoffman and Tomlin [20] compute the expectation over the measurements only for pairs of sensors, thus decreasing the dimension of the required integration. Instead of MI, in this work, we approximate the MI gradient. A related work which uses the MI gradient is Ref. [21], in which the computational complexity is reduced by integrating over binary sensor measurements and only for sensors whose fields of view overlap. A fully distributed approach based on belief consensus is proposed in Ref. [22]. This paper is also related to consensus control, which seeks agreement in the states of multi-agent dynamical systems. Recent results [23,24] address switching topologies, nontrivial delays, and asynchronous estimation but with the main difference that the sensors agree on their own states, while in this work they need to agree on the exogenous state of the source.

1.2 Contributions. We develop a *distributed* approach for stochastic source seeking using a mobile sensor network, which *does not rely on a model of the signal field or global localization*.

¹Corresponding author.

Contributed by the Dynamic Systems Division of ASME for publication in the JOURNAL OF DYNAMIC SYSTEMS, MEASUREMENT, AND CONTROL. Manuscript received January 31, 2014; final manuscript received June 10, 2014; published online October 21, 2014. Assoc. Editor: Dejan Milutinovic.

Our method uses a finite-difference (FD) scheme to estimate the signal gradient correctly, *even when the sensor formation is not maintained well*. In the model-based case, we show that a SA to the MI gradient *using only a few predicted signal measurements* is enough to provide good control performance. This is in contrast with existing work which insists on improving the quality of the gradient estimate as much as possible.

1.3 Paper Organization. In Sec. 2, we describe the considered source-seeking scenarios precisely. Our model-free and model-based approaches are discussed in detail in Secs. 3 and 4, respectively, assuming all-to-all communication among the sensors. Distributed versions are presented and analyzed in Sec. 5. Finally, in Sec. 6, we present an application to wireless radio source localization and compare the performance of the two methods.

2 Problem Formulation

Consider a team of n sensing robots with states $\{x_{1,t}, \dots, x_{n,t}\} \subset \mathcal{X} \cong \mathbb{R}^{d_x}$ at time t . The states typically comprise pose and velocity information but might include other operational parameters too. At a high-level planning stage, we suppose that the vehicles have discrete single-integrator dynamics $x_{i,t+1} = x_{i,t} + u_{i,t}$, where $u_{i,t} \in \mathcal{U}$ is the control input to sensor i . The task is to localize a *static* signal source, whose unknown state is $y \in \mathcal{Y} \cong \mathbb{R}^{d_y}$. The state captures the source position and other observable properties of interest. At time t , each sensor i has access to a noisy measurement $z_{i,t} \in \mathcal{Z} \cong \mathbb{R}^{d_z}$ of the signal generated by y

$$z_{i,t} = h(x_{i,t}, y) + v_{i,t} \quad (1)$$

where $v_{i,t}$ is the measurement noise, whose values are independent at any pair of times and among sensors. The noise depends on the states of the sensor and the source, i.e., $v_{i,t}(x_{i,t}, y)$, but to simplify notation we do not make it explicit. We assume that the noise is zero mean and has a finite second moment, i.e., $\mathbb{E}v_{i,t} = 0, \forall i, t, x_{i,t}$ and $\text{tr}(\mathbb{E}[v_{i,t}v_{i,t}^T]) < \infty$. In the remainder, we use the notation $x_t := [x_{1,t}^T, \dots, x_{n,t}^T]^T, u_t := [u_{1,t}^T, \dots, u_{n,t}^T]^T, z_t := [z_{1,t}^T, \dots, z_{n,t}^T]^T$, and $v_t := [v_{1,t}^T, \dots, v_{n,t}^T]^T$.

In the model-free scenario, the sensors simply receive measurements without knowing the signal model $h(\cdot, \cdot)$. We suppose that the team adopts some arbitrary formation, with center of mass $m_t := \sum_{i=1}^n x_{i,t}/n$ at time t , which can be enforced using potential fields [3] or convex optimization [25]. The sensors use the centroid m_t as the estimate of the source state y at time t and try to lead it toward the true source location based on the received measurements. Let $f: \mathcal{X} \rightarrow \mathcal{Y}$ be a known transformation, which maps the team centroid to a source estimate. For example, if the robot state space captures both position and orientation, e.g., $\mathcal{X} = SE(2)$, but we are interested only in position estimates for the source, e.g., $\mathcal{Y} = \mathbb{R}^2$, then f will be the projection which extracts the position components from the centroid $m_t \in \mathcal{X}$. We consider the following problem.

Problem 2.1 (Model-Free Source Seeking). *Assume that the measurement signal in Eq. (1) is scalar² and its expectation is maximized at the true state y of the source*

$$f^{-1}(y) \in \arg \max_{x \in \mathcal{X}} h(x, y) \quad (2)$$

Generate a sequence of control inputs u_0, u_1, \dots for the team of sensors in order to drive its centroid m_t toward a maximum of the signal field $h(\cdot, y)$.

²The assumption is made only to simplify the presentation of the gradient ascent approach in the model-free case. The approach generalizes to signals of higher dimension.

In the model-based case, the sensors have accurate knowledge of $h(\cdot, \cdot)$ which can be exploited to maximize the information that future measurements provide about the signal source. We choose MI as a measure of informativeness and formulate the following problem.

Problem 2.2 (Model-Based Source Seeking). *Given the sensor poses $x_{t-1} \in \mathcal{X}^n$ and a prior distribution of the source state y at time $t-1$, choose the control input $u_t \in \mathcal{U}^n$, which optimizes the following:*

$$\begin{aligned} & \max_{u_{1,t}, \dots, u_{n,t}} I(y; z_t | x_t) \\ \text{s.t.} \quad & x_{i,t} = x_{i,t-1} + u_{i,t}, \quad i = 1, \dots, n \\ & z_{i,t} = h(x_{i,t}, y) + v_{i,t}, \quad i = 1, \dots, n \end{aligned} \quad (3)$$

We resort to SA methods in both scenarios and emphasize their usefulness in simplifying the algorithms, while providing theoretic guarantees about the performance.

3 Model-Free Source Seeking

3.1 Model-Free Algorithm. Our model-free approach is to design an iterative optimization scheme, which causes the centroid m_t of the robot formation to ascend the gradient $g(x, y) := \nabla_x h(x, y)$ of the measurement signal. The gradient ascent leads m_t to a (often local) maximum of the signal field, which is appropriate in view of Assumption (2). In detail, the desired dynamics for the centroid are

$$m_{t+1} = m_t + \gamma_t g(m_t, y) \quad (4)$$

A complication arises because the sensors do not have access to $g(\cdot, y)$ and can only measure a noisy version of $h(\cdot, y)$ at their current positions. Supposing noise-free measurements for now, the sensors can approximate the signal gradient at the formation centroid via a FD scheme

$$g(m_t, y) = \nabla_x h(m_t, y) = W(x_t) \begin{pmatrix} h(x_{1,t}, y) \\ \vdots \\ h(x_{n,t}, y) \end{pmatrix} - b_t \quad (5)$$

where $W(x_t) \in \mathbb{R}^{d_x \times n}$ is a matrix of FD weights, which depends on the sensor states x_t , and $b_t \in \mathbb{R}^{d_x}$ captures the error in the approximation. The most natural way to obtain the FD weights is to require that the approximation is exact for a set of test functions $\psi_i, i = 1, \dots, n$, commonly polynomials, which can represent the shape of $g(\cdot, y)$. In particular, the following relation needs to hold:

$$\begin{bmatrix} \psi_1(x_{1,t}) & \cdots & \psi_1(x_{n,t}) \\ \vdots & & \vdots \\ \psi_n(x_{1,t}) & \cdots & \psi_n(x_{n,t}) \end{bmatrix} W(x_t)^T = \begin{bmatrix} \frac{\partial}{\partial x} \psi_1(m_t) \\ \vdots \\ \frac{\partial}{\partial x} \psi_n(m_t) \end{bmatrix} \quad (6)$$

where $(\partial/\partial x)\psi_i(x)$ is a row vector of partial derivatives. When $x_{i,t} \in \mathbb{R}$, the most common set of test functions are the monomials $\psi_i(x) = x^{i-1}$, in which case (6) becomes a Vandermonde system. The standard (monomial) FD approach is problematic when the states $x_{i,t}$ are high-dimensional and not in a lattice configuration because the system in Eq. (6) becomes ill-conditioned. These difficulties are alleviated by using RBFs $\psi_i(x) := \phi(\|x - x_{i,t}\|)$ as test functions. In particular, using Gaussian RBFs, $\phi(d) := e^{-(\delta d)^2}$, with *shape parameter* $\delta > 0$, guarantees that Eq. (6) is nonsingular [26]. Then, the FD weights obtained from Eq. (6) as a function of x_t are

$$W(x_t) = R(x_t)^T \Phi(x_t)^{-T} \quad (7)$$

where, for $x \in \mathcal{X}^n$, we let $\Phi_{ij}(x) := e^{-\delta^2 \|x_j - x_i\|_2^2}$ and

$$R(x) := \begin{bmatrix} 2\delta^2 e^{-\delta^2 \left\| x_1 - \sum_{i=1}^n x_i/n \right\|_2^2} \left(x_1 - \sum_{i=1}^n x_i/n \right)^T \\ \vdots \\ 2\delta^2 e^{-\delta^2 \left\| x_n - \sum_{i=1}^n x_i/n \right\|_2^2} \left(x_n - \sum_{i=1}^n x_i/n \right)^T \end{bmatrix} \quad (8)$$

Since the measurements are noisy, sensor i can observe only $z_{i,t}$ rather than $h(x_{i,t}, y)$. As a result, the gradient ascent (4) can be implemented only approximately via $g(m_t, y) \approx W(x_t)z_t$ instead of Eq. (5) and with the additional complication that the measurement noise makes the iterates m_t random. Our stochastic model-free source-seeking algorithm is

$$m_{t+1} = m_t + \gamma_t W(x_t) z_t \quad (9)$$

The convergence of similar source-seeking schemes is often studied in a deterministic framework [3] by assuming that the noise can be neglected, which is difficult to justify. In Sec. 3.2, we show that the center of mass m_t , following the dynamics (9) with appropriately chosen step-sizes γ_t , converges to a neighborhood of a local maximum of $h(\cdot, y)$. Assuming all-to-all communication or a centralized location, which receives all state and measurement information from the sensors, the stochastic gradient ascent (9) can be implemented as is. It requires that the sensors are localized relative to one another, i.e., in the inertial frame of one sensor, but not globally, in the world frame. Notably, it is also not important to maintain a rigid sensor formation as the correct FD weights necessary to combine the observations are recomputed at every measurement location. Section 3.2 shows that the only requirement is that the sensor team is not contained in a subspace of \mathbb{R}^{d_x} when measuring (e.g., at least three noncollinear sensors are needed for $d_x = 2$).

3.2 Convergence Analysis. To carry out the convergence analysis of the stochastic gradient ascent in Eq. (9), we resort to the theory of SAs [27,28]. It is sufficient to consider the following SA algorithm:

$$m_{t+1} = m_t + \gamma_t (g(m_t) + b_t + D_t) \quad (10)$$

where b_t is a bias term, D_t is a random zero-mean perturbation, γ_t is a small step-size, and m_t is a random sequence whose asymptotic behavior is of interest. The main result is that the iterates m_t in Eq. (10) asymptotically follow the integral curves of the ordinary differential equation (ODE) $\dot{m} = g(m)$. Since in our case with a fixed source state y , $g(m) := \nabla_x h(m, y)$, the ODE method (Ref. [28] Chap. 2 and Ref. [29]) shows that the iterates $\{m_t\}$ almost surely (a.s.) converge to the set $\{x | \nabla_x h(x, y) = 0\}$ of critical points of $h(\cdot, y)$ under the following Assumptions:³

- (A1) The map g is Lipschitz continuous;⁴
- (A2) Step-sizes $\{\gamma_t\}$ are positive scalars satisfying

$$\sum_{t=0}^{\infty} \gamma_t = \infty \text{ and } \sum_{t=0}^{\infty} \gamma_t^2 < \infty;$$

(A3) $\{D_t\}$ is martingale difference sequence with respect to the family of σ -algebras $\mathcal{F}_t := \sigma(m_0, D_s, 0 \leq s \leq t)$, i.e., D_t is measurable with respect to \mathcal{F}_t , $\mathbb{E}[\|D_t\|] < \infty$, and $\mathbb{E}[D_{t+1} | \mathcal{F}_t] = 0$ a.s. for all $t \geq 0$. Also, D_t is square integrable with $\mathbb{E}[\|D_{t+1}\|^2 | \mathcal{F}_t] \leq K(1 + \|m_t\|^2)$ a.s. for $t \geq 0$ and some constant $K > 0$;

(A4) $\{m_t\}$ is bounded, i.e., $\sup_t \|m_t\| < \infty$ a.s.;

(A5) $\{b_t\}$ is bounded and $b_t \rightarrow 0$ a.s. as $t \rightarrow \infty$.

The proposed source-seeking algorithm (9) can be converted to the SA form (10) as follows:

$$\begin{aligned} m_{t+1} &= m_t + \gamma_t W(x_t) z_t = m_t + \gamma_t W(x_t) \begin{pmatrix} h(x_{1,t}, y) + v_{1,t} \\ \vdots \\ h(x_{n,t}, y) + v_{n,t} \end{pmatrix} \\ &= m_t + \gamma_t (g(m_t, y) + b_t + W(x_t)v_t) \end{aligned}$$

where the second equality follows from Eq. (5). Assumption (A1) ensures that $\dot{m} = g(m, y)$ has a unique solution for any initial condition and any fixed source state y . Assumption (A2) can be satisfied by an appropriate choice of the step-size, e.g., $\gamma_t = 1/(t+1)$. The selection of proper step-sizes is an important practical issue that is not emphasized in this paper but is discussed at length in Refs. [18], [27], and [30]. We can satisfy (A4) by requiring that the environment \mathcal{X} of the sensors is bounded and if necessary use a projected version of the gradient ascent [28, Chap. 5.4]. This also ensures that the FD weights are bounded and in turn (A3) is satisfied

$$\begin{aligned} (\mathbb{E}\|D_t\|_2)^2 &\leq \mathbb{E}\|D_t\|_2^2 = \mathbb{E}[\|D_t\|_2^2 | \mathcal{F}_{t-1}] = \mathbb{E}[\|W(x_t)v_t\|_2^2] \\ &\leq \|W(x_t)\|_2^2 \mathbb{E}\|v_t\|_2^2 = \|W(x_t)\|_2^2 \sum_{i=1}^n \text{tr}(\mathbb{E}[v_{i,t}v_{i,t}^T]) < \infty \\ \mathbb{E}[D_t | \mathcal{F}_{t-1}] &= \mathbb{E}[W(x_t)v_t] = W(x_t)\mathbb{E}v_t = 0 \end{aligned}$$

since the measurement noise in Eq. (1) is uncorrelated in time and has zero mean and a finite second moment. Note that the error term in Eq. (5) violates (A5) because it does not converge to 0. However, if we ensure that the sensor formation is not contained in a subspace of \mathbb{R}^{d_x} , then b_t remains bounded by some $\varepsilon_0 > 0$, i.e., $\sup_t \|b_t\| \leq \varepsilon_0$. Then, the argument in Ref. [28, Chap. 5, Theorem 6] shows that the iterates m_t converge a.s. to a small neighborhood of a local maximum, whose size depends on ε_0 . The result is summarized below:

THEOREM 1. Suppose that the gradient $g(x, y) = \nabla_x h(x, y)$ of the measurement signal is Lipschitz continuous³ in x , the step-sizes γ_t in Eq. (9) satisfy (A2), the sensor state space \mathcal{X} is bounded, and the sensor formation is not contained in a subspace of \mathbb{R}^{d_x} at the measurement locations. Then, algorithm (9) converges to a small neighborhood around a local maximum of the signal field $h(\cdot, y)$.

4 Model-Based Source Seeking

4.1 Model-Based Algorithm. In this section, we address Problem 2.2 assuming all-to-all communication. The sensors can follow the gradient of the cost function in Eq. (3) to reach a local maximum:

$$x_{t+1} = x_t + \gamma_t \nabla_x I(y; z_t | x) \Big|_{x=x_t} \quad (11)$$

where γ_t is the step-size at time t . Let $p(z|y, x)$ denote the probability density function (pdf) of the measurement signal in Eq. (1).

³While Assumptions (A1)–(A5) are sufficient to prove the convergence in our application, they are by no means the weakest possible. If necessary some can be relaxed using the results in stochastic approximation [27,28].

⁴Given two metric spaces (\mathcal{X}, d_x) and (\mathcal{G}, d_g) , a function $g: \mathcal{X} \rightarrow \mathcal{G}$ is Lipschitz continuous if there exists a real constant $0 \leq L < \infty$ such that: $d_g(g(x_1), g(x_2)) \leq L d_x(x_1, x_2), \forall x_1, x_2 \in \mathcal{X}$.

Let $p_t(y)$ be the pdf used by the sensors at time t to estimate the state of the source, which is assumed independent of x_t . The following theorem gives an expression for the MI gradient provided that $p(z|y, x)$ is differentiable with respect to the sensor configurations.

THEOREM 2 (Ref. [31]). *Let random vectors Y and Z be jointly distributed with pdf $p(y, z|x)$, which is differentiable with respect to the parameter $x \in \mathcal{X}$. Suppose that the support of $p(y, z|x)$ does not depend on x . Then, the gradient with respect to x of the MI between Y and Z is*

$$\nabla_x I(Y; Z|x) = \iint (\nabla_x p(y, z|x)) \log \frac{p(z|y, x)}{p(z|x)} dy dz$$

where $p(z|y, x)$ and $p(z|x)$ are the conditional and the marginal pdfs of Z , respectively.

Obtaining the MI gradient is computationally very demanding for two reasons. First, an approximate representation is needed for the continuous pdfs in the integral. Second, at time t , the integration is over the collection of all sensor measurements $z_t = [z_{1,t}^T, \dots, z_{n,t}^T]^T$, which can have a very high dimension in practice. As mentioned in Sec. 1, most existing work has focused on accurate approximations. However, Theorem 2 allows us to make a key observation

$$\nabla_x I(y; z_t|x_t) = \mathbb{E}[\pi_t(z_t, x_t) |x_t] = \int \pi_t(z_t, x_t) p_t(z_t|x_t) dz_t$$

where

$$\begin{aligned} \pi_t(z, x) &:= \int \frac{\nabla_x p(z|y, x)}{p_t(z|x)} p_t(y) \log \frac{p(z|y, x)}{p_t(z|x)} dy \\ p_t(z|x) &:= \int p(z|y, x) p_t(y) dy \end{aligned} \quad (12)$$

where the independence between y and x_t is used for the decomposition: $p_t(y, z|x) = p(z|y, x) p_t(y)$. Relying on the signal model, the sensors can simulate realizations of the random variable z_t , iid with pdf $p_t(z_t|x_t)$. Instead of computing the integral in Eq. (12) needed for the gradient ascent (11), we propose the following stochastic algorithm:

$$x_{t+1} = x_t + \gamma_t \pi_t(z_t, x_t) \quad (13)$$

It can be written in the SA form (10) as follows:

$$\begin{aligned} x_{t+1} &= x_t + \gamma_t \mathbb{E}_{z_t}[\pi_t(z_t, x_t) |x_t] + \gamma_t D_t \\ &= x_t + \gamma_t (\nabla_x I(y; z_t|x_t) + D_t) \end{aligned}$$

where $D_t := \pi_t(z_t, x_t) - \mathbb{E}[\pi_t(z_t, x_t) |x_t]$. To evaluate the convergence, we consider Assumptions (A1)–(A5) again. As before, satisfaction of (A2) is achieved by a proper step-size choice, while (A4) holds due to the bounded workspace \mathcal{X} . Assumption (A5) is satisfied because in this case the bias term is zero. To verify (A3), note that D_t is measurable with respect to $\mathcal{F}_t = \sigma(x_0, D_s, 0 \leq s \leq t)$ and for $t \geq 1$

$$\begin{aligned} \mathbb{E}[D_t | \mathcal{F}_{t-1}] &= \mathbb{E}[\pi_t(z_t, x_t) - \mathbb{E}[\pi_t(z_t, x_t) |x_t] | \mathcal{F}_{t-1}] \\ &= \mathbb{E}[\pi_t(z_t, x_t) | \mathcal{F}_{t-1}] - \mathbb{E}[\pi_t(z_t, x_t) |x_t] = 0 \end{aligned}$$

Finally, if $p(z|y, x)$ and its gradient $\nabla_x p(z|y, x)$ are sufficiently regular (e.g., the former is bounded away from zero and the latter is Lipschitz continuous and bounded), the square integrability condition on D_t is satisfied.

This analysis demonstrates that even if a single z_t sample is used to approximate the MI gradient (instead of the integration in Eq. (12)), the stochastic gradient ascent (13) will converge to a local maximum of the MI between the source state and the future sensor measurements.

4.2 Implementation Details. To implement the stochastic gradient ascent in Eq. (13), the sensors need to propagate $p_t(y)$ over time and sample from $p_t(z_t|x_t)$. We achieve the first requirement by a particle filter [32, Chap. 4], which approximates $p_t(y)$ by a set of weighted samples $\{w_t^m, y_t^m\}_{m=1}^{N_p}$ as follows: $p_t(y) \approx \sum_{m=1}^{N_p} w_t^m \delta(y - y_t^m)$, where $\delta(\cdot)$ is a Dirac delta function. Using the particle set, we can write π and the measurement pdf as follows:

$$\begin{aligned} \pi_t(z, x) &\approx \sum_{m=1}^{N_p} w_t^m \frac{\nabla_x p(z|y_t^m, x)}{p(z|x)} \log \frac{p(z|y_t^m, x)}{p(z|x)} \\ p_t(z|x) &\approx \sum_{m=1}^{N_p} w_t^m p(z|y_t^m, x) \end{aligned}$$

where $p(z|y, x)$ and its gradient can be decomposed further

$$\begin{aligned} p(z|y, x) &= \prod_{j=1}^n p(z_j|y, x_j) \\ \frac{\partial p(z|y, x)}{\partial x_k} &= \frac{\partial p(z_k|y, x_k)}{\partial x_k} \prod_{j \neq k} p(z_j|y, x_j) \end{aligned}$$

due to the independence of the observations in Eq. (1). In practice, there is a trade-off between moving the sensors and spending time approximating the gradient of the MI (12). The SA in Eq. (13) uses a single sample from $p_t(\cdot|x_t)$ but if sampling is fast compared to the time needed to relocate the sensors, more samples can be used to get a better estimate of the gradient. We use Monte Carlo integration, which proceeds as follows:

- (1) Sample $\bar{m}(l)$ from the discrete distribution $w_t^1, \dots, w_t^{N_p}$.
- (2) Sample $\bar{z}(l)$ from the pdf $p(\cdot|y_{\bar{m}(l)}^m, x_t)$.
- (3) Repeat steps (1) and (2) to obtain N_z samples $\{\bar{z}(l)\}_{l=1}^{N_z}$.
- (4) Approximate: $\nabla_x I(y; z_t|x_t) \approx \frac{1}{N_z} \sum_{l=1}^{N_z} \pi_t(\bar{z}(l), x_t)$.

Note that the advantage of improving the gradient estimate is not clear and should not necessarily be prioritized over the sensor motion. The SA techniques show that even an approximation with a single sample is sufficient to make progress. In contrast, the related approaches mentioned in the Introduction insist on improving the quality of the gradient estimate as much as possible. Depending on the application, this can slow down the robot motion and possibly make the algorithms impractical. Our more flexible approach adds an extra degree of freedom by allowing a trade-off between the gradient estimation quality and the motion speed of the sensors.

5 Distributed Algorithms

In many scenarios, all-to-all communication is either infeasible or prone to failures. In this section, we present distributed versions of the model-free and the model-based algorithms. Since the model-free algorithm should be applicable to light-weight platforms with no global localization capabilities, the sensors use noisy relative measurements of their neighbors' locations to estimate the collective formation state. In the model-based case, the sensors may spread around the environment and we are forced to assume that each agent is capable of estimating its own state $x_{i,t}$. We begin with preliminaries on distributed estimation.

5.1 Preliminaries on Distributed Estimation. Let the communication network of the n sensors be represented by an undirected graph $G = (\{1, \dots, n\}, E)$. Suppose that the sensors need to estimate an unknown static parameter $\theta^* \in \Theta$, where $\Theta \subseteq \mathbb{R}^{d_\theta}$ is a convex space. At discrete times $k \in \mathbb{N}$, each sensor i receives a random sequence of iid signals $s_i(k) \in \mathbb{R}^{d_i}$, which are drawn from

a distribution with conditional pdf $l_i(\cdot|\theta)$ and are independent of those received by the other sensors. These individual signals, although potentially informative, do not reveal the parameter completely, i.e., each sensor faces a local identification problem. We assume, however, that θ^* is identifiable if one has access to the signals received by all sensors. Thus, individual sensors need to supplement their local observations with information communicated with their neighbors. In particular, each sensor i propagates a pdf $p_{i,k} : \Theta \rightarrow \mathbb{R}_{\geq 0}$ over the parameter space

$$p_{i,k+1}(\theta) = \eta_{i,k} l_i(s_i(k+1)|\theta) \prod_{j \in N_i \cup \{i\}} (p_{j,k}(\theta))^{a_{ij}}$$

$$\hat{\theta}_i(k) \in \arg \max_{\theta \in \Theta} p_{i,k}(\theta) \quad (14)$$

where $\eta_{i,k}$ is a normalization constant ensuring that $p_{i,k+1}$ is a proper pdf, N_i is the set of nodes (neighbors) connected to sensor i , and a_{ij} are weights such that $\sum_{j \in N_i \cup \{i\}} a_{ij} = 1$. The update is the same as the standard Bayes rule with the exception that sensor i does not just use its own prior but a *geometric average* of its neighbors' priors. Given that G is connected, Rad and Tahbaz-Salehi [33] show that the distributed estimator (14) is weakly consistent⁵ under broad assumptions on the signal models $l_i(\cdot|\theta)$. The results in Refs. [34] and [35] suggest that this algorithm is even applicable to a time-varying graph topology with asynchronous communication.

5.1.1 Specialization to Gaussian Distributions. We now specialize the general scheme of Ref. [33] to Gaussian distributions. To our knowledge, this specialization is new and the theorem obtained below (Theorem 3) shows that the resulting distributed linear Gaussian estimator is mean square consistent,⁶ which is stronger than the weak consistency⁴ shown in Ref. [33, Theorem 1]. Suppose that the agents' measurement signals are *linear* in the parameter θ^* and perturbed by Gaussian noise

$$s_i(k) = H_i \theta^* + \varepsilon_i(k), \quad \varepsilon_i(k) \sim \mathcal{N}(0, E_i), \quad \forall i \quad (15)$$

Let $\mathcal{G}(\omega, \Omega)$ denote a Gaussian distribution (in information space) with mean $\Omega^{-1}\omega$ and covariance matrix Ω^{-1} . Since the private observations (15) are linear Gaussian, without the loss of generality, the pdf $p_{i,k}$ of agent i is the pdf of a Gaussian $\mathcal{G}(\omega_{i,k}, \Omega_{i,k})$. Exploiting that the parameter θ^* is static, the update equation of the distributed filter in Eq. (14), specialized to Gaussian distributions, is

$$\omega_{i,k+1} = \sum_{j \in N_i \cup \{i\}} a_{ij} \omega_{j,k} + H_i^T E_i^{-1} s_i(k)$$

$$\Omega_{i,k+1} = \sum_{j \in N_i \cup \{i\}} a_{ij} \Omega_{j,k} + H_i^T E_i^{-1} H_i \quad (16)$$

$$\hat{\theta}_i(k) := \Omega_{i,k}^{-1} \omega_{i,k}$$

In this linear Gaussian case, we prove (in Ref. [36]) a strong result about the quality of the estimates in Eq. (16).

THEOREM 3. *Suppose that the communication graph G is connected and the matrix $[H_1^T \dots H_n^T]^T$ has rank d_θ . Then, the estimates (16) of all agents converge in mean square to θ^* , i.e.,*

$$\lim_{k \rightarrow \infty} \mathbb{E} \left[\left\| \hat{\theta}_i(k) - \theta^* \right\|^2 \right] = 0, \quad \forall i$$

5.1.2 Specialization to Particle Distributions. Suppose that the pdf $p_{i,k}$ is represented by a set of particles $\{w_{i,k}^m, \theta_{i,k}^m\}_{m=1}^{N_p}$, which are *identical* for all sensors initially (at $k=0$). Since the

⁵Weak consistency means that the estimates $\hat{\theta}_i(k)$ converge in probability to θ^* , i.e. $\lim_{k \rightarrow \infty} \mathbb{P}(\|\hat{\theta}_i(k) - \theta^*\| \geq \varepsilon) = 0$ for any $\varepsilon > 0$ and all i .

⁶Mean-square consistency means that the estimates $\hat{\theta}_i(k)$ converge in mean-square to θ^* , i.e. $\lim_{k \rightarrow \infty} \mathbb{E}[\|\hat{\theta}_i(k) - \theta^*\|^2] = 0$ for all i .

parameter θ^* is stationary, the particle positions $\theta_{i,k}^m$ will remain the same across the sensors for all time. The update equation of the distributed filter in Eq. (14), specialized to particle distributions, only needs to propagate the particle importance weights $w_{i,k}^m$ and is summarized in Algorithm 1.

Algorithm 1: Distributed particle filter at sensor i

- 1: **Input:** Particle sets $\{w_{j,k}^m, \theta_{j,k}^m\}$ for $m=1, \dots, N_p$ and $j \in N_i \cup \{i\}$, private signal $s_i(k+1)$, and pdf $l_i(\cdot|\cdot)$
- 2: **Output:** Particle weights $\{w_{i,k+1}^m\}$ for $m=1, \dots, N_p$
- 3: Average priors: $\bar{w}_{i,k}^m \leftarrow \exp\left(\sum_{j \in N_i \cup \{i\}} a_{ij} \log(w_{j,k}^m)\right)$
- 4: Update: $w_{i,k+1}^m \leftarrow \bar{w}_{i,k}^m l_i(s_i(k+1)|\theta_{i,k}^m)$ for $m=1, \dots, N_p$
- 5: Normalize the weights
- 6: **Return** $\{w_{i,k+1}^m\}$ for $m=1, \dots, N_p$

5.2 Distributed Model-Free Algorithm. To distribute the model-free algorithm (9), the sensor formation needs to estimate its configuration x_t , the centroid m_t , and the SA to the signal gradient $W(x_t)z_t$ at each measurement location (i.e., at each time t) using only local information. We introduce a fast time-scale $k=0, 1, \dots$, which will be used for the estimation procedure at each time t . During this, the sensors remain stationary and we drop the index t to simplify the notation. As mentioned earlier, we suppose that each sensor i receives a relative measurement of the state of each of its neighbors $j \in N_i$

$$s_{ij}(k) = x_j - x_i + \varepsilon_{ij}(k), \quad \varepsilon_{ij}(k) \sim \mathcal{N}(0, E_{ij}) \quad (17)$$

where $\varepsilon_{ij}(k)$ is the measurement noise, which is independent at any pair of times on the fast time-scale and across sensor pairs. If each sensor manages to estimate the states of the whole sensor formation using the measurements $\{s_{ij}(k)\}$, then each can compute the FD weights in Eq. (7) on its own.

The distributed linear Gaussian estimator (16) can be employed to estimate the sensor states x . Notice that it is sufficient to estimate x in a local frame because neither the FD computation (7) nor the gradient ascent (9) requires global state information. Assume that all sensors know that sensor 1 is the origin at every measurement location. Let $x^* := [0^T \ (x_2 - x_1)^T \ \dots \ (x_n - x_1)^T]^T$ denote the true sensor states in the frame of sensor 1. Let $\hat{x}^i(k)$ denote the estimate that sensor i has of x^* at time k on the fast time-scale. The vector form of the measurement equations (17) is

$$s(k) = (B \otimes I_{d_x})^T x^* + \varepsilon(k) \quad (18)$$

where B is the incidence matrix of the communication graph G . The measurements (18) fit the linear Gaussian model in Eq. (15). Since the first element of x^* is always 0, only $(n-1)d_x$ components need to be estimated. As the rank of $B \otimes I_{d_x}$ is also $(n-1)d_x$, Theorem 3 allows us to use the distributed estimator (16) to update $\hat{x}^i(k)$.

Concurrently with the state estimation, sensor i would be obtaining observations $z_{i,t}(k)$ of the signal field for $k=0, 1, \dots$ ⁷ In the centralized case (Sec. 3), each sensor uses the following gradient approximation:

$$g(m_t, y) \approx W(x_t)z_t = \sum_{i=1}^n \mathbf{col}_i(W(x_t))z_{i,t} \quad (19)$$

where $\mathbf{col}_i(W(x_t))$ denotes the i th column of the FD-weight matrix. Since x_t and z_t are not available in the distributed setting, each sensor can use its local measurements $z_{i,t}(k)$ and its estimate

⁷The time-scales of the relative state measurements and the signal measurements might be different but for simplicity we keep them the same.

$\hat{x}_i^j(k)$ of the sensor states to form its own local estimate of the signal gradient

$$\hat{g}_{i,t}(k) := \text{col}_i(W(\hat{x}_i^j(k))) \frac{1}{k+1} \sum_{\tau=0}^k z_{i,t}(\tau) \quad (20)$$

In order to obtain an approximation to $g(m_t, y)$ as in Eq. (19) in a distributed manner, we use a high-pass dynamic consensus filter [37] to have the sensors agree on the value of the sum

$$\hat{g}_i(k) := n \left(\frac{1}{n} \sum_{i=1}^n \hat{g}_{i,t}(k) \right)$$

Each node maintains a state $q_{i,k}$, receives an input $\mu_{i,k}$, and provides an output $r_{i,k}$ with the following dynamics:

$$q_{i,k+1} = q_{i,k} + \beta \sum_{j \in N_i} (q_{j,k} - q_{i,k}) + \beta \sum_{j \in N_i} (\mu_{j,k} - \mu_{i,k}) \quad (21)$$

$$r_{i,k} = q_{i,k} + \mu_{i,k}$$

where $\beta > 0$ is a step-size. For a connected network, Ref. [37, Theorem 1] guarantees that $r_{i,k}$ converges to $1/n \sum_i \mu_{i,k}$ as $k \rightarrow \infty$. The following result can be shown by letting $\mu_{i,k} := \hat{g}_{i,t}(k)$ and is proved in the Appendix.

THEOREM 4. Suppose that the communication graph G is strongly connected. If the sensor nodes estimate their states x^* from the relative measurements (18) using algorithm (16), compute the FD weights (7) using the state estimates, and run the dynamic consensus filter (21) with input $\mu_{i,k} := \hat{g}_{i,t}(k)$, which was defined in Eq. (20), then the output $r_{i,k}$ of the consensus filter satisfies

$$n \left(\lim_{k \rightarrow \infty} \mathbb{E}[r_{i,k}] \right) = g(m^*, y) + b, \quad \forall i \in \{1, \dots, n\}$$

where $g(m^*, y)$ is the true signal gradient at $m^* := \sum_{i=1}^n x_i^*/n$ and b is the error in the FD approximation (5).

After this procedure, the agents agree on a centroid for the formation and a gradient estimate, which can be used to compute the next formation centroid according to Eq. (9). Since the FD weights are recomputed at every t , the formation need not be maintained accurately. This allows the sensors to avoid obstacles and takes care of the motion uncertainty.

5.3 Distributed Model-Based Algorithm. In this section, we aim to distribute the model-based source-seeking algorithm (13). We assume that sensors which are sufficiently far from each other receive independent information. This is justified because when the sensing footprints of two sensors do not overlap, their sensed signals (if any) will not be coming from the same source. As a result, computing the MI gradient in Eq. (12) with respect to x_i is decoupled from the states of the distant sensors.

THEOREM 5. Let \mathcal{V}_i denote the set of sensors (excluding i) whose sensing footprints overlap with that of sensor i . Let $\bar{\mathcal{V}}_i$ denote the rest of the sensors. Suppose that sensor i 's measurements, z_i , are independent (not conditionally on y , as before) of the measurements, $z_{\bar{\mathcal{V}}_i}$, obtained by the sensors $\bar{\mathcal{V}}_i$, i.e., $p_t(z_i, z_{\bar{\mathcal{V}}_i} | x_i, x_{\bar{\mathcal{V}}_i}) = p_t(z_i | x_i) p_t(z_{\bar{\mathcal{V}}_i} | x_{\bar{\mathcal{V}}_i})$. Then,

$$\frac{\partial}{\partial x_i} I(y; z_i, z_{\mathcal{V}_i}, z_{\bar{\mathcal{V}}_i} | x_i, x_{\mathcal{V}_i}, x_{\bar{\mathcal{V}}_i}) = \frac{\partial}{\partial x_i} I(y; z_i, z_{\mathcal{V}_i} | x_i, x_{\mathcal{V}_i})$$

Proof. By the chain rule of MI and then the independence of z_i and $z_{\bar{\mathcal{V}}_i}$,

$$\begin{aligned} I(y; z_i, z_{\mathcal{V}_i}, z_{\bar{\mathcal{V}}_i} | x_i, x_{\mathcal{V}_i}, x_{\bar{\mathcal{V}}_i}) \\ &= I(y; z_i, z_{\mathcal{V}_i} | x_i, x_{\mathcal{V}_i}) + I(y; z_{\bar{\mathcal{V}}_i} | z_i, z_{\mathcal{V}_i}, x_i, x_{\mathcal{V}_i}, x_{\bar{\mathcal{V}}_i}) \\ &= I(y; z_i, z_{\mathcal{V}_i} | x_i, x_{\mathcal{V}_i}) + I(y; z_{\bar{\mathcal{V}}_i} | z_{\mathcal{V}_i}, x_{\mathcal{V}_i}, x_{\bar{\mathcal{V}}_i}) \end{aligned}$$

The second term above is constant with respect to x_i ■

As a result of Theorem 5 and the SA algorithm in Eq. (13), sensor i updates its pose as follows:

$$x_{i,t+1} = x_{i,t} + \gamma_t \pi_t(z_{\{\mathcal{V}_i\} \cup \mathcal{V}_i, t}, x_{\{\mathcal{V}_i\} \cup \mathcal{V}_i, t}) \quad (22)$$

This update is still not completely distributed as it requires knowledge of $x_{\mathcal{V}_i, t}$ and the pdf p_t .⁸ We propose to distribute the computation of p_t via the distributed particle filter (Algorithm 1). Then, each sensor maintains its own estimate of the source pdf, $p_{i,t}$, represented by a particle set $\{w_{i,t}^m, y_{i,t}^m\}$. Given a new measurement, $z_{i,t+1}$, sensor i averages its prior, $p_{i,t}$, with the priors of its neighbors and updates it using Bayes rule. Finally, to obtain $x_{\mathcal{V}_i, t}$ we use a flooding algorithm (Algorithm 2). The convergence analysis of the gradient ascent scheme in the distributed case (22) remains the same as in Sec. 4 because each sensor i computes the complete MI gradient. This is possible because due to Theorem 5 the states and measurements of distant sensors are not needed, while Algorithm 2 provides the information from the nearby sensors.

Algorithm 2: State exchange algorithm at sensor i

- 1: **Input:** Communication radius r_c , sensing radius r_s , state x_i
- 2: **Output:** Array a_i with $a_i[j] = x_j$ if $j \in \mathcal{V}_i \cup \{i\}$ and $a_i[j] = \text{empty}$ else
- 3: $a_i[i] \leftarrow x_i$, $a_i[j] \leftarrow \text{empty}$, $j \neq i$ ▷ Holds the required sensor states
- 4: $b \leftarrow \min\{\text{ceil}(2r_s/r_c), n\}$ ▷ Number of rounds needed
- 5: **for** $k = 1 \dots b$ **do**
- 6: Send a_i to neighbors N_i , receive $\{a_j\}$ from $j \in N_i$
- 7: **for** $j \in N_i$ **do**
- 8: **for** $l = 1 \dots n$ **do**
- 9: **if** $(a_i[l] = \text{empty}) \&\& (a_j[l] \neq \text{empty})$ **then** $a_i[l] \leftarrow a_j[l]$

6 Applications

The performance of the source-seeking algorithms is demonstrated in simulation using a team of ten sensors to localize the source of a wireless radio signal. A radio signal is suitable for comparing the two algorithms because it is very noisy and difficult to model and yet most approaches for wireless source seeking are model-based. We begin by modeling the received signal strength (RSS), which is needed for the model-based algorithm.

6.1 RSS Model. Let the positions of a wireless source and receiver in 2D be y and x , respectively. The RSS (dBm) at x is modeled as

$$P_{rx}(x, y) = P_{tx} + G_{tx} - L_{tx} + G_{rx} - L_{rx} - L_{fs}(x, y) - L_m(x, y) - R(x, y)$$

where P_{tx} is the transmitter output power (18 dBm in our experiments), G_{tx} is the transmitter antenna gain (1.5 dBi), L_{tx} is the transmitter loss (0 dB), G_{rx} is the receiver antenna gain (1.5 dBi), L_{rx} is the receiver loss (0 dB), L_{fs} is the free space loss (dB), L_m is the multipath loss (dB), and R is the noise. The free space loss is modeled as

$$L_{fs}(x, y) = -27.55 + 20 \log_{10}(\nu) + 20 \log_{10}(\|x - y\|_2)$$

where ν is the frequency (2400 MHz). The model from Ref. [38] is used for the multipath loss

$$L_m(x, y) = \begin{cases} \alpha + \beta \lambda(x, y), & \text{if } \lambda(x, y) > 0 \\ 0, & \text{else} \end{cases}$$

⁸Since all sensors have the same observation model $h(\cdot, \cdot)$, each sensor can simulate measurements $z_{\mathcal{V}_i, t}$ as long as it knows the configurations $z_{\mathcal{V}_i, t}$.

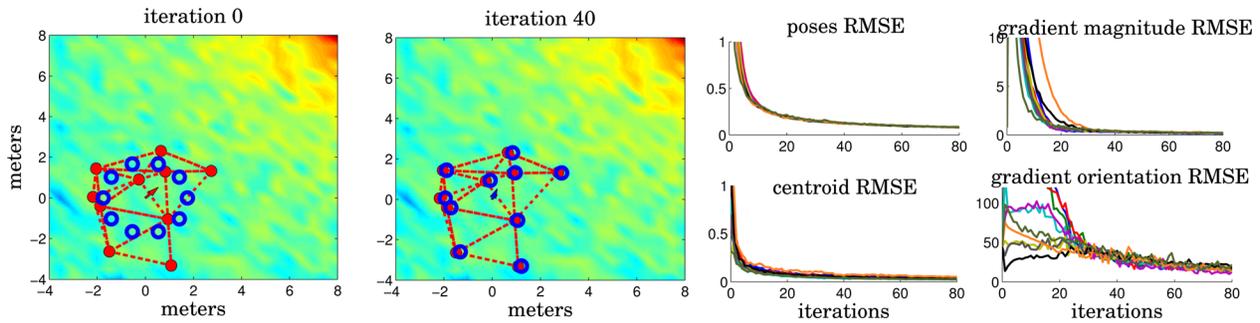


Fig. 1 Joint position and gradient estimation at a single measurement location (on the fast time-scale). The first plot shows the true sensor positions (red circles), initial position estimates (blue circles), and the true gradient of the signal field (red arrow). The second plot shows the position estimates after 40 iterations (blue circles) and the gradient estimate of sensor 1 (blue arrow). The third column shows the root mean squared error (RMSE) of the position (top) and centroid (bottom) estimates of all sensors averaged over 50 independent repetitions. The fourth column shows the RMSE of the gradient magnitude and orientation estimates.

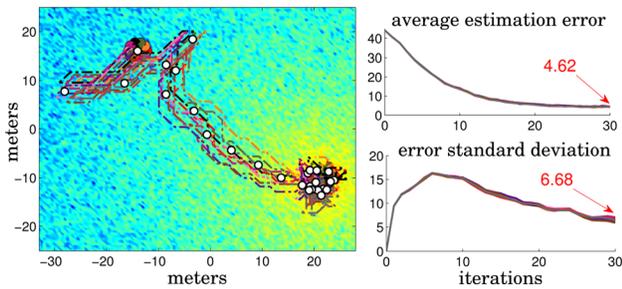


Fig. 2 The paths followed by the sensors after 30 iterations of the model-free source-seeking algorithm in an obstacle-free environment. The white circles indicate sensor 1's estimates of the source position over time. The plots on the right show the average error of the source position estimates and its standard deviation averaged over 50 independent repetitions.

where α is a multiwall constant (30 dB), β is a wall attenuation factor (15 dB/m), and $\lambda(x, y)$ denotes the distance traveled by the ray from y to x through occupied cells in the environment (represented as an occupancy grid). Finally, if the measurement is line-of-sight, i.e., $\lambda(x, y) = 0$, the fading $R(x, y)$ is Rician (μ, σ); otherwise it is Rayleigh (σ). We used $\mu = 4$ dB and $\sigma = 20$ dB in the simulations.

6.2 Simulation Results. The first experiment aims at verifying the conclusions of Theorem 4 when the sensor formation is not maintained well, namely that the distributed relative pose estimation and the consensus on the local FD gradient estimates converge asymptotically to an unbiased (up to the error in the FD approximation) gradient estimate. Ten sensors were arranged in a distorted “circular” formation (see Fig. 1) and were held stationary during the estimation procedure (on the fast time-scale). Initially, the sensors assumed that they were in a perfect circular formation of radius of 1.75 m. Relative measurements (17) with noise covariance $E_{ij} = 0.4I_2$ were exchanged to estimate the sensor states. At each time k , sensor i used its estimate $\hat{x}^i(k)$ to compute the FD weights via Eq. (7). Wireless signal measurements obtained according to the RSS model were combined with the FD weights to form the local gradient estimates (20), which were used to update the state of the consensus filter according to Eq. (21). Figure 1 shows that the errors in the pose and the gradient estimates tend to zero after 80 iterations on the fast time-scale.

Next, we demonstrate the ability of our algorithms to localize the source of a wireless signal obtained according to the RSS model of Sec. 6.1. The performance of the model-free algorithm is illustrated in Fig. 2. A circular formation with radius of 1.75 m consisting of ten sensors was maintained. The communication radius was 6 m, while the sensing radius was infinite. The sensors did not coordinate to maintain the formation. They were kept together by the agreement on the centroid and the signal gradient,

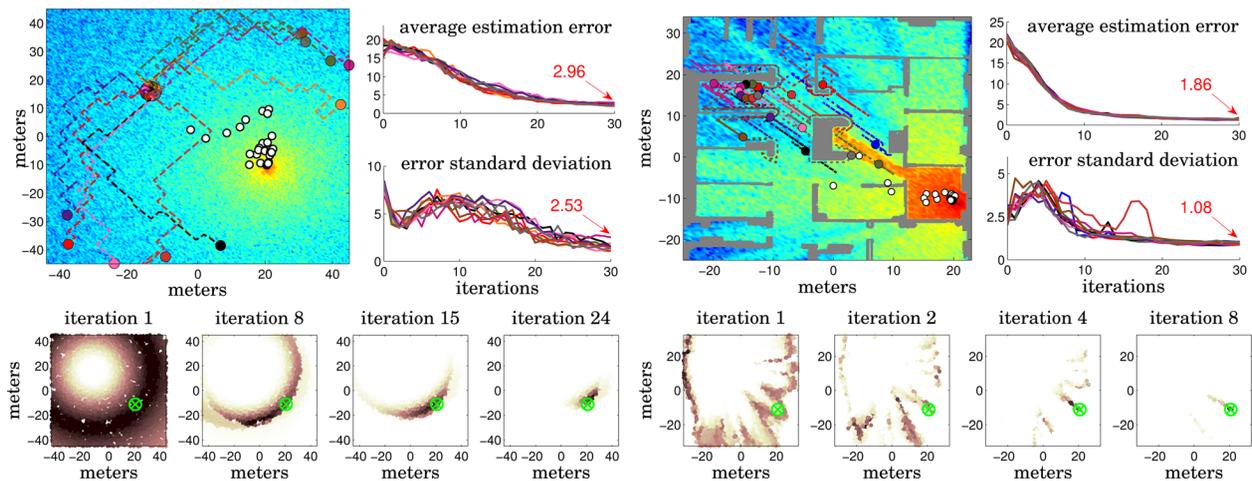


Fig. 3 The paths followed by the sensors after 30 iterations of the model-based source-seeking algorithm in an environment without obstacles (left) and with obstacles (right). The white circles indicate sensor 1's estimates of the source position over time. The plots show the average error of the source position estimates and its standard deviation averaged over 50 independent repetitions. The evolution of sensor 1's distributed particle filter is shown in each scenario (bottom row).

achieved via the distributed state estimation and the consensus filter. At time t , each sensor i applied the control $u_{i,t} = \gamma_i \hat{g}_{i,t}(K_{\max})$, where $\hat{g}_{i,t}(K_{\max})$ is the gradient estimate after $K_{\max} = 50$ iterations on the fast time-scale and γ_i is the step-size. Unlike the persistent measurements in Fig. 1, the sensors measured their relative states and the wireless signal only 10 times and stopped updating their local gradient estimates to enable faster convergence of the consensus filter. The initial distance between the signal source and the centroid of the sensor formation was 44.2 m. Averaged over 50 independent repetitions, the sensors managed to estimate the source location within 4.62 m in 30 iterations.

The same initial source and sensor positions were used to set up the model-based experiments. Figure 3 shows the performance in environments with and without obstacles. The communication radius was 10 m, while the sensing radius was infinite. The sensors maintained distributed particle filters with 4000 particles and used five signal measurements to update the filters before moving (unlike the ten used in the model-free case). A stochastic MI gradient was obtained via ten simulated signal measurements only. Averaged over 50 independent repetitions, the sensors managed to estimate the source location within 2.96 m in the obstacle-free case and 1.86 m in the obstacle case after 30 iterations of the algorithm. It is interesting to note that the performance of the model-based algorithm is better when obstacles are present in the environment. When the model is good and the environment is known, the wall attenuation of the signal helps discount many hypothetical source locations, which would not be possible in the obstacle-free case (see the filter evolution in Fig. 3). As a result, the model-based algorithm outperforms the model-free one. However, we expect that as the quality of the model degrades so would the performance of the model-based approach and the model-free algorithm would become more attractive.

7 Conclusion

Distributed model-free and model-based approaches for source seeking with a mobile sensor network were developed. The stochastic gradient ascent approach in the model-free case does not need global localization and is robust to deformations in the geometry of the sensor team. The SA simplifies the algorithm and provides convergence guarantees. The model-based method has the sensors follow a stochastic gradient of the MI between their expected measurements and source estimates. In this case, the SA enables a key trade-off between time spent moving the sensors and time spent planning the most informative move. The experiments show that the model-based algorithm outperforms the model-free one when an accurate model of the signal is available. Its drawbacks are that it relies on knowledge of the environment, global localization, and a flooding algorithm to exchange the sensor states, which can be demanding for the network. If computation is limited, the environment is unknown, the signal is difficult to model, or global localization is not available, the model-free algorithm would be the natural choice. In the future work, we plan to apply the algorithms to other signals, compare their performance to other approaches in the literature, and carry out real-world experiments.

Acknowledgment

This work was supported by ONR-HUNT Grant No. N00014-08-1-0696 and by TerraSwarm, one of six centers of STARnet, a Semiconductor Research Corporation program sponsored by MARCO and DARPA.

Appendix: Proof of Theorem 4

From Theorem 3, $\hat{x}^i(k) \xrightarrow{L^2} x^*, \forall i$, which implies convergence in L^1 and in probability. Convergence in L^1 implies that the sequence $\{\hat{x}^i(k)\}$ is uniformly integrable (UI) for all i [39, Theorem 5.5.2].

We claim that this implies that the sequence of FD weights $W(\hat{x}^i(k))$ computed in Eq. (7) is UI for each i . The matrix Φ in Eq. (7) is a bounded continuous function of $\hat{x}^i(k)$, which means that there exists a constant $K_i^\Phi \leq \infty$ for each i such that $\|\Phi(\hat{x}^i(k))^{-T}\|_1 \leq K_i^\Phi$. Define $\alpha_i(k) := \hat{x}^i(k) - \sum_{j=1}^n \hat{x}_j^i(k)/n$. From Eq. (8),

$$\begin{aligned} \|W(\hat{x}^i(k))\|_1 &\leq \left\| \begin{bmatrix} 2\delta^2 e^{-\delta^2 \|\alpha_1(k)\|_2^2} \alpha_1^T(k) \\ \vdots \\ 2\delta^2 e^{-\delta^2 \|\alpha_n(k)\|_2^2} \alpha_n^T(k) \end{bmatrix}^T \right\|_1 \|\Phi(\hat{x}^i(k))^{-T}\|_1 \\ &\leq 2\delta^2 K_i^\Phi \sum_{j=1}^n e^{-\delta^2 \|\alpha_j(k)\|_2^2} \|\alpha_j(k)\|_1 \\ &\leq 2\delta^2 K_i^\Phi \sum_{j=1}^n \left\| \hat{x}_j^i(k) - \frac{1}{n} \sum_{l=1}^n \hat{x}_l^i(k) \right\|_1 \\ &\leq 4\delta^2 K_i^\Phi \sum_{j=1}^n \|\hat{x}_j^i(k)\|_1 = 4\delta^2 K_i^\Phi \|\hat{x}^i(k)\|_1 \end{aligned}$$

By UI of $\{\hat{x}^i(k)\}$, for any $\varepsilon > 0$, there exist $K_i \in [0, \infty)$ such that $\mathbb{E}[\|\hat{x}^i(k)\|_1 \mathbf{1}_{\{\|\hat{x}^i(k)\|_1 \geq K_i\}}] \leq \varepsilon$ for all k . Then for all i, k

$$\begin{aligned} &\mathbb{E} \left[\|W(\hat{x}^i(k))\|_1 \mathbf{1}_{\{\|W(\hat{x}^i(k))\|_1 \geq 4\delta^2 K_i^\Phi K_i\}} \right] \\ &\leq 4\delta^2 K_i^\Phi \mathbb{E} \left[\|\hat{x}^i(k)\|_1 \mathbf{1}_{\{4\delta^2 K_i^\Phi \|\hat{x}^i(k)\|_1 \geq 4\delta^2 K_i^\Phi K_i\}} \right] \leq 4\delta^2 K_i^\Phi \varepsilon \end{aligned}$$

Since $W(\hat{x}^i(k))$ is a continuous function of $\hat{x}_i(k)$ by the continuous mapping theorem, $W(\hat{x}^i(k)) \xrightarrow{p} W(x^*), \forall i$. This, coupled with the UI of $\{W(\hat{x}^i(k))\}$ for all i implies that $W(\hat{x}^i(k)) \xrightarrow{L^1} W(x^*), \forall i$. The signal measurements $z_i(\tau)$ in Eq. (20) are independent of the estimates $W(\hat{x}^i(k))$ because the latter are based on the relative measurements in Eq. (17). Therefore,

$$\begin{aligned} \mathbb{E} \hat{g}_i(k) &= \mathbb{E}[\text{col}_i(W(\hat{x}^i(k)))] \frac{1}{k+1} \sum_{\tau=0}^k \mathbb{E} z_i(\tau) \\ &= \mathbb{E}[\text{col}_i(W(\hat{x}^i(k)))] h(x_i^*, y) \rightarrow \text{col}_i(W(x^*)) h(x_i^*, y) \end{aligned} \quad (23)$$

Now, consider the behavior of the consensus filter in Eq. (21) with $\mu_{i,k} = \hat{g}_i(k)$. Eliminating the state $q_{i,k}$ and writing the equations in matrix form gives

$$r_{k+1} = (I_{nd_c} - \beta(L \otimes I_{d_x})) r_k + (\mu_{k+1} - \mu_k)$$

where L is the Laplacian of the communication graph G . Taking expectations above results in a deterministic linear time-invariant system, which was analyzed in Ref. [37]. In light of Eq. (23), Proposition 1 in Ref. [37] shows that for all i

$$\lim_{k \rightarrow \infty} \left(\mathbb{E}[r_{i,k}] - \frac{1}{n} \sum_{i=1}^n \text{col}_i(W(x^*)) h(x_i^*, y) \right) = 0$$

Finally, the FD approximation in Eq. (5) shows that

$$\lim_{k \rightarrow \infty} \mathbb{E} r_{ik} = \frac{1}{n} W(x^*) \begin{pmatrix} h(x_1^*, y) \\ \vdots \\ h(x_n^*, y) \end{pmatrix} = \frac{1}{n} (g(m^*, y) + b), \quad \forall i$$

References

- [1] Lux, R., and Shi, W., 2004, "Chemotaxis-Guided Movements in Bacteria," *Crit. Rev. Oral. Biol. Med.*, **15**(4), pp. 207–220.

- [2] Frankel, R., Bazylinski, D., Johnson, M., and Taylor, B., 1997, "Magneto-Aerotaxis in Marine Coccoid Bacteria," *Biophys. J.*, **73**(2), pp. 994–1000.
- [3] Ögren, P., Fiorelli, E., and Leonard, N., 2004, "Cooperative Control of Mobile Sensor Networks," *IEEE Trans. Autom. Control*, **49**(8), pp. 1292–1302.
- [4] Sukhatme, G., Dhariwal, A., Zhang, B., Oberg, C., Stauffer, B., and Caron, D., 2007, "Design and Development of a Wireless Robotic Networked Aquatic Microbial Observing System," *Environ. Eng. Sci.*, **24**(2), pp. 205–215.
- [5] Rybski, P., Stoeter, S., Erickson, M., Gini, M., Hougen, D., and Papanikolopoulos, N., 2000, "A Team of Robotic Agents for Surveillance," International Conference on Autonomous Agents, Barcelona, Spain, *ACM*, New York, pp. 9–16.
- [6] Kumar, V., Rus, D., and Singh, S., 2004, "Robot and Sensor Networks for First Responders," *IEEE Pervasive Comput.*, **3**(4), pp. 24–33.
- [7] Wu, W., and Zhang, F., 2011, "Experimental Validation of Source Seeking With a Switching Strategy," *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 3835–3840.
- [8] Li, S., and Guo, Y., 2012, "Distributed Source Seeking by Cooperative Robots: All-to-All and Limited Communications," *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1107–1112.
- [9] Brinón-Arranz, L., and Schenato, L., 2013, "Consensus-Based Source-Seeking With a Circular Formation of Agents," *European Control Conference*, pp. 2831–2836.
- [10] Zhang, F., and Leonard, N., 2010, "Cooperative Filters and Control for Cooperative Exploration," *IEEE Trans. Autom. Control*, **55**(3), pp. 650–663.
- [11] Choi, J., Oh, S., and Horowitz, R., 2009, "Distributed Learning and Cooperative Control for Multi-Agent Systems," *Automatica*, **45**(12), pp. 2802–2814.
- [12] Jadaliha, M., Lee, J., and Choi, J., 2012, "Adaptive Control of Multiagent Systems for Finding Peaks of Uncertain Static Fields," *ASME J. Dyn. Syst. Meas. Control*, **134**(5), p. 051007.
- [13] Azuma, S., Sakar, M., and Pappas, G., 2012, "Stochastic Source Seeking by Mobile Robots," *IEEE Trans. Autom. Control*, **57**(9), pp. 2308–2321.
- [14] Zhang, C., Arnold, D., Ghods, N., Siranosian, A., and Krstić, M., 2007, "Source Seeking With Non-Holonomic Unicycle Without Position Measurement and With Tuning of Forward Velocity," *Syst. Control Lett.*, **56**(3), pp. 245–252.
- [15] Liu, S., and Krstić, M., 2010, "Stochastic Source Seeking for Nonholonomic Unicycle," *Automatica*, **46**(9), pp. 1443–1453.
- [16] Stanković, M., and Stipanović, D., 2010, "Extremum Seeking Under Stochastic Noise and Applications to Mobile Sensors," *Automatica*, **46**(8), pp. 1243–1251.
- [17] Ghods, N., and Krstić, M., 2011, "Source Seeking With Very Slow or Drifting Sensors," *ASME J. Dyn. Syst. Meas. Control*, **133**(4), p. 044504.
- [18] Atanasov, N., Le Ny, J., Michael, N., and Pappas, G., 2012, "Stochastic Source Seeking in Complex Environments," *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 3013–3018.
- [19] Charrow, B., Kumar, V., and Michael, N., 2013, "Approximate Representations for Multi-Robot Control Policies That Maximize Mutual Information," *Robotics: Science and Systems (RSS)*, Berlin, Germany.
- [20] Hoffmann, G., and Tomlin, C., 2010, "Mobile Sensor Network Control Using Mutual Information Methods and Particle Filters," *IEEE Trans. Autom. Control*, **55**(1), pp. 32–47.
- [21] Dames, P., Schwager, M., Kumar, V., and Rus, D., 2012, "A Decentralized Control Policy for Adaptive Information Gathering in Hazardous Environments," *IEEE Conference on Decision and Control (CDC)*, Dec., pp. 2807–2813.
- [22] Julian, B., Angermann, M., Schwager, M., and Rus, D., 2012, "Distributed Robotic Sensor Networks: An Information-Theoretic Approach," *Int. J. Rob. Res.*, **31**(10), pp. 1134–1154.
- [23] Yin, G., Yuan, Q., and Wang, L., 2013, "Asynchronous Stochastic Approximation Algorithms for Networked Systems: Regime-Switching Topologies and Multiscale Structure," *SIAM Multiscale Model. Simul.*, **11**(3), pp. 813–839.
- [24] Yu, W., Zheng, W., Chen, G., Ren, W., and Cao, J., 2011, "Second-Order Consensus in Multi-Agent Dynamical Systems With Sampled Position Data," *Automatica*, **47**(7), pp. 1496–1503.
- [25] Derenick, J., and Spletzer, J., 2007, "Convex Optimization Strategies for Coordinating Large-Scale Robot Formations," *IEEE Trans. Rob.*, **23**(6), pp. 1252–1259.
- [26] Fornberg, B., Lehto, E., and Powell, C., 2013, "Stable Calculation of Gaussian-Based RBF-FD Stencils," *Comput. Math. Appl.*, **65**(4), pp. 627–637.
- [27] Kushner, H., and Yin, G., 2003, *Stochastic Approximation and Recursive Algorithms and Applications*, 2 ed., Springer-Verlag, New York.
- [28] Borkar, V., 2008, *Stochastic Approximation: A Dynamical Systems Viewpoint*, Cambridge University Press, Cambridge, UK.
- [29] Ljung, L., 1977, "Analysis of Recursive Stochastic Algorithms," *IEEE Trans. Autom. Control*, **22**(4), pp. 551–575.
- [30] Spall, J., 2003, *Introduction to Stochastic Search and Optimization*, John Wiley & Sons, Hoboken, NJ.
- [31] Schwager, M., Dames, P., Rus, D., and Kumar, V., 2011, "A Multi-Robot Control Policy for Information Gathering in the Presence of Unknown Hazards," *Proceedings of International Symposium on Robotics Research*, Aug.
- [32] Thrun, S., Burgard, W., and Fox, D., 2005, *Probabilistic Robotics*, MIT, Cambridge, MA.
- [33] Rad, K., and Tahbaz-Salehi, A., 2010, "Distributed Parameter Estimation in Networks," *IEEE Conference on Decision and Control (CDC)*, pp. 5050–5055.
- [34] Shahrampour, S., and Jadbabaie, A., 2013, "Exponentially Fast Parameter Estimation in Networks Using Distributed Dual Averaging," *IEEE Conference on Decision and Control (CDC)*, pp. 6196–6201.
- [35] Tahbaz-Salehi, A., and Jadbabaie, A., 2010, "Consensus Over Ergodic Stationary Graph Processes," *IEEE Trans. Autom. Control*, **55**(1), pp. 225–230.
- [36] Atanasov, N., Le Ny, J., and Pappas, G., 2014, "Distributed Algorithms for Stochastic Source Seeking With Mobile Robot Networks: Technical Report," preprint arXiv: 1402.0051.
- [37] Spanos, D., Olfati-Saber, R., and Murray, R., 2005, "Dynamic Consensus on Mobile Networks," 16th IFAC World Congress, International Federation of Automatic Control, Prague, Czech Republic.
- [38] Capulli, F., Monti, C., Vari, M., and Mazzenga, F., 2006, "Path Loss Models for IEEE 802.11a Wireless Local Area Networks," *3rd International Symposium on Wireless Communication Systems*, pp. 621–624.
- [39] Durrett, R., 2010, *Probability: Theory and Examples*, Vol. 4, Cambridge University, Cambridge University Press, New York.