2011 50th IEEE Conference on Decision and Control and
European Control Conference (CDC-ECC)
Orlando, FL, USA, December 12-15, 2011

Nonuniform coverage control on the line

Naomi Ehrich Leonard and Alex Olshevsky

Abstract— **This paper investigates control laws allowing mobile, autonomous agents to optimally position themselves on the line for distributed sensing in a nonuniform field. We show that a simple static control law, based only on local measurements of the field by each agent, drives the agents to the optimal positions in time that is quadratic in the number of agents. Further, we exhibit a dynamic control law which, under slightly stronger assumptions on the capabilities and knowledge of each agent, drives the agents to the optimal positions an order of magnitude faster, namely in time that is linear in the number of agents. Both algorithms are fully distributed and robust to unpredictable loss and addition of agents.**

I. INTRODUCTION

Widespread deployment of networks of sensors and autonomous vehicles is expected to revolutionize our ability to monitor and control physical environments from remote locations. However, for such networks to achieve their full range of applicability, they must be capable of operating in uncertain and unstructured environments without centralized supervision. Realizing the full potential of such systems will require the development of protocols that are fully autonomous, distributed, and adaptive in the face of rapidly changing environments.

An important problem in this context is the coverage problem. A collection of mobile sensors need to determine how to distribute themselves over a region given an observation field they can measure; the sensors should be positioned so that the likelihood of detecting an event of interest is maximized. If the probability distribution of the event is uniform over the area, then the optimal solution will involve a uniform spacing out of the agents. On the other hand if this probability distribution is not uniform, then the sensors should be more densely positioned in the subregions that have higher event probability.

There is a considerable literature on coverage algorithms for groups of dynamic agents, which we do not survey in its entirety here; we will refer the reader to [12], [6], [5], [9], [2], [1], [3], [4], [11] and the references therein. In [5], uniform coverage algorithms are derived using Voronoi cells and gradient laws for distributed dynamical systems. Uniform constrained coverage control is addressed in [12] where the constraint is a minimum limit on node degree. Virtual potentials enable repulsion between agents to maximize coverage and attraction between agents to enforce the constraint. In [9], gradient control laws are proposed to move sensors to a configuration that maximizes expected event detection frequency. Local rules are enforced by defining a sensing radius for each agent, which also makes computations simpler. The approach is demonstrated for a nonuniform but symmetric density field with and without communication constraints. Further results for distributed coverage control are presented in [6] for a coverage metric defined in terms of the Euclidean metric with a weighting factor that allows for nonuniformity. As in [6], the methodology makes use of Voronoi cells and Lloyd descent algorithms. The papers [3], [4] identified a class of non-convex regions for which the coverage problem may be solved by reduction to the convex case through a well-chosen transformation of the region. The papers [1], [2] explored an optimization-based approach to some variations of the covering problem.

The paper [11] considered the general nonuniform coverage problem with a non-Euclidean distance, and it proposed and proved the correctness of a coverage control law in the plane. However, the control law of [11] is only partially distributed, in that it relies on a "cartogram computation" step which requires some global knowledge of the domain.

Our work builds on the results of [11] to design a control law for the nonuniform coverage problem in the one-dimensional case when the agents are positioned on the line. We develop fully distributed coverage control laws for a nonuniform field in this setting, and moreover, we prove quantitative convergence bounds on the performance of these algorithms. Interestingly, we find that relatively modest increases in the capabilities and knowledge of each agent can translate into considerable improvements of the global performance. These improvements are obtained by implementing more sophisticated distributed algorithms; in particular, we make heavy use of the technique of lifting Markov chains, introduced in [7], and subsequently used to accelerate distributed computation in [8].

We begin with an introduction to the nonuniform coverage problem in Section II. In Section III, we present our first fully distributed control law for the coverage problem. The execution of this control law only requires the agents to be able to measure distances to their neighbors and measure the field around their location. The main result of this section is Theorem 1, which demonstrates the correctness of the algorithm and gives a quantitative bound on its performance. We show that it takes n agents on the order of n^2 discrete-time updates to come close to the optimal configuration regardless of the initial conditions.

In Section IV, we present another fully distributed control law for coverage. The execution of this control law requires more capabilities on the part of the agents: they store

The authors are with Department of Mechanical and Aerospace Engineering, Princeton University, Princeton, NJ 08544, USA, naomi@princeton,edu, aolshevs@princeton.edu

This research was supported in part by AFOSR grant FA9550-07-1-0-0528 and ONR grant N00014-09-1-1074.

several numbers in memory, communicate these numbers to their neighbors at every round, and moreover, they know approximately (within a constant factor) how many agents there are in total. Subject to these assumptions, we derive a considerable speedup over the simple static control law of Section III. The main result of this section is Theorem 6, which demonstrates the correctness of the algorithm and proves that it takes a network of n agents on the order of n discrete-time updates to come close to the optimal positions. This is an order of magnitude improvement over the control law of Section III.

II. NONUNIFORM COVERAGE

We introduce the nonuniform coverage problem in this section; our exposition closely follows the expositions of [11], [5]. We consider n mobile agents initially situated at arbitrary positions $x_1(0), x_2(0), \ldots, x_n(0)$ which, for simplicity, we will henceforth assume to be located in the interval $[0,1]$. There is a strictly positive, piecewise-continuous function $\rho : [0,1] \to (0,\infty)$, which measures the density of information or resource at each point. The goal is to bring the agents from their initial configuration to a static configuration that allows them to optimally sense in the density field ρ. Intuitively, we would like more agents to be positioned in areas where ρ is high, and fewer agents positioned in areas where ρ is low.

More formally, for $a, b \in [0,1]$ we define the metric

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

It is easy to see this defines a valid metric between points in $[0,1]$. Relative to the ordinary distance $|a-b|$, this metric expands regions where ρ is large and shrinks regions where ρ is small. We will find it convenient later to refer to the quantities $\rho_{\max} = \sup_{z\in[0,1]}\rho(z)$ and the similarly-defined $\rho_{\min}$.

Following [11], we define the coverage of a set of points $x_1, \ldots, x_n$ relative to the density field ρ as

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

Given the positions $x_1, \ldots, x_n$ of the agents and the density field ρ, computing Φ requires computing the distance d_ρ from any point in $[0,1]$ to the closest x_i. The coverage metric Φ is then the largest of these distances. A smaller Φ implies the vehicles achieve better coverage of the domain $[0,1]$. We use Φ^* to denote the best (smallest) possible coverage

$$\Phi^* = \inf_{(x_1,\ldots,x_n)\in[0,1]^n} \Phi(x_1, \ldots, x_n, \rho).$$

In this paper, we are concerned with designing control laws which drive agents towards positions with coverage Φ^*. As pointed out in [11], the problem of optimal positioning with a nonuniform distance is closely related to information gathering and sensor array optimization problems. A typical problem is to minimize shortest response time from a collection of vehicles to any point in a terrain of varying roughness. In that case, the non-Euclidean distance d_ρ appears because rougher bits of terrain take longer to traverse. Another such problem is the detection of acoustic signals; the objective is to place sources so they can detect a source anywhere in an inhomogeneous medium. In that case, the non-Euclidean distance d_ρ appears as a result of the spatially varying refractive index of the environment.

III. A STATIC COVERAGE CONTROL LAW

We now describe and analyze a simple distributed control law that drives the vehicles towards optimal coverage. First, we need to define the notion of a ρ-weighted median between points.

Definition The α-median $m_\rho^\alpha(a,b)$ is defined as the point $c \in (a,b)$ which satisfies

$$\int_a^c \rho(z)dz = \alpha \int_c^b \rho(z)dz.$$

Due to the strict positivity of ρ, it is easy to see that a unique such point exists for any $\alpha \geq 0$.

We can now state the coverage control law. We will assume for convenience that agents are labeled $1, \ldots, n$ from left to right. This makes it easier to state what follows; however, the actual implementation of the algorithm does not require the use of these labels.

A static coverage control law: the agents iterate as

$$\begin{aligned} x_1(t+1) &= m_\rho^{1/2}(0, x_2(t)) \\ x_i(t+1) &= m_\rho^1(x_{i-1}(t), x_{i+1}(t)), \quad i = 2, \ldots, n-1 \\ x_n(t+1) &= m_\rho^2(x_{n-1}(t), 1) \end{aligned}$$

We first briefly outline how this scheme may be implemented without knowledge of the labels $1, \ldots, n$ by the nodes. A node i will initially check whether it has a left neighbor and a right neighbor, or whether it is a "border agent" with a single neighbor. Suppose it has two neighbors. Then, it will measure the distance d^{L} to its left neighbor and d^{R} to its right neighbor, and denoting its position (which it does not know) by x_i, will measure ρ in the interval $[x_i - d^{\mathrm{L}}, x_i + d^{\mathrm{R}}]$. This gives it enough information to compute the 1-median of the positions of its neighbors, and it moves to this location. "Border agents" can similarly implement this control law without knowledge of their labels.

Next, we remark that this scheme may be interpreted as a distributed implementation of the cartogram approach introduced in [11] specialized to the line. Intuitively, each of the middle nodes $2, \ldots, n-1$ seeks to position itself "in the middle" of its neighboring agents while stretching areas with high ρ and shrinking areas with low ρ; this is precisely the distributed computation of the cartograms used in [11].

Our goal in this section is to prove that this iteration solves the coverage control problem and to provide quantitative bounds on its performance. The main result of this section is the following theorem.

Theorem 1: Each $x_i(t)$ has a limit, and the limiting set of positions have coverage Φ^*. Moreover, after $O(n^2 \log(\frac{n}{\epsilon}\frac{\rho_{\max}}{\rho_{\min}}))$ rounds, each agent is within ϵ of its final limit.

We next turn to the proof of this theorem. We first write down the optimality conditions for achieving Φ^*.

Lemma 2. The equations

$$2d_\rho(0,x_1) = d_\rho(x_1,x_2) = \cdots d_\rho(x_{n-1},x_n) = 2d_\rho(x_n,1) \quad (1)$$

have a unique solution which achieves coverage Φ^*. Moreover,

$$\Phi^* = \frac{1}{2n} d_\rho(0,1).$$

We next introduce a change of variables which makes our static control law easier to analyze. We define

$$F(x) = \int_0^x \rho(z)dz,$$

and note that $F(1) = d_\rho(0,1)$. Moreover, for any two points $a < b$ in $[0,1]$,

$$d_\rho(a,b) = F(b) - F(a),$$

and

$$F(m_\rho^\alpha(a,b)) = \frac{F(a) + \alpha F(b)}{1+\alpha}. \quad (2)$$

The next lemma restates our coverage control law in a particularly convenient form.

Lemma 4: Assuming there are at least two agents, let us define

$$\begin{aligned} d_0(t) &= 2F(x_1(t)) \\ d_i(t) &= F(x_{i+1}(t)) - F(x_i(t)), \quad i = 2, \ldots, n-1 \\ d_n(t) &= 2(F(1) - F(x_n(t))) \end{aligned}$$

and let $d(t)$ be the vector in $\mathbb{R}^{n+1}$ which stacks the variables d_i. Finally, for $n \geq 5$ we define

$$U_n = \begin{pmatrix} -4 & 4 & & & & & \\ 2 & -5 & 3 & & & & \\ & 3 & -6 & 3 & & & \\ & & \ddots & \ddots & \ddots & & \\ & & & 3 & -6 & 3 & \\ & & & & 3 & -5 & 2 \\ & & & & & 4 & -4 \end{pmatrix}$$

Then $d(t)$ follows the dynamics

$$d(t+1) = (I + \frac{1}{6}U_{n+1})d(t).$$

Lemma 5: Let $k \geq 5$ and let $P_k = I + U_k/6$. Then the spectrum of P_k is real. Labeling it from smallest to largest as $\lambda_k(P) \leq \cdots \leq \lambda_2(P) \leq \lambda_1(P) = 1$, we have

$$\max(|\lambda_k(P)|, |\lambda_2(P)|) \leq 1 - \frac{1}{3k^2}.$$

Proof: Consider an undirected line graph on k nodes with self loops at each node, as in Figure 1; we will assign weights to the edges as shown in that figure. Moreover, we will define w_i be the sum of all the weights incident on node i, i.e. $w_i = \sum_{j=1}^n w_{ij}$. Clearly, $(w_1, w_2, \ldots, w_n, w_{n+1}) = (3, 6, \ldots, 6, 3)$. With these definitions in place, we observe that for all i,j, $P_{ij} = w_{ij}/w_i$.

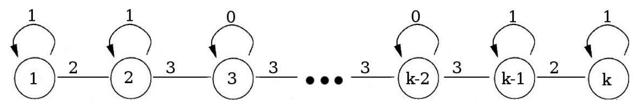

Fig. 1. The weighted graph capturing the transition matrix P_k for $k \geq 5$: $P_{ij} = w_{ij}/w_i$ where w_{ij} is the weight of the edge (i,j) and w_i is the sum of all the weights incident on node i.

Define the inner product $\langle x, y \rangle = \sum_{i=1}^k w_i x_i y_i$. Then,

$$\langle x, Py \rangle = \sum_{i,j=1}^n w_{ij} x_i y_j = \sum_{i,j=1}^n w_{ji} x_i y_j = \langle Px, y \rangle,$$

so that P is self-adjoint in this inner product, and in particular its spectrum must be real. Observing that the largest eigenvalue of P is 1 with the eigenvector of all ones, some straightforward manipulations with the Courant-Fisher theorem give

$$\begin{aligned} \lambda_2(P) &= \max_{\langle x, \mathbf{1} \rangle = 0, \langle x, x \rangle = 1} \langle x, Px \rangle \\ &= 1 - \min_{\sum_i w_i x_i = 0, \sum_i w_i x_i^2 = 1} \sum_{i=1}^{k-1} w_{i,i+1}(x_{i+1} - x_i)^2 \end{aligned}$$

We now lower bound the last term on the right hand side using a variation of the argument from [10]. The minimum is achieved at some vector x (since we are minimizing a continuous function over a compact set); use the notation y for the minimizer. Consider the vector with i'th entry is $w_i y_i$. Let m denote the index of its smallest entry and M the index of its largest entry; without loss of generality, we may assume $m < M$. Observe that the constraint $\sum_i w_i y_i = 0$ implies that $y_{\mathrm{m}} < 0$ while the constraint $\sum_i w_i y_i^2 = 1$ implies that the average value of $w_i y_i^2$ is $1/k$, which means $w_M y_M^2 \geq 1/k$ or $y_M \geq 1/\sqrt{w_M k} \geq 1/\sqrt{6k}$. Thus $\frac{1}{\sqrt{6k}} \leq y_M - y_m = \sum_{i=m}^{M-1} y_{i+1} - y_i$. Applying Cauchy-Schwarz, $\frac{1}{6k} \leq k \sum_{i=m}^{M-1} (y_{i+1} - y_i)^2$, and therefore, $\sum_{i=1}^{k-1} w_{i,i+1}(y_{i+1} - y_i)^2 \geq \frac{1}{2k^2}$. Putting it all together, this implies the desired bound on λ_2.

A similar argument (which is also a variation of an argument from [10]) proves the bound for λ_k. **q.e.d.**

Theorem 1 now follows as a consequence of the previous lemma.

Remark: Observe that the coverage control law we have presented in this section is naturally robust to addition and deletion of agents as well as changes in ρ. Indeed, as long as the density and the number of agents stop changing at some point, the algorithm is guaranteed to converge to the optimal configuration. An open problem is to prove performance bounds for this algorithm in the scenario when the number of agents and the density are continually changing.

IV. A DYNAMIC COVERAGE CONTROL LAW

In this section, we propose another control law for the nonuniform coverage problem on the line. We draw heavily on the paper [7], which described a fast "non-reversible Markov sampler" for sampling a uniform random number from $\{1, \ldots, n\}$. We show that it is possible to use this sampler as the basis for a coverage control law which works an order of magnitude faster than the static control law we have described in the previous section.

Here, we make stronger stronger assumptions on the capabilities and knowledge of each agent. We now assume that agents can store numbers in memory, transmit numbers to their neighbors, and can detect when their neighbors move to a new location. However, every node will only be storing and transmitting/sensing two numbers per every step of the control law, so that the additional effort expended is not excessively onerous. Finally, we will assume that each agent has an estimate U of the total number of agents, and that this estimate is accurate within a constant factor c:

$$\frac{n}{c} \leq U \leq cn. \tag{3}$$

We refer to the control law of this section as the "dynamic coverage control law," since in contrast with the control law of the previous section, the feedback law has dynamics of its own. This control law follows two steps: an initial measurement phase, and the subsequent communication/measurement/movement stages.

Dynamic control law: Nodes keep track of the variables $z_i(t), z_i'(t)$, initialized in the first step as

$$\begin{aligned} z_1(t) &= \frac{1}{2}\int_0^{m_\rho^1(x_1(0),x_2(0))} \rho(z)dz \\ z_i(t) &= \frac{1}{2}\int_{m_\rho^1(x_{i-1}(0),x_i(0))}^{m_\rho^1(x_i(0),x_{i+1}(0))} \rho(z)dz \quad i=2,\ldots,n-1 \\ z_n(t) &= \frac{1}{2}\int_{m_\rho^1(x_{n-1}(0),x_n(0))}^{1} \rho(z)dz \end{aligned}$$

and $z_i'(0) = z_i(0)$ for each i. At each step, nodes transmit their variables $z_i(t), z_i'(t)$ to their neighbors, and then set their values $z_i(t+1), z_{i'}(t+1)$ to be linear combinations of their previous values and the values they have just received. The linear combination taken by each agent is derived from a rule based on Figure 2. Note that this figure contains nodes labeled $1, \ldots, n$ and $1', \ldots, n'$. Node i sets $z_i(t+1)$ to be a linear combination of those values $z_k(t)$ which have edges going from k to i; the coefficient it puts in front of $z_k(t)$ is the label on the edge. The value of $z_{i'}(t+1)$ is determined in the same way. For example, agent 1 updates as

$$\begin{aligned} z_1(t+1) &= (1-\frac{1}{U})z_1'(t) + \frac{1}{2U}z_2'(t) \\ z_1'(t+1) &= \frac{1}{2}(1-\frac{1}{U})z_2'(t) + \frac{1}{U}z_1'(t) \end{aligned}$$

.

After updating their variables $z_i, z_{i'}$, the agents move as follows. Agent 1 moves to the position c that satisfies

$$\int_0^c \rho(z)dz = z_1(t) + z_{1'}(t).$$

Each other agent $i > 0$ waits for the agent to the left of it to move to the position $x_{i-1}(t+1)$, and then moves to the position c that satisfies

$$\int_{x_{i-1}(t+1)}^c \rho(z)dz = z_i(t) + z_{i'}(t).$$

While the above update rule is somewhat involved, it has a simple interpretation. Consider the Markov chain of Figure 2. If $z_i(t)$ is the probability of being at node i at time t, and $z_i'(t)$ is the probability of being at node i' at time t, then the variables $z_i(t), z_{i'}(t)$ satisfy the above recursion.

Indeed, this recursion is an adaptation of the "non-reversible Markov chain sampler" from [7]. It was observed in that paper that while an ordinary "diffusive" Markov chain on the line graph which, say, moves to the right and left each with probability $1/2$ takes on the order of n^2 steps to come close to the uniform distribution, the "guided" Markov chain of Figure 2 takes on the order of n steps to come close to the uniform distribution. The dynamic coverage control law of this section is an attempt to harness this insight for the coverage problem.

Our goal in this section is to prove that this iteration solves the coverage control problem and to demonstrate that it is an order of magnitude faster than the static control law of the previous section. The main result of this section is the following theorem.

Theorem 6: Each $x_i(t)$ has a limit, and the limiting set of positions have coverage Φ^*. Moreover, after $O(n\log(\frac{\rho_{\max}n}{\rho_{\min}\epsilon}))$ rounds of updates, each agent is within ϵ of its final limit.

The proof closely mimics the proof of the related results from [7]. Let $z(t)$ denote the row vector $\begin{pmatrix} z_1(t) & \ldots & z_n(t) & z_{1'}(t) & \ldots z_{n'}(t) \end{pmatrix}$. Let K to denote the matrix that maps $z(t)$ to $z(t+1)$ through right-multiplication:

$$z(t+1) = z(t)K.$$

Observe that K is a nonnegative, irreducible stochastic matrix. Standard results in Markov chain theory imply that the above iteration converges to a scaled multiple of the stationary probability vector of the chain. Thus we can immediately conclude that each $z_i(t)$ has a limit, and consequently, the positions of the agents under the dynamic coverage control law have limits as well. Moreover, observe that the vector of all ones is a right eigenvector of K, which implies that $\sum_{i=1}^n z_i(t) + z_{i'}(t)$ does not change after the execution of each update. Since $\sum_{i=1}^n z_i(0) + z_{i'}(0) = F(1)$, we can immediately conclude that no agent ever moves outside of $[0,1]$.

Lemma 7: The stationary probability of the Markov chain in Figure 2 is

$$\pi_1 = \pi_{1'} = \pi_n = \pi_{n'} = \frac{1}{4(n-1)}$$

and for all other nodes,

$$\pi_i = \pi_{i'} = \frac{1}{2(n-1)}.$$

This lemma implies that

$$\lim_t z_1(t) = \lim_t z_{1'}(t) = \lim_t z_n(t) = \lim_t z_{n'}(t) = \frac{F(1)}{4(n-1)}, \tag{4}$$

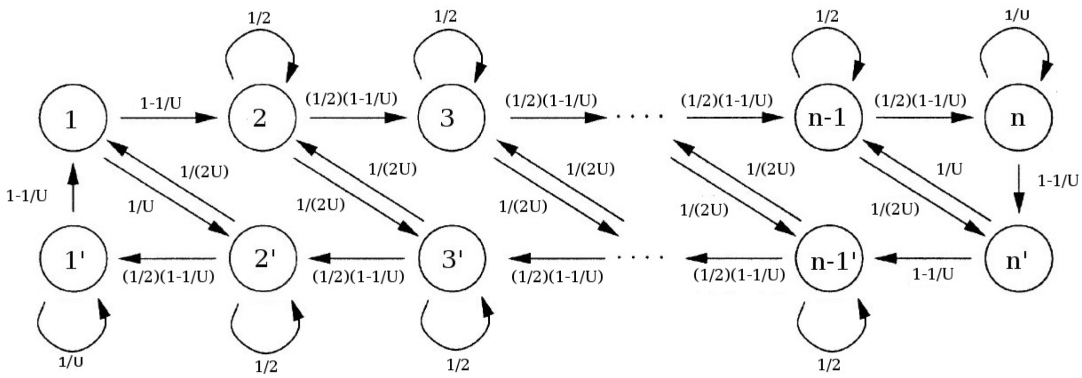

Fig. 2. The Markov chain representing the update equations of the dynamic coverage control law.

and

$$\lim_t z_i(t) = \lim_t z_{i'}(t) = \frac{F(1)}{2(n-1)}, \tag{5}$$

which implies by Lemma 2 that the limiting set of positions do have optimal coverage Φ^*. It remains to bound the time until the agents approach these positions.

We find it convenient to use the largest l^1 distance between the rows of K^t and their ultimate limit as a measure of convergence. In particular, let $(K^t)_i$ mean the i'th row of K^t, and let

$$v(t) = \max_i ||(K^t)_i - \pi||_1.$$

We use $t(n, \epsilon)$ to denote the time until $v(t)$ permanently sinks below ϵ.

Lemma 8 $t(n, \epsilon) = O(n \log 1/\epsilon)$.

The proof follows closely the proof of the main result of [7], and we omit it here.

Lemma 9 After $O(\log n(nF(1)/\epsilon))$ steps, we will always have

$$|z_i(t) - \pi_i F(1)| \leq \frac{\epsilon}{2n}.$$

Proof: We show that after $O(n \log n/\epsilon')$ iterations, the inequality

$$|z_i(t) - \pi_i F(1)| \leq \frac{\epsilon'}{2n} F(1) \tag{6}$$

is satisfied. Taking $\epsilon' = \epsilon/F(1)$ in this statement yields the lemma.

To prove Eq. (6), observe first that scaling the density ρ multiplies both sides by the same number so that we may assume without loss of generality that $F(1) = 1$. In this case, $z_i(t)$ is the probability that the random walk that starts at node k with probability $z_k(0)$ is at node i at time t. This is a convex combination of the entries of the i'th column of K^t. By the previous Lemma, the i'th entry of each row is not more than $\epsilon'/(2n)$ from π_i after $O(n \log n/\epsilon')$ steps, and the convex combination of these entries must have the same property. **q.e.d.**

Theorem 6 now follows by a straightforward application of Lemma 9 to the density $\rho/\rho_{\min}$.

Remark: We remark that it is possible to describe a modification of this algorithm that is robust to unpredictable loss and deletion of agents, and we briefly sketch this modification here. When a node is added, it sets its $z_i(t), z_{i'}(t)$ to zero, and the algorithm proceeds as before. If an agent k dies, then its neighboring agents can infer its values $z_k(t), z_{k'}(t)$ from either the previous transmission of agent k or agent k's position when it dies; they can then increase their own values $z_{k-1}, z_{(k-1)'}, z_{k+1}, z_{(k+1)'}$ in such a way as to keep $\sum_i z_i(t) + z_{i'}(t)$ the same. As long as the number of agents stabilizes eventually, the control law will converge to the correct answer. An open problem is to prove a guarantee on performance if the number of agents is continually changing.

V. SIMULATIONS

We report here on several simulations of our coverage control laws. We are able to observe that quite often the performance is considerably better than the theoretical upper bounds derived in this paper, and that the dynamic control law of Section IV gives a considerable practical speedup over the static control law of Section III.

Figure 3 shows the results from a simulation with random initial conditions. In this case, $x_i(0)$ is the i'th largest value of n random variables, all uniform on $[0, 1]$. The density ρ was uniform on $[0, 1]$. Moreover, we assumed that each agent knows the total number of agents in the system, i.e. $U = n$. The top figure shows some snapshots from the progress of both algorithms, while the bottom figure shows the time until the stopping condition $\sum_{i=1}^n (x_i(t) - \lim_t x_i(t))^2 \leq 10^{-4}$ holds for the first time. The randomness of the initial conditions seems to result in a reasonably quick convergence. We see that for the range of parameters in the graph the static control law has a convergence time which grows slower than the quadratic growth proved in Theorem 1, while the dynamic control law has convergence time which appears to grow somewhat slower than the linear upper bound of Theorem 6.

On the other hand, it is not hard to find examples for which the bounds on growth rates from Theorems 1 and 6 do occur. In Figure 4, we see the performance when every agent starts with $x_i(0) = 1$; every other aspect is the same as the simulation in Figure 3. We see that in this case the convergence times do seem to grow quadratically for the static control law, and linearly for the dynamic control law.

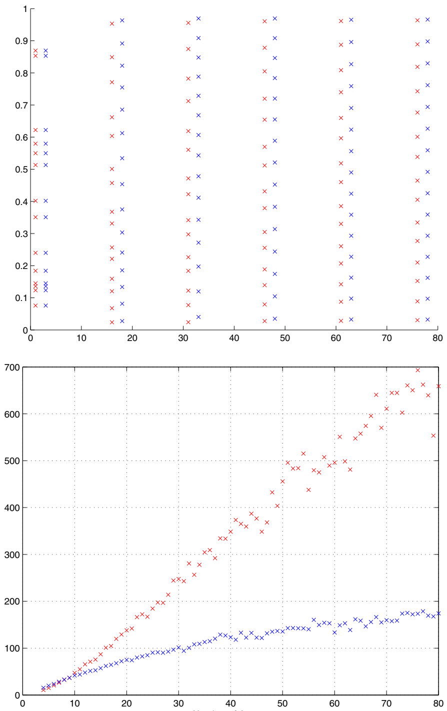

Fig. 3. Top figure displays the progress of both the static (red) and dynamic (blue) coverage control laws. In both cases, the initial condition consists of 15 nodes with the same randomly generated positions. The update number is represented on the x-axis, the distribution of agents in $[0,1]$ on the y-axis. The bottom shows the number of agents on the x-axis, and the number of update rounds until agents are close to their final values on the y-axis. Initial conditions are random as well, and each point on the bottom graph is an average of 80 runs.

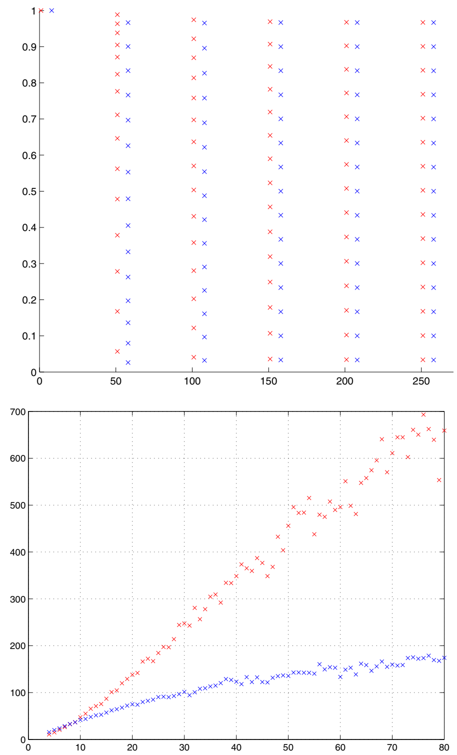

Fig. 4. Top figure displays the progress of both the static (red) and dynamic (blue) coverage control laws. The initial conditions consist of 15 nodes all of which start with $x_i(0) = 1$. The update number is represented on the x-axis, the distribution of agents in $[0,1]$ on the y-axis. The bottom figure shows the number of agents on the x-axis, and the number of update rounds until agents are close to their final values on the y-axis; just as in the top figure, every node starts out with $x_i(0) = 1$.

VI. CONCLUSIONS

We have investigated distributed control laws for mobile, autonomous agents to position themselves on the line for optimal coverage in a nonuniform field. Our main results are stated in Theorems 1 and 6. Theorem 1 gives a quantitative upper bound on the convergence time of a simple control law for coverage. Theorem 6 discusses a more complicated control law which, while making stronger assumptions on the capabilities of each agent, manages to accomplish the coverage task an order of magnitude faster in the worst case.

Our work suggests a number of open questions. It is of interest to understand whether the increased capabilities of the agents in Section IV are really necessary to achieve better performance. In addition, it would be interesting to explore whether the results described here extend to two and higher dimensions, and in particular, whether a dynamic control law such as the one in Theorem 6 might be useful for speeding up performance in more general settings.

REFERENCES

[1] A. Ahmadzadeh, A. Jadbabaie, V. Kumar, G.J. Pappas, Dynamic coverage using receeding horizon control, *Proceedings of the European Control Conference,* 2007.

[2] A. Ahmadzadeh, J. Keller, G. J. Pappas, A. Jadbabaie, V. Kumar, An optimization-based approach to time-critical cooperative surveillance and coverage with unmanned aerial vehicles, *Proceedings of the International Symposium on Experimental Robotics,* 2006.

[3] C. H. Caicedo-Nunez, M. Zefran, Performing coverage on nonconvex domains, *Proceedings 17th IEEE International Conferene on Control Applications,* 2008.

[4] C. H. Caicedo-Nunez, M. Zefran, A coverage algorithm for a class of non-convex regions, *Proceedings of the 47th IEEE Conference on Decision and Control,* 2008.

[5] J. Cortes, F. Bullo, Coordination and geometric optimization via distributed dynamical systems, *SIAM Journal on Control and Optimization,* vol 44, pp. 1543–1574, 2005.

[6] J. Cortes, S. Martinez, T. Karatas, F. Bullo, Coverage control for mobile sensing networks, *IEEE Transactions on Robotics and Automation,* vol. 20, pp. 243-255, 2004.

[7] P. Diaconis, S. Holmes, R.M. Neal, Analysis of a nonreversible Markov chain sampler, *Annals of Applied Probability,* vol. 10, no. 3, pp. 726–752, 2000.

[8] K. Jung, D. Shah, J. Shin, Distributed averaging via lifted Markov chains, *IEEE Transactions on Information Theory,* vol.56, no. 1, 2010.

[9] W. Li, C. G. Cassandras, Distributed cooperative coverage control of sensor networks, *Proceedings of the 44th IEEE Conference on Decision and Control*, pp. 2542-2547, 2005.

[10] H.J. Landau, A.M. Odlyzko, Bounds for eigenvalues of certain stochastic matrices, *Linear Algebra and Its Applications*, vol. 38, pp. 5–15, 1981.

[11] F. Lekien, N.E. Leonard, Nonuniform coverage and cartograms, *SIAM Journal on Control and Optimization*, vol. 48, no. 1, pp. 351–372, 2009.

[12] S. Poduri, G. S. Sukhatme, Constrained coverage for mobile sensor networks, *Proceedings of the IEEE International Conference on Robotics and Automation,* pp. 165-172., 2004.