Distributed Algorithms for Environment Partitioning in Mobile Robotic Networks

Marco Pavone, Alessandro Arsie, Emilio Frazzoli, Francesco Bullo

***Abstract*—A widely applied strategy for workload sharing is to equalize the workload assigned to each resource. In mobile multi-agent systems, this principle directly leads to equitable partitioning policies whereby (i) the environment is equitably divided into subregions of equal measure, (ii) one agent is assigned to each subregion, and (iii) each agent is responsible for service requests originating within its own subregion. The current lack of distributed algorithms for the computation of equitable partitions limits the applicability of equitable partitioning policies to limited-size multi-agent systems operating in known, static environments. In this paper, first, we design provably correct and spatially distributed algorithms that allow a team of agents to compute a convex and equitable partition of a convex environment. Second, we discuss how these algorithms can be extended so that a team of agents can compute, in a spatially distributed fashion, convex and equitable partitions with additional features, e.g., equitable and median Voronoi diagrams. Finally, we discuss two application domains for our algorithms, namely dynamic vehicle routing for mobile robotic networks and wireless ad hoc networks. Through these examples we show how one can couple the algorithms presented in this paper with equitable partitioning policies to make these amenable to distributed implementation; more in general, we illustrate a systematic approach to devise spatially distributed control policies for a large variety of multi-agent coordination problems. Our approach is related to the classic Lloyd algorithm, and exploits the unique features of power diagrams.**

I. Introduction

In the near future, large groups of autonomous agents will be used to perform complex tasks including transportation and distribution, logistics, surveillance, search and rescue operations, humanitarian demining, environmental monitoring, and planetary exploration. The potential advantages of multi-agent systems are, in fact, numerous. For instance, the intrinsic parallelism of a multi-agent system provides robustness to failures of single agents, and in many cases can guarantee better time efficiency. Moreover, it is possible to reduce the total implementation and operation cost, increase reactivity and system reliability, and add flexibility and modularity to monolithic approaches.

In essence, agents can be interpreted as *resources* to be *shared* among *customers*. In surveillance and exploration missions, customers are points of interest to be visited; in transportation and distribution applications, customers are people demanding some goods or services (e.g., utility repair); in logistics tasks, customers could be troops on the battlefield.

A widely applied strategy for workload sharing is to equalize the total workload assigned to each resource. In mobile multi-agent systems, this strategy naturally leads to *equitable partitioning policies* [1]–[4]. An equitable partitioning policy is a workload sharing policy whereby the environment $\mathcal{Q} \subset \mathbb{R}^d$ is *equitably* partitioned into m openly disjoint subregions $\mathcal{Q}_i$ ($i \in \{1, \dots, m\}$) whose union is $\mathcal{Q}$, where m is the number of available agents; then, each agent i is assigned to subregion $\mathcal{Q}_i$, and each customer in $\mathcal{Q}_i$ receives service by the agent assigned to $\mathcal{Q}_i$. In this paper, equitability is in the following sense: If we model the workload for subregion $\mathcal{T} \subseteq \mathcal{Q}$ as $\lambda_{\mathcal{T}} \doteq \int_{\mathcal{T}} \lambda(x)\, dx$, where $\lambda(\cdot)$ is a measure over $\mathcal{Q}$, the workload for agent i is $\lambda_{\mathcal{Q}_i}$ (the measure λ can represent, for example, the density of customers over $\mathcal{Q}$, or, in a stochastic setting, their arrival rate). Then, an equitable partition (i.e., a partition that guarantees equitable workload sharing) is a partition where $\lambda_{\mathcal{Q}_i} = \lambda_{\mathcal{Q}}/m$, for all i.

Equitable partitioning policies are predominant for three main reasons: (i) efficiency, (ii) ease of design and (iii) ease of analysis. Equitable partitioning policies are, therefore, ubiquitous in applications involving multi-agent systems. To date, nevertheless, to the best of our knowledge, all equitable partitioning policies inherently assume a centralized computation of the partition of the environment (henceforth, we will refer to algorithms for the computation of partitions as *partitioning algorithms*). This fact is in sharp contrast with the desire of a fully distributed architecture for a multi-agent system. The lack of a fully distributed architecture limits the applicability of equitable partitioning policies to limited-size multi-agent systems operating in a known, static environment.

The contribution of this paper is threefold. First, we design provably correct and spatially distributed algorithms that allow a team of agents to compute a convex and equitable partition of a convex environment. Our approach is related to the classic Lloyd algorithm from vector quantization theory [5], [6], and exploits the unique features of power diagrams, a generalization of Voronoi diagrams (a similar approach is studied in [7] in the context of sensor networks performing static coverage optimization with area constraints; see also [8] for another interesting application of power diagrams in the context of power-constrained mobile sensor networks). Second, we discuss how these algorithms can be extended so that a team of agents can compute, in a spatially distributed fashion, convex and equitable partitions with additional features critical to applications. For example, we consider equi-

Submitted December 6, 2008, revised version October 5, 2010.

Marco Pavone is with the Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109 (Marco.Pavone@jpl.nasa.gov).

Alessandro Arsie is with the Department of Mathematics, University of Toledo, Toledo, OH 43606 (alessandro.arsie@utoledo.edu).

Emilio Frazzoli is with the Laboratory for Information and Decision Systems, Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA 02139 (frazzoli@mit.edu).

Francesco Bullo is with the Center for Control Engineering and Computation, University of California at Santa Barbara, CA 93106 (bullo@engineering.ucsb.edu).

table and median Voronoi diagrams, which play a key role in several application domains: to the best of our knowledge, no algorithm (centralized or distributed) is currently available for their computation, but they are within the scope of our analysis. Third, we discuss two important application domains for our algorithms, namely, dynamic vehicle routing for mobile robotic networks and wireless ad hoc networks. Through these examples we illustrate a systematic approach to devise spatially distributed control policies for the class of multi-agent coordination problems that admit equitable partitioning policies as a solution. This approach consists in combining the partitioning algorithms presented in this paper with suitable single-agent control laws. In other words, the partitioning algorithms we devise in this paper are a "building block" instrumental to design spatially distributed control policies for a large variety of multi-agent coordination problems.

We mention that our algorithms, although motivated in the context of multi-agent systems, are a novel contribution to the field of computational geometry. In particular we address, using a dynamical system framework, the well-studied equitable convex partition problem (see [9] and references therein); moreover, our analysis provides new insights in the geometry of Voronoi diagrams and power diagrams (including some existence and impossibility results).

The paper is organized as follows. In Section II we provide the necessary tools from calculus, algebraic topology, and computational geometry. In Section III we first prove some existence results for power diagrams, and then we design provably correct and spatially distributed algorithms for the computation of equitable partitions. In Section IV we discuss how one can extend these algorithms to enable the spatially distributed computation of convex and equitable partitions with additional features (e.g., equitable and median Voronoi diagrams). In Section V we describe two application domains for the algorithms developed in this paper, namely dynamic vehicle routing and deployment of wireless ad hoc networks, and in Section VI we present results from numerical experiments. Finally, in Section VII, we draw our conclusions.

II. Background

In this section we introduce some notation and briefly review some concepts from calculus, algebraic topology, and computational geometry, on which we will rely extensively later in the paper.

A. Notation

Let $\|\cdot\|$ denote the Euclidean norm. Let $\mathcal{Q}$ be a compact, convex subset of $\mathbb{R}^d$. We denote the boundary of $\mathcal{Q}$ as $\partial\mathcal{Q}$ and the Lebesgue measure of $\mathcal{Q}$ as $|\mathcal{Q}|$. We define the diameter of $\mathcal{Q}$ as: $\operatorname{diam}(\mathcal{Q}) \doteq \max\{||p-q|| \,|\, p, q \in \mathcal{Q}\}$. The distance from a point x to a set M is defined as $\operatorname{dist}(x, M) \doteq \inf_{p\in M} \|x - p\|$. We define $I_m \doteq \{1, 2, \ldots, m\}$. Let $G = (g_1, \ldots, g_m) \subset \mathcal{Q}^m$ denote the location of m points in $\mathcal{Q}$. A *partition* (or tessellation) of $\mathcal{Q}$ is a collection of m closed subsets $\{\mathcal{Q}_i\}_{i=1}^m$ with disjoint interiors whose union is $\mathcal{Q}$. A partition $\{\mathcal{Q}_i\}_{i=1}^m$ is *convex* if each $\mathcal{Q}_i$, $i \in I_m$, is convex. Let $\lambda : \mathcal{Q} \to \mathbb{R}_{>0}$ be a measure over $\mathcal{Q}$, absolutely continuous with respect to the Lebesgue measure. Define $\lambda_{\mathcal{T}} \doteq \int_{\mathcal{T}} \lambda(x)\, dx$ for any $\mathcal{T} \subseteq \mathcal{Q}$; a partition $\{\mathcal{Q}_i\}_{i=1}^m$ of the environment $\mathcal{Q}$ is *equitable* with respect to λ if $\lambda_{\mathcal{Q}_i} = \lambda_{\mathcal{Q}_j}$ for all $i, j \in I_m$.

Finally, we define the saturation function $\operatorname{sat}_{a,b}(x)$, with $a < b$, as:

$$\operatorname{sat}_{a,b}(x) = \begin{cases} 1, & \text{if } x > b, \\ (x-a)/(b-a), & \text{if } a \le x \le b, \\ 0, & \text{if } x < a. \end{cases}$$

B. Variation of an Integral Function due to a Domain Change

The following result is related to classic divergence theorems [10]. Let $Q = Q(y) \subset \mathcal{Q}$ be a region that depends smoothly on a real parameter $y \in \mathbb{R}$ and that has a well-defined boundary $\partial Q(y)$ for all y. Let h be a density function over $\mathcal{Q}$. Then

$$\frac{d}{dy}\int_{Q(y)} h(x)\, dx = \int_{\partial Q(y)} \left(\frac{dx}{dy}\cdot n(x)\right) h(x)\, dx, \quad (1)$$

where $v \cdot w$ denotes the scalar product between vectors v and w, $n(x)$ is the unit outward normal to $\partial Q(y)$ at x, and dx/dy is the derivative of the boundary points with respect to y.

C. A Sufficient Condition for the Surjectivity of a Map

The following two results will be fundamental to prove some existence theorems and are a direct consequence of the theory of degree of continuous maps between spheres (see the Appendix for the definition of degree of a map).

Theorem 2.1 (Surjectivity of continuous maps): Let B^m be a closed m-dimensional ball and let S^{m-1} be its boundary, namely an $(m-1)$-dimensional sphere. Let $f : B^m \to B^m$ be a continuous map and assume that its restriction $f_{S^{m-1}} : S^{m-1} \to S^{m-1}$ has degree different from 0. Then f is *onto* B^m.

Proof: See Appendix. ■

In the sequel we will also need the following result.

Lemma 2.2 (Degree of continuous bijective maps): Let $f : S^m \to S^m$, with $m \ge 1$, be a continuous bijective map from an m-dimensional sphere to itself. Then the degree of f is equal to ± 1.

Proof: See Appendix. ■

D. Voronoi Diagrams and Power Diagrams

We refer the reader to [11] and [12] for comprehensive treatments, respectively, of Voronoi diagrams and power diagrams. Assume that $G = (g_1, \ldots, g_m)$ is an ordered set of *distinct* points. The *Voronoi diagram* $\mathcal{V}(G) = (V_1(G), \ldots, V_m(G))$ of $\mathcal{Q}$ generated by points G is defined by

$$V_i(G) = \{x \in \mathcal{Q} |\ \|x - g_i\| \le \|x - g_j\|,\ \forall j \ne i,\ j \in I_m\}.$$

We refer to G as the set of *generators* of $\mathcal{V}(G)$, and to $V_i(G)$ as the Voronoi cell or region of dominance of the i-th generator. For $g_i, g_j \in G$, $i \ne j$, we define the *bisector* between g_i and g_j as $b(g_i, g_j) = \{x \in \mathcal{Q} |\ \|x - g_i\| = \|x - g_j\|\}$. The face $b(g_i, g_j)$ bisects the line segment joining g_i and g_j, and this line segment is orthogonal to the face (*Perpendicular Bisector*

Property). One can easily show that each Voronoi cell is a convex set, and thus a Voronoi diagram of $\mathcal{Q}$ is a convex partition of $\mathcal{Q}$ (see Figure 1(a)). The Voronoi diagram of an ordered set of possibly *coincident* points is not well defined.

Assume, now, that each point $g_i \in G$ has assigned an individual weight $w_i \in \mathbb{R}$, $i \in I_m$; let $W = (w_1, \dots, w_m)$. We define the *power distance* as

$$d_P(x, g_i; w_i) \doteq \|x - g_i\|^2 - w_i.$$

We refer to the pair (g_i, w_i) as a *power point* and define $G_W = \Big((g_1, w_1), \dots, (g_m, w_m)\Big)$. Two power points (g_i, w_i) and (g_j, w_j) are coincident if $g_i = g_j$ *and* $w_i = w_j$. Assume that G_W is an ordered set of distinct power points. Similarly as before, the *power diagram* $\mathcal{V}(G_W) = (V_1(G_W), \dots, V_m(G_W))$ of $\mathcal{Q}$ generated by power points G_W is defined by

$$V_i(G_W) = \{x \in \mathcal{Q} |\ \|x - g_i\|^2 - w_i \leq \|x - g_j\|^2 - w_j, \quad \forall j \neq i,\ j \in I_m\}.$$

We refer to G_W as the set of *power generators* of $\mathcal{V}(G_W)$, and to $V_i(G_W)$ as the power cell or region of dominance of the i-th power generator; moreover we call g_i and w_i, respectively, the position and the weight of the power generator (g_i, w_i). One can easily show that a power diagram is a convex partition of $\mathcal{Q}$. Notice that, when all weights are the same, the power diagram of $\mathcal{Q}$ coincides with the Voronoi diagram of $\mathcal{Q}$. Indeed, power diagrams are the generalized Voronoi diagrams that have the strongest similarities to the original diagrams [13]. There are some differences, though. First, a power cell might be empty. Second, g_i might not be in its power cell (see Figure 1(b), where each weight is positive, and each power generator (g_i, w_i) is represented by a circle whose center is g_i and whose radius is $\sqrt{w_i}$). Finally, the bisector of (g_i, w_i) and (g_j, w_j), $i \neq j$, is

$$b\Big((g_i, w_i), (g_j, w_j)\Big) = \{x \in \mathcal{Q} |\ (g_j - g_i)^{\mathrm{T}} x = \frac{1}{2}(\|g_j\|^2 - \|g_i\|^2 + w_i - w_j)\}. \tag{2}$$

Hence, $b\Big((g_i, w_i), (g_j, w_j)\Big)$ is a face orthogonal to the line segment $\overline{g_i\, g_j}$ and passing through the point g^*_{ij} given by

$$g^*_{ij} = \frac{\|g_j\|^2 - \|g_i\|^2 + w_i - w_j}{2\|g_j - g_i\|^2}(g_j - g_i);$$

this last property will be crucial in the remainder of the paper: it means that, by changing the values of the weights, it is possible to arbitrarily move the bisector between the positions of the two corresponding power generators, while still preserving the orthogonality constraint. The power diagram of an ordered set of possibly *coincident* power points is not well defined.

For simplicity, we will refer to $V_i(G)$ ($V_i(G_W)$) as V_i. When the two Voronoi (power) cells V_i and V_j are adjacent (i.e., they share a face), g_i ((g_i, w_i)) is called a *Voronoi* (*power*) *neighbor* of g_j ((g_j, w_j)), and vice-versa. The set of indices of the Voronoi (power) neighbors of g_i ((g_i, w_i)) is denoted by N_i. We also denote the (i, j)-face as $\Delta_{ij} \doteq V_i \cap V_j$.

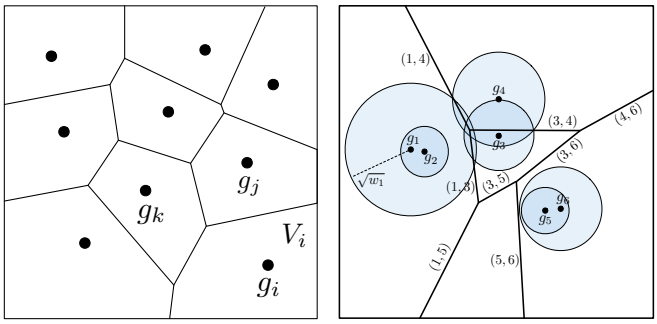

(a) A Voronoi Diagram.

(b) A power diagram [13]. The weights w_i are all positive. Power generator (g_2, w_2) has an empty cell. Power generator (g_5, w_5) is outside its region of dominance.

Fig. 1. Examples of Voronoi diagrams and power diagrams.

III. A Spatially-Distributed Algorithm to Compute Equitable Partitions

In this section we develop a provably correct and spatially distributed algorithm for the computation of a convex and equitable partition of a convex environment $\mathcal{Q}$ (see, e.g., [14] for a rigorous definition of spatially distributed algorithms). In the next section we will present an extension of this algorithm, which enables the spatially distributed computation of convex and equitable partitions with additional features (e.g., convex and equitable partitions that are approximations of equitable Voronoi diagrams).

To develop our algorithms, we restrict our attention to a specific class of partitions, namely the class of *power diagrams*. The reason for focusing on power diagrams is threefold. First, power diagrams can be viewed as a map between sets of weighted points and regions of dominance; as it will become apparent in Section III-B, where we give an overview of the proposed algorithm, this property simplifies considerably the task of designing spatially distributed algorithms for environment partitioning. Second, several well-known and practically important convex partitions, such as median Voronoi diagrams, are particular types of power diagrams; hence, power diagrams are rather general. Finally, equitable power diagrams are always guaranteed to exist, as we show next.

A. On the Existence of Equitable Power Diagrams

An important property of power diagrams is that an equitable power diagram always exists for any λ (notice that in general, when λ is non-uniform, an equitable Voronoi diagram may *fail* to exist, as we will show in section III-E). Indeed, as shown in the next theorem, an equitable power diagram (with respect to any given λ) exists for *any* vector of *distinct* points $G = (g_1, \dots, g_m)$ in $\mathcal{Q}$.

Theorem 3.1 (Existence of equitable power diagrams): Let $G = (g_1, \dots, g_m)$ be the positions of $1 \leq m < \infty$ distinct points in $\mathcal{Q}$. Then, there exist weights w_i, $i \in I_m$, such that the power points $\Big((g_1, w_1), \dots, (g_m, w_m)\Big)$ generate a power diagram that is equitable with respect to λ. Moreover, given a vector of weights W^* that yields an equitable power diagram,

the set of all vectors of weights yielding an equitable power diagram is $\mathcal{W}^* \doteq \{W^* + t[1,\dots,1] \mid t \in \mathbb{R}\}$.

Proof: It is not restrictive to assume that $\lambda_{\mathcal{Q}} = 1$ (i.e., we normalize the measure of $\mathcal{Q}$), since $\mathcal{Q}$ is bounded. The strategy of the proof is to use a topological argument to force existence. Specifically, we view a power diagram as a "map" that maps vectors of weights into vectors of measures of power cells, and we show that this map is surjective by applying Theorem 2.1. The surjectivity of the "power diagram map", in turn, *implies* that it must exist a vector of weights that realizes an equitable power diagram.

We begin the proof by constructing a weight space. Let $D = \mathrm{diam}(\mathcal{Q})^2/2$, and consider the cube $\mathcal{C} \doteq [-D, D]^m$ (see Figure 2). This is the weight space and we consider weight vectors W taking value in $\mathcal{C}$; this is not restrictive, since if the difference between two weights is larger than $2D = \mathrm{diam}(\mathcal{Q})^2$, at least one cell has measure zero, and thus the corresponding power diagram can not be equitable. Second, consider the standard *m-simplex* of measures $\lambda_{\mathcal{Q}_1}, \dots, \lambda_{\mathcal{Q}_m}$ (where $\mathcal{Q}_1, \dots, \mathcal{Q}_m$ are the power cells). This can be realized in $\mathbb{R}^m$ as the subset of defined by $\sum_{i=1}^m \lambda_{\mathcal{Q}_i} = 1$ with the condition $\lambda_{\mathcal{Q}_i} \geq 0$ (see again Figure 2). Let us call this set "the measure simplex $\mathcal{A}$" (notice that it is $(m-1)$-dimensional).

We call $f : \mathcal{C} \to \mathcal{A}$ the map associating, according to the power distance, a weight vector W with the corresponding vector of measures $(\lambda_{\mathcal{Q}_1}, \dots, \lambda_{\mathcal{Q}_m})$. Since the points in G are assumed to be distinct, this map is continuous.

To prove the first statement of the theorem, we will prove that $f : \mathcal{C} \to \mathcal{A}$ is surjective for every $m \geq 1$ by using induction on m, starting with the base case $m = 3$ (the statement for $m = 1$ and $m = 2$ is trivially checked). We consider as base case $m = 3$ since its study, which can be aided by visualization, contains most of the ideas involved in the inductive step and makes the corresponding proof more transparent. When $m = 3$, the weight space $\mathcal{C}$ is a three-dimensional cube with vertices $v_0 = [-D, -D, -D]$, $v_1 = [D, -D, -D]$, $v_2 = [-D, D, -D]$, $v_3 = [-D, -D, D]$, $v_4 = [D, -D, D]$, $v_5 = [-D, D, D]$, $v_6 = [D, D, -D]$, and $v_7 = [D, D, D]$. The measure simplex $\mathcal{A}$ is a triangle with vertices u_1, u_2, and u_3 that correspond to the cases 1) $\lambda_{\mathcal{Q}_1} = 1, \lambda_{\mathcal{Q}_2} = 0, \lambda_{\mathcal{Q}_3} = 0$, 2) $\lambda_{\mathcal{Q}_1} = 0, \lambda_{\mathcal{Q}_2} = 1, \lambda_{\mathcal{Q}_3} = 0$, and 3) $\lambda_{\mathcal{Q}_1} = 0, \lambda_{\mathcal{Q}_2} = 0, \lambda_{\mathcal{Q}_3} = 1$, respectively. Moreover, call e_1, e_2 and e_3 the edges opposite to the vertices u_1, u_2, and u_3, respectively. The edges e_i are, therefore, given by the condition $\{\lambda_{\mathcal{Q}_i} = 0\}$ (see Figure 2).

Let us return to the map $f : \mathcal{C} \to \mathcal{A}$. It is easy to see that f is constant on sets of the form $\mathcal{W} \doteq \{\{W + t(1,1,1)\} \cap \mathcal{C}, \quad t \in \mathbb{R}\}$, where W is a weight vector in $\mathcal{C}$; in other words, whenever two sets of weights differ by a common quantity, they are mapped to the same point in $\mathcal{A}$. Moreover, fixing a point $p \in \mathcal{A}$, we have that $f^{-1}(p)$ is simply given by a set of the form $\mathcal{W}$ for a suitable W (a proof of this fact is provided, for any $m \geq 3$, within the proof of the inductive step, which is presented in the Appendix). Hence, the "fibers" of f, i.e., the loci where f is constant, are straight lines *parallel* to the main diagonal v_0v_7 (the second statement in Theorem 3.1 is an immediate consequence of this fact). Note that the image of the diagonal v_0v_7 is exactly the point p_0 of $\mathcal{A}$ for which the measures are those of a Voronoi partition, since the weights are all equal.

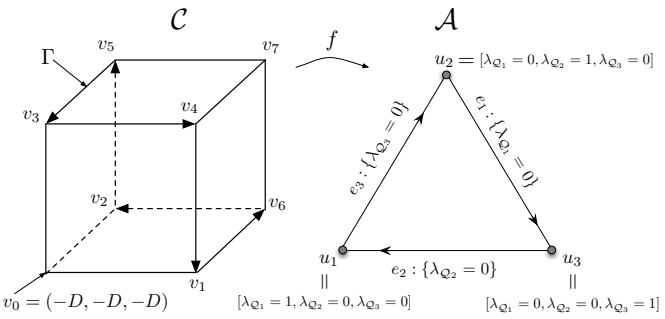

Fig. 2. Weight space $\mathcal{C}$, measure simplex $\mathcal{A}$, and the power diagram map $f : \mathcal{C} \to \mathcal{A}$ for $m = 3$.

On the weight space $\mathcal{C}$ let us define the following equivalence relation: $W^1 \equiv W^2$ if and only if they are on a line parallel to the main diagonal v_0v_7. The map $f : \mathcal{C} \to \mathcal{A}$ induces a continuous map (still called f by abuse of notation) from $\mathcal{C}/\equiv$ to $\mathcal{A}$ having the same image. Let us identify $\mathcal{C}/\equiv$ with a simpler topological space. Since any line in the cube parallel to the main diagonal v_0v_7 is entirely determined by its intersections with the three faces $F_3 = \mathcal{C} \cap \{w_3 = -D\}$, $F_2 = \mathcal{C} \cap \{w_2 = -D\}$, and $F_1 = \mathcal{C} \cap \{w_1 = -D\}$, we can identify $\mathcal{C}/\equiv$ with the union of these faces and we call this union $\mathcal{F}$. We therefore have a continuous map $f : \mathcal{F} \to \mathcal{A}$ that has the same image of the original f; besides, the induced map $f : \mathcal{F} \to \mathcal{A}$ is *injective* by construction, since each fiber intersects $\mathcal{F}$ in only one point.

Observe that $\mathcal{F}$ is homeomorphic to B^2, the 2-dimensional ball, like $\mathcal{A}$ itself. Up to homeomorphisms, therefore, the map $f : \mathcal{F} \to \mathcal{A}$ can be viewed as a map (again called f by abuse of notation) $f : B^2 \to B^2$. Consider the closed loop Γ given by v_2v_5, v_5v_3, v_3v_4, v_4v_1, v_1v_6, v_6v_2 with this orientation (see Figure 2). This loop is the boundary of $\mathcal{F}$ and we think of it also as the boundary of B^2. Taking into account the continuity of f, it is easy to see that f maps Γ onto the boundary of $\mathcal{A}$. For example, while we move on the edges v_2v_5 and v_5v_3, that are characterized by having $w_1 = -D$, the corresponding point on the measure simplex moves on the edge e_1.

Moreover, since f is injective by construction, the inverse image of any point on the boundary of $\mathcal{A}$ is just one element of Γ. Identifying the boundary of $\mathcal{A}$ with S^1 (up to homeomorphisms) and the loop Γ with S^1 (up to homeomorphisms) we have a bijective continuous map $f_{S^1} : S^1 \to S^1$. By Lemma 2.2 the degree of f is equal to ± 1 and, therefore, f is onto $\mathcal{A}$, using Theorem 2.1. This proves the base case $m = 3$. The proof of the inductive step is provided in the Appendix. ■

Some remarks are in order.

Remark 3.2 (General measure assignment): In the proof of the above theorem, we actually proved that for *any* measure vector $(\lambda_{\mathcal{Q}_1}, \dots, \lambda_{\mathcal{Q}_m})$ in $\mathcal{A}$ there exists a weight vector $W \in \mathcal{C}$ realizing it through the map f. This could be useful in some applications.

Remark 3.3 (Uniqueness of equitable power diagrams): Since all vectors of weights in $\mathcal{W}^*$ yield exactly the *same* power diagram, we conclude that the positions of the

generators *uniquely* induce an equitable power diagram.

B. Overview of the Algorithm

Henceforth, we assume that $\mathcal{Q}$ is a compact, convex subset of $\mathbb{R}^2$ (we will discuss more general environments in Remark 4.5). Each agent locally controls a power generator $(g_i(t), w_i(t)) \in \mathcal{Q} \times \mathbb{R}$, where $t \in \mathbb{R}_{\geq 0}$ denotes dependence on time. We will refer to the power cell $V_i = V_i(G_W)$, where $G_W = \Big((g_1, w_1), \ldots, (g_m, w_m)\Big)$, as the region of dominance of agent i, and to the partition into regions of dominance induced by the set of generators[1] G_W as $\mathcal{V}(G_W)$.

The key idea is to construct an energy function with the properties that (1) it depends on the weights of the generators, and (2) all its critical points correspond to vectors of weights yielding an equitable power diagram (whose existence is guaranteed by Theorem 3.1). Then, the agents update their weights according to a spatially distributed gradient-descent law (while maintaining the positions of the generators *fixed*) until a critical point of the energy function (and hence an equitable power diagram) is reached.

Assume, henceforth, that the positions of the generators are *distinct*, i.e., $g_i \neq g_j$ for $i \neq j$, and define the set

$$S \doteq \Big\{(w_1, \ldots, w_m) \in \mathbb{R}^m \,|\, \lambda_{V_i} > 0, \;\; \forall i \in I_m\Big\}.$$

In other words, set S contains all vectors of weights such that no region of dominance has measure equal to *zero*.

We introduce the energy function $H_{\mathcal{V}} : S \to \mathbb{R}_{>0}$:

$$H_{\mathcal{V}}(W) \doteq \sum_{i=1}^{m} \left(\int_{V_i(W)} \lambda(x) dx \right)^{-1} = \sum_{i=1}^{m} \lambda_{V_i(W)}^{-1}, \qquad (3)$$

where $W = (w_1, \ldots, w_m)$.

C. Smoothness and Gradient of $H_{\mathcal{V}}$

We now analyze the smoothness properties of $H_{\mathcal{V}}$. In the following, let $\gamma_{ij} \doteq \|g_j - g_i\|$.

Theorem 3.4 (Smoothness of $H_{\mathcal{V}}$): Assume the generators' positions are *distinct*, i.e., $g_i \neq g_j$ for $i \neq j$. The following statements hold:

1) the function $H_{\mathcal{V}}$ is continuously differentiable on S, where for each $i \in \{1, \ldots, m\}$

$$\frac{\partial H_{\mathcal{V}}}{\partial\, w_i} = \sum_{j \in N_i} \frac{1}{2\gamma_{ij}} \left(\frac{1}{\lambda_{V_j}^2} - \frac{1}{\lambda_{V_i}^2} \right) \int_{\Delta_{ij}} \lambda(x)\, dx, \qquad (4)$$

2) all critical points of $H_{\mathcal{V}}$ are weight vectors that yield an equitable power diagram.

Proof: By assumption, $g_i \neq g_j$ for $i \neq j$, thus the power diagram is well defined. Since the motion of a weight w_i only affects power cell V_i and its neighboring cells V_j for $j \in N_i$, we have

$$\frac{\partial H_{\mathcal{V}}}{\partial w_i} = -\frac{1}{\lambda_{V_i}^2} \frac{\partial \lambda_{V_i}}{\partial w_i} - \sum_{j \in N_i} \frac{1}{\lambda_{V_j}^2} \frac{\partial \lambda_{V_j}}{\partial w_i}.$$

[1]For brevity, we will often refer to a power generator simply as a generator.

Now, the result in equation (1) provides the means to analyze the variation of an integral function due to a domain change. Since the boundary of V_i satisfies $\partial V_i = \cup_j \Delta_{ij} \cup B_i$, where $\Delta_{ij} = \Delta_{ji}$ is the edge between V_i and V_j, and B_i is the boundary between V_i and $\mathcal{Q}$ (if any, otherwise $B_i = \emptyset$), we have

$$\frac{\partial \lambda_{V_i}}{\partial w_i} = \sum_{j \in N_i} \int_{\Delta_{ij}} \left(\frac{\partial x}{\partial w_i} \cdot n_{ij}(x) \right) \lambda(x)\, dx + \underbrace{\int_{B_i} \left(\frac{\partial x}{\partial w_i} \cdot n_{ij}(x) \right) \lambda(x)\, dx}_{=0}, \qquad (5)$$

where we defined n_{ij} as the unit normal to Δ_{ij}, outward of V_i (thus $n_{ji} = -n_{ij}$). The second term is trivially equal to zero if $B_i = \emptyset$; it is also equal to zero if $B_i \neq \emptyset$, since the integrand is zero almost everywhere. Similarly,

$$\frac{\partial \lambda_{V_j}}{\partial w_i} = \int_{\Delta_{ij}} \left(\frac{\partial x}{\partial w_i} \cdot n_{ji}(x) \right) \lambda(x)\, dx. \qquad (6)$$

To evaluate the scalar product between the derivative of the boundary points and the unit normal to the boundary in equations (5) and (6), we differentiate equation (2) with respect to w_i at every point $x \in \Delta_{ij}$; we get

$$\frac{\partial x}{\partial w_i} \cdot (g_j - g_i) = \frac{1}{2}.$$

From equation (2) we have $n_{ij} = (g_j - g_i) / \|g_j - g_i\|$, and the desired explicit expressions for the scalar products in equation (5) and in equation (6) follow immediately (recalling that $n_{ji} = -n_{ij}$). Collecting the above results, we obtain the partial derivative with respect to w_i.

The proof of the characterization of the critical points (i.e., the proof of the second statement) is an immediate consequence of the expression for the gradient of $H_{\mathcal{V}}$; we omit it in the interest of brevity. ∎

Remark 3.5 (Spatially distributed gradient computation): The computation of the gradient in Theorem 3.4 is spatially distributed over the dual graph of the power diagram (we call such graph the power-Delaunay graph), since the summation in equation (4) only runs through the indices of generators with neighboring power cells.

Example 3.6 (Gradient of $H_{\mathcal{V}}$ for uniform measure): The gradient of $H_{\mathcal{V}}$ simplifies considerably when λ is uniform. In this case, it is straightforward to verify that

$$\frac{\partial H_{\mathcal{V}}}{\partial\, w_i} = \frac{|\mathcal{Q}|}{2\,\lambda_{\mathcal{Q}}} \sum_{j \in N_i} \frac{\delta_{ij}}{\gamma_{ij}} \left(\frac{1}{|V_j|^2} - \frac{1}{|V_i|^2} \right),$$

where δ_{ij} is the length of the boundary segment Δ_{ij}.

D. Spatially Distributed Algorithm for Equitable Partitioning

Each agent updates its own weight according to the following control law defined over the set S:

$$\dot{w}_i(t) = -\frac{\partial H_{\mathcal{V}}}{\partial w_i}(W(t)), \quad t \geq 0, \qquad (7)$$

where the dot represents differentiation with respect to time, and where we assume that the partition $\mathcal{V}(W) = \{V_1, \ldots, V_m\}$ is continuously updated. One can prove the following result.

Theorem 3.7 (Correctness of algorithm (7)): Assume that the positions of the generators are *distinct*, i.e., $g_i \neq g_j$ for $i \neq j$. Consider the gradient vector field on S defined by equation (7). Then generators' weights starting at $t = 0$ at $W(0) \in S$ and evolving under (7) remain in S and converge asymptotically to a critical point of $H_{\mathcal{V}}$, i.e., to a vector of weights yielding an equitable power diagram.

Proof: Define the set

$$U \doteq \Big\{(w_1, \ldots, w_m) \in \mathbb{R}^m \mid \sum_{i=1}^{m} w_i = c\Big\},$$

where $c \in \mathbb{R}$ is an arbitrary constant. Let $\Omega \doteq S \cap U$. We next prove that generators' weights starting in Ω and evolving under (7) converge to a vector of weights yielding an equitable power diagram.

By assumption, $g_i \neq g_j$ for $i \neq j$, thus the power diagram is well defined. First, we prove that set Ω is positively invariant with respect to (7). Noticing that control law (7) is a gradient descent law, we have for all trajectories starting in Ω

$$\lambda_{V_i(W(t))}^{-1} \leq H_{\mathcal{V}}(W(t)) \leq H_{\mathcal{V}}(W(0)), \quad i \in I_m,\ t \geq 0.$$

Since the measures of the power cells depend continuously on the weights, we conclude that the measures of all power cells will be bounded away from zero; thus, the weights will belong to S for all $t \geq 0$, that is, $W(t) \in S\ \forall t \geq 0$. Moreover, the sum of the weights is invariant, in fact

$$\frac{\partial \sum_{i=1}^{m} w_i}{\partial t} = -\sum_{i=1}^{m} \frac{\partial H_{\mathcal{V}}}{\partial w_i} = -\sum_{i=1}^{m} \sum_{j \in N_i} \frac{1}{2\gamma_{ij}} \Big(\frac{1}{\lambda_{V_j}^2} - \frac{1}{\lambda_{V_i}^2}\Big) \int_{\Delta_{ij}} \lambda(x)\, dx = 0,$$

since $\gamma_{ij} = \gamma_{ji}$, $\Delta_{ij} = \Delta_{ji}$, and $j \in N_i \Leftrightarrow i \in N_j$. Thus, we have $W(0) \in \Omega \Rightarrow W(t) \in U\ \forall t \geq 0$. Since for all trajectories starting in Ω one has $W(t) \in S\ \forall t \geq 0$ and $W(t) \in U\ \forall t \geq 0$, we conclude that $W(0) \in \Omega \Rightarrow W(t) \in S \cap U = \Omega\ \forall t \geq 0$, that is, set Ω is positively invariant.

Second, function $H_{\mathcal{V}}$ is clearly non-increasing along system trajectories starting in Ω, that is, $\dot{H}_{\mathcal{V}} \leq 0$ in Ω.

Third, all trajectories with initial conditions in Ω are bounded. Indeed, we have already shown that $\sum_{i=1}^{m} w_i(t) = c$ along system trajectories starting in Ω. This implies that weights remain within a bounded set: If, by contradiction, a weight could become arbitrarily positive large, another weight would become arbitrarily negative large (since the sum of weights is constant), and the measure of at least one power cell would vanish, which contradicts the fact that Ω is positively invariant.

Finally, by Theorem 3.4, $H_{\mathcal{V}}$ is continuously differentiable in Ω. Hence, by invoking the LaSalle invariance principle (see, for instance, [6]), generators' weights with initial conditions in Ω and evolving under (7) will converge asymptotically to the set of critical points of $H_{\mathcal{V}}$ in Ω, which is not empty as confirmed by Theorem 3.1. Indeed, by Theorem 3.1, we know that all vectors of weights yielding an equitable power diagram differ by a common translation. Thus, the largest invariant set of $H_{\mathcal{V}}$ in Ω contains only one point. This implies that for all $W(0) \in \Omega$ the limit $\lim_{t\to\infty} W(t)$ exists and it is equal to a vector of weights that yields an equitable power diagram.

The theorem then follows since c was chosen arbitrarily. ■

Some remarks are in order.

Remark 3.8 (Global convergence): By Theorem 3.7, convergence to an equitable power diagram is *global* with respect to S for *any* set of generators' distinct positions. Indeed, there is a very natural choice for the initial values of the weights. Assuming that at $t = 0$ agents are in $\mathcal{Q}$ and in distinct positions, each agent initializes the position of its generator to its physical position, and the corresponding weight to zero. Then, the initial partition is a Voronoi tessellation; since λ is positive on $\mathcal{Q}$, each initial cell has nonzero measure, and therefore $W(0) \in S$.

Remark 3.9 (Spatially distributed algorithm): The computation of the partial derivative of $H_{\mathcal{V}}$ with respect to the i-th weight only requires information from the agents with neighboring power cells. Therefore, the gradient descent law (7) is indeed a *spatially distributed* algorithm over the power-Delaunay graph. We mention that, in a power diagram, each power generator has an average number of neighbors less than or equal to six [13]; therefore, the computation of gradient (7) is scalable (on average) with the number of agents.

Remark 3.10 (Partitions with general measure assignment): The focus of this paper is on equitable partitions. Notice, however, that it is easy to extend the previous algorithm to obtain a spatially distributed (again, over the power-Delaunay graph) control law that provides *any* desired measure vector $(\lambda_{V_1}, \ldots, \lambda_{V_m})$. In particular, assume that we desire a partition such that $\lambda_{V_i} = \beta_i \lambda_{\mathcal{Q}}$, where $\beta_i \in (0, 1)$, $\sum_{i=1}^{m} \beta_i = 1$. If we redefine $H_{\mathcal{V}} : S \to \mathbb{R}_{>0}$ as

$$H_{\mathcal{V}}(W) \doteq \sum_{i=1}^{m} \frac{\beta_i^2}{\lambda_{V_i(W)}},$$

then, following the same steps as before, it is possible to show that under control law

$$\dot{w}_i = -\frac{\partial H_{\mathcal{V}}}{\partial w_i}(W) = \sum_{j \in N_i} \frac{1}{2\gamma_{ij}} \left(\frac{\beta_j^2}{\lambda_{V_j}^2} - \frac{\beta_i^2}{\lambda_{V_i}^2}\right) \int_{\Delta_{ij}} \lambda(x)\, dx,$$

the solution converges to a vector of weights that yields a power diagram with the property $\lambda_{V_i} = \beta_i \lambda_{\mathcal{Q}}$ (whose existence is guaranteed by Remark 3.2).

E. On the Use of Power Diagrams instead of Voronoi Diagrams

A natural question that arises is whether a similar result can be obtained by using Voronoi diagrams (of which power diagrams are a generalization). The answer is positive if we constrain λ to be uniform over $\mathcal{Q}$, but it is negative for general measures λ, as we briefly discuss next.

Indeed, when λ is *uniform* over $\mathcal{Q}$, an equitable Voronoi diagram always exists. We prove this result in a slightly more general setup.

Definition 3.11 (Unimodal Property): Let $\mathcal{Q} \subset \mathbb{R}^d$ be a bounded, measurable set (not necessarily convex). We say that $\mathcal{Q}$ enjoys the Unimodal Property if there exists a unit vector $v \in \mathbb{R}^d$ such that the following holds. For each $s \in \mathbb{R}$, define the slice $\mathcal{Q}^s \doteq \{x \in \mathcal{Q}, \quad v \cdot x = s\}$, and call $\psi(s) \doteq m_{d-1}(\mathcal{Q}^s)$ the $(d-1)$-dimensional Lebesgue measure of the slice. Then, the function ψ is unimodal. In other words, ψ attains its global maximum at a point $\bar{s}$, is increasing on $(-\infty, \bar{s}]$, and decreasing on $[\bar{s}, \infty)$.

When λ is uniform, the Unimodal Property (notice that every compact, convex set enjoys this property) turns out to be a sufficient condition for the existence of equitable Voronoi diagrams, as stated in the following theorem.

Theorem 3.12 (Existence of equitable Voronoi diagrams): If $\mathcal{Q} \subset \mathbb{R}^d$ is a bounded, measurable set satisfying the Unimodal Property and λ is uniform over $\mathcal{Q}$, then for every $m \geq 1$ there exists an equitable Voronoi diagram with m (Voronoi) generators.

Proof: See Appendix. ∎

Then, an equitable Voronoi diagram can be achieved by using a gradient descent law conceptually similar to the one discussed previously (the details are presented in [15]). We emphasize that the above existence result on equitable Voronoi diagrams seems to be new in the rich literature on Voronoi tessellations.

While an equitable Voronoi diagram always exists when λ is *uniform* over $\mathcal{Q}$, in general, for non-uniform λ, an equitable Voronoi diagram fails to exist, as the following counterexample shows.

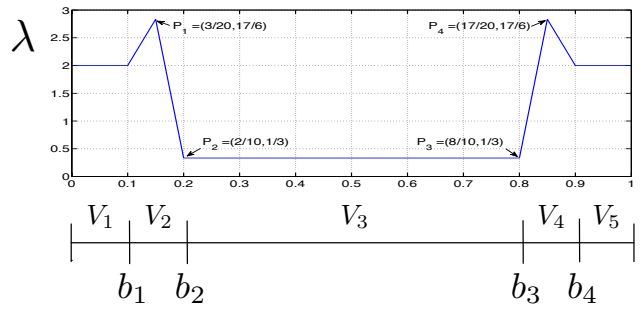

Fig. 3. Example of non-existence of an equitable Voronoi diagram on a line. The above tessellation is an equitable partition, but not a Voronoi diagram.

Example 3.13 (Existence problem on a line): Consider a one-dimensional Voronoi diagram. In this case a Voronoi cell is a half line or a line segment (called a Voronoi line). It is easy to notice that the boundary point between two adjacent Voronoi lines is the mid-point of the generators of those Voronoi lines. Consider the measure λ in Figure 3, whose support is the interval $[0, 1]$. Assume $m = 5$. Let b_i $(i = 1, \ldots, 4)$ be the position of the i-th leftmost boundary point and g_i be the position of the i-th leftmost generator $(i = 1, \ldots, 5)$. It is easy to verify that the only admissible configuration for the boundary points in order to obtain an equitable Voronoi diagram is the one depicted in Figure 3 (the V_i's represent the Voronoi lines). Now, by the Perpendicular Bisector Property, it must hold:

$$\begin{cases} g_3 - b_2 = b_2 - g_2 \\ g_4 - b_3 = b_3 - g_3 \end{cases}$$

Therefore, we would require $g_4 - g_2 = 2(b_3 - b_2) = 1.2$; this is impossible, since g_2 must belong to the interval $[0.1, 0.2]$ and g_4 must belong to the interval $[0.8, 0.9]$.

IV. Distributed Algorithms for Equitable Partitions with Additional Features

In this section we devise spatially distributed algorithms to compute convex, equitable partitions with additional features (e.g., convex and equitable partitions that are approximations of equitable Voronoi diagrams). In the next section we will provide in-depth motivations for the partitions we study here and two applications for the algorithms we devise.

From an algorithmic standpoint, the key idea we exploit is that an equitable power diagram can be obtained by just changing the values of the weights (while keeping the generators's positions *fixed*), as shown in Theorem 3.7. Thus, one can use the degrees of freedom given by the positions of the generators to achieve *additional objectives*. Specifically, we now assume that both generators' weights and positions obey a first order dynamical behavior

$$\begin{cases} \dot{w}_i = u_i^w, \\ \dot{g}_i = u_i^g. \end{cases}$$

Define the set

$$\tilde{S} \doteq \Big\{ \big((g_1, w_1), \ldots, (g_m, w_m) \big) \in (\mathcal{Q} \times \mathbb{R})^m \,|\, g_i \neq g_j \text{ for all } i \neq j, \text{and } \lambda_{V_i} > 0 \ \ \forall i \in I_m \Big\}.$$

The *primary* objective is to achieve a convex and equitable partition, and is captured, similarly as before, by the energy function $\tilde{H}_{\mathcal{V}} : \tilde{S} \to \mathbb{R}_{>0}$

$$\tilde{H}_{\mathcal{V}}(G_W) \doteq \sum_{i=1}^{m} \lambda_{V_i(G_W)}^{-1}.$$

Theorem 4.1 (Smoothness of $\tilde{H}_{\mathcal{V}}$): The following statements hold:

1) the function $\tilde{H}_{\mathcal{V}}$ is continuously differentiable on $\tilde{S}$, where for each $i \in \{1, \ldots, m\}$

$$\frac{\partial \tilde{H}_{\mathcal{V}}}{\partial g_i} = \sum_{j \in N_i} \Big(\frac{1}{\lambda_{V_j}^2} - \frac{1}{\lambda_{V_i}^2} \Big) \int_{\Delta_{ij}} \frac{(x - g_i)}{\gamma_{ij}} \lambda(x)\, dx,$$

$$\frac{\partial \tilde{H}_{\mathcal{V}}}{\partial w_i} = \sum_{j \in N_i} \Big(\frac{1}{\lambda_{V_j}^2} - \frac{1}{\lambda_{V_i}^2} \Big) \int_{\Delta_{ij}} \frac{1}{2\gamma_{ij}} \lambda(x)\, dx,$$

2) all critical points of $\tilde{H}_{\mathcal{V}}$ are generators' positions and weights that yield an equitable power diagram.

Proof: The proof of this theorem is very similar to the proof of Theorem 3.4; we omit it in the interest of brevity (the derivation of the partial derivative $\frac{\partial \tilde{H}_{\mathcal{V}}}{\partial g_i}$ can be found in [16]). ∎

Notice that the computation of the gradient in Theorem 4.1 is spatially distributed over the power-Delaunay graph. For brevity, we denote the vectors $\pm \frac{\partial \tilde{H}_{\mathcal{V}}}{\partial g_i}$ with $v_{\pm \partial \tilde{H}_i}$, respectively.

Three possible *additional* objectives are discussed in the remainder of this section.

A. On Approximating Equitable and Median Power Diagrams

We call a power diagram $\mathcal{V}(G_W)$ a *median power diagram* of $\mathcal{Q}$ with respect to the measure λ if the ordered set of generators' positions G is equal to the ordered set of generalized medians of the sets in $\mathcal{V}(G_W)$ with respect to λ, i.e., if for all $i \in \{1, \dots, m\}$

$$g_i = \arg\min_{g \in \mathbb{R}^2} \int_{V_i(G_W)} \|g - x\| \lambda(x). \tag{8}$$

If all weights are equal, a median power diagram is referred to as a median Voronoi diagram (since when all weights are equal a power diagram reduces to a Voronoi diagram). It is possible to show that a median power diagram always exists for any compact domain $\mathcal{Q}$ and density λ [17]. The minimization problem in equation (8) is a strictly convex optimization problem, and we denote its (unique) solution with g_i^*; the point g_i^* can be readily computed by using iterative methods, e.g., gradient-descent methods where the gradient is given by the formula

$$\frac{\partial \int_{V_i(G_W)} \|g - x\| \lambda(x)\, dx}{\partial g} = \int_{V_i(G_W)} \frac{g - x}{\|g - x\|} \lambda(x)\, dx.$$

The main motivation to study equitable and median power diagrams is that, as it will be discussed in Section IV-C, they can be used to approximate equitable and median Voronoi diagrams. Such diagrams play a special role in the context of dynamic vehicle routing for robotic networks (see Section V for further details). Remarkably, in a median Voronoi diagram the shape of the cells, under certain conditions, resembles that of regular polygons [17].

A natural candidate control law for the computation of an equitable and median power diagram (or at least for the computation of an approximation of it) is to let the positions of the generators move toward the medians of the corresponding regions of dominance, when this motion does not increase the disagreement between the measures of the cells (i.e., it does not make the time derivative of $\tilde{H}_{\mathcal{V}}$ positive). Accordingly, we introduce the following C^∞ saturation function

$$\Theta(x) \doteq \begin{cases} 0 & \text{for } x \leq 0\ , \\ \exp\left(-\frac{1}{(\beta x)^2}\right) & \text{for } x > 0, \quad \beta \in \mathbb{R}_{>0}; \end{cases}$$

moreover, we denote the vector $g_i^* - g_i$ as $v_{g_i^*, g_i}$. Then, each agent updates its own power generator according to the following control law defined over the set $\tilde{S}$:

$$\begin{aligned} \dot{w}_i &= -\frac{\partial \tilde{H}_{\mathcal{V}}}{\partial w_i}, \\ \dot{g}_i &= \alpha\, \Theta(v_{g_i^*, g_i} \cdot v_{-\partial \tilde{H}_i})\ v_{g_i^*, g_i}, \end{aligned} \tag{9}$$

where we assume that the partition $\mathcal{V}(G_W) = \{V_1, \dots, V_m\}$ is continuously updated, and where $\alpha \in \mathbb{R}_{>0}$. The term $\Theta(v_{g_i^*, g_i} \cdot v_{-\partial \tilde{H}_i})$ is needed to make the right-hand side of (9) compatible with the minimization of $\tilde{H}_{\mathcal{V}}$; in fact, due to the presence of $\Theta(\cdot)$, $\dot{g}_i = 0$ whenever $v_{g_i^*, g_i} \cdot v_{-\partial \tilde{H}_i} \leq 0$. In other words, g_i moves toward the median of its cell *if and only if* this motion is compatible with the minimization of $\tilde{H}_{\mathcal{V}}$. Note that the vector field in equation (9) is Lipschitz continuous and that the computation of the right-hand side of (9) is spatially distributed over the power-Delaunay graph.

As in many algorithms that involve the update of generators of Voronoi diagrams, it is possible (even though simulations show that this is "highly unlikely") that under control law (9) there exists a time t^* and $i, j \in I_m$ such that $g_i(t^*) = g_j(t^*)$. In such a case, either the power diagram is not defined (when $w_i(t^*) = w_j(t^*)$), or there is an empty cell (when $w_i(t^*) \neq w_j(t^*)$), and there is no obvious way to specify the behavior of control law (9) for these singularity points. Then, to make the set $\tilde{S}$ positively invariant, we have to make a *technical* modification to the update equation for the positions of the generators. The idea is to stop the positions of two generators when they are *close* and are on a *collision course*.

Define, for $\Delta \in \mathbb{R}_{>0}$, the set $M_i(G, \Delta) \doteq \{g_j \in G \,|\, \|g_j - g_i\| \leq \Delta,\ g_j \neq g_i\}$. In other words, M_i is the set of generators' positions within an (Euclidean) distance Δ from g_i. For $\delta \in \mathbb{R}_{>0}$, $\delta < \Delta$, define the gain function $\Psi(\rho, \vartheta) : [0, \Delta] \times [0, 2\pi] \mapsto \mathbb{R}_{\geq 0}$ (see Figure 4) as follows:

$$\begin{cases} \frac{\rho - \delta}{\Delta - \delta} & \text{if } \delta < \rho \leq \Delta \quad \text{and } 0 \leq \vartheta < \pi, \\ \frac{\rho - \delta}{\Delta - \delta}(1 + \sin\vartheta) - \sin\vartheta & \text{if } \delta < \rho \leq \Delta \quad \text{and } \pi \leq \vartheta \leq 2\pi, \\ 0 & \text{if } \rho \leq \delta \quad \text{and } 0 \leq \vartheta < \pi, \\ -\frac{\rho}{\delta} \sin\vartheta & \text{if } \rho \leq \delta \quad \text{and } \pi \leq \vartheta \leq 2\pi. \end{cases}$$

It is easy to see that $\Psi(\cdot, \cdot)$ is a continuous function on $[0, \Delta] \times [0, 2\pi]$ and it is globally Lipschitz there. Function $\Psi(\cdot, \cdot)$ has the following motivation. Let ρ be equal to $\|g_j - g_i\|$ (for some $g_j \in M_i(G, \Delta)$), and let v_x be a vector such that the tern $\{v_x, (g_j - g_i), v_x \times (g_j - g_i)\}$ is an orthogonal basis of $\mathbb{R}^3$, co-oriented with the standard basis. In Figure 4, v_x corresponds to the x axis and $g_j - g_i$ corresponds to the y axis. Finally, let ϑ be the angle between v_x and $v_{g_i^*, g_i}$, where $0 \leq \vartheta \leq 2\pi$. If $\rho \leq \delta$ and $0 \leq \vartheta < \pi$, then g_i is *close* to g_j and it is on a *collision course*, thus we set the gain to zero. Similar considerations hold for the other three cases; for example, if $\rho \leq \delta$ and $\pi \leq \vartheta \leq 2\pi$, the positions of the generators are *close*, but they are not on a collision course, thus the gain is positive. In practice, one should choose *small* values for the constants δ and Δ (e.g., in our simulations, they are set to values in the order of $\mathrm{diam}(\mathcal{Q}) \cdot 10^{-5}$).

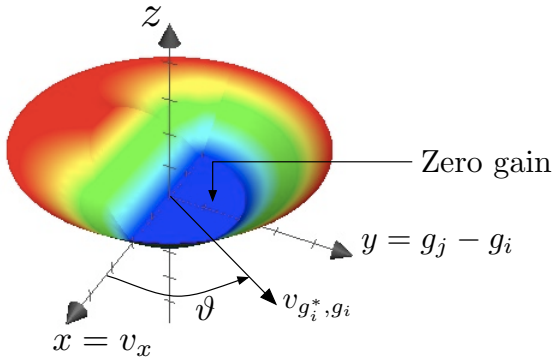

Fig. 4. Gain function used to avoid that the positions of two power generators can coincide.

Accordingly, we modify control law (9) as follows:

$$\begin{aligned}\dot{w}_i &= -\frac{\partial \tilde{H}_{\mathcal{V}}}{\partial w_i} \doteq u_i^{\text{med,w}},\\ \dot{g}_i &= \alpha\, \Theta(v_{g_i^*,g_i} \cdot v_{-\partial \tilde{H}_i})\; v_{g_i^*,g_i} \cdot \\ &\qquad \prod_{g_j \in M_i(G,\Delta)} \Psi\Big(\|g_j - g_i\|,\ \vartheta_{ij}\Big) \doteq u_i^{\text{med,g}},\end{aligned} \tag{10}$$

where ϑ_{ij} is the angle between v_x and $v_{g_i^*,g_i}$ (see the above discussion for the definition of v_x). If $M_i(G,\Delta)$ is the empty set, then we have an empty product, whose numerical value is 1. Note that the right-hand side of (10) is Lipschitz continuous, since it is a product of C^1 functions and Lipschitz continuous functions, and it can be still computed in a spatially distributed way (in fact, it only requires information from the agents with neighboring power cells, and whose generators' positions are within a distance Δ). Next theorem shows that algorithm (10) is still guaranteed to provide an equitable power diagram, whose closeness to an equitable and median power diagram will be discussed in section VI.

Theorem 4.2 (Correctness of algorithm (10)*):* Consider the vector field on $\tilde{S}$ defined by equation (10). Then generators' positions and weights starting at $t = 0$ at $G_W(0) \in \tilde{S}$ and evolving under (10) remain in $\tilde{S}$ and converge asymptotically to the set of critical points of the objective function $\tilde{H}_{\mathcal{V}}$ (i.e., to the set of vectors of generators' positions and weights that yield an equitable power diagram).

Proof: The proof is virtually identical to the one of Theorem 3.7, and we omit it in the interest of brevity. We only notice that $\tilde{H}_{\mathcal{V}}$ is non-increasing along system trajectories

$$\dot{\tilde{H}}_{\mathcal{V}} = \sum_{i=1}^m \frac{\partial \tilde{H}_{\mathcal{V}}}{\partial g_i}\cdot \dot{g}_i + \frac{\partial \tilde{H}_{\mathcal{V}}}{\partial w_i}\dot{w}_i = \sum_{i=1}^m \underbrace{\frac{\partial \tilde{H}_{\mathcal{V}}}{\partial g_i}\cdot \dot{g}_i}_{\leq 0} - \Big(\frac{\partial \tilde{H}_{\mathcal{V}}}{\partial w_i}\Big)^2 \leq 0.$$

Moreover, the components of vector field (10) for the position of each generator are either zero or point toward $\mathcal{Q}$ (since the median of a cell must be within $\mathcal{Q}$); therefore, each generator will remain within the compact set $\mathcal{Q}$. ∎

B. On Approximating Equitable Voronoi Diagrams

As we will show in Section V, in some applications it could be preferable to have power diagrams as close as possible to Voronoi diagrams (recall that when λ is not uniform, an equitable Voronoi diagram could fail to exist). The objective of obtaining a power diagram close to a Voronoi diagram can be translated in the minimization of the function $K : \mathbb{R}^m \to \mathbb{R}_{\geq 0}$:

$$K(W) \doteq \frac{1}{2}\sum_{i=1}^m w_i^2;$$

when $w_i = 0$ for all $i \in I_m$, one has $K(W) = 0$ and the corresponding power diagram coincides with a Voronoi diagram. To include the minimization of the additional objective K it is natural to consider, instead of control law (7), the following update law for the weights:

$$\dot{w}_i = -\frac{\partial H_{\mathcal{V}}}{\partial w_i} - \frac{\partial K}{\partial w_i} = -\frac{\partial H_{\mathcal{V}}}{\partial w_i} - w_i. \tag{11}$$

However, $H_{\mathcal{V}}$ (defined in equation (3)) is no longer a valid Lyapunov function for control law (11). The idea, then, is to let the positions of the generators move so that $\frac{\partial \tilde{H}_{\mathcal{V}}}{\partial g_i}\cdot \dot{g}_i - \frac{\partial \tilde{H}_{\mathcal{V}}}{\partial w_i} w_i = 0$. In other words, the dynamics of generators' positions are used to compensate the effect of the term $-w_i$ (present in the weights' dynamics) on the time derivative of $\tilde{H}_{\mathcal{V}}$.

Thus, we set up the following control law, with ε_1, ε_2 and ε_3 positive *small* constants, $\varepsilon_2 > \varepsilon_1$,

$$\begin{aligned}\dot{w}_i &= -\frac{\partial \tilde{H}_{\mathcal{V}}}{\partial w_i} - w_i\, \text{sat}_{\varepsilon_1,\varepsilon_2}\Big(\|v_{\partial \tilde{H}_i}\|\Big)\, \text{sat}_{0,\varepsilon_3}\Big(\text{dist}(g_i, \partial V_i)\Big),\\ \dot{g}_i &= w_i \frac{\partial \tilde{H}_{\mathcal{V}}}{\partial w_i}\frac{v_{\partial \tilde{H}_i}}{\|\ v_{\partial \tilde{H}_i}\|^2}\, \text{sat}_{\varepsilon_1,\varepsilon_2}\Big(\|v_{\partial \tilde{H}_i}\|\Big)\, \text{sat}_{0,\varepsilon_3}\Big(\text{dist}(g_i, \partial V_i)\Big).\end{aligned} \tag{12}$$

The gain $\text{sat}_{\varepsilon_1,\varepsilon_2}\Big(\|v_{\partial \tilde{H}_i}\|\Big)$ is needed to make the right-hand side of (12) Lipschitz continuous, while the gain $\text{sat}_{0,\varepsilon_3}\Big(\text{dist}(g_i, \partial V_i)\Big)$ ensures that generators' positions stay within $\mathcal{Q}$. Notice that the computation of the right-hand side of (12) is spatially distributed over the power-Delaunay graph.

As before, it is possible (even though simulations show that this is "highly unlikely") that under control law (12) there exists a time t^* and $i, j \in I_m$ such that $g_i(t^*) = g_j(t^*)$. Thus, similarly as before, we modify the update equations (12) as follows

$$\begin{aligned}\dot{w}_i &= -\frac{\partial \tilde{H}_{\mathcal{V}}}{\partial w_i} - w_i\, \text{sat}_{\varepsilon_1,\varepsilon_2}\Big(\|v_{\partial \tilde{H}_i}\|\Big)\, \text{sat}_{0,\varepsilon_3}\Big(\text{dist}(g_i, \partial V_i)\Big)\cdot \\ &\qquad \prod_{g_j \in M_i(G,\Delta)} \Psi\Big(\|g_j - g_i\|,\ \vartheta_{ij}\Big) \doteq u_i^{\text{vor,w}},\\ \dot{g}_i &= w_i \frac{\partial \tilde{H}_{\mathcal{V}}}{\partial w_i}\frac{v_{\partial \tilde{H}_i}}{\|\ v_{\partial \tilde{H}_i}\|^2}\, \text{sat}_{\varepsilon_1,\varepsilon_2}\Big(\|v_{\partial \tilde{H}_i}\|\Big)\, \text{sat}_{0,\varepsilon_3}\Big(\text{dist}(g_i, \partial V_i)\Big)\cdot \\ &\qquad \prod_{g_j \in M_i(G,\Delta)} \Psi\Big(\|g_j - g_i\|,\ \vartheta_{ij}\Big) \doteq u_i^{\text{vor,g}},\end{aligned} \tag{13}$$

where ϑ_{ij} is defined as in Section IV-A, with $w_i \frac{\partial \tilde{H}_{\mathcal{V}}}{\partial w_i} v_{\partial \tilde{H}_i}$ playing the role of $v_{g_i^*,g_i}$.

Next theorem shows that algorithm (13) is still guaranteed to provide an equitable power diagram, whose closeness to an equitable Voronoi diagram will be discussed in section VI.

Theorem 4.3 (Correctness of algorithm (13)*):* Consider the vector field on $\tilde{S}$ defined by equation (13). Then generators' positions and weights starting at $t = 0$ at $G_W(0) \in \tilde{S}$ and evolving under (13) remain in $\tilde{S}$ and converge asymptotically to the set of critical points of the objective function $\tilde{H}_{\mathcal{V}}$ (i.e., to the set of vectors of generators' positions and weights that yield an equitable power diagram).

Proof: Consider $\tilde{H}_{\mathcal{V}}$ as a Lyapunov function candidate. First, we prove that set $\tilde{S}$ is positively invariant with respect to (13). Indeed, by definition of (13), we have $g_i \neq g_j$ for $i \neq j$ for all $t \geq 0$ (therefore, the power diagram is always well defined). Moreover, it is straightforward to see that $\dot{\tilde{H}}_{\mathcal{V}} \leq 0$. Therefore, it holds

$$\lambda_{V_i(G_W(t))}^{-1} \leq \tilde{H}_{\mathcal{V}}(G_W(t)) \leq \tilde{H}_{\mathcal{V}}(G_W(0)), \quad i \in I_m,\ t \geq 0.$$

Since the measures of the power cells depend continuously on the generators' positions and weights, we conclude that the measures of all power cells will be bounded away from zero. Furthermore, since $\dot{g}_i = 0$ on the boundary of $\mathcal{Q}$ for all $i \in I_m$, each generator will remain within the compact set $\mathcal{Q}$. Thus, the generators' positions and weights will belong to $\tilde{S}$ for all $t \geq 0$, that is, $G_W(t) \in \tilde{S}\ \forall t \geq 0$.

Second, as stated before, $\tilde{H}_{\mathcal{V}} \colon \tilde{S} \to \mathbb{R}_{>0}$ is non-increasing along system trajectories, i.e., $\dot{\tilde{H}}_{\mathcal{V}} \leq 0$ in $\tilde{S}$.

Third, all trajectories with initial conditions in $\tilde{S}$ are bounded. Indeed, we have already shown that each generator remains within the compact set $\mathcal{Q}$ under control law (13). As far as the weights are concerned, we start by noticing that the time derivative of the sum of the weights is

$$\frac{\partial \sum_{i=1}^m w_i}{\partial t} = -\sum_{i=1}^m w_i \mathrm{sat}_{\varepsilon_1,\varepsilon_2}\Big(\|v_{\partial \tilde{H}_i}\|\Big)\, \mathrm{sat}_{0,\varepsilon_3}\Big(\mathrm{dist}(g_i, \partial V_i)\Big) \cdot \prod_{g_j \in M_i(G,\Delta)} \Psi\Big(\|g_j - g_i\|,\ \vartheta_{ij}\Big),$$

since, similarly as in the proof of Theorem 3.7, $\sum_{i=1}^m \frac{\partial \tilde{H}_{\mathcal{V}}}{\partial w_i} = 0$. Moreover, the magnitude of the difference between any two weights is bounded by a constant $B \in \mathbb{R}_{>0}$, that is,

$$|w_i - w_j| \leq B \quad \text{for all } i,\, j \in I_m. \tag{14}$$

In fact, if, by contradiction, the magnitude of the difference between any two weights could become arbitrarily large, the measure of at least one power cell would vanish, since the positions of the generators are confined within $\mathcal{Q}$. Assume, for the sake of contradiction, that weights' trajectories are unbounded. This means that

$$\forall R > 0 \quad \exists t \geq 0 \text{ and } \exists\, j \in I_m \quad \text{such that} \quad |w_j(t)| \geq R.$$

For simplicity, assume that $w_i(0) = 0$ for all $i \in I_m$ (the extension to arbitrary initial conditions in $\tilde{S}$ is straightforward). Choose $R = 2mB$ and let t_2 be the time instant such that $|w_j(t_2)| = R$, for some $j \in I_m$. Without loss of generality, assume that $w_j(t_2) > 0$. Because of constraint (14), we have $\sum_{i=1}^m w_i(t_2) \geq \frac{B}{2} m(3m+1)$. Let t_1 be the *last* time before t_2 such that $w_j(t) = mB$; because of continuity of trajectories, t_1 is well defined. Then, because of constraint (14), we have (i) $\sum_{i=1}^m w_i(t_1) \leq \frac{B}{2} m(3m-1) < \sum_{i=1}^m w_i(t_2)$, *and* (ii) $\frac{\partial \sum_{i=1}^m w_i(t)}{\partial t} \leq 0$ for $t \in [t_1, t_2]$ (since $w_j(t) \geq mB$ for all $t \in [t_1, t_2]$ and equation (14) implies $\min_{i \in I_m} w_i(t) > 0$ for all $t \in [t_1, t_2]$); thus, we get a contradiction.

Finally, by Theorem 4.1, $\tilde{H}_{\mathcal{V}}$ is continuously differentiable in $\tilde{S}$. Hence, by the LaSalle invariance principle, under the descent flow (13) the generators' positions and weights will converge asymptotically to the set of critical points of $\tilde{H}_{\mathcal{V}}$, which is not empty as confirmed by Theorem 3.1. ■

C. On Approximating Equitable and Median Voronoi Diagrams

In many applications, it is desirable to obtain approximations of equitable and median Voronoi diagrams. For example, such diagrams are intimately related to the solution of the well-known dynamic vehicle routing problem, where the objective is to plan optimal multi-vehicle routes to perform tasks that are generated over time by an exogenous process (see Section V-A). Moreover, as the number of generators increases, equitable and median Voronoi diagrams assume an hexagonal honeycomb structure where each cell has the same area (assuming that λ is uniform) [17]. This fact has interesting applications in the context of wireless ad hoc networks (see Section V-B). In general, equitable and median Voronoi diagrams provide subregions having the *same measure* and whose shapes show *circular symmetry*.

In light of Theorems 4.2 and 4.3, it is possible to obtain a power diagram approximating an equitable and median Voronoi diagram by combining control laws (10) and (13). In particular, we set up the following spatially distributed control law:

$$\begin{aligned} \dot{w}_i &= u_i^{\text{med,w}} + u_i^{\text{vor,w}}, \\ \dot{g}_i &= u_i^{\text{med,g}} + u_i^{\text{vor,g}}. \end{aligned} \tag{15}$$

Next theorem shows that algorithm (15) is still guaranteed to provide an equitable power diagram, whose closeness to an equitable and median Voronoi diagram will be discussed in section VI.

Theorem 4.4 (Correctness of algorithm (15)*):* Consider the vector field on $\tilde{S}$ defined by equation (15). Then generators' positions and weights starting at $t = 0$ at $G_W(0) \in \tilde{S}$ and evolving under (15) remain in $\tilde{S}$ and converge asymptotically to the set of critical points of the objective function $\tilde{H}_{\mathcal{V}}$ (i.e., to the set of vectors of generators' positions and weights that yield an equitable power diagram).

Proof: The proof of this theorem is a straightforward combination of the proofs of Theorems 4.2 and 4.3. ■

To the best of our knowledge, this is the first algorithm to compute approximations of equitable and median Voronoi diagrams. We observe that one can obtain a power diagram approximating an equitable and *centroidal* Voronoi diagram by simply replacing the motion toward the median with a motion toward the centroid (see [18] for an introduction to centroidal Voronoi diagrams and for a discussion on their practical importance).

We conclude this section with a remark about the validity of our algorithms when the environment is more general than a compact, convex subset of $\mathbb{R}^2$.

Remark 4.5 (General environments): Both the existence theorem 3.1 and the convergence theorems 3.7, 4.2, 4.3, and 4.4 indeed hold for any compact, connected environment in $\mathbb{R}^d$. In particular, the existence theorem 3.1 and the convergence theorems 3.7, 4.2, 4.3, and 4.4 hold even if the environment $\mathcal{Q}$ has "holes", namely it is not simply connected or it has nontrivial homology. In fact, in the presence of "holes", the map associating weight vectors with vectors of measures is still surjective to the boundary of the measure simplex; then, since in the "weight space" there are no "holes" and for the measure simplex one needs to prove surjectivity only on the boundary, the argument relying on the topological degree carries over. Of course, if the environment is non-convex, some of the power cells might be non-convex.

V. Applications

In this section we present two application domains for the partitioning algorithms presented in Sections III and IV. The discussion of these applications serves also to illustrate a systematic approach to devise spatially distributed control policies for the class of multi-robot coordination problems that admit equitable partitioning policies as a solution.

A. Dynamic Vehicle Routing

An important application of the algorithms presented in this paper lies in the context of dynamic vehicle routing problems, where the objective is to plan optimal multi-vehicle routes to perform tasks that are generated over time by an exogenous process.

Specifically, we consider the following general model of dynamic vehicle routing problem, know in the literature as the m-vehicle Dynamic Traveling Repairman Problem (m-DTRP) [1]: m vehicles operating in a bounded environment $\mathcal{Q}$ and traveling with bounded velocity must service demands whose time of arrival, location and on-site service are stochastic. The objective is to find a routing policy to service demands over an infinite horizon that minimizes the expected system time (wait plus service) of the demands. There are many practical settings in which such problem arises; e.g., any distribution system which receives orders in real time and makes deliveries based on these orders (e.g., courier services) is a clear candidate. Surveillance missions where a team of unmanned aerial vehicles must visit locations of events dynamically originating within a protected environment is a second important example.

The key concept linking the algorithms presented in this paper with routing policies for the m-DTRP is that of π-*partitioning policy*. Given a single-vehicle routing policy π for the 1-DTRP (e.g., a first-come first-served policy) and m vehicles, a π-partitioning policy is a multi-vehicle policy such that 1) the environment $\mathcal{Q}$ is partitioned according to some partitioning scheme into m openly disjoint subregions $\mathcal{Q}_i$, $i \in \{1, \dots, m\}$, whose union is $\mathcal{Q}$, 2) one vehicle is assigned to each subregion (thus, there is a one-to-one correspondence between vehicles and subregions), and 3) each vehicle executes the single-vehicle policy π to service demands that fall within its own subregion.

The following two results, valid under the assumption that the measure λ is uniform, characterize the optimality of two types of π-partitioning policies [19].

Theorem 5.1 (Optimality of π-partitioning policies): Assume π^* is a single-vehicle optimal policy for the 1-DTRP. For m vehicles,

1) a π-partitioning policy using a partitioning scheme whereby $\{Q_i\}_{i=1}^m$ is an equitable and median Voronoi diagram is an almost optimal policy in light load (i.e., when the arrival rate of demands is "small") and an optimal policy in heavy load (i.e., when the arrival rate of demands is "large");
2) a π-partitioning policy using a partitioning scheme whereby $\{Q_i\}_{i=1}^m$ is an equitable partition is an optimal policy in heavy load.

Light and heavy load can be defined more formally in terms of *load factor*; we refer the interested reader to [19]. The *almost* optimality in light load is to be understood as follows: in light load, the average system time becomes a function of the loitering locations of the vehicles [1], and the generalized median locations that give rise to a median Voronoi diagram correspond to local minima or saddle points of this function. One can state a similar set of results for the general case where the measure λ is *not* uniform; the details are omitted in the interest of brevity and can be found in [19].

In light of Theorem 5.1, a *systematic* approach to obtain multi-vehicle routing policies with provable performance guarantees and amenable to distributed implementation is to combine the partitioning algorithms presented in this work with the optimal single-vehicle routing policies developed in [19]. Note that an equitable power diagram guarantees optimal performance in heavy load, while an equitable and median Voronoi diagram provides almost optimal performance in light load and optimal performance in heavy load.

Accordingly, the first step to obtain a spatially distributed multi-vehicle routing policy is to associate each vehicle i with a *power generator* (g_i, w_i), which is an artificial variable locally controlled by the i-th vehicle. We define the region of dominance for vehicle i as the power cell $V_i = V_i(G_W)$, where $G_W = \Big((g_1, w_1), \dots, (g_m, w_m)\Big)$ (see Figure 5). Then, each vehicle applies to its generator one of the previous partitioning algorithms (e.g. control law (15), if one desires performance guarantees in both light and heavy load), while simultaneously performing within its own region of dominance the optimal single-vehicle routing policies described in [19] (see Figure 5).

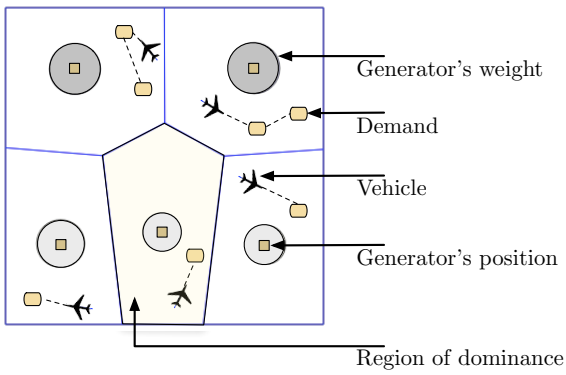

Fig. 5. Vehicles, demands, power generators, and regions of dominance. Radii of light (dark) grey circles represent the magnitudes of positive (negative) weights.

B. Hybrid Networks

A wireless ad-hoc network consists of a group of nodes which communicate with each other over a wireless channel without any centralized control; in situations where there is no fixed infrastructure, for example, battlefields, catastrophe control, etc., wireless ad hoc networks become valuable alternatives to fixed infrastructure networks for nodes to communicate with each other. To improve throughput capacity, one

can add a sparse network of more sophisticated nodes (supernodes) providing long-distance communication. Assuming that normal nodes are independently and uniformly located in the environment, supernodes should divide the area according to a hexagonal tessellation [3], where all hexagonal cells have the same area. One can design a spatially distributed algorithm to deploy the supernodes into an equitable partition with hexagonal cells as follows. Each supernode is associated with a power generator, and we let the physical position of each supernode coincide with the position of its corresponding power generator. Then, setting $\lambda \equiv 1$, each supernode updates its power generator (and, hence, its physical position) according to algorithm (15). Since, when λ is uniform, algorithm (15) provides equitable partitions with almost-hexagonal polygonal cells (see Section IV-C), the supernodes will deploy themselves into a near optimal configuration.

VI. Simulations and Discussion

The algorithms presented in Section IV are designed to provide approximations of equitable power diagrams with additional features (note that equitability of the partition is always guaranteed). In this section we study by simulation the *quality* of such approximations. Due to space constraints, we focus on algorithm (15), which is designed to provide approximations of equitable and median Voronoi diagrams (again, equitability is a guaranteed property).

We introduce three criteria to judge, respectively, *closeness* to a median power diagram, *closeness* to a Voronoi diagram, and circular symmetry of a partition (in particular, closeness to partitions with hexagonal cells).

A. Closeness to Median Power Diagrams

Consider a power diagram $\mathcal{V}(G_W) = (V_1(G_W), \dots, V_m(G_W))$, and let g_i^* be the median of power cell V_i, $i \in I_m$. We consider the following metric to measure closeness to a median power diagram:

$$d_{\text{med}} \doteq \frac{1}{m} \sum_{i=1}^{m} \frac{\|g_i^* - g_i\|}{\text{diam}(V_i(G_W))}.$$

Clearly, d_{med} equals zero when a power diagram coincides with a median power diagram. We will also refer to d_{med} as the *median defect* of a power diagram.

B. Closeness to Voronoi Diagrams

In a Voronoi diagram, the intersection between the bisector of two neighboring generators g_i and g_j and the line segment joining g_i and g_j is the midpoint $g_{ij}^{\text{vor}} \doteq (g_i + g_j)/2$. Then, if we define g_{ij}^{pow} as the intersection, in a power diagram, between the bisector of two neighboring generators (g_i, w_i) and (g_j, w_j) and the line segment joining their positions g_i and g_j, a possible way to measure the *distance* d_{vor} of a power diagram from a Voronoi diagram is the following:

$$d_{\text{vor}} \doteq \frac{1}{2N} \sum_{i=1}^{m} \sum_{j \in N_i} \frac{\|g_{ij}^{\text{pow}} - g_{ij}^{\text{vor}}\|}{0.5\, \gamma_{ij}},$$

where N is the number of neighboring relationships and, as before, $\gamma_{ij} = \|g_j - g_i\|$. Clearly, if a power diagram is also a Voronoi diagram (i.e., if all weights are equal), $d_{\text{vor}} = 0$. We will also refer to d_{vor} as the *Voronoi defect* of a power diagram.

C. Circular Symmetry of a Partition

A quantitative manifestation of circular symmetry is the well-known *isoperimetric inequality*, which states that among all planar objects of a given perimeter the circle encloses the largest area. More precisely, given a planar region $\mathcal{Q}$ with perimeter $p_{\mathcal{Q}}$ and area $|\mathcal{Q}|$, then $p_{\mathcal{Q}}^2 - 4\pi|\mathcal{Q}| \geq 0$, and equality holds if and only if $\mathcal{Q}$ is a circle. Then, we can define the *isoperimetric ratio* as follows: $R_{\mathcal{Q}} = \frac{4\pi|\mathcal{Q}|}{p_{\mathcal{Q}}^2}$; by the isoperimetric inequality $R_{\mathcal{Q}} \leq 1$, with equality only for circles. Interestingly, for a regular n-gon the isoperimetric ratio R_n is $R_n = \frac{\pi}{n \tan \frac{\pi}{n}}$, which converges to 1 for $n \to \infty$. Accordingly, given a partition $\{\mathcal{Q}_i\}_{i=1}^m$, we consider as a measure for the circular symmetry of a partition the "average isoperimetric ratio" $R_{\{\mathcal{Q}_i\}_{i=1}^m} \doteq \frac{1}{m} \sum R_{\mathcal{Q}_i}$.

D. Simulation Results

All simulations are performed on a machine with a 2.4GHz Intel Core Duo processor and 4GB of RAM. The code is written in C++ and makes use of the C++ Computational Geometry Algorithms Library `CGAL`[2].

We apply algorithm (15) to ten power generators, whose initial positions are independently and uniformly distributed in the unit square $\mathcal{Q}$, and whose weights are initialized to zero. Time is discretized with a step $dt = 0.01$, and each simulation run consists of 600 iterations (thus, the final time is $T = 6$). Define the area error ϵ as $\epsilon \doteq (\lambda_{i_{\text{max}}} - \lambda_{i_{\text{min}}})/(\lambda_{\mathcal{Q}}/m)$, evaluated at time $T = 6$; in the definition of ϵ, $\lambda_{i_{\text{max}}}$ is the measure of the power cell with maximum measure and $\lambda_{i_{\text{min}}}$ is the measure of the power cell with minimum measure. We perform two sets of simulations. In the first set of simulations we consider a measure λ *uniform* over $\mathcal{Q}$, i.e., $\lambda \equiv 1$, while in the second set of simulations we consider a measure λ that follows a gaussian distribution, namely $\lambda(x, y) = e^{-5((x-0.8)^2+(y-0.8)^2)}$, $(x, y) \in \mathcal{Q}$, whose peak is at the top-right corner of the unit square. Each set of simulations consists of 50 simulation runs.

Tables I and II summarize simulation results for the uniform λ (λ=unif) case and for the gaussian λ (λ=gauss) case. In both cases, average and worst-case values of the area error ϵ, median defect d_{med}, Voronoi defect d_{vor}, and average isoperimetric ratio $R_{\{\mathcal{Q}_i\}_{i=1}^m}$ are with respect to 50 simulation runs. Notice that for both measures, after 600 iterations, (i) the worst-case area error is below 8%, (ii) the worst-case values of d_{vor} and d_{med} are very small, and, finally, (iii) cells have, approximately, the circular symmetry of squares (since $R_4 \approx 0.78$). Hence, simulation results show that algorithm (15) consistently provides very good approximations of equitable and median Voronoi diagrams. In both cases, the positions of the generators always stayed within their corresponding power

[2] `CGAL` is freely available for academic research use at `http://www.cgal.org/`.

TABLE I
AVERAGE PERFORMANCE OF CONTROL LAW (15).

λ	$\mathbb{E}[\epsilon]$	$\mathbb{E}[d_{\text{med}}]$	$\mathbb{E}[d_{\text{vor}}]$	$\mathbb{E}[R_{\{\mathcal{Q}_i\}}]$
unif	0.1%	2.8%	0.2%	0.75
gauss	1.4%	2.2%	1%	0.74

TABLE II
WORST-CASE PERFORMANCE OF CONTROL LAW (15).

λ	$\max \epsilon$	$\max d_{\text{med}}$	$\max d_{\text{vor}}$	$\min R_{\{\mathcal{Q}_i\}}$
unif	0.5%	3.4%	0.4%	0.7
gauss	7.7%	3.8%	2.7%	0.7

cells. Figure 6 shows some typical equitable partitions that are achieved with control law (15) (the number of generators is 10 and $\mathcal{Q}$ is an irregular convex polygon).

We have also performed extensive simulations of algorithms (7), (10), and (13). In general, these algorithms provide equitable power diagrams with "long and skinny" power cells (i.e., with a low value of the average isoperimetric ratio $R_{\{\mathcal{Q}_i\}_{i=1}^m}$). Moreover, they sometimes lead to partitions where some of the generators' positions are outside their corresponding power cells.

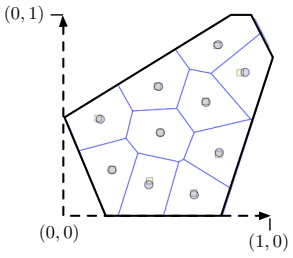

(a) Typical equitable partition of $\mathcal{Q}$ for $\lambda(x, y) = 1$ ($\epsilon = 0.01\%$, $d_{\text{med}} = 2.3\%$, $d_{\text{vor}} = 0.4\%$, $R_{\{\mathcal{Q}_i\}}$=0.76).

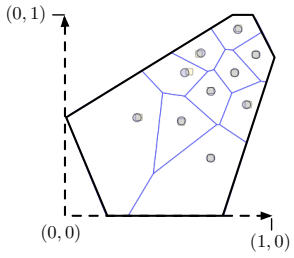

(b) Typical equitable partition of $\mathcal{Q}$ for $\lambda(x, y) = e^{-5((x-0.8)^2+(y-0.8)^2)}$ ($\epsilon = 0.7\%$, $d_{\text{med}} = 2.1\%$, $d_{\text{vor}} = 2.2\%$, $R_{\{\mathcal{Q}_i\}} = 0.74$).

Fig. 6. Typical equitable partitions achieved by using control law (15) after 1500 iterations. The squares represent the positions of the generators, while the circles represent the medians. Notice how each bisector intersects the line segment joining the two corresponding power neighbors almost at the midpoint; hence both partitions are very close to Voronoi partitions.

VII. CONCLUSION

We have presented provably correct, spatially distributed algorithms for the computation of a convex and equitable partition of a convex environment. We have also considered the issue of computing convex and equitable partitions with additional features (e.g., convex and equitable partitions that are approximations of equitable and median Voronoi diagrams). Finally, we have discussed how the algorithms devised in this paper represent a building block instrumental to design spatially distributed control policies for several multi-agent coordination problems, including dynamic vehicle routing, and deployment of wireless networks.

This paper leaves numerous important extensions open for further research. First, all the algorithms we proposed are synchronous: we plan to devise algorithms that are amenable to asynchronous implementation. Second, it is of interest to consider the setting where the measure λ evolves in time. This would require a characterization of the convergence rate of our algorithms. Finally, to assess the closed-loop robustness and the feasibility of our algorithms, we plan to implement them on a network of unmanned aerial vehicles.

ACKNOWLEDGMENTS

We gratefully acknowledge Professor A. Bressan's contribution in deriving the proof of Theorem 3.12. The research leading to this work was partially supported by the National Science Foundation through grants #0705451 and #0705453 and by the Office of Naval Research through grant # N00014-07-1-0721.

REFERENCES

[1] D. J. Bertsimas and G. J. van Ryzin, "Stochastic and dynamic vehicle routing in the Euclidean plane with multiple capacitated vehicles," *Operations Research*, vol. 41, no. 1, pp. 60–76, 1993.

[2] O. Baron, O. Berman, D. Krass, and Q. Wang, "The equitable location problem on the plane," *European Journal of Operational Research*, vol. 183, no. 2, pp. 578–590, 2007.

[3] B. Liu, Z. Liu, and D. Towsley, "On the capacity of hybrid wireless networks," in *IEEE INFOCOM 2003*, San Francisco, CA, Apr. 2003, pp. 1543–1552.

[4] J. Carlsson, D. Ge, A. Subramaniam, A. Wu, and Y. Ye, "Solving min-max multi-depot vehicle routing problem," in *Lectures on Global Optimization*, P. Pardalos and T. Coleman, Eds. American Mathematical Society and Field Institute, 2007, pp. 31–46.

[5] S. P. Lloyd, "Least-squares quantization in PCM," *IEEE Transactions on Information Theory*, vol. 28, no. 2, pp. 129–137, 1982.

[6] F. Bullo, J. Cortés, and S. Martínez, *Distributed Control of Robotic Networks*, ser. Applied Mathematics Series. Princeton University Press, 2009, available at `http://www.coordinationbook.info`.

[7] J. Cortés, "Coverage optimization and spatial load balancing by robotic sensor networks," *IEEE Transactions on Automatic Control*, vol. 55, no. 3, pp. 749–754, 2010.

[8] A. Kwok and S. Martínez, "Deployment algorithms for a power-constrained mobile sensor network," *International Journal of Robust and Nonlinear Control*, vol. 20, no. 7, pp. 725–842, 2010.

[9] J. Carlsson, B. Armbruster, and Y. Ye, "Finding equitable convex partitions of points in a polygon efficiently," *ACM Transactions on Algorithms*, no. 6, 2010.

[10] I. Chavel, *Eigenvalues in Riemannian Geometry*. New York, NY: Academic Press, 1984.

[11] A. Okabe, B. Boots, K. Sugihara, and S. N. Chiu, *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*. New York, NY: John Wiley & Sons, 2000.

[12] H. Imai, M. Iri, and K. Murota, "Voronoi diagram in the Laguerre geometry and its applications," *SIAM Journal on Computing*, vol. 14, no. 1, pp. 93–105, 1985.

[13] F. Aurenhammer, "Power diagrams: properties, algorithms and applications," *SIAM Journal on Computing*, vol. 16, no. 1, pp. 78–96, 1987.

[14] J. Cortés, S. Martínez, and F. Bullo, "Spatially-distributed coverage optimization and control with limited-range interactions," *ESAIM. Control, Optimisation & Calculus of Variations*, vol. 11, pp. 691–719, 2005.

[15] M. Pavone, E. Frazzoli, and F. Bullo, "Decentralized algorithms for stochastic and dynamic vehicle routing with general demand distribution," in *Proc. IEEE Conf. on Decision and Control*, New Orleans, LA, Dec. 2007, pp. 4869–4874.

[16] ——, "Distributed policies for equitable partitioning: Theory and applications," in *Proc. IEEE Conf. on Decision and Control*, Cancun, Mexico, Dec. 2008, pp. 4191–4197.

[17] E. Zemel, "Probabilistic analysis of geometric location problems," *Annals of Operations Research*, vol. 1, no. 3, pp. 215–238, October 1984.

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

[19] M. Pavone, E. Frazzoli, and F. Bullo, "Adaptive and distributed algorithms for vehicle routing in a stochastic and dynamic environment," *IEEE Transactions on Automatic Control*, to appear in the 2011 June issue, available at http://arxiv.org/abs/0903.3624.
[20] V. Guillemin and A. Pollack, *Differential Topology*. Prentice-Hall, 1974, republished by the American Mathematical Society 2010.
[21] A. Hatcher, *Algebraic Topology*. Cambridge, U.K.: Cambridge University Press, 2002, available at www.math.cornell.edu/∼hatcher/AT/ATpage.html.
[22] A. Bressan, *Personal Communication*, 2008.

APPENDIX

A. Definition of Degree of a Map

We start with the simplest definition of degree of a map. Let $f : X \to Y$ be a smooth map between connected compact manifolds X and Y of the same dimension, and let $p \in Y$ be a regular value for f (regular values abound due to Sard's lemma [20]). Since X is compact, $f^{-1}(p) = \{x_1, \dots, x_m\}$ is a finite set of points, and since p is a regular value it means that $f_{U_i} : U_i \mapsto f(U_i)$ is a local diffeomorphism, where U_i is a suitable open neighborhood of x_i. Diffeomorphisms can be either orientation preserving or orientation reversing. Let d^+ be the number of points x_i in $f^{-1}(p)$ at which f is orientation preserving (i.e., $\det(\mathrm{Jac}(f)) > 0$, where $\mathrm{Jac}(f)$ is the Jacobian matrix of f) and d^- be the number of points in $f^{-1}(p)$ at which f is orientation reversing (i.e., $\det(\mathrm{Jac}(f)) < 0$). Since X is connected, it can be proved that the number $d^+ - d^-$ is independent of the choice of $p \in Y$ and one defines the degree of f as $d^+ - d^-$. The degree can be also defined for a *continuous* map $f : X \to Y$ among connected oriented topological manifolds of the same dimensions, this time using homology groups or the local homology groups at each point in $f^{-1}(p)$ whenever the set $f^{-1}(p)$ is finite. For more details see [21].

B. Proof of Theorem 2.1

Proof: Since f as a map from S^{m-1} to S^{m-1} is different from zero, then the map $f_{S^{m-1}}$ is onto the sphere. If f is not onto B^m, then it is homotopic to a map $B^m \to S^{m-1}$, and then $f_{S^{m-1}} : S^{m-1} \to S^{m-1}$ is homotopic to the trivial map (since it extends to the ball). Therefore $f_{S^{m-1}} : S^{m-1} \to S^{m-1}$ has zero degree, contrary to the assumption that it has degree different from zero. ∎

C. Proof of Lemma 2.2

Proof: The map f is a continuous bijective map from a compact space to a Hausdorff space, and therefore it is a homeomorphism. By observing that a homeomorphism $f : S^m \to S^m$ has degree ± 1 (see, for instance, [21, page 136]), we obtain the claim. ∎

D. Proof of Inductive Step in Theorem 3.1

Proof: Here we suppose that we have proved that the map f is surjective for $m-1$ power generators and we show how to use this to prove that the map f is surjective for m power generators.

If we have m power generators, the weight space is given by an m dimensional cube $\mathcal{C} = [-D, D]^m$, in complete analogy with the case of 3 generators. The m-simplex of measures is again defined as a set $\mathcal{A} \doteq \{(\lambda_{\mathcal{Q}_1}, \dots, \lambda_{\mathcal{Q}_m}) \in \mathbb{R}^m\}$ such that $\lambda_{\mathcal{Q}_i} \geq 0$ for $i \in \{1, \dots, m\}$ and $\sum_{i=1}^m \lambda_{\mathcal{Q}_i} = 1$. Note that $\mathcal{A}$ is homeomorphic to the $(m-1)$-dimensional ball B^{m-1}. As before, a power diagram can be viewed as a continuous map $f : \mathcal{C} \to \mathcal{A}$. It is easy to see that f is constant on sets of the form $\mathcal{W} \doteq \{\{W + t(1, \dots, 1)\} \cap \mathcal{C}, \quad t \in \mathbb{R}\}$, where W is a weight vector in $\mathcal{C}$. Moreover, fixing a point $p \in \mathcal{A}$ we have that $f^{-1}(p)$ is given by a set of the form $\mathcal{W}$ for a suitable W. Indeed, assume this is not the case, then the vector of measures $(\lambda_{\mathcal{Q}_1}, \dots, \lambda_{\mathcal{Q}_m})$ is obtained via f using two sets of weights: $W^1 \doteq (w_1^1, \dots, w_m^1)$ and $W^2 \doteq (w_1^2, \dots, w_m^2)$, and W^1 and W^2 do not belong to the same $\mathcal{W}$, namely it is not possible to obtain W^2 as $W^1 + t(1, \dots, 1)$ for a suitable t. This means that the vector difference $W^2 - W^1$ is not a multiple of $(1, \dots, 1)$. Therefore, there exists a nonempty set of indices J such that $w_j^2 - w_j^1 \geq w_k^2 - w_k^1$, whenever $j \in J$ and for all $k \in \{1, \dots m\}$, and such that the previous inequality is strict for at least one $k^* \in \{1, \dots m\}$. Now, among the indices in J there exists at least one of them, call it j^*, such that the generator j^* is a neighbor of generator k^*, due to the fact that the domain $\mathcal{Q}$ is connected. Indeed, if for all neighbors of k^* the inequality is not strict, choose as new k^* (denoted with a slight abuse of notation $k^{*,\prime}$ - in general their number is larger than one) each neighbor of k^* and look for neighbors of $k^{*,\prime}$ in the set J. If the search provides no result, repeat taking as new k^* each neighbor of $k^{*,\prime}$ and keep looking for a j^*. Since there are finitely many generators and the set is connected, eventually with this procedure one explores the entire set, and if no j^* is found then W^2 is expressible as $W^1 + t(1, \dots, 1)$ for a suitable t, contrary to the current assumption. Therefore, without loss of generality, we can assume that among the indices in J there exists at least one of them, call it j^*, such that the generator j^* is a neighbor of generator k^*. However, since $w_{j^*}^2 - w_{j^*}^1 > w_{k^*}^2 - w_{k^*}^1$, and $w_{j^*}^2 - w_{j^*}^1 \geq w_k^2 - w_k^1$ for all $k \in \{1, \dots, m\}$, then the measure $\lambda_{\mathcal{Q}_{j^*}}$ corresponding to the choice of weights W^2 is strictly larger than the measure $\lambda_{\mathcal{Q}_{j^*}}$ corresponding to the choice of weights W^1. This proves that $f^{-1}(p)$ is given only by sets of the form $\mathcal{W}$.

Accordingly, we introduce an equivalence relation on $\mathcal{C}$, declaring that two sets of weights W^1 and W^2 are equivalent if and only if they belong to the same $\mathcal{W}$. Let us call $\equiv$ this equivalence relation. It is immediate to see that f descends to a map $f : \mathcal{C}/\equiv \to \mathcal{A}$ (still called f by abuse of notation), and that f is now injective. It is easy to identify $\mathcal{C}/\equiv$ with the union of the $(m-1)$-dimensional faces of $\mathcal{C}$ given by $\mathcal{F} = \cup_{i=1}^m (\mathcal{C} \cap \{w_i = -D\})$ (see Figure 7). In this way we get a continuous injective map $f : \mathcal{F} \to \mathcal{A}$ that has the same image as the original f. Notice also that $\mathcal{F}$ is homeomorphic to the closed $(m-1)$-dimensional ball, thus, up to homeomorphisms, f can be viewed as a map $f : B^{m-1} \to B^{m-1}$.

We want to prove that the map $f_{\partial\mathcal{F}}$, given by the restriction of f to $\partial\mathcal{F}$, is onto $\partial\mathcal{A}$. To see this, consider one of the $(m-2)$-dimensional faces $\partial\mathcal{A}_i$ of $\partial\mathcal{A}$, which are identified by the condition $\lambda_{\mathcal{Q}_i} = 0$ (see Figure 7). Consider the face F_i in $\mathcal{F}$, where F_i is given by $F_i \doteq \mathcal{C} \cap \{w_i = -D\}$. We claim that the set $S_i \doteq \partial F_i \cap \partial\mathcal{F}$ is mapped onto $\partial\mathcal{A}_i$ by f. Observe that set S_i is described by the following equation: $S_i = \cup_{j \neq i} (\{w_i =$

$-D, w_j = D\} \cap \mathcal{F}$), so S_i is exactly equivalent to a set of type $\mathcal{F}$ for $m-1$ generators. Moreover, observe that $\partial\mathcal{A}_i$ can also be identified with the measure simplex for $m-1$ generators. By inductive hypothesis the map $f : S_i \to \partial\mathcal{A}_i$ is surjective, and therefore also the map $f_{\partial\mathcal{F}}$ is onto $\partial\mathcal{A}$.

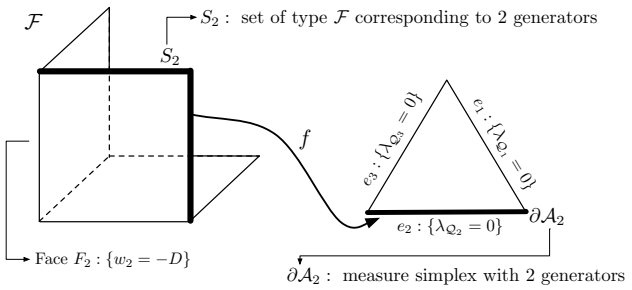

Fig. 7. The set S_2 and $f : S_2 \to \partial\mathcal{A}_2$ for the case $m = 3$. Unfortunately, for $m > 3$, visualization is not possible. Notice that S_2 is indeed a set of type $\mathcal{F}$ for $m = 2$ generators.

Since $f_{\partial\mathcal{F}}$ is a bijective continuous map among $(m-2)$-dimensional spheres (up to homeomorphisms), it has degree ± 1 by Lemma 2.2. Finally, we conclude that f is onto $\mathcal{A}$, using Theorem 2.1. ■

E. Proof of Theorem 3.12

Proof: The proof mainly relies on [22]. Let v be the unit vector introduced in the definition of the Unimodal Property. Then, there exist unique values $s_0 < s_1 < \ldots < s_m$ such that $s_0 = \inf\{s; \mathcal{Q}^s \neq \emptyset\}$, $s_m = \sup\{s; \mathcal{Q}^s \neq \emptyset\}$, and

$$\lambda_{\{x\in\mathcal{Q};\, v\cdot x \leq s_k\}} = \frac{k}{m}\lambda_{\mathcal{Q}}, \quad k = 1, \ldots, m-1. \tag{16}$$

Consider the intervals $L_i \doteq [s_{i-1}, s_i]$, $i \in I_m$. We claim that one can choose points $g_i = t_i v \in \mathbb{R}^d$, $i \in I_m$, such that $t_i \in L_i$ and the corresponding Voronoi diagram is

$$\begin{aligned} \mathcal{Q}_i &= \{x \in \mathcal{Q} \mid \|x - g_i\| = \min_k \|x - g_k\|\} \\ &= \{x \in \mathcal{Q} \mid v \cdot x \in [s_{i-1}, s_i]\}. \end{aligned} \tag{17}$$

Together, equation (16) and equation (17) yield the desired result.

Since, by assumption, $\mathcal{Q}$ enjoys the Unimodal Property, there exists an index $\bar{i} \in \{1, \ldots, m\}$ such that the length of the intervals $L_i = [s_{i-1}, s_i]$ decreases as i ranges from 1 to $\bar{i}$, then increases as i ranges from $\bar{i}$ to m. Let $L_{\bar{i}} = [s_{\bar{i}-1}, s_{\bar{i}}]$ be the smallest of these intervals, and define $t_{\bar{i}} \doteq \frac{s_{\bar{i}-1}+s_{\bar{i}}}{2} \in L_{\bar{i}}$. By induction, for i increasing from $\bar{i}$ to $m-1$, define t_{i+1} as the symmetric to t_i with respect to s_i, so that $t_{i+1} = 2s_i - t_i$, $i = \bar{i}, \bar{i}+1, \ldots, m-1$. Since the length of L_{i+1} is larger than the length of L_i, we have

$$t_i \in L_i \Rightarrow t_{i+1} \in L_{i+1}. \tag{18}$$

Similarly, for i decreasing from $\bar{i}$ to 2, we define $t_{i-1} = 2s_{i-1} - t_i$, $i = \bar{i}, \bar{i}-1, \ldots, 2$. Since the interval L_{i-1} is now larger than the interval L_i, we have

$$t_i \in L_i \Rightarrow t_{i-1} \in L_{i-1}. \tag{19}$$

Equations (18)-(19) imply $t_i \in L_i$ for all $i = 1, \ldots, m$. Hence the second equality in equation (17) holds. ■

Marco Pavone is a Staff Engineer/Technologist within the Advanced Robotic Controls Group at the Jet Propulsion Laboratory, California Institute of Technology. He received the Laurea degree in Systems and Control Engineering from the University of Catania, Italy, in 2004, the Diploma Engineer degree from Scuola Superiore di Catania, Italy, in 2005, and the Ph.D. degree in Aeronautics and Astronautics from the Massachusetts Institute of Technology, in 2010. His main research interests lie in the area of design and control of autonomous systems, with a focus on distributed coordination of multi-agent networks, robotic motion planning, transportation networks, and formation flying.

Alessandro Arsie received the B.Sc. degree in physics from the University of Padova, Padova, Italy, in 1997, the Ph.D. degree in mathematical physics from the International School for Advanced Studies, Trieste, Italy, in 2001, and the M.Sc. degree in quantitative finance and risk management from the Business University Bocconi, Milan, Italy, in 2005. He has been a Postdoctoral Associate in the Department of Mathematics, University of Bologna, Bologna, Italy, a Postdoctoral Associate in the department of Mechanical and Aerospace Engineering at UCLA, in the Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, and in the Department of Mathematics, Penn State University, University Park. He is currently assistant professor of mathematics at the University of Toledo, Toledo. His past research interests were focused on algebraic geometry while his current research interests are centered around the study of geometric and analytic aspects of dynamical systems and their applications.

Emilio Frazzoli received the Laurea degree in aerospace engineering from the University of Rome La Sapienza, Rome, Italy, in 1994 and the Ph.D. degree in navigation and control systems from the Department of Aeronautics and Astronautics, Massachusetts Institute of Technology (MIT), Cambridge, in 2001. From 2001 to 2004, he was an Assistant Professor of aerospace engineering at the University of Illinois at Urbana-Champaign. From 2004 to 2006, he was an Assistant Professor of mechanical and aerospace engineering at the University of California, Los Angeles. He is currently an Associate Professor of aeronautics and astronautics with the Laboratory for Information and Decision Systems, Department of Aeronautics and Astronautics, MIT. His main research interests lie in the general area of planning and control for mobile cyber-physical systems, with a particular emphasis on autonomous vehicles, mobile robotics, and transportation networks. Dr. Frazzoli received the National Science Foundation (NSF) CAREER Award in 2002.

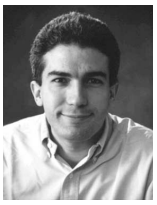

Francesco Bullo received the Laurea degree "summa cum laude" in Electrical Engineering from the University of Padova, Italy, in 1994, and the Ph.D. degree in Control and Dynamical Systems from the California Institute of Technology in 1999. From 1998 to 2004, he was an Assistant Professor with the Coordinated Science Laboratory at the University of Illinois at Urbana-Champaign. He is currently a Professor with the Mechanical Engineering Department and the Center for Control, Dynamical Systems and Computation at the University of California, Santa Barbara.

His main research interest is multi-agent networks with application to robotic coordination, distributed computing and power networks; he has worked on problems in vehicle routing, geometric control, and motion planning. He is the coauthor of "Geometric Control of Mechanical Systems" (Springer, 2004, 0-387-22195-6) and "Distributed Control of Robotic Networks" (Princeton, 2009, 978-0-691-14195-4). His students' papers were finalists for the Best Student Paper Award at the IEEE Conference on Decision and Control (2002, 2005, 2007), and the American Control Conference (2005, 2006, 2010). He has published more than 150 papers in international journals, books, and refereed conferences. He has served, or is serving, on the editorial boards of the "IEEE Transactions on Automatic Control", the "ESAIM: Control, Optimization, and the Calculus of Variations" and the "SIAM Journal of Control and Optimization".