Coordinated deployment of mobile sensing networks with limited-range interactions

Jorge Cortés   Sonia Martínez   Francesco Bullo
Coordinated Science Laboratory
University of Illinois at Urbana-Champaign
Urbana, Illinois 61801, USA
{jcortes,smartine,bullo}@uiuc.edu

***Abstract*— This paper presents coordination algorithms for groups of mobile agents performing deployment and coverage tasks. As an important modeling constraint, we assume that each mobile agent has a limited sensing or communication radius. We focus on (1) a comprehensive smoothness analysis of a class of locational optimization functions (including a generalized statement of the Conservation-of-Mass Law), and (2) a discrete-time convergence result based on a recently-developed generalized statement of LaSalle Invariance Principle. These coordination algorithms have convergence guarantees and are spatially distributed with respect to appropriate proximity graphs. Numerical simulations illustrate the results.**

I. Introduction

The current technological development of relatively inexpensive communication, computation, and sensing devices has lead to an intense research activity devoted to the distributed control and coordination of networked systems. The potential advantages of networked robotic systems are their versatility and robustness in the realization of multiple tasks such as manipulation in hazardous environments, pollution detection, estimation and map-building of unknown environments. A fundamental problem in the distributed coordination of mobile robots is that of providing stable and decentralized control laws that are scalable with the number of network agents. Indeed, there have been various efforts to provide rigorous procedures with convergence guarantees using a combination of potential energy shaping methods, gyroscopic forces, and graph theory (see [1], [2], [3], [4], [5], [6], [7] and references therein). In our previous work [8], we studied distributed algorithms for deployment and optimal coverage using tools from computational geometry, nonsmooth analysis and geometric optimization.

In devising coordination algorithms it is important to progressively account for the various restrictions that real-world systems impose. Building on [8], this paper develops spatially-distributed algorithms for coverage control amenable to implementation on (more) realistic network models; we do this by considering the following new aspects. First, we enforce the communication or sensing capacity of an agent to be restricted to a bounded region, typically much smaller than the environment of interest. We refer to these information exchanges between agents as "limited-range interactions." Second, we design discrete-time gradient ascent control laws, motivated by the argument that discrete-time feedback algorithms are indeed the ones truly amenable to implementation in a group of agents exchanging information over a communication network. We deal with these problems via a combination of tools from graph theory, locational optimization, and systems theory.

The contributions of the paper are the following: First, we formalize the concept of spatially-distributed algorithms via the notion of proximity graphs. Second, we analyze the smoothness properties of an important class of objective functions, called multi-center functions, common in locational optimization, quantization theory, and geometric optimization. Our analysis supersedes the results in [9], [10], [11]. As an important outcome, we determine the extent in which certain multi-center functions are spatially distributed and with respect to which proximity graphs. Finally, we design novel spatially-distributed discrete-time control laws for groups of robots with the objective of steering the location of a group of robots to local maxima of the multi-center function. We formally analyze their performance and illustrate their behavior in simulations. In our analysis, we make use of useful extensions of the Conservation-of-Mass Law and of the discrete-time LaSalle Invariance Principle. These extensions are, to the best of our knowledge, not present in classical texts on the subject.

One fundamental scientific problem in the study of coordination algorithms is scalability with respect to communication complexity. It is important to design algorithms with communication requirements that scale nicely with the number of network agents. However, it is impossible to quantify the communication complexity of any algorithm without a detailed communication model. Adopting a computational geometric approach, this paper classifies the complexity of coordination algorithms via the proximity graphs with respect to which the algorithms are spatially distributed. The underlying assumption is that low complexity proximity graphs (e.g., graphs with a low number of edges) will require limited communication in a realistic implementation.

Throughout the paper we consider only extremely simple models for the dynamics of each individual agent (the state of each agent is a point in $\mathbb{R}^2$ and its dynamical model is an integrator). This feature is a natural consequence of our focus on network-wide coordination aspects. We shall also interchangeably refer to agent as location or point.

The organization of the paper is as follows. In Section II we review various preliminary mathematical concepts and we develop useful extensions of classical analysis results.

In Section III we study the smoothness of the multi-center function and show in what sense its gradient is spatially distributed. In Section IV we design spatially-distributed coverage algorithms in discrete time. The numerical outcomes of the algorithms' implementation are reported in Section V. Finally, we discuss possible avenues of future research in Section VI. We refer the interested reader to the journal version of this paper [12] for a detailed discussion of the computational geometric and algorithmic aspects of this work, and the proofs of all statements in this manuscript.

II. Preliminary developments

In this section we present a variety of preliminary concepts. Proximity graphs from computational geometry will allow us to introduce the notion of spatially-distributed algorithms. We refer the reader to [13, Chapter 1] for an exposition of standard graph-theoretical notions.

A. Voronoi partitions and proximity graphs

We start by reviewing the notion of Voronoi partition; we refer the reader to [14], [10] for comprehensive treatments. A *covering* of $\mathbb{R}^2$ is a collection of subsets of $\mathbb{R}^2$ whose union is $\mathbb{R}^2$; a *partition* of $\mathbb{R}^2$ is a covering whose subsets have disjoint interiors. Let $\mathcal{P}$ be a set of n distinct points $\{p_1, \dots, p_n\}$ in $\mathbb{R}^2$. The *Voronoi partition of* $\mathbb{R}^2$ generated by $\mathcal{P}$ is the collection of sets $\{V_i(\mathcal{P})\}_{i\in\{1,\dots,n\}}$ defined by $V_i(\mathcal{P}) = \{q \in \mathbb{R}^2 \mid \|q - p_i\| \leq \|q - p_j\|, \text{ for all } p_j \in \mathcal{P}\}$. Here, $\|\cdot\|$ denotes the standard Euclidean norm. It is customary to refer to $V_i(\mathcal{P})$ as V_i.

Next, we present some relevant concepts on proximity graph functions [15]. Let $\mathbb{F}(\mathbb{R}^2)$ be the collection of finite point sets in $\mathbb{R}^2$; an element of $\mathbb{F}(\mathbb{R}^2)$ is denoted by $\mathcal{P} = \{p_1, \dots, p_n\} \subset \mathbb{R}^2$, where $p_1, \dots, p_n$ are distinct. Let $\mathbb{G}(\mathbb{R}^2)$ be the set of undirected graphs whose vertex set is an element of $\mathbb{F}(\mathbb{R}^2)$. A *proximity graph function* $\mathcal{G} : \mathbb{F}(\mathbb{R}^2) \to \mathbb{G}(\mathbb{R}^2)$ associates to a point set $\mathcal{P}$ an undirected graph with vertex set $\mathcal{P}$ and edge set $\mathcal{E}_{\mathcal{G}}(\mathcal{P})$, where $\mathcal{E}_{\mathcal{G}} : \mathbb{F}(\mathbb{R}^2) \to \mathbb{F}(\mathbb{R}^2 \times \mathbb{R}^2)$ is such that $\mathcal{E}_{\mathcal{G}}(\mathcal{P}) \subseteq \mathcal{P} \times \mathcal{P} \setminus \operatorname{diag}(\mathcal{P} \times \mathcal{P})$ for any $\mathcal{P}$. Here, $\operatorname{diag}(\mathcal{P}\times\mathcal{P}) = \{(p, p) \in \mathcal{P} \times \mathcal{P} \mid p \in \mathcal{P}\}$. In other words, the edge set of a proximity graph depends on the location of its vertices. For $p \in \mathbb{R}^2$ and $r \in \mathbb{R}_+ = [0, +\infty)$, let $B_r(p) = \{q \in \mathbb{R}^2 \mid \|q - p\| \leq r\}$. Consider the following proximity graphs (see Figure 1):

(i) the *Delaunay graph* $\mathcal{G}_{\mathrm{D}}$ with $(p_i, p_j) \in \mathcal{E}_{\mathrm{D}}(\mathcal{P})$ if and only if $V_i(\mathcal{P}) \cap V_j(\mathcal{P}) \neq \emptyset$;

(ii) the *r-disk graph* $\mathcal{G}_{\text{disk}}(r)$ with $(p_i, p_j) \in \mathcal{E}_{\text{disk}}(\mathcal{P}, r)$ if and only if $\|p_i - p_j\| \leq r$;

(iii) the *r-Delaunay graph* $\mathcal{G}_{\text{disk}\cap\mathrm{D}}(r)$ is the intersection of $\mathcal{G}_{\text{disk}}(r)$ and $\mathcal{G}_{\mathrm{D}}$;

(iv) the *r-limited Delaunay graph* $\mathcal{G}_{\mathrm{LD}}(r)$ with $(p_i, p_j) \in \mathcal{E}_{\mathrm{LD}}(\mathcal{P}, r)$ if and only if $\Delta_{ij}(\mathcal{P}, r) \triangleq \big(V_i(\mathcal{P}) \cap B_{\frac{r}{2}}(p_i)\big) \cap \big(V_j(\mathcal{P}) \cap B_{\frac{r}{2}}(p_j)\big) \neq \emptyset$.

Remark 2.1: Let $\mathcal{P}$ be a set of n distinct points $\{p_1, \dots, p_n\} \subset \mathbb{R}^2$, and $r \in \mathbb{R}_+$. The collection $\{V_i(\mathcal{P}) \cap B_{\frac{r}{2}}(p_i)\}_{i\in\{1,\dots,n\}}$ is a partition of the set $\cup_i B_{\frac{r}{2}}(p_i) \subset \mathbb{R}^2$. The boundary of $V_i(\mathcal{P}) \cap B_{\frac{r}{2}}(p_i)$ is the union of a finite number of segments and arcs; see Figure 2. Therefore, at fixed $\mathcal{P}$, there exist n numbers $M_i(r) \geq 0$, $i \in \{1, \dots, n\}$, of distinct arcs $\operatorname{arc}_{i,1}(r), \dots, \operatorname{arc}_{i,M_i(r)}(r)$ of radius $\frac{r}{2}$ in $\partial(V_i(\mathcal{P}) \cap B_{\frac{r}{2}}(p_i))$ with the property that

$$\partial\big(V_i(\mathcal{P}) \cap B_{\frac{r}{2}}(p_i)\big) = \big(\cup_{p_j\in\mathcal{N}_{\mathcal{G}_{\mathrm{LD}}(r),p_i}(\mathcal{P})} \Delta_{ij}(r)\big) \cup \big(\cup_{l\in\{1,\dots,M_i(r)\}} \operatorname{arc}_{i,l}(r)\big),$$

where $\mathcal{N}_{\mathcal{G}_{\mathrm{LD}}(r),p_i}(\mathcal{P})$ denotes the set of neighbors in $\mathcal{G}_{\mathrm{LD}}(\mathcal{P}, r)$ of the vertex p_i.

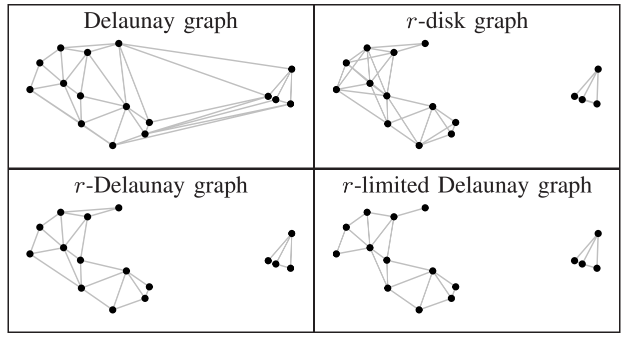

Fig. 1. Delaunay, r-disk, r-Delaunay, and r-limited Delaunay graphs for a random configuration of 16 generators with coordinates contained in the rectangle $[0, 1.9] \times [0, .75]$. The parameter r is taken equal to $.45$.

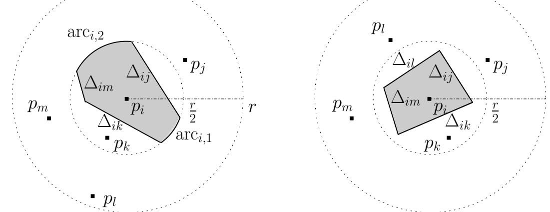

Fig. 2. The shaded regions are examples of sets $V_i(\mathcal{P}) \cap B_{\frac{r}{2}}(p_i)$. In the figure on the right, $\partial V_i(\mathcal{P}) \cap B_{\frac{r}{2}}(p_i)$ does not contain any arc.

B. Spatially-distributed functions and set-valued maps

The notion of proximity graph is defined for sets of distinct points $\mathcal{P} = \{p_1, \dots, p_n\}$. However, we will often consider tuples of elements of $\mathbb{R}^d$ of the form $P = (p_1, \dots, p_n)$, i.e., ordered sets of possibly coincident points. Let $i_{\mathbb{F}} : (\mathbb{R}^d)^n \to \mathbb{F}(\mathbb{R}^d)$ be the natural immersion, i.e., $\mathcal{P} = i_{\mathbb{F}}(P)$ is the point set that contains only the distinct points in P. Note that $i_{\mathbb{F}}$ is invariant under permutations of its arguments and that the cardinality of $i_{\mathbb{F}}(P)$ is in general less than or equal to n. More precisely, if

$$\mathcal{S} = \{P \in (\mathbb{R}^2)^n \mid p_i = p_j \text{ for some } i \neq j \in \{1, \dots, n\}\},$$

then $\#\mathcal{P} < n$ if $P \in \mathcal{S}$ and $\#\mathcal{P} = n$ if $P \notin \mathcal{S}$. The *Voronoi covering* $\mathcal{V}(p_1, \dots, p_n) = \{V_i(p_1, \dots, p_n)\}_{i\in\{1,\dots,n\}}$ *generated by the tuple* $(p_1, \dots, p_n)$ is defined by assigning to each point p_i its corresponding Voronoi cell in the Voronoi partition generated by $\mathcal{P}$. Coincident points in the tuple $(p_1, \dots, p_n)$ have the same Voronoi cell.

We can now extend the notion of proximity graphs to this setting. Given a proximity graph function $\mathcal{G}$ with edge set function $\mathcal{E}_{\mathcal{G}}$, we define (with a slight abuse of notation)

$$\mathcal{G} = \mathcal{G} \circ i_{\mathbb{F}} : (\mathbb{R}^2)^n \to \mathbb{G}(\mathbb{R}^2),$$
$$\mathcal{E}_{\mathcal{G}} = \mathcal{E}_{\mathcal{G}} \circ i_{\mathbb{F}} : (\mathbb{R}^2)^n \to \mathbb{F}(\mathbb{R}^2 \times \mathbb{R}^2).$$

Moreover, the map $\mathcal{N}_{\mathcal{G}} : (\mathbb{R}^2)^n \to (\mathbb{F}(\mathbb{R}^2))^n$ is the function whose jth component is $\mathcal{N}_{\mathcal{G},j}(p_1,\dots,p_n) = \mathcal{N}_{\mathcal{G},p_j}(i_{\mathbb{F}}(p_1,\dots,p_n))$. Coincident points in the tuple $(p_1,\dots,p_n)$ have the same set of neighbors.

Given a set Y and a proximity graph function $\mathcal{G}$, a map $T : (\mathbb{R}^2)^n \to Y^n$ is *spatially distributed over* $\mathcal{G}$ if there exist a map $\tilde{T} : \mathbb{R}^2 \times \mathbb{F}(\mathbb{R}^2) \to Y$, with the property that, for all $(p_1,\dots,p_n) \in (\mathbb{R}^d)^n$ and for all $j \in \{1,\dots,n\}$,

$$T_j(p_1,\dots,p_n) = \tilde{T}(p_j, \mathcal{N}_{\mathcal{G},j}(p_1,\dots,p_n)),$$

where T_j denotes the jth-component of T. In other words, the jth component of a spatially distributed map at $(p_1,\dots,p_n)$ can be computed with only the knowledge of the vertex p_j and the neighboring vertices in the undirected graph $\mathcal{G}(\{p_1,\dots,p_n\})$. We are now in a position to state an important property of the r-limited Delaunay graph.

Lemma 2.2: Let $r \in \mathbb{R}_+$. The set of neighbors map $\mathcal{N}_{\mathcal{G}_{\text{LD}(r)}}$ of $\mathcal{G}_{\text{LD}}(r)$ is spatially distributed over $\mathcal{G}_{\text{disk}}(r)$.

Loosely speaking, this result means that, to compute $\mathcal{G}_{\text{LD}}$, each agent needs to know only the location of all the agents in a disk of radius r. This property is to be contrasted with the centralized computation required to determine $\mathcal{G}_{\text{disk}\cap\text{D}}$.

C. Piecewise smooth sets and Conservation-of-Mass Law

A set $S \subset \mathbb{R}^2$ is *strictly star-shaped* if there exists a point $p \in S$ such that for all $s \in \partial S$ and all $\lambda \in (0,1]$, one has $\lambda p + (1-\lambda)s \in \text{int}(S)$. A *curve* C in $\mathbb{R}^2$ is the image of a map $\gamma : [a,b] \to \mathbb{R}^2$. The map γ is a *parameterization of* C. A curve $\gamma : [a,b] \to \mathbb{R}^2$ is *simple* if it is not self-intersecting, i.e., if γ is injective on (a,b). A curve is *closed* if $\gamma(a) = \gamma(b)$. A set $\Omega \subset \mathbb{R}^2$ is *piecewise smooth* if its boundary, $\partial\Omega$, is a simple closed curve that admits a continuous and piecewise smooth parameterization $\gamma : \mathbb{S}^1 \to \mathbb{R}^2$. Likewise, a collection of sets $\{\Omega(x) \subset \mathbb{R}^2 \mid x \in (a,b)\}$ is a *piecewise smooth family* if $\Omega(x)$ is piecewise smooth for all $x \in (a,b)$, and there exists $\gamma : \mathbb{S}^1 \times (a,b) \to \mathbb{R}^2$, $(\theta,x) \mapsto \gamma(\theta,x)$, differentiable with respect to its second argument such that for each $x \in (a,b)$, the map $\theta \mapsto \gamma_x(\theta) = \gamma(\theta,x)$ is a continuous and piecewise smooth parameterization of $\partial\Omega(x)$. We refer to γ as a *parameterization for the family* $\{\Omega(x) \subset \mathbb{R}^2 \mid x \in (a,b)\}$. The following result is an extension of the integral form of the Conservation-of-Mass Law in fluid mechanics [16]. Given a curve C parameterized by a piecewise smooth map $\gamma : [a,b] \to C$, recall that $\int_C f = \int_a^b f(\gamma(t))\,|\dot{\gamma}(t)|\,dt$ denotes the line integral over C of $f : C \subset \mathbb{R}^2 \to \mathbb{R}$, and is independent of the selected parameterization.

Proposition 2.3: Let $\{\Omega(x) \subset Q \,|\, x \in (a,b)\}$ be a piecewise smooth family with $\Omega(x)$ strictly star-shaped for all $x \in (a,b)$. Let $\phi : Q \times (a,b) \to \mathbb{R}$ be continuously differentiable with respect to its second argument for all $x \in (a,b)$ and almost all $q \in \Omega(x)$, and such that for each $x \in (a,b)$, the maps $q \mapsto \phi(q,x)$ and $q \mapsto \frac{\partial\phi}{\partial x}(q,x)$ are integrable on $\Omega(x)$. Then, the function $(a,b) \ni x \mapsto \int_{\Omega(x)} \phi(q,x)dq$ is continuously differentiable and

$$\frac{d}{dx}\int_{\Omega(x)} \phi(q,x)dq = \int_{\Omega(x)} \frac{\partial\phi}{\partial x}(q,x)dq + \int_{\partial\Omega(x)} n^t \frac{\partial\gamma}{\partial x}\,\phi(\cdot,x)\,,$$

where $n : \partial\Omega(x) \to \mathbb{R}^2$, $q \mapsto n(q)$, is the unit outward normal to $\partial\Omega(x)$ at $q \in \partial\Omega(x)$, and $\gamma : \mathbb{S}^1 \times (a,b) \to \mathbb{R}^2$ is a parameterization for the family $\{\Omega(x) \subset \mathbb{R}^2 \mid x \in (a,b)\}$.

D. Discrete-time LaSalle Invariance Principle

We review here some notions on discrete-time algorithms [17]. An *algorithm on* $S \subset \mathbb{R}^N$ is a set-valued map $T : S \to 2^S \setminus \{\emptyset\}$. A map from S to S is simply a singleton-valued map. For any $x_0 \in S$, T generates feasible sequences as follows: given $x_n \in S$, the map T yields $T(x_n) \subset S$. From this set, an arbitrary element x_{n+1} may be selected,

$$x_{n+1} \in T(x_n)\,, \quad n \in \mathbb{N} \cup \{0\}\,. \tag{1}$$

An algorithm T is *closed at* $x \in S$ if for all convergent sequences $x_k \to x$, $x'_k \to x'$ with $x'_k \in T(x_k)$, one has $x' \in T(x)$. An algorithm is *closed on* $W \subset S$ if it is closed at x, for all $x \in W$. In particular, every continuous map $T : S \to S$ is closed on S. A set $C \subset S$ is *weakly positively invariant with respect to* T if for any $x_0 \in C$ there exists $x \in T(x_0)$ such that $x \in C$. The function $U : S \to \mathbb{R}$ is *non-increasing along* T *on* W if $U(x') \leq U(x)$ for all $x' \in T(x)$ and all $x \in W$.

Theorem 2.4 (Discrete-time LaSalle Invariance Principle): Let T be a closed algorithm on $W \subset \mathbb{R}^N$ and let U be a continuous non-increasing function along T on W. Let $x_0 \in W$ and assume the sequence $\{x_n \mid n \in \mathbb{N} \cup \{0\}\}$ defined via $x_{n+1} \in T(x_n)$ belongs to W and is bounded. Then there exists $c \in \mathbb{R}$ such that

$$x_n \longrightarrow M \cap U^{-1}(c)\,,$$

where M is the largest weakly positively invariant set contained in $\{x \in \overline{W} \mid \exists y \in T(x) \text{ such that } U(y) = U(x)\}$.

III. Limited-range locational optimization

Let Q be a simple convex polygon in $\mathbb{R}^2$ including its interior. Let $\text{diam}(Q) = \max_{q,p \in Q} \|q-p\|$. For $\delta < \epsilon \in \mathbb{R}_+$, let $D_{[\delta,\epsilon]}(p) = \{q \in \mathbb{R}^2 \mid \delta \leq \|q-p\| \leq \epsilon\}$ denote the annulus in $\mathbb{R}^2$ centered at p of radius δ and ϵ; we also define $D_{[0,\epsilon]}(p) = B_\epsilon(p) = \{q \in \mathbb{R}^2 \mid \|q-p\| \leq \epsilon\}$ and $D_{[\delta,+\infty)}(p) = \{q \in \mathbb{R}^2 \mid \delta \leq \|q-p\|\}$. Let $n_{B_\epsilon(p)}(q)$ denote the unit outward normal to $B_\epsilon(p)$ at $q \in \partial B_\epsilon(p)$. Given $S \subset Q$, let 1_S denote the indicator function, $1_S(q) = 1$ if $q \in S$, and $1_S(q) = 0$ if $q \notin S$. Throughout the paper, given a point set $\mathcal{P} = \{p_1,\dots,p_n\}$, we consider the restriction of the Voronoi partition $\mathcal{V}(\mathcal{P})$ generated by $\mathcal{P}$ to the polygon Q, $\{V_i(\mathcal{P}) \cap Q\}_{i \in \{1,\dots,n\}}$. For ease of exposition, we denote this restriction in the same way as the standard Voronoi partition. Given a tuple $P = (p_1,\dots,p_n) \in Q^n$, recall that $\mathcal{P} = i_{\mathbb{F}}(P)$ denotes the associated point set that only contains the corresponding distinct points.

A *density function* $\phi : Q \to \overline{\mathbb{R}}_+$ is a bounded function on Q. Given a set $S \subset Q$, let $\text{area}_\phi(S) = \int_S \phi(q)dq$. A *performance function* $f : \overline{\mathbb{R}}_+ \to \mathbb{R}$ is a non-increasing and piecewise differentiable function with finite jump discontinuities at $R_1, \dots, R_m \in \mathbb{R}_+$, with $R_1 < \cdots < R_m$. For convenience, we set $R_0 = 0$ and $R_{m+1} = +\infty$, and write

$$f(x) = \sum_{\alpha=1}^{m+1} f_\alpha(x)\, 1_{[R_{\alpha-1}, R_\alpha)}(x)\,,$$

where $f_\alpha : [R_{\alpha-1}, R_\alpha] \to \mathbb{R}$, $\alpha \in \{1, \dots, m+1\}$ are non-increasing and differentiable with $f_\alpha(R_\alpha) > f_{\alpha+1}(R_\alpha)$ for $\alpha \in \{1, \dots, m\}$. Given a density function ϕ and a performance function f, we consider the *multi-center function* $\mathcal{H} : Q^n \to \mathbb{R}$ defined by

$$\mathcal{H}(P) = \int_Q \max_{i \in \{1,\dots,n\}} f(\|q - p_i\|)\phi(q) dq. \tag{2}$$

Note that $\mathcal{H}$ is an aggregate objective function since it depends on all the locations $p_1, \dots, p_n$. It will be of interest to find local maxima for $\mathcal{H}$.

Remark 3.1: In the optimal placement problem of large numbers of spatially-distributed sensors, (1) $\mathcal{H}$ provides the *expected value of the sensing performance* provided by the group of sensors over any point in the environment Q, where (2) ϕ is the *distribution density function* representing a measure of information or probability that some event take place over Q, and (3) f describes the *sensing performance* of the sensors. Because of noise and loss of resolution, the sensing performance at point q taken from the ith sensor at the position p_i degrades with the distance $\|q - p_i\|$. •

Theorem 3.2: Given a density function ϕ and a performance function f, the multi-center function $\mathcal{H}$ is globally Lipschitz on Q^n, and continuously differentiable on $Q^n \setminus \mathcal{S}$, where for each $i \in \{1, \dots, n\}$

$$\begin{aligned}\frac{\partial \mathcal{H}}{\partial p_i}(P) = &\int_{V_i(P)} \frac{\partial}{\partial p_i} f(\|q - p_i\|)\phi(q) dq + \sum_{\alpha=1}^{m} (f_\alpha(R_\alpha) - \\ &f_{\alpha+1}(R_\alpha)) \cdot \Big(\sum_{k=1}^{M_i(2R_\alpha)} \int_{\mathrm{arc}_{i,k}(2R_\alpha)} n_{B_{R_\alpha}(p_i)}(q)\phi(q) dq \Big),\end{aligned} \tag{3}$$

with $\mathrm{arc}_{i,k}(2R_\alpha)$, $k \in \{1, \dots, M_i(2R_\alpha)\}$ the arcs in the boundary of $V_i(P) \cap B_{R_\alpha}(p_i)$.

This result is a consequence of Proposition 2.3; we refer the interested reader to [12] for a detailed proof. For particular choices of performance function, the corresponding multi-center function and its gradient have different features.

a) Centroid problem: If f is piecewise differentiable with no jump discontinuities, then

$$\frac{\partial \mathcal{H}}{\partial p_i}(P) = \int_{V_i(P)} \frac{\partial}{\partial p_i} f(\|q - p_i\|)\phi(q) dq\,.$$

This is the result known in the locational optimization literature [10], [9]. For $f(x) = -x^2$, $\mathcal{H}$ reads

$$\mathcal{H}(P) = -\sum_{i=1}^{n} \int_{V_i(P)} \|q - p_i\|^2 \phi(q) dq \triangleq -\sum_{i=1}^{n} J_{V_i, p_i}\,,$$

where $J_{W,p}$ denotes the polar moment of inertia of the set $W \subset Q$ about the point p. In addition, the gradient of $\mathcal{H}$ is

$$\frac{\partial \mathcal{H}}{\partial p_i}(P) = 2\int_{V_i(P)} (q - p_i)\phi(q) dq = 2\,\mathrm{M}_{V_i(P)}(\mathrm{CM}_{V_i(P)} - p_i).$$

Here M_W and CM_W denote, respectively, the mass and the center of mass with respect to ϕ of $W \subset Q$. The critical points of $\mathcal{H}$ are $P \in Q^n$ with $p_i = \mathrm{CM}_{V_i(P)}$ for all $i \in \{1, \dots, n\}$, i.e. *centroidal Voronoi configurations*, cf. [9].

b) Area problem: If $f(x) = 1_{[0,R]}(x)$, then $\mathcal{H}$ corresponds to the area, measured according to ϕ, covered by the union of the n balls $B_R(p_1), \dots, B_R(p_n)$, that is,

$$\mathcal{H}(P) = \mathrm{area}_\phi(\cup_{i=1}^n B_R(p_i))\,.$$

In this case, the first term in equation (3) vanishes and then

$$\frac{\partial \mathcal{H}}{\partial p_i}(P) = \sum_{k=1}^{M_i(2R)} \int_{\mathrm{arc}_{i,k}(2R)} n_{B_R(p_i)}\, \phi\,.$$

Given $P \in Q^n$, if the ith agent is surrounded by neighbors in the graph $\mathcal{G}_{\mathrm{LD}}(2R)$ in such a way that $M_i(2R) = 0$, then the $\mathcal{H}$ does not depend on p_i. This situation is depicted in Figure 2 (see the example on the right) and captures the fact that the total area covered by the agents is not affected by an infinitesimal displacement of the ith agent.

c) Mixed centroid-area problem: If f is given by $x \mapsto f(x) = -x^2\, 1_{[0,R)}(x) + b \cdot 1_{[R,+\infty)}(x)$, for $b \le -R^2$, the multi-center function takes the form

$$\mathcal{H}(P) = -\sum_{i=1}^{n} J_{V_i(P) \cap B_R(p_i), p_i} + b\ \mathrm{area}_\phi(Q \setminus \cup_{i=1}^n B_R(p_i))\,,$$

and its partial derivative with respect p_i is

$$\begin{aligned}\frac{\partial \mathcal{H}}{\partial p_i}(P) &= 2\,\mathrm{M}_{V_i(P)\cap B_R(p_i)}(\mathrm{CM}_{V_i(P)\cap B_R(p_i)} - p_i) \\ &\quad - (R^2 + b) \sum_{k=1}^{M_i(2R)} \int_{\mathrm{arc}_{i,k}(2R)} n_{B_R(p_i)}\, \phi\,.\end{aligned}$$

In the particular case when $b = -R^2$, the function f is continuous and therefore the gradient of $\mathcal{H}$ takes the form

$$\frac{\partial \mathcal{H}}{\partial p_i}(P) = 2\,\mathrm{M}_{V_i(P)\cap B_R(p_i)}(\mathrm{CM}_{V_i(P)\cap B_R(p_i)} - p_i)\,.$$

In this case, the critical points of $\mathcal{H}$ are $P \in Q^n$ such that $p_i = \mathrm{CM}_{V_i(P)\cap B_R(p_i)}$ for all $i \in \{1, \dots, n\}$. We refer to them as *R-centroidal Voronoi configurations*. For $R \ge \mathrm{diam}(Q)$, they coincide with the standard centroidal Voronoi configurations over Q.

Corollary 3.3: Let ϕ and f be a density and a performance function, respectively. The gradient of $\mathcal{H}$ with respect to the agents' location $P \in Q^n$ is spatially distributed over $\mathcal{G}_{\mathrm{D}}$. Furthermore, if $f(x) = b$ for all $x \ge R$, then the gradient of $\mathcal{H}$ with respect to the agents' location is spatially distributed over $\mathcal{G}_{\mathrm{LD}}(2R)$.

IV. Design of spatially-distributed algorithms for coverage control

Here, we design discrete-time implementations of the gradient ascent for a general aggregate objective function $\mathcal{H}$. We start by extending the definition of $\mathcal{H}$ to consider general partitions $\mathcal{W}$ of Q. Let $P \in Q^n$ and let $\mathcal{W} = \{W_i \subset Q\}_{i=1}^n$

be a partition of Q such that W_i is a convex polygon and $p_i \in W_i$, for $i \in \{1, \ldots, n\}$. Define

$$\mathcal{H}_e(P, \mathcal{W}) = \sum_{i=1}^{n} \int_{W_i} f(\|q - p_i\|)\phi(q)dq\,.$$

The function $\mathcal{H}_e$ is differentiable with respect to its first variable on the whole Q^n. Note that, if $\mathcal{H}_1(p, W) = \int_W f(\|q - p\|)\phi(q)dq$, then we can also write

$$\mathcal{H}_e(P, \mathcal{W}) = \sum_{i=1}^{n} \mathcal{H}_1(p_i, W_i)\,.$$

The following two equalities hold

$$\mathcal{H}(P) = \mathcal{H}_e(P, \mathcal{V}(P))\,, \quad \text{for all } P \in Q^n\,, \quad (4)$$

$$\frac{\partial \mathcal{H}_e}{\partial p_i}(P, \mathcal{V}(P)) = \frac{\partial \mathcal{H}}{\partial p_i}(P)\,, \quad \text{for all } P \in Q^n \setminus \mathcal{S}\,. \quad (5)$$

Let $P \in \mathcal{S}$ and consider a partition $\mathcal{W} = \{W_i \subset Q\}_{i=1}^n$ of Q such that W_i is a convex polygon and $p_i \in W_i$, for $i \in \{1, \ldots, n\}$. Let $i_0, j_0 \in \{1, \ldots, n\}$, $i_0 \neq j_0$ such that $p_{i_0} = p_{j_0}$. Then, $V_{i_0}(P) = V_{j_0}(P)$, and $\mathcal{V}(P)$ is no longer a partition of Q, but a covering. Nevertheless, one could consider the line determined by the edge $W_{i_0} \cap W_{j_0}$ and the associated hyperplanes $H_{i_0, W_{i_0} \cap W_{j_0}}$ and $H_{j_0, W_{i_0} \cap W_{j_0}}$ such that $W_{i_0} \subset H_{i_0, W_{i_0} \cap W_{j_0}}$ and $W_{j_0} \subset H_{j_0, W_{i_0} \cap W_{j_0}}$. With a slight abuse of notation, redefining

$$V_{i_0}(P) = V_{i_0}(P) \cap H_{i_0, W_{i_0} \cap W_{j_0}}\,,$$
$$V_{j_0}(P) = V_{j_0}(P) \cap H_{j_0, W_{i_0} \cap W_{j_0}}\,,$$

the collection $\mathcal{V}(P)$ can be seen a partition of Q. This procedure can be extended if there are more than two coincident agents $\{i_1, \ldots, i_s\}$ at a point $p \in Q$ by defining

$$V_{i_\mu}(P) = V_{i_\mu}(P) \cap \left(\cap_{\nu \in \{1,\ldots,s\} \setminus \{\mu\}} H_{i_\mu, W_{i_\mu} \cap W_{i_\nu}} \right),$$

for $\mu \in \{1, \ldots, s\}$. In the following, such a construction will be tacitly performed for $P \in \mathcal{S}$ and a partition $\mathcal{W}$ of Q. One can show (cf. [12]) that the Voronoi partition is optimal within the set of partitions of Q, that is, $\mathcal{H}_e(P, \mathcal{W}) \leq \mathcal{H}_e(P, \mathcal{V}(P))$ for all $P \in Q^n$. Moreover, the inequality is strict if f is strictly decreasing and the partitions $\mathcal{V}(P)$ and $\mathcal{W}$ differ by a set of non-zero measure.

We are now ready to characterize a class of algorithms with guaranteed convergence to the set of critical points of the aggregate objective function $\mathcal{H}$.

Proposition 4.1 (Discrete-time ascent): Let $T : Q^n \to 2^{Q^n}$ be a closed algorithm on $Q^n \setminus \mathcal{S}$ such that:

(a) for all $P \in Q^n$, $T(P) \cap \mathcal{S} = \emptyset$;
(b) for all $P \in Q^n \setminus \mathcal{S}$, $P' \in T(P)$ and $i \in \{1, \ldots, n\}$, $\mathcal{H}_1(p'_i, V_i(P)) \geq \mathcal{H}_1(p_i, V_i(P))$;
(c) for all $P \in \mathcal{S}$ and $P' \in T(P)$, $\mathcal{H}(P') > \mathcal{H}(P)$;
(d) if $P \in Q^n \setminus \mathcal{S}$ is not a critical point of $\mathcal{H}$, then for all $P' \in T(P)$, there exists $j \in \{1, \ldots, n\}$ such that $\mathcal{H}_1(p'_j, V_j(P)) > \mathcal{H}_1(p_j, V_j(P))$.

Let $P_0 \in Q^n$ denote the initial agents' location. Then, any sequence $\{P_n \mid n \in \mathbb{N} \cup \{0\}\}$ generated according to equation (1) converges to the set of critical points of $\mathcal{H}$.

This result is a consequence of Theorem 2.4; we refer the interested reader to [12] for a detailed proof. In what follows, we devise a general algorithm $T : Q^n \to 2^{Q^n}$ verifying properties (a)-(d) in Proposition 4.1. One can show that if $p_0 \in \partial W$, then $\frac{\partial \mathcal{H}_1(\cdot, W)}{\partial p}(p_0) \neq 0$ points toward the interior of W. If $p_0 \in \operatorname{int}(W)$ is not a critical point, then one also has that $\frac{\partial \mathcal{H}_1(\cdot, W)}{\partial p}(p_0) \neq 0$. For both cases, there exists $\epsilon = \epsilon(p_0, W) > 0$ such that the point p_δ defined by

$$p_\delta = p_0 + \delta\, \frac{\partial \mathcal{H}_1(\cdot, W)}{\partial p}(p_0) \in W$$

has the property that $\mathcal{H}_1(p_\delta) > \mathcal{H}_1(p_0)$, for all $\delta \in (0, \epsilon)$, and $\mathcal{H}_1(p_\epsilon) = \mathcal{H}_1(p_0)$. As it is usually done in nonlinear programming [17], the computation of the step-size ϵ can be implemented numerically via a "line search". With this discussion in mind, let us define the *line search algorithm* $T_{\text{ls}} : Q^n \to 2^{Q^n}$ as follows:

Given $P \in Q^n$, let $P' \in T_{\text{ls}}(P)$ if, for $i \in \{1, \ldots, n\}$ with the property that $p_i \neq p_j, j \in \{1, \ldots, n\} \setminus \{i\}$,

$$p'_i = p_i + \delta\, \frac{\partial \mathcal{H}_1(\cdot, V_i(P))}{\partial p}(p_i)\,, \quad (6)$$

with $\delta \in \left[\frac{\epsilon(p_i, V_i(P))}{3}, \frac{\epsilon(p_i, V_i(P))}{2}\right]$, and, for each set $\{i_1, \ldots, i_s\}$ of coincident indexes at $p \in Q$,

$$p'_{i_\mu} = p_{i_\mu} + \delta\, \frac{\partial \mathcal{H}_1(\cdot, Y_{i_\mu})}{\partial p}(p_{i_\mu})\,, \quad (7)$$

with $\delta \in \left[\frac{\epsilon(p_i, Y_i)}{3}, \frac{\epsilon(p_i, Y_i)}{2}\right]$, where $\{Y_{i_1}, \ldots, Y_{i_s}\}$ is a partition of $V_{i_1}(P) = \cdots = V_{i_s}(P)$ verifying $p \in Y_{i_\mu}$, for $\mu \in \{1, \ldots, s\}$.

Proposition 4.2: The algorithm $T_{\text{ls}} : Q^n \to 2^{Q^n}$ defined by equations (6)-(7) is closed on $Q^n \setminus \mathcal{S}$, and verifies properties (a)-(d) in Proposition 4.1.

Corollary 4.3: Consider the algorithm $T_{\text{ls}} : Q^n \to 2^{Q^n}$ defined by equations (6)-(7). Then

(i) For a general performance function f, the algorithm T_{ls} is spatially distributed over $\mathcal{G}_{\text{D}}$. If, in addition, $f(x) = b$ for all $x \geq R$, then T_{ls} is spatially distributed over $\mathcal{G}_{\text{LD}}(2R)$;
(ii) The sequence of agents' locations generated by T_{ls} according to equation (1) starting at $P_0 \in Q^n$, converges asymptotically to the set of critical points of the aggregate objective function $\mathcal{H}$.

Remark 4.4: In a practical setting, the sensing and/or communication capabilities of each network agent are restricted to a bounded region specified by a finite radius $r > 0$. Therefore, as a way of approximating the gradient ascent of the multi-center function $\mathcal{H}$ for an arbitrary performance function f, one could implement the discrete-time algorithm T_{ls} corresponding to the multi-center function induced by the performance function $f_{\frac{r}{2}} : \overline{\mathbb{R}}_+ \to \mathbb{R}$ given by $f_{\frac{r}{2}}(x) = f(x)$ for $x < \frac{r}{2}$ and $f_{\frac{r}{2}}(x) = f(\operatorname{diam}(Q))$ for $x \geq \frac{r}{2}$. This latter algorithm has the advantage of being spatially distributed over $\mathcal{G}_{\text{LD}}(r)$. We do not follow this path here, and instead refer the interested reader to [12] for a complete exposition, including a detailed characterization of the approximation properties of this procedure. •

V. Simulations

To illustrate the performance of the coordination algorithms, we include some simulation results. The algorithms are implemented in `Mathematica` as a library of routines and a main program running the simulation. We show executions of the discrete-time algorithm T_{ls} (cf. equations (6)-(7)) for the area problem and the mixed centroid-area problem with discontinuous sensing performance in Figs. 3 and 4, respectively. Measuring displacements in meters, we consider the domain Q with vertices

$$\{(0,0),(2.125,0),(2.9325,1.5),(2.975,1.6),\\ (2.9325,1.7),(2.295,2.1),(0.85,2.3),(0.17,1.2)\}.$$

The diameter of this domain is $\operatorname{diam}(Q) = 3.37796$. In all examples, the density function ϕ (represented by means of its contour plot) is the sum of five Gaussian functions of the form $5\,\exp(6(-(x-x_{\text{center}})^2-(y-y_{\text{center}})^2))$ with centers $(x_{\text{center}}, y_{\text{center}})$ at $(2,.25)$, $(1,2.25)$, $(1.9,1.9)$, $(2.35,1.25)$ and $(.1,.1)$. Measured with respect to ϕ, the area of the domain is $\operatorname{area}_\phi(Q) = 8.61656$. Each agent operates with a finite sensing/communication radius equal to $r = .45$.

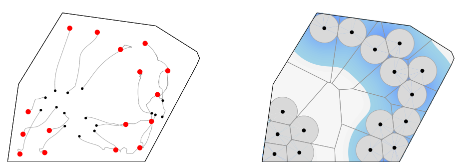

Fig. 3. Area problem (with agent performance $f(x) = 1_{[0,\frac{r}{2}]}(x)$): discrete-time algorithm T_{ls} for 16 agents on a convex polygonal environment. The left figure illustrates the gradient ascent flow and the right figure represents the final location and Voronoi partition. For each agent i, the intersection $V_i \cap B_{\frac{r}{2}}(p_i)$ is plotted in light gray. After 18 seconds, the value of the multi-center function is approximately 6.28977.

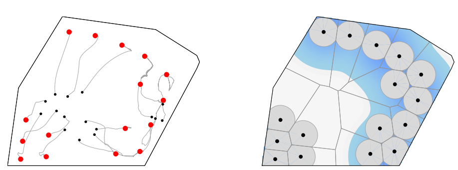

Fig. 4. Mixed centroid-area problem (with discontinuous agent performance $f(x) = -x^2\, 1_{[0,\frac{r}{2})}(x) - \operatorname{diam}(Q)^2 \cdot 1_{[\frac{r}{2},+\infty)}(x)$): discrete-time algorithm T_{ls} for 16 agents on a convex polygonal environment. The left figure illustrates the gradient ascent flow and the right figure represents the final location and Voronoi partition. For each agent i, the intersection $V_i \cap B_{\frac{r}{2}}(p_i)$ is plotted in light gray. After 13.5 seconds, the value of the multi-center function is approximately -6.803.

As mentioned in Remark 4.4, the execution of the coordination algorithm in Figure 4 can be regarded as a limited-range implementation of the gradient ascent of the multi-center function $\mathcal{H}$ corresponding to the agent performance $f(x) = -x^2$; this performance function does not have any range limitation. To compare both executions, we computed the percentage error in the value of $\mathcal{H}$ at their final configurations, which is approximately 30.7%. We also observed that the percentage error of the performance of the limited-range implementation improves with higher values of the ratio $r/\operatorname{diam}(Q)$.

VI. Conclusions and future work

We have presented novel spatially-distributed discrete-time algorithms for coordinated motion of groups of agents. Future research lines include (1) distributed implementation of deterministic annealing techniques (which promise to overcome local maxima), (2) visibility-based algorithms for coverage in non-convex environments, and (3) distributed algorithms for other cooperative behaviors and sensing tasks, e.g., detection, estimation, and map-building.

Acknowledgements

This material is based upon work supported in part by ARO Grant DAAD 190110716, ONR YIP Award N00014-03-1-0512, and NSF SENSORS Award IIS-0330008. Sonia Martínez's work was partially supported by a Fulbright PostDoctoral Fellowship from the Spanish Ministery of Education and Culture.

References

[1] H. Ando, Y. Oasa, I. Suzuki, and M. Yamashita, "Distributed memoryless point convergence algorithm for mobile robots with limited visibility," *IEEE Transactions on Robotics and Automation*, vol. 15, no. 5, pp. 818–828, 1999.

[2] A. Jadbabaie, J. Lin, and A. S. Morse, "Coordination of groups of mobile autonomous agents using nearest neighbor rules," *IEEE Transactions on Automatic Control*, vol. 48, no. 6, pp. 988–1001, 2003.

[3] R. Olfati-Saber and R. M. Murray, "Consensus problems in networks of agents with switching topology and time-delays," *IEEE Transactions on Automatic Control*, Apr. 2003, to appear.

[4] P. Ögren, E. Fiorelli, and N. E. Leonard, "Cooperative control of mobile sensor networks: adaptive gradient climbing in a distributed environment," *IEEE Transactions on Automatic Control*, Aug. 2004, to appear.

[5] K. M. Passino, *Biomimicry for Optimization, Control, and Automation*. New York, NY: Springer Verlag, 2004, in print.

[6] J. Lin, A. S. Morse, and B. D. O. Anderson, "The multi-agent rendezvous problem," in *IEEE Conf. on Decision and Control*, Maui, Hawaii, Dec. 2003, pp. 1508–1513.

[7] Z. Lin, M. Broucke, and B. Francis, "Local control strategies for groups of mobile autonomous agents," *IEEE Transactions on Automatic Control*, vol. 49, no. 4, pp. 622–629, 2004.

[8] J. Cortés, S. Martínez, T. Karatas, and F. Bullo, "Coverage control for mobile sensing networks," *IEEE Transactions on Robotics and Automation*, vol. 20, no. 2, pp. 243–255, 2004.

[9] Q. Du, V. Faber, and M. Gunzburger, "Centroidal Voronoi tessellations: applications and algorithms," *SIAM Review*, vol. 41, no. 4, pp. 637–676, 1999.

[10] A. Okabe, B. Boots, K. Sugihara, and S. N. Chiu, *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*, 2nd ed., ser. Wiley Series in Probability and Statistics. New York, NY: John Wiley & Sons, 2000.

[11] R. M. Gray and D. L. Neuhoff, "Quantization," *IEEE Transactions on Information Theory*, vol. 44, no. 6, pp. 2325–2383, 1998, Commemorative Issue 1948-1998.

[12] J. Cortés, S. Martínez, and F. Bullo, "Spatially-distributed coverage optimization and control with limited-range interactions," *ESAIM. Control, Optimisation & Calculus of Variations*, Jan. 2004, submitted.

[13] R. Diestel, *Graph Theory*, 2nd ed., ser. Graduate Texts in Mathematics. New York, NY: Springer Verlag, 2000, vol. 173.

[14] M. de Berg, M. van Kreveld, and M. Overmars, *Computational Geometry: Algorithms and Applications*. New York, NY: Springer Verlag, 1997.

[15] J. W. Jaromczyk and G. T. Toussaint, "Relative neighborhood graphs and their relatives," *Proceedings of the IEEE*, vol. 80, no. 9, pp. 1502–1517, 1992.

[16] A. J. Chorin and J. E. Marsden, *A Mathematical Introduction to Fluid Mechanics*, 3rd ed., ser. Texts in Applied Mathematics. New York, NY: Springer Verlag, 1994, vol. 4.

[17] D. G. Luenberger, *Linear and Nonlinear Programming*, 2nd ed. Reading, MA: Addison-Wesley, 1984.