

Nonuniform Line Coverage From Noisy Scalar Measurements

Peter Davison, Naomi Ehrich Leonard, Alex Olshevsky, and Michael Schwemmer

Abstract—We study the problem of distributed coverage control in a network of mobile agents arranged on a line. The goal is to design distributed dynamics for the agents to achieve optimal coverage positions with respect to a scalar density field that measures the relative importance of each point on the line. Unlike previous work, which implicitly assumed the agents know this density field, we only assume that each agent can access noisy samples of the field at points close to its current location. We provide a simple randomized protocol wherein every agent samples the scalar field at three nearby points at each step and which guarantees convergence to the optimal positions. We further analyze the convergence time of this protocol and show that, under suitable assumptions, the squared distance to the optimal coverage configuration decays as $O(1/t)$ with the number of iterations t , where the constant scales polynomially with the number of agents n . We illustrate these results with simulations.

Index Terms—Decentralized control, networked control systems.

I. INTRODUCTION

As technological advances have improved the capabilities, reliability, and cost of robotic sensing platforms, their potential for deployment in autonomous, cooperative networks has gained significant attention. The emergent capabilities of mobile sensor networks promise to revolutionize complex tasks such as surveillance, exploration and environmental monitoring. However, development of high-level capabilities requires solutions to lower-level problems such as formation control and coverage control, and these solutions should be distributed, adaptive to changing environments, and robust to uncertainty and changes in network topology.

The present work focuses on coverage control, where the goal is to optimally locate the nodes, or agents, to maximize the so-called coverage metric, which measures the largest distance from a point in a domain of interest to the closest node. The coverage problem may be thought of in terms of interception time: the optimal coverage locations in a domain minimize the largest “response time” from the node locations to a point in the domain.

Manuscript received October 15, 2013; revised May 25, 2014 and October 16, 2014; accepted December 2, 2014. Date of publication December 2, 2014; date of current version June 24, 2015. This work was supported in part by ONR Grant N00014-09-1-1074, by the MBI, and by the NSF under Grant DMS 0931642. Recommended by Associate Editor A. Papachristodoulou.

P. Davison is with the Department of Aeronautics and Astronautics, MIT, Cambridge, MA 02139 USA (e-mail: pdavison@mit.edu).

N. E. Leonard is with the Department of Mechanical and Aerospace Engineering, Princeton University, Princeton, NJ 08544 USA (e-mail: naomi@princeton.edu).

A. Olshevsky is with the Department of Industrial and Enterprise Systems Engineering, University of Illinois at Urbana-Champaign, Urbana, IL 01536 USA (e-mail: aolshev2@illinois.edu).

M. Schwemmer is with the Mathematical Biosciences Institute, Ohio State University, Columbus OH 43210 USA (e-mail: schwemmer.2@mbi.osu.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TAC.2014.2375732

The coverage problem often involves distances which differ from the ordinary Euclidean metric by a weight factor which puts a heavier weight on some regions relative to others; in this case, the problem is referred to as the *nonuniform* coverage control problem, whereas if the distances are Euclidean then it is standard to refer to the problem as the *uniform* coverage problem. The weight factor of each point is usually referred to as the density field. Intuitively, one may think of minimizing interception time in a terrain of varying roughness, where different regions of the terrain take longer to traverse compared to others. In a nonuniform field, agents should be closer together in regions of higher density and more spread out in regions of lower density.

We will consider a particular case of the nonuniform coverage problem when the agents are arranged on a line. Our interest in the line coverage problem is motivated by two distinct considerations.

First, the line coverage problem is the natural model for the “border patrol” problem, wherein we must position n nodes on a curve in \mathbb{R}^2 such that the line distance from any point on the line to the closest node is minimized. After a reparametrization of the curve, this is exactly the line coverage problem. In this context, the curve usually represents a physical border; the n nodes will usually be autonomous vehicles; the distance metric being optimized is the largest interception time from any point on the border to the closest vehicle; and the density of a location represents difficulty of travel, determined by the roughness of the terrain in that location.

The terrain roughness can be measured by each vehicle at its current location using a laser stripe generator (see [21]). Correspondingly, we seek an adaptive protocol by means of which n vehicles can explore the border and optimally position themselves without knowing the terrain roughness in advance. Since the optimal position will depend on the terrain roughness at all points on the border, any protocol for this problem will need to ensure that the vehicles do not neglect any part of the border in their sampling. Consequently, we will study algorithms wherein each vehicle repeatedly samples the terrain density in the region of the border closer to it than to other agents (thus ensuring no part of the border is neglected) and, based on these samples and the positions of its neighbors, moves to a new location.

A related application is adaptive sampling applications in two and three dimensions. There is often one dimension where nonuniformity dominates and a protocol is needed for coverage in the dominant direction; e.g., in the ocean autonomous vehicles measuring temperature would use the protocol to optimally position themselves along the thermocline in the vertical water column.

Secondly, the general coverage control problem has been the subject of much recent interest within the control community; we refer the reader to the recent papers [1]–[7], [9]–[16] [17], [19], [20] and the references therein. The nonuniform line coverage problem is the one-dimensional version of the general coverage control problem, and consequently, it provides a simplified setting to make advances in addressing outstanding questions in coverage control.

Due to the large amount of related literature, we will only attempt to survey the works most directly relevant to our problem and approach. In [5] and [6], a distributed uniform coverage control law is developed which makes use of Voronoi partitions and gradient descent laws. The nonuniform case is treated in part by using density-dependent gradient descent laws with the Voronoi partitions computed for the uniform case. Coverage with communication constraints is treated in [13], [17].

The nonuniform coverage control problem is addressed in [11] where a density-dependent distance metric is defined that stretches and shrinks subregions of high and low density. A cartogram transformation (which needs to be known by the agents) is then used to compute the Voronoi partitions and convergence to optimal nonuniform coverage is proved in the case of a static or slowly time-varying density field. The necessity of knowing the cartogram transformation is relaxed in the case of nonuniform coverage control on the line in [12] where fully distributed nonuniform coverage protocols are derived. A follow-up work [7] considers the line coverage problem where only samples of ρ are available and explores the performance guarantees associated with some simple strategies. Some recent work [4], [8] considers strategies for coverage on the line when some proportion of the sensors are expected to randomly fail.

In this work, our focus is on deriving rigorously correct protocols for optimal line coverage under the assumption that every node has access to noisy measurements of the field ρ . This is a considerably weaker assumption compared to the previous literature (e.g., [5], [6], [11], [12]) where convergence to optimal coverage was established with update rules that involved exact computation of integrals of ρ , thus implicitly assuming that the density is known to all the nodes. Our main result is a randomized protocol that drives all the agents to the optimal configuration from only three noisy samples of ρ at each step. Moreover, we derive upper bounds on the convergence time of our protocol that scale polynomially with the number of nodes n . We perform simulations that show the convergence times of our protocol are quite reasonable, scaling considerably faster than our worst-case bounds.

The paper is structured as follows. In Section II we provide the necessary background to state the problem formally and briefly summarize our main results. Section III then contains our protocol and a proof of its convergence and convergence rate. We illustrate these results with simulations in Section IV.

II. BACKGROUND AND SUMMARY OF OUR RESULTS

We will consider a network of n mobile agents whose locations are assumed, for simplicity, to lie within the interval $[0,1]$. We will denote the positions of these agents, or nodes, at time step t by $x_1^{(t)}, x_2^{(t)}, \dots, x_n^{(t)}$. We adopt the convention that the labeling of agents from 1 to n matches their initial order at time step 0 along the line from left to right, i.e., $x_1^{(0)} \leq x_2^{(0)} \leq \dots \leq x_n^{(0)}$. We will use $x_0^{(t)}$ and $x_{n+1}^{(t)}$ to denote 0 and 1 for all t , respectively; this will simplify notation.

The information density field, $\rho : \mathbb{R} \rightarrow (0, \infty)$ is assumed to be a positive differentiable function which is bounded above and bounded away from zero from below. Following [11], we define the distance between two points a, b as

$$d_\rho(a, b) = \int_{\min(a,b)}^{\max(a,b)} \rho(z) dz.$$

Intuitively, this distance function stretches regions of high ρ relative to regions of lower ρ . Using this notion of distance, the coverage metric Φ is then the largest distance from any point in the domain $[0,1]$ to the agent that is nearest to it:

$$\Phi(x_1, \dots, x_n, \rho) = \max_{y \in [0,1]} \left[\min_{i=1, \dots, n} d_\rho(y, x_i) \right].$$

The optimum coverage Φ^* is the minimum of Φ over all possible agent configurations in the interval $[0,1]$. It is not hard to see (and was proven in [12]) that for any ρ satisfying the above assumptions, Φ is minimized at a unique vector $x = x^*$ among vectors with nondecreasing entries, i.e., those satisfying $x_1 \leq x_2 \leq \dots \leq x_n$.

The nonuniform line coverage problem is to design a protocol that drives all the agents to a configuration achieving coverage Φ^* . Furthermore, this protocol must be distributed, meaning that each agent is limited to repeatedly updating its position based only on the positions of its closest neighbors on the left and right as well as measurements of the density ρ taken near its location. The special case when ρ is identically equal to 1 is referred to as the uniform line coverage problem.

It is easy to see that the optimal coverage configuration is invariant under any scaling of the density field ρ . We will therefore assume henceforth, without loss of generality, that the range of ρ is $[1, \rho_{\max}]$ for some finite positive $\rho_{\max} \geq 1$.

A variety of protocols for the line coverage problems are available (see [5], [6], [11] and the recent paper [12] focused on the line coverage problem). However, all of these works implicitly assume that the agents know ρ exactly because they include integrals of ρ in the update equations used by each agent. By contrast, in this paper we instead assume that each agent only has access to *noisy* samples of ρ . Specifically, we will assume that every agent can take samples of the form $\hat{\rho}(z) = \rho(z) + w$, where w is noise and z is a point between the measuring agent's left and right neighbors. We remark that this assumption may involve some physical travel on the part of each node at every step; for example, if it chooses to sample ρ at a point outside its physical sensing radius (but still between its left and right neighbors), it will need to move closer to that point. The noises w are assumed to be independent, have zero mean and bounded support, and an upper bound on this support, which we will denote by M , is known to all the nodes. Furthermore, we will assume that $\hat{\rho}(z)$ is nonnegative with probability 1; this occurs, for example, if the noise support is not too large. Intuitively, since the density $\rho(z)$ represents difficulty of travel (and is therefore nonnegative) we require noisy estimation of it to result in nonnegative samples.

Our main contribution in this paper is a protocol for the nonuniform line coverage problem wherein each node uses only three samples of the density at each step. Moreover, we obtain precise bounds on the convergence speed of our protocol under the additional assumption that every node knows a rough estimate of the total number of nodes in the network. We will show that the per-node expected square distance from optimal coverage configuration decays as $O(n^5/t)$, where the constant depends on the quantity ρ_{\max} as well as on the support of the noise distribution M .

III. NONUNIFORM LINE COVERAGE FROM NOISY SCALAR MEASUREMENTS

We next describe the line coverage protocol we present and analyze here. We begin with an informal sketch. We will show that optimal coverage configurations can be characterized as the minima of a certain Lyapunov function, and the first-order conditions for optimality prescribe that the ρ -weighted distances between each agent and each of its two closest neighbors be in a certain proportion. This naturally suggests an algorithm wherein each agent repeatedly moves to put these distances in the right proportion. However, without exact knowledge of ρ , the nodes estimate the current ρ -weighted distances through random sampling. In particular, each node will sample ρ at three points, one between itself and its right neighbor, one between itself and its left neighbor, and one at its current location. These samples of ρ allow it to compute an unbiased estimate of where its next location should be in order to put the two distances into the desired proportion.¹

¹In line with previous protocols for coverage control which relied on computation of Voronoi partitions (e.g., [5], [6], [11]), this process may be thought of in terms of each node computing its Voronoi cell (interval) from a few noisy samples at each stage.

However, because all the estimates obtained by the nodes in this way are noisy, we will introduce a stepsize; thus nodes will move only part-way towards their new positions, and the size of the move at each step will (slowly) decay to zero with time. Intuitively, the position of each node is influenced by all previous samples of ρ collected by it and neighboring nodes, and consequently as time goes on and the node's position reflects more and more past samples, the node will need to move less in response to each new sample.

A. A Formal Statement of the Protocol

At time step t , node k will sample the density ρ at a random point $r_k^{(t)}$ uniformly in the interval $[x_k^{(t)}, x_{k+1}^{(t)}]$, i.e., between itself and its right neighbor; and at a random point $l_k^{(t)}$ uniformly in the interval $[x_{k-1}^{(t)}, x_k^{(t)}]$, i.e., between itself and its left neighbor (recall our convention that $x_0^{(t)} = 0, x_{n+1}^{(t)} = 1$). Finally, node k samples the density at its own location $x_k^{(t)}$. After obtaining these samples node k proceeds to set²

$$\begin{aligned} R_k^{(t)} &= \widehat{\rho}\left(r_k^{(t)}\right)\left(x_{k+1}^{(t)} - x_k^{(t)}\right) \\ L_k^{(t)} &= \widehat{\rho}\left(l_k^{(t)}\right)\left(x_k^{(t)} - x_{k-1}^{(t)}\right) \end{aligned}$$

and then to update

$$\begin{aligned} x_1^{(t+1)} &= x_1^{(t)} - \frac{\alpha(t)\widehat{\rho}\left(x_1^{(t)}\right)}{8(\rho_{\max} + M)^2}\left(2L_1^{(t)} - R_1^{(t)}\right) \\ x_k^{(t+1)} &= x_k^{(t)} - \frac{\alpha(t)\widehat{\rho}\left(x_k^{(t)}\right)}{8(\rho_{\max} + M)^2}\left(L_k^{(t)} - R_k^{(t)}\right) \quad (1) \\ &\text{when } k = 2, \dots, n-1 \end{aligned}$$

$$x_n^{(t+1)} = x_n^{(t)} - \frac{\alpha(t)\widehat{\rho}\left(x_n^{(t)}\right)}{8(\rho_{\max} + M)^2}\left(L_n^{(t)} - 2R_n^{(t)}\right) \quad (2)$$

where $\alpha(t)$ is a stepsize which satisfies

$$0 \leq \alpha(t) \leq 1 \text{ for all } t, \quad \sum_{t=0}^{\infty} \alpha(t) = +\infty, \quad \sum_{t=0}^{\infty} \alpha^2(t) < \infty.$$

We remark that choosing the stepsize $\alpha(t) = 1/t^p$ for any exponent $p \in (1/2, 1]$ satisfies all three of these conditions. Finally, we will refer to this protocol as the *randomized scalar coverage protocol*.

B. Main Result

The main result of this paper is the following theorem.

Theorem 1: The positions $x^{(t)} = (x_1^{(t)}, \dots, x_n^{(t)})^T$ converge to the unique minimizer configuration $x^* = (x_1^*, \dots, x_n^*)^T$ of the coverage metric Φ with probability one. Moreover, if every node knows an upper bound U on the total number of nodes n and chooses the stepsize $\alpha(t) = 8U^2(\rho_{\max} + M)^2 / (8U^2(\rho_{\max} + M)^2 + t)$ then we will have the expected error bound

$$E\left[\frac{1}{n}\|x^{(t)} - x^*\|_2^2\right] \leq \frac{16nU^4(\rho_{\max} + M)^4(4\rho_{\max}^2 + 2\|\rho'\|_{\infty}\rho_{\max})}{8U^2(\rho_{\max} + M)^2 + t}$$

where $\|\rho'\|_{\infty} = \sup_{z \in [0,1]} |\rho'(z)|$.

Under the assumption that every node approximately knows the total number of nodes, e.g., if we have $n \leq U \leq 2n$, an implication of this theorem is that we need to wait $O(n^5/\epsilon)$ iterations until the average square error $E[\frac{1}{n}\|x^{(t)} - x^*\|_2^2]$ is below ϵ in expectation, where the constant within the O -notation depends on the density ρ

²Recall our notation: $\widehat{\rho}(z)$ is $\rho(z)$ plus a zero-mean random variable with support in $[-M, M]$, and all these random variables are jointly independent.

and the noise support M . Thus, the main result of this paper is that this decay is linear in the number of iterations t and the constant in front of this decay scales polynomially with the number of nodes n .

C. Proofs

We will shortly turn to the proof of Theorem 1. Before doing so, however, we need to demonstrate something more basic: that the protocol preserves the ordering of the agents i.e., that $x_1^{(t)} \leq x_2^{(t)} \leq \dots \leq x_n^{(t)}$ for all t . If this were not the case, our protocol would not be truly *distributed*: since the next position of node k is affected by the positions of nodes $k-1$ and $k+1$ at each step, it is crucial that these three nodes continue to be each other's closest neighbors.

Proposition 2: Under the randomized scalar coverage protocol, we have that with probability 1

$$0 = x_0^{(t)} \leq x_1^{(t)} \leq x_2^{(t)} \leq \dots \leq x_n^{(t)} \leq x_{n+1}^{(t)} = 1$$

for all integers t .

Proof of Proposition 2: By assumption the statement is true at time $t=0$, and we prove it by induction. Suppose that the statement holds at time t and consider node i . As a consequence of the update rule (2) and the fact that $R_i^{(t)}, L_i^{(t)} \geq 0$, we have

$$\begin{aligned} x_i^{(t+1)} - x_i^{(t)} &\leq \frac{\alpha(t)}{8(\rho_{\max} + M)} 2R_i^{(t)} \frac{\widehat{\rho}\left(x_i^{(t)}\right)}{\rho_{\max} + M} \\ &\leq \frac{R_i^{(t)}}{4(\rho_{\max} + M)} \\ &\leq \frac{x_{i+1}^{(t)} - x_i^{(t)}}{4}. \end{aligned}$$

A similar argument establishes that

$$x_i^{(t+1)} - x_i^{(t)} \geq \frac{1}{4}\left(x_{i-1}^{(t)} - x_i^{(t)}\right)$$

and these two inequalities imply the proposition. \blacksquare

Now that we have established that the protocol remains distributed by preserving the ordering of the nodes, we turn to the proof of Theorem 1. First we will argue that the optimal coverage point is the minimum of a certain Lyapunov function; this is Lemma 3 below, which defines a function $Q(x_1, \dots, x_n)$ minimized at optimal coverage. Next, we bound how much $Q(x_1, \dots, x_n)$ decreases at each step. This appears to be difficult to do directly. However, relying on the key idea that Q becomes convex after a position-dependent stretching of the coordinate space, we will prove a number of inequalities which will allow us to argue that Q decreases at every step by a certain fraction of the full distance to the optimal value.

We begin now with a series of lemmas executing this plan which will culminate in the proof of Theorem 1. Our first lemma introduces the Lyapunov function Q :

Lemma 3: Define

$$\begin{aligned} Q(x_1, \dots, x_n) &= 2\left(\int_0^{x_1} \rho(z) dz\right)^2 + \left(\int_{x_1}^{x_2} \rho(z) dz\right)^2 \\ &\quad + \dots + \left(\int_{x_{n-1}}^{x_n} \rho(z) dz\right)^2 + 2\left(\int_{x_n}^1 \rho(z) dz\right)^2. \end{aligned}$$

Then $Q(x_1, \dots, x_n)$ has a unique global minimizer $x^* = (x_1^*, x_2^*, \dots, x_n^*)^T$; this minimizer x^* is also a minimizer of Φ and satisfies $0 \leq x_1^* \leq x_2^* \leq \dots \leq x_n^* \leq 1$.

Proof: It is easy to see that $Q(x_1, \dots, x_n)$ must have a global minimum, since it is continuous and blows up if at least one of the variables goes off to infinity. Setting all partial derivatives of Q to zero and rearranging, we obtain the set of equations

$$\begin{aligned} 2 \int_0^{x_1} \rho(z) dz &= \int_{x_1}^{x_2} \rho(z) dz \\ \int_{x_{i-1}}^{x_i} \rho(z) dz &= \int_{x_i}^{x_{i+1}} \rho(z) dz \quad \text{for } i = 2, \dots, n-1 \\ \int_{x_{n-1}}^{x_n} \rho(z) dz &= 2 \int_{x_n}^1 \rho(z) dz. \end{aligned} \tag{3}$$

Any global minimum of Q must satisfy these equations. It is in Lemma 2 in [12] that these equations have a unique solution which minimizes Φ , and the proof of that lemma establishes that the entries of this solution are monotonic and lie in $[0, 1]$. ■

For convenience of notation, we will define F to be the function which maps $[0, 1]$ into $[0, \rho_{\max}]$ by $F(x) = \int_0^x \rho(z) dz$. Note that (3) can be conveniently restated in terms of F , e.g., the equality for $i = 2, \dots, n-1$ in (3) is simply $F(x_i) - F(x_{i-1}) = F(x_{i+1}) - F(x_i)$.

Next, we will need a technical estimate on the smallest value assumed by a certain quadratic form on the unit sphere.

Lemma 4:

$$\min_{\|x\|_2=1} x_1^2 + (x_2 - x_1)^2 + \dots + (x_n - x_{n-1})^2 + x_n^2 \geq \frac{1}{n^2}.$$

Proof: Since $\|x\|_2 = 1$ we have that at least one component x_i satisfies $|x_i| \geq 1/\sqrt{n}$. Without loss of generality, let us assume $x_i > 0$; else, we can simply replace x with $-x$. We then have that

$$\frac{1}{\sqrt{n}} \leq x_i = (x_1 - 0) + (x_2 - x_1) + \dots + (x_i - x_{i-1}).$$

Applying Cauchy-Schwarz yields the lemma. ■

Next we prove a technical lemma which lower bounds the size of the gradient of Q at each step as a fraction of the distance to the optimal value. The proof uses the previous Lemma 4 and proceeds by relying on a coordinate-depending stretching of the space making Q convex.

Lemma 5:

$$\frac{\|Q'(x)\|_2^2}{Q(x) - Q(x^*)} \geq \frac{4}{n^2}.$$

Proof: Define

$$\begin{aligned} G(y_1, \dots, y_n) &= 2y_1^2 + (y_2 - y_1)^2 \\ &+ \dots + (y_n - y_{n-1})^2 + 2 \left(\int_0^1 \rho(z) dz - y_n \right)^2. \end{aligned}$$

Observe that

$$Q(x_1, \dots, x_n) = G(F(x_1), F(x_2), \dots, F(x_n)). \tag{4}$$

As a consequence of this and Lemma 3, we can conclude that $G(y)$ has a unique minimum y^* satisfying $y_i^* = \int_0^{x_i^*} \rho(z) dz$. The Hessian

of $G(y_1, \dots, y_n)$ is easily computed; it is

$$G''(y) = \begin{pmatrix} 6 & -2 & 0 & 0 & \dots & 0 \\ -2 & 4 & -2 & 0 & \dots & 0 \\ 0 & -2 & 4 & -2 & \dots & 0 \\ \vdots & \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & 0 & 0 & -2 & 4 & -2 \\ 0 & 0 & 0 & 0 & -2 & 6 \end{pmatrix}.$$

We will use the standard notation of e_i to mean the column vector with a 1 in the i 'th component and zero elsewhere; moreover, we will use $e_{i,j}$ to denote the vector with a 1 in the i 'th component, a -1 in the j 'th component, and zeros elsewhere. Then it is easy to verify that

$$G'' = 4e_1 e_1^T + 2 \sum_{i=1}^{n-1} e_{i,i+1} e_{i,i+1}^T + 4e_n e_n^T$$

so that its smallest eigenvalue satisfies

$$\begin{aligned} \lambda_{\min} &= \min_{\|y\|_2=1} y^T G'' y \\ &= \min_{\|y\|_2=1} 4y_1^2 + 2 \sum_{i=1}^{n-1} (y_i - y_{i+1})^2 + 4y_n^2 \geq \frac{2}{n^2} \end{aligned} \tag{5}$$

where the last inequality follows from Lemma 4. Thus G is a strongly convex function, and a standard bound on the norm of its gradient (see Lemma 3 in [18, Ch. 1.4]) is

$$\frac{\|G'(y)\|_2^2}{G(y) - G(y^*)} \geq 2\lambda_{\min}.$$

Further, for any $x \in \mathbb{R}^n$ choosing $y \in \mathbb{R}^n$ defined by $y_i = F(x_i)$, we have by (4) and our assumption that $\rho \geq 1$ everywhere, that

$$\frac{\|Q'(x)\|_2^2}{Q(x) - Q(x^*)} \geq \frac{\|G'(y)\|_2^2}{G(y) - G(y^*)}.$$

Putting together the last two inequalities and plugging in the bound on λ_{\min} from (5) we obtain the current lemma. ■

Having now established the lower bound on the norm of the gradient of $Q(x)$, we now proceed to the proof of Theorem 1. We proceed by arguing that our protocol is a randomized version of gradient descent on the function $Q(x)$, which can be shown to converge despite the nonconvexity of Q .

Proof of Theorem 1: We begin by rewriting the update equation in more convenient form. Specifically, comparing the randomized control law, namely (2), with the definition of the function $Q(x_1, \dots, x_n)$, we observe that we can write

$$x^{(t+1)} = x^{(t)} - \frac{\alpha^{(t)}}{16(\rho_{\max} + M)^2} \bar{g}^{(t)}$$

where $E[\bar{g}^{(t)}] = Q'(x^{(t)})$ and

$$\begin{aligned} \|\bar{g}^{(t)}\|_2^2 &= \left(\hat{\rho}(x_1^{(t)}) \left(4L_1^{(t)} - 2R_1^{(t)} \right) \right)^2 \\ &+ \sum_{i=2}^{n-1} \left(\hat{\rho}(x_i^{(t)}) \left(2L_i^{(t)} - 2R_i^{(t)} \right) \right)^2 \\ &+ \left(\hat{\rho}(x_n^{(t)}) \left(2L_n^{(t)} - 4R_n^{(t)} \right) \right)^2 \\ &\leq (\rho_{\max} + M)^2 \sum_{i=1}^n \left(\max \left(4L_i^{(t)}, 4R_i^{(t)} \right) \right)^2 \\ &\leq 64(\rho_{\max} + M)^4 \end{aligned}$$

and all the above inequalities hold with probability 1. Next, we observe that Q is a twice differentiable function, so that we may expand it in a Taylor series. Since

$$\frac{\partial^2 Q}{\partial x_1^2}(x) = 6\rho^2(x_1) + 2\rho'(x_1) \left(2 \int_0^{x_1} \rho(z) dz - \int_{x_1}^{x_2} \rho(z) dz \right)$$

$$\frac{\partial^2 Q}{\partial x_n^2}(x) = 6\rho^2(x_n) + 2\rho'(x_n) \left(\int_{x_{n-1}}^{x_n} \rho(z) dz - 2 \int_{x_n}^1 \rho(z) dz \right)$$

$$\frac{\partial^2 Q}{\partial x_i^2}(x) = 4\rho^2(x_i) + 2\rho'(x_i) \left(\int_{x_{i-1}}^{x_i} \rho(z) dz - \int_{x_i}^{x_{i+1}} \rho(z) dz \right)$$

when $i = 2, \dots, n-1$

$$\frac{\partial^2 Q}{\partial x_i \partial x_j}(x) = -2\rho(x_i)\rho(x_j) \text{ when } j \in \{i-1, i+1\}$$

$$\frac{\partial^2 Q}{\partial x_i \partial x_j}(x) = 0 \text{ when } j \notin \{i-1, i, i+1\}$$

by Gershgorin circles, it follows that as long as $0 \leq x_1^{(t)} \leq \dots \leq x_n^{(t)} \leq 1$, the largest eigenvalue of $Q''(x)$ is never more than $8\rho_{\max}^2 + 4\|\rho'\|_{\infty}\rho_{\max}$ in magnitude. Thus as long as both $x^{(t)}$ and $x^{(t+1)}$ have entries between 0 and 1 and monotonically nondecreasing (note that by Proposition 2 this holds at every time t), we may bound $Q(x^{(t+1)})$ via the Taylor expansion with Lagrange remainder form as

$$Q(x^{(t+1)}) \leq Q(x^{(t)}) + \nabla Q(x^{(t)})^T (x^{(t+1)} - x^{(t)}) + (4\rho_{\max}^2 + 2\|\rho'\|_{\infty}\rho_{\max}) \|x^{(t+1)} - x^{(t)}\|_2^2$$

or

$$Q(x^{(t+1)}) \leq Q(x^{(t)}) - \frac{\alpha(t)}{16(\rho_{\max} + M)^2} \sum_{i=1}^n [Q'(x^{(t)})]_i \bar{g}_i^{(t)} + \frac{(4\rho_{\max}^2 + 2\|\rho'\|_{\infty}\rho_{\max}) \alpha^2(t)}{16^2(\rho_{\max} + M)^4} \|\bar{g}^{(t)}\|_2^2.$$

Taking expectations and using Lemma 5

$$E [Q(x^{(t+1)}) - Q(x^*) | x^{(t)}] \leq \left(1 - \frac{\alpha(t)}{4n^2(\rho_{\max} + M)^2} \right) \times (Q(x^{(t)}) - Q(x^*)) + \frac{(4\rho_{\max}^2 + 2\|\rho'\|_{\infty}\rho_{\max}) \alpha^2(t)}{4}.$$

By Lemma 10 of [18, Ch. 1.4] it follows that $Q(x^{(t)}) - Q(x^*)$ approaches zero with probability 1, which by the uniqueness of the minimizer x^* implies that $x^{(t)} \rightarrow x^*$ with probability 1. Now under the additional assumption that every node knows an upper bound U and chooses the stepsize of $\alpha(t) = 8U^2(\rho_{\max} + M)^2 / (8U^2(\rho_{\max} + M)^2 + t)$, we have that

$$E [Q(x^{(t+1)}) - Q(x^*)] \leq \left(1 - \frac{2}{8U^2(\rho_{\max} + M)^2 + t} \right) E [Q(x^{(t)}) - Q(x^*)] + \frac{16U^4(\rho_{\max} + M)^4 (4\rho_{\max}^2 + 2\|\rho'\|_{\infty}\rho_{\max})}{(8U^2(\rho_{\max} + M)^2 + t)^2}.$$

We now claim that for all $t \geq 0$, we have that

$$E [Q(x^{(t)}) - Q(x^*)] \leq \frac{16U^4(\rho_{\max} + M)^4 (4\rho_{\max}^2 + 2\|\rho'\|_{\infty}\rho_{\max})}{8U^2(\rho_{\max} + M)^2 + t}.$$

We prove this claim by induction. Indeed, at $t = 0$, since the initial positions are in the interval $[0, 1]$, it is immediate that $Q(x(0)) \leq 2\rho_{\max}^2$, which proves the statement at $t = 0$. Now suppose that the inequality holds at time t . For simplicity of notation, let us adopt the shorthands $C = 16U^4(\rho_{\max} + M)^4 (4\rho_{\max}^2 + 2\|\rho'\|_{\infty}\rho_{\max})$ and $U' = 8U^2(\rho_{\max} + M)^2$. We then have that

$$\begin{aligned} E [Q(x^{(t+1)}) - Q(x^*)] &\leq \left(1 - \frac{2}{U' + t} \right) E [Q(x^{(t)}) - Q(x^*)] \\ &\quad + \frac{C}{(U' + t)^2} \\ &\leq \left(1 - \frac{2}{U' + t} \right) \frac{C}{(U' + t)} + \frac{C}{(U' + t)^2} \\ &\leq C \left(\frac{1}{U' + t} - \frac{1}{(U' + t)^2} \right) \\ &\leq \frac{C}{(U' + t + 1)} \end{aligned}$$

which proves the claim. Finally, we show this implies the theorem. Observe that $G(y) - G(y^*) \geq (1/n^2) \|y - y^*\|_2^2$ due to the fact that $\lambda_{\min}(G''(y)) \geq 2/n^2$ from (5). Thus

$$Q(x) - Q(x^*) \geq \frac{1}{n^2} \|F(x) - F(x^*)\|_2^2 \geq \frac{1}{n^2} \|x - x^*\|_2^2$$

where the last step follows from the fact that $\rho \geq 1$. \blacksquare

IV. SIMULATIONS

We briefly report on a simulation intended to gauge the practical convergence time of our protocol. Fig. 1 shows simulation for a system of 20 nodes for two starting points: one chosen uniformly at random and one which initially places all the nodes at one corner. All noises are uniform in $[-1/2, 1/2]$. In both cases, the density is uniform, which allows convergence to the optimal configuration to be ‘‘read off’’ from the graph by watching the spacings equalize. We chose the stepsize $\alpha(t)$ by setting $\alpha(t) = 1$ for the first half of the iterations and setting $\alpha(t) = 1/\sqrt{t}$ for the latter half; this decays more slowly as compared to the stepsize we used to obtain Theorem 1 but appears to be advantageous for a network of 20 nodes executing several thousand iterations.

Our simulations confirm our theoretical convergence results. Furthermore, they suggest that our error bounds are conservative; indeed, the system of twenty nodes reaches close to the optimal configuration after thousands of iterations, while the upper bounds of Theorem 1 are several orders of magnitude larger. An open question is to obtain improved theoretical guarantees that better match the faster speed we observe.

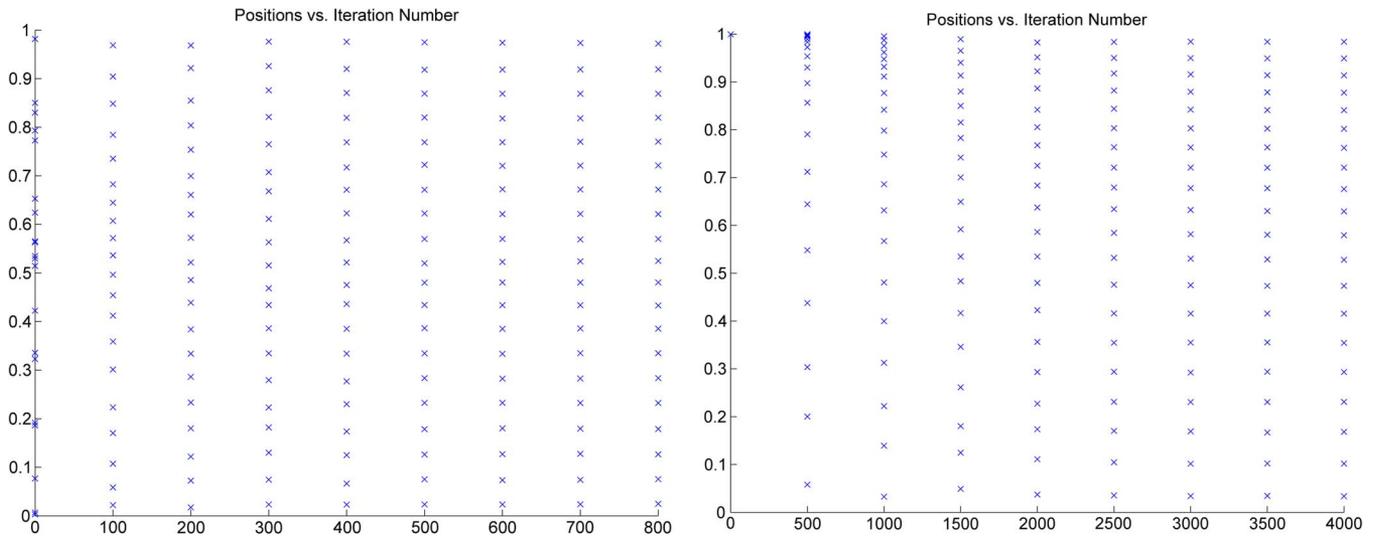


Fig. 1. Positions of the nodes as a function of iteration number for a network of 20 nodes. On the left, the starting positions are uniformly random in $[0,1]$, while on the right, all nodes begin with $x_i(0) = 1$.

REFERENCES

- [1] F. Bullo, R. Carli, and P. Frasca, "Gossip coverage control for robotic networks: Dynamical systems on the space of partitions," *SIAM J. Control Optimiz.*, vol. 50, no. 1, pp. 419–447, 2012.
- [2] C. H. Caicedo-Nunez and M. Zefran, "A coverage algorithm for a class of non-convex regions," in *Proc. 47th IEEE Conf. Decision Control*, 2008.
- [3] J. Cortes, "Coverage optimization and spatial load balancing by robotic sensor networks," *IEEE Trans. Autom. Control*, vol. 55, no. 3, pp. 749–754, Mar. 2010.
- [4] J. Cortes, "Deployment of an unreliable robotic networks for spatial estimation," *Syst. Control Lett.*, vol. 61, no. 1, pp. 41–49, 2012.
- [5] J. Cortes and F. Bullo, "Coordination and geometric optimization via distributed dynamical systems," *SIAM J. Control Optimiz.*, vol. 44, pp. 1543–1574, 2005.
- [6] J. Cortes, S. Martinez, T. Karatas, and F. Bullo, "Coverage control for mobile sensing networks," *IEEE Trans. Robot. Autom.*, vol. 20, no. 2, pp. 243–255, 2004.
- [7] P. Davison, M. Schwemmer, and N. E. Leonard, "Distributed nonuniform coverage with limited scalar measurements," in *Proc. 50th Allerton Conf. Communication, Control, Signal Process.*, 2012.
- [8] P. Frasca, F. Garin, B. Gerencser, and J. M. Hendrickx, "One-Dimensional Coverage by Unreliable Sensors," [Online]. Available: <http://arxiv.org/abs/1404.7711>
- [9] C. Gao, J. Cortes, and F. Bullo, "Notes on averaging over acyclic digraphs and discrete coverage control," *Automatica*, vol. 44, no. 8, pp. 2120–2127, 2008.
- [10] K. Laventall and J. Cortes, "Coverage control by multi-robot networks with limited range anisotropic sensory," *Int. J. Control*, vol. 82, no. 6, pp. 113–1121, 2009.
- [11] F. Lekien and N. E. Leonard, "Nonuniform coverage and cartograms," *SIAM J. Control Optimiz.*, vol. 48, no. 1, pp. 351–372, 2009.
- [12] N. Leonard and A. Olshevsky, "Nonuniform coverage control on the line," *IEEE Trans. Autom. Control*, vol. 58, no. 11, pp. 2743–2755, 2013.
- [13] W. Li and C. G. Cassandras, "Distributed cooperative coverage control of sensor networks," in *Proc. IEEE Conf. Decision Control*, 2005, pp. 2542–2547.
- [14] S. Martinez, F. Bullo, J. Cortes, and E. Frazzoli, "On synchronous robotic networks—Part II: Time complexity of rendezvous and deployment algorithms," *IEEE Trans. Autom. Control*, vol. 52, no. 12, pp. 2214–2226, Dec. 2007.
- [15] K. J. Obermeyer, A. Ganguli, and F. Bullo, "Multi-agent deployment for visibility coverage in polygonal environments with holes," *Int. J. Robust Nonlin. Control*, vol. 21, no. 12, pp. 1467–1492, 2011.
- [16] M. Pavone, A. Arsie, E. Frazzoli, and F. Bullo, "Distributed algorithms for environmental partitioning in mobile robotic networks," *IEEE Trans. Autom. Control*, vol. 56, no. 8, pp. 1834–1848, Aug. 2011.
- [17] S. Poduri and G. S. Sukhatme, "Constrained coverage for mobile sensor networks," in *Proc. Int. Conf. Robotics Automation*, 2004, pp. 165–171.
- [18] B. Polyak, *Introduction to Optimization*. New York: Optimization Software Inc., 1987.
- [19] S. Susca, S. Martinez, and F. Bullo, "Monitoring environmental boundaries with a robotic sensor network," *IEEE Trans. Control Syst. Technol.*, vol. 16, no. 2, pp. 288–296, 2008.
- [20] N. Sydney and D. A. Paley, "Optimal sampling of nonstationary spatiotemporal fields using a mobile sensory network," in *Proc. 3rd IFAC Workshop Distributed Estimation Control Netw. Syst.*, 2012.
- [21] A. M. Zhang and R. A. Russell, "Surface roughness measurement for outdoor mobile robotic applications," in *Proc. 2004 Australasian Conf. Robotics Automation*, 2004.