

Distributed Policies for Equitable Partitioning: Theory and Applications

Marco Pavone, Emilio Frazzoli, Francesco Bullo

Abstract—The most widely applied resource allocation strategy is to balance, or equalize, the total workload assigned to each resource. In mobile multi-agent systems, this principle directly leads to equitable partitioning policies in which (i) the workspace is divided into subregions of equal measure, (ii) each agent is assigned to a unique subregion, and (iii) each agent is responsible for service requests originating within its own subregion. In this paper, we design distributed and adaptive policies to allow a team of agents to achieve a convex and equitable partition of a convex workspace. Our approach is related to the classic Lloyd algorithm, and exploits the unique features of Power Diagrams. We discuss possible applications to routing of vehicles in stochastic and dynamic environments, and to wireless networks. Simulation results are presented and discussed.

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 allocated to *customers*. In surveillance and exploration missions, customers are points of interest to be visited; in transportation and distribution applications, customers are people demanding some service (e.g., utility repair) or goods; in logistics tasks, customers could be troops in the battlefield.

The most widely applied resource allocation strategy is to balance, or equalize, the total workload assigned to each resource. While, in principle, several strategies are able to guarantee workload-balancing in multi-agent systems, *equitable partitioning policies* are predominant [1]–[4]. A *partitioning policy* is an algorithm that, as a function of the number m of agents and, possibly, of their position and other information, partitions a bounded workspace A into subregions A_i , for $i \in \{1, \dots, m\}$. Then, each agent i is assigned to subregion A_i , and each customer in A_i receives service from the agent

assigned to A_i . Accordingly, if we model the *workload* for subregion $S \subseteq A$ as $\lambda_S \doteq \int_S \lambda(x) dx$, where $\lambda(x)$ is a measure over A , then the workload for agent i is λ_{A_i} . Then, load balancing calls for equalizing the workload λ_{A_i} in the m subregions or, in equivalent words, requires to compute an *equitable partition* of the workspace A (i.e., a partition in subregions with the same measure).

Equitable partitioning policies are predominant for three main reasons: (i) efficiency, (ii) ease of design and (iii) ease of analysis; they are, therefore, ubiquitous in multi-agent system applications. To date, nevertheless, to the best of our knowledge, all equitable partitioning policies inherently assume a *centralized* computation of the workspace's tessellation. 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 three-fold. First, utilizing appropriate partitioning policies, we design distributed and adaptive policies to allow a team of agents to achieve an equitable partition. Our approach is related to the classic Lloyd algorithm from quantization theory [5], and exploits the unique features of Power Diagrams, a generalization of Voronoi Diagrams (see [6] for another interesting application of Power Diagrams in mobile sensor networks). A remarkable feature of our algorithms is that they guarantee *convex* A_i subregions (provided that the workspace is convex). Second, we design heuristic distributed algorithms that not only yield equitable partitions, but also provide “fat” (i.e., with a small diameter for a given area) subregions. Fat subregions, in general, improve overall performance. Third, we discuss some applications of our algorithms; we focus, in particular, on the Dynamic Traveling Repairman Problem (DTRP) [1], and on hybrid networks.

We, finally, 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 [7] and references therein); moreover, our results provide new insights in the geometry of Voronoi Diagrams and Power Diagrams.

II. BACKGROUND

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

Marco Pavone and Emilio Frazzoli are with the Laboratory for Information and Decision Systems, Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, {pavone, frazzoli}@mit.edu.

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

A. Notation

Let $\|\cdot\|$ denote the Euclidean norm. Let A be a compact, convex subset of \mathbb{R}^d . We denote the boundary of A as ∂A . The distance from a point x to a set M is defined as $\text{dist}(x, M) \doteq \inf_{p \in M} \|x - p\|$. We define $I_m \doteq \{1, 2, \dots, m\}$. Let $G = (g_1, \dots, g_m) \subset (\mathbb{R}^d)^m$ denote the location of m points. A *partition* (or tessellation) of A is a collection of m closed subsets $\mathcal{A} = \{A_1, \dots, A_m\}$ with disjoint interiors whose union is A . The partition of A is *convex*, if each subset A_i , $i \in I_m$, is convex.

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

The following result is related to classic divergence theorems [8]. Let $\Omega = \Omega(y) \subset A$ be a region that depends smoothly on a real parameter $y \in \mathbb{R}$ and that has a well-defined boundary $\partial\Omega(y)$ for all y . Let h be a density function over A . Then

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

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

C. Voronoi Diagrams and Power Diagrams

We refer the reader to [9] and [10] for comprehensive treatments, respectively, of Voronoi diagrams and Power Diagrams. Assume that G is an ordered set of distinct points. The *Voronoi Diagram* $\mathcal{V}(G) = (V_1(G), \dots, V_m(G))$ generated by points $G = (g_1, \dots, g_m)$ is defined by

$$V_i(G) = \{x \mid \|x - g_i\| \leq \|x - g_j\|, \forall j \neq i, j \in I_m\}. \quad (2)$$

We refer to G as the set of *generators* of $\mathcal{V}(G)$, and to $V_i(G)$ as the Voronoi cell of the i -th generator. For $g_i, g_j \in G$, $i \neq j$, let

$$b(g_i, g_j) = \{x \mid \|x - g_i\| = \|x - g_j\|\} \quad (3)$$

be the bisector of g_i and g_j ; 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*). It is easy to verify that each Voronoi cell is a *convex* set.

Assume, now, that each generator $g_i \in G$ has assigned an individual weight $w_i \in \mathbb{R}$, $i \in I_m$. We define $W = (w_1, \dots, w_m)$. In some sense, w_i measures the capability of g_i to influence its neighborhood. This is expressed by the power distance

$$d_P(x, g_i; w_i) \doteq \|x - g_i\|^2 - w_i. \quad (4)$$

We refer to the pair (g_i, w_i) as a *power point*. We define $G_W = ((g_1, w_1), \dots, (g_m, w_m)) \in (\mathbb{R}^d \times \mathbb{R})^m$. 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* $\mathcal{V}(G_W) =$

$(V_1(G_W), \dots, V_m(G_W))$ generated by power points $G_W = ((g_1, w_1), \dots, (g_m, w_m))$ is defined by

$$V_i(G_W) = \{x \mid \|x - g_i\|^2 - w_i \leq \|x - g_j\|^2 - w_j, \forall j \neq i, j \in I_m\}. \quad (5)$$

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 of the i -th power generator; moreover we call g_i and w_i , respectively, the position and the weight of power generator (g_i, w_i) . Notice that, when all weights are the same, the Power Diagram coincides with the Voronoi Diagram. Each power cell is, as well, a convex set (as it can be easily verified). Indeed, Power Diagrams are the generalized Voronoi Diagrams that have the strongest similarities to the original diagrams [11]. There are some differences, though. First, a power cell might be empty. Second, g_i might not be in its power cell. Finally, the bisector of (g_i, w_i) and (g_j, w_j) , $i \neq j$, is

$$b((g_i, w_i), (g_j, w_j)) = \{x \mid (g_j - g_i)^T x = \frac{1}{2}(\|g_j\|^2 - \|g_i\|^2 + w_i - w_j)\}. \quad (6)$$

Hence, $b((g_i, w_i), (g_j, w_j))$ is a face still 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 remaining of the paper: it means that, by changing weights, it is possible to arbitrarily move the bisector, while still preserving the orthogonality constraint. Notice that the Power diagram of an ordered set of possibly *coincident* power points is not well-defined. We define

$$\Gamma_{\text{coinc}} = \left\{ G_W \mid g_i = g_j \text{ and } w_i = w_j \text{ for some } i \neq j \in I_m \right\}. \quad (7)$$

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 define the (i, j) -face as $\Delta_{ij} \doteq V_i \cap V_j$.

III. PROBLEM FORMULATION

A total of m identical mobile agents provide service in a compact, convex service region $A \subseteq \mathbb{R}^d$. Let λ be a measure whose bounded support is A (in equivalent words, λ is not zero only on A); for any set S , we define the *workload* for region S as $\lambda_S \doteq \int_S \lambda(x) dx$. The measure λ models service requests, and can represent, for example, the density of customers over A , or, in a stochastic setting, their arrival rate. Given the measure λ , a partition $\{A_i\}_i$ of the workspace A is *equitable* if $\lambda_{A_i} = \lambda_{A_j}$ for all $i, j \in I_m$.

A *partitioning policy* is an algorithm that, as a function of the number m of agents and, possibly, of their position and other information, partitions a bounded workspace A into subregions A_i , $i \in I_m$. Then, each agent i is assigned to subregion A_i , and each service request in A_i receives service from the agent assigned to A_i . We refer to subregion A_i as the *region of dominance* of agent i . Given a measure λ and a partitioning policy, m agents are in a *convex equipartition configuration* with respect to λ if the associated partition is equitable and convex.

In this paper we study the following problem: find a distributed equitable partitioning policy that allows m mobile agents to reach a convex equipartition configuration (with respect to λ).

IV. FROM CONVEX EQUIPARTITIONS TO POWER DIAGRAMS

In [12], the authors present a distributed algorithm for the local computation of Voronoi cells. Therefore, it is tempting to consider the set of agents as a set of Voronoi generators; then, each agent i computes, using the algorithm in [12], its Voronoi cell V_i , and the set $A_i \doteq V_i \cap A$ becomes its region of dominance. By the properties of Voronoi Diagrams, the resulting regions of dominance are convex, tessellate A , but, in general, the resulting partition is not equitable. To overcome this problem, in [13] we introduced the idea of enabling the generators (in this setting the agents) to move, according to a distributed control law, along the gradient of a locational optimization function toward an *equitable* Voronoi Diagram, where each Voronoi cell has the same measure with respect to λ (i.e., $\lambda_{V_i} = \lambda_A/m$ for all $i \in I_m$); in fact, an equitable Voronoi Diagram yields an equitable partition $\{A_i\}_i$ of A , since $\lambda_{V_i} = \lambda_{V_i \cap A} = \lambda_{A_i}$ for all $i \in I_m$. However, this approach assumes that an equitable Voronoi Diagram exists.

Indeed, while an equitable Voronoi Diagram always exists when λ is *constant* over A [14], in general, for non-constant λ , an equitable Voronoi Diagram fails to exist, as the following counterexample shows.

Example 4.1 (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 Fig. 1, whose support is the interval $[0, 1]$. Assume $m = 5$. Let b_i ($i = 1, \dots, 4$) be the position of the i -th rightmost boundary point and g_i be the position of the i -th rightmost generator ($i = 1, \dots, 5$). 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 Fig. 1. Now, by the Perpendicular Bisector Property, we require:

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

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

Thus, for non-constant λ , in general, an equitable Voronoi Diagram fails to exist. A possible solution is to use Power Diagrams. On the one hand, Power Diagrams are the generalized Voronoi Diagrams that have the strongest similarities to the original diagrams [11], on the other hand, since the bisector is not required to pass through the midpoint of the line joining two neighbor generators, they are much more flexible. Notice that, following the ideas in [12], it is possible to compute Power Diagrams in a distributed way. The problem of existence of equitable Power Diagrams is studied next.




Fig. 1. Example of non-existence of an equitable Voronoi Diagram on a line. The above tessellation is an equitable partition, but not a Voronoi Diagram.

Theorem 4.2 (Existence of equitable Power Diagrams):

Let $A \subset \mathbb{R}^d$ be a compact, convex set, and λ be a measure whose support is A . Then, for every $m \geq 1$, there exist distinct points g_i , $i \in I_m$, all in the interior of A , and weights w_i , $i \in I_m$, such that the corresponding Power Diagram is equitable with respect to λ , i.e., $\lambda_{V_i} = \lambda_{V_j}$ for all $i, j \in I_m$.

Proof: Notice that the power distance is preserved under roto-translation. By compactness, there exist points $a, b \in A$ such that $\|b - a\| = \max_{z, z' \in A} \|z - z'\|$. By a translation of coordinates, we can assume $a = 0$. Define $v \doteq b/\|b\|$; by a rotation of coordinates, we can assume, without loss of generality, that v coincides with the first vector of the canonical basis e_1 . For each $s \in \mathbb{R}$, define the slice $V^s \doteq \{x \in A \mid e_1 \cdot x = s\}$. Then, there exist unique values $s_0 < s_1 < \dots < s_m$ such that $s_0 = \inf\{s \mid V^s \neq \emptyset\}$, $s_m = \sup\{s \mid V^s \neq \emptyset\}$, and

$$\lambda_{\{x \mid e_1 \cdot x \leq s_i\}} = \frac{i}{m} \lambda_A, \quad i = 1, \dots, m-1. \quad (8)$$

Define: $g_i = e_1(s_{i-1} + s_i)/2$, $i \in I_m$. We want, now, to choose the weights in such a way that $s_i e_1 \in b((g_i, w_i), (g_{i+1}, w_{i+1}))$ for $i = 1, \dots, m-1$ (see Fig. 2 for the case $A \subset \mathbb{R}^2$). This is, indeed, always possible. Recalling, in fact, Eq. (6), we set

$$w_{i+1} = \|g_{i+1}\|^2 - \|g_i\|^2 + w_i - 2(g_{i+1} - g_i)^T s_i e_1.$$

By setting $w_1 = 0$, the above recursive equation yields the weights w_i , $\forall i$.

The last step is to show that, for all $i \in I_m$,

$$\begin{aligned} \tilde{V}_i &\doteq \{x \mid e_1 \cdot x \in [s_{i-1}, s_i]\} = \\ &\{x \mid d_P(x, g_i; w_i) \leq d_P(x, g_j; w_j), \forall j \neq i\} = V_i. \end{aligned} \quad (9)$$




Fig. 2. Construction used for the proof of Theorem 4.2.

Together, Eq. (8) and Eq. (9) yield the desired result.

Given the weights thus computed, we have that $s_i e_1 \in b((g_i, w_i), (g_{i+1}, w_{i+1}))$. As a consequence, we have (with obvious modifications for g_1 and g_m):

$$\begin{aligned} \|x - g_i\|^2 - w_i &\leq \|x - g_{i+1}\|^2 - w_{i+1}, \quad \forall x \in \tilde{V}_i, \\ \|x - g_i\|^2 - w_i &\leq \|x - g_{i-1}\|^2 - w_{i-1}, \quad \forall x \in \tilde{V}_i. \end{aligned} \quad (10)$$

First, we want to show that

$$\|x - g_i\|^2 - w_i \leq \|x - g_{i+2}\|^2 - w_{i+2}, \quad \forall x \in \tilde{V}_i.$$

Assume, by contradiction, that there exists $\bar{x} \in \tilde{V}_i$ such that

$$\|\bar{x} - g_i\|^2 - w_i > \|\bar{x} - g_{i+2}\|^2 - w_{i+2}.$$

We can assume, without loss of generality, that $\bar{x} \cdot e_1 \in [g_i \cdot e_1, s_i]$. Define $\bar{x}_b \doteq \bar{x} + (s_i - \bar{x} \cdot e_1)e_1$. Clearly, $\bar{x}_b \in b((g_i, w_i), (g_{i+1}, w_{i+1}))$ and it belongs to both \tilde{V}_i and \tilde{V}_{i+1} . Since $\|\bar{x} - g_i\|^2 \leq \|\bar{x}_b - g_i\|^2$ and $\|\bar{x} - g_{i+2}\|^2 \geq \|\bar{x}_b - g_{i+2}\|^2$, we get

$$\|\bar{x}_b - g_{i+2}\|^2 - w_{i+2} < \|\bar{x}_b - g_i\|^2 - w_i = \|\bar{x}_b - g_{i+1}\|^2 - w_{i+1}.$$

This is a contradiction with respect to (10), since $\bar{x}_b \in \tilde{V}_{i+1}$.

It is, then, easy to show that for all $r \in \mathbb{N}_+$ such that $i + r \leq m$

$$\|x - g_i\|^2 - w_i \leq \|x - g_{i+r}\|^2 - w_{i+r}, \quad \forall x \in \tilde{V}_i.$$

By identical arguments, for all $r \in \mathbb{N}_+$ such that $i - r \geq 1$

$$\|x - g_i\|^2 - w_i \leq \|x - g_{i-r}\|^2 - w_{i-r}, \quad \forall x \in \tilde{V}_i.$$

Therefore, the set \tilde{V}_i is a subset of the power cell of generator i . Finally, we have to show that every point not belonging to \tilde{V}_i can not belong to the power cell of generator i . Assume, by contradiction, that there exists $\bar{x} \notin \tilde{V}_i$ such that \bar{x} belongs to the power cell of generator i . Since a power cell is a convex set, then the convex set $S \doteq \{x | x = ag_i + (1 - a)\bar{x}, a \in [0, 1]\}$ is entirely contained in the power cell of generator i . Then, there exists a point in S that belongs to the interior either of \tilde{V}_{i-1} or of \tilde{V}_{i+1} . Therefore, because of Eq. (10), such point can not belong to the power cell of generator i , a contradiction. Hence, the second equality in Eq. (9) holds. ■

V. GRADIENT DESCENT LAW FOR EQUITABLE PARTITIONING

In this section, we design distributed policies to allow a team of agents to achieve a convex equipartition configuration.

A. Virtual Generators

The first step is to associate to each agent i a *virtual power generator* (virtual generator for short) (g_i, w_i) . We define the region of dominance for agent i as the set $A_i \doteq V_i(G_W) \cap A$, where $G_W = ((g_1, w_1), \dots, (g_m, w_m))$, and $V_i(G_W)$ is the power cell of the i -th generator in the Power Diagram $\mathcal{V}(G_W)$. By the properties of Power Diagrams, the resulting regions of dominance are convex and tessellate A . Moreover, if the Power Diagram $\mathcal{V}(G_W)$ is equitable, i.e., each power cell has the same measure with respect to λ (or, in equivalent words, $\lambda_{V_i} = \lambda_A/m$ for all $i \in I_m$), then the resulting partition $\{A_i\}_i$ of A is equitable, since $\lambda_{V_i} = \lambda_{V_i \cap A} = \lambda_{A_i}$, for all $i \in I_m$. Henceforth, we assume that $G_W \in (\mathbb{R}^2 \times \mathbb{R})^m$, although extensions to other dimensions (i.e., to $G_W \in (\mathbb{R}^d \times \mathbb{R})^m$ with $d \neq 2$) are possible. A virtual generator (g_i, w_i) is simply an artificial variable locally controlled by the i -th agent; in particular, g_i is a virtual point and w_i is its weight (see Fig. 3). Virtual generators allow us to decouple the problem of achieving an equitable partition into regions of dominance from that of positioning an agent inside its own region of dominance.

We shall assume that each agent has sufficient information available to determine: (1) its power cell, and (2) the locations of all outstanding events in its power cell. A control policy that relies on information (1) and (2) is *distributed* in the sense that the behavior of each agent depends only on the location of the other agents with contiguous power cells. A spatially distributed algorithm for the local computation and maintenance of power cells can be designed following the ideas in [12].

B. Locational Optimization

The key idea is to enable virtual generators to follow a (distributed) gradient descent law such that an equitable Power Diagram is reached. Define the set

$$S \doteq \{G_W \in (\mathbb{R}^2 \times \mathbb{R})^m | G_W \neq \Gamma_{\text{coinc}} \text{ and } \lambda_{V_i} > 0 \forall i \in I_m\}.$$

We introduce the locational optimization function $H_V : S \rightarrow \mathbb{R}_{>0}$

$$H_V(G_W) \doteq \sum_{i=1}^m \lambda_{V_i}^{-1} + \frac{1}{2} \sum_{i=1}^m \text{dist}^2(g_i, A). \quad (11)$$

Notice that $\text{dist}(g_i, A) = 0$ if $g_i \in A$.

C. The gradient of H_V

The gradient of H_V is computed in the following theorem. We point out that this gradient can be computed in a distributed way, since it depends only on the location of the other agents with contiguous power cells. In the following,




Fig. 3. Agents, virtual generators and regions of dominance.

let $\gamma_{ij} \doteq \|g_j - g_i\|$. Moreover, let v_{A,g_i} be the vector from the point closest to g_i in A to g_i ; $v_{A,g_i} = 0$ if $g_i \in A$.

Theorem 5.1: Given a measure λ , the locational optimization function H_V is continuously differentiable on S , where for each $i \in \{1, \dots, m\}$

$$\begin{aligned} \frac{\partial H_V}{\partial g_i} &= \sum_{j \in N_i} \left(\frac{1}{\lambda_{V_j}^2} - \frac{1}{\lambda_{V_i}^2} \right) \int_{\Delta_{ij}} \frac{(x - g_i)}{\gamma_{ij}} \lambda(x) dx + v_{A,g_i}, \\ \frac{\partial H_V}{\partial w_i} &= \sum_{j \in N_i} \left(\frac{1}{\lambda_{V_j}^2} - \frac{1}{\lambda_{V_i}^2} \right) \int_{\Delta_{ij}} \frac{1}{2\gamma_{ij}} \lambda(x) dx, \end{aligned} \quad (12)$$

Furthermore, the critical configurations of H_V are generators' locations and weights with the property that all power cells have measure equal to λ_A/m , and all generators' locations are within A .

Proof: We first consider the partial derivative with respect to g_i . Let g_i^k be the k -th component of g_i ($k = 1, 2$). Similarly, let v_{A,g_i}^k be the k -th component of v_{A,g_i} ($k = 1, 2$). Since the motion of a generator g_i affects only power cell V_i and its neighboring cells V_j for $j \in N_i$, we have

$$\frac{\partial H_V}{\partial g_i^k} = -\frac{1}{\lambda_{V_i}^2} \frac{\partial \lambda_{V_i}}{\partial g_i^k} - \sum_{j \in N_i} \frac{1}{\lambda_{V_j}^2} \frac{\partial \lambda_{V_j}}{\partial g_i^k} + v_{A,g_i}^k. \quad (13)$$

Now, the result in Eq. (1) provides the means to analyze the variation of an integral function due to a domain change. Since the boundary of V_i satisfies $\partial V_i = \cup_j \Delta_{ij}$, where $\Delta_{ij} = \Delta_{ji}$ is the boundary between V_i and V_j , we have

$$\frac{\partial \lambda_{V_i}}{\partial g_i^k} = \sum_{j \in N_i} \int_{\Delta_{ij}} \left(\frac{dx}{dg_i^k} \cdot n_{ij}(x) \right) \lambda(x) dx \quad (14)$$

where we defined n_{ij} as the unit normal to Δ_{ij} outward of V_i (therefore we have $n_{ji} = -n_{ij}$). Similarly,

$$\frac{\partial \lambda_{V_j}}{\partial g_i^k} = \int_{\Delta_{ij}} \left(\frac{dx}{dg_i^k} \cdot n_{ji}(x) \right) \lambda(x) dx. \quad (15)$$

To evaluate the scalar product between the derivative of boundary points and the unit outward normal to the border in Eq. (14) and in Eq. (15), we differentiate Eq. (6) with respect to g_i^k at every point $x \in \Delta_{ij}$; we get

$$\frac{\partial x}{\partial g_i^k} \cdot (g_j - g_i) = e_k^T \cdot (x - g_i), \quad (16)$$

where e_k is the k -th vector of the canonical basis ($k = 1, 2$) in \mathbb{R}^2 . From Eq. (6) we have $n_{ij} = (g_j - g_i) / \|g_j - g_i\|$,

and the desired explicit expressions for the scalar products in Eq. (14) and in Eq. (15) follow immediately (recalling that $n_{ji} = -n_{ij}$).

Collecting the above results, we get the partial derivative with respect to g_i . The proof for the partial derivative with respect to w_i is similar and is omitted. The proof of the characterization of the critical points is an immediate consequence of the expression for the gradient of H_V ; we omit it in the interest of brevity. ■

D. Gradient Descent Law

Assume that generators' weights and positions obey a first order dynamical behavior described by $\dot{g}_i = u_i^g$ and $\dot{w}_i = u_i^w$. Consider H_V an aggregate objective function to be minimized and impose that the weight w_i and the position g_i follow the gradient descent given by (12). In more precise terms, we set up the following control law defined over the set S

$$u_i^g = -\frac{\partial H_V}{\partial g_i}(G_W), \quad u_i^w = -\frac{\partial H_V}{\partial w_i}(G_W), \quad (17)$$

where we assume that the partition $\mathcal{V}(G_W) = \{V_1, \dots, V_m\}$ is continuously updated. Let $\Omega \subseteq S$ be the set such that generators' positions and weights starting at $t = 0$ at $G_W(0) \in \Omega$ and evolving under (17) do not reach Γ_{coinc} . Clearly, Ω is non-empty. One can prove the following result.

Theorem 5.2: Consider the gradient vector field on Ω defined by equation (17). Then generators' positions and weights starting at $t = 0$ at $G_W(0) \in \Omega$, and evolving under (17) remain in Ω and converge asymptotically to the set of critical points of H_V (i.e., to the set of vectors of generators' positions and weights that yield an equitable Power Diagram).

Proof: First, we prove that set Ω is positively invariant with respect to (17), i.e., that $G_W(t) \neq \Gamma_{\text{coinc}}$, $t \geq 0$, and $\lambda_{V_i}(t) > 0$, $t \geq 0$, $i \in I_m$. Indeed, by definition of Ω , we have $G_W(t) \neq \Gamma_{\text{coinc}}$ for all $t \geq 0$ (therefore, the Power Diagram is always well-defined). Moreover, it is straightforward to see that $\dot{H}_V \leq 0$. Therefore, it holds

$$\lambda_{V_i}^{-1}(t) \leq H_V(G_W(t)) \leq H_V(G_W(0)), \quad i \in I_m, t \geq 0.$$

We conclude that the measures of all power cells are bounded away from zero. Thus, the generators' positions and weights will belong to Ω for all $t \geq 0$, i.e., $G_W(t) \in \Omega \forall t \geq 0$.

Second, as discussed before, $H_V : \Omega \rightarrow \mathbb{R}_{>0}$ is non-increasing along system trajectories, i.e., $\dot{H}_V(G_W) \leq 0$ in Ω .

Third, all trajectories with initial conditions in Ω are bounded. Indeed, $(1/2)\text{dist}^2(g_i(t), A) \leq H_V(G_W(t)) \leq H_V(G_W(0))$, $i \in I_m$, $t \geq 0$; therefore the generators' positions will remain within a bounded set. Moreover, it is easy to see that the sum of the weights is invariant under control law (17), i.e., $\frac{\partial \sum_{i=1}^m w_i}{\partial t} = 0$; thus we have $\sum_{i=1}^m w_i(t) = \sum_{i=1}^m w_i(0) \doteq w(0)$ along system trajectories. 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 the weights is constant), and the measure of at least one power cell would vanish (since the positions of the generators remain within a bounded set), which contradicts the fact that Ω is positively invariant.

Finally, by Theorem 5.1, H_V is continuously differentiable in Ω . Hence, by invoking the LaSalle invariance principle, under the descent flow (17), the generators' positions and weights will converge asymptotically to the set of critical points of H_V within Ω , which is non-empty by Theorem 4.2. ■

Remark 5.3: A spatially distributed algorithm for the local computation and maintenance of power cells can be designed following the ideas in [12]. Moreover, the partial derivative of H_V with respect to the i -th virtual generator only depends on the virtual generators with neighboring power cells. Therefore, the gradient descent law (17) is indeed a distributed control law. We mention that, in a Power Diagram, the average number of neighbors of a generator is less than or equal to six [11]; therefore, the computation of gradient (17) is scalable with the number of agents.

Remark 5.4: The analytical characterization of set Ω is an open problem. However, in all our simulations, the solution of the gradient flow (17) starting at general initial conditions $G_W(0) \in S$ never reached Γ_{coinc} and always converged to the set of critical points of H_V , i.e. to an equitable partition with generators' locations inside A (for an in-depth discussion of simulation results, see Section VII).

VI. LLOYD DESCENT FOR EQUITABLE PARTITIONING

The previous gradient descent law, although effective in providing a convex equitable partition, sometimes yields long and "skinny" subregions, whereas in some applications "fat" subregions (i.e., with a small diameter for a given area) are desirable. In this section, we introduce an heuristic distributed algorithm to achieve an equitable partition into convex and "fat" (indeed hexagon-like) subregions. As before, such algorithm is designed to be implemented by a distributed network of agents.

The idea is to extend the continuous-time Lloyd algorithm presented in [15]. As before, we associate to each agent i a virtual generator (g_i, w_i) . The mass and the centroid of the power cell V_i , $i \in I_m$, are then defined as

$$M_{V_i} = \int_{V_i} \lambda(x) dx, \quad C_{V_i} = \frac{1}{M_{V_i}} \int_{V_i} x \lambda(x) dx.$$

Then, each agent $i \in I_m$ updates its own virtual generator according to the following Lloyd descent

$$\dot{g}_i = -k_{\text{prop}}(g_i - C_{V_i}), \quad \dot{w}_i = \frac{1}{|N_i|} \left(\sum_{j \in N_i} \lambda_{V_j} \right) - \lambda_{V_i}, \quad (18)$$

where k_{prop} is a positive gain, and $|N_i|$ is the number of neighbors of virtual generator i . If all weights are initialized to the same value \bar{w} , and if \dot{w}_i is set to zero, then the control law (18) reduces to the continuous-time Lloyd algorithm in [15]; in particular, the generators will converge to the set of centroidal Voronoi configurations, where all cells are

TABLE I
PERFORMANCE OF CONTROL LAWS 1 AND 2.


Error ϵ	Law 1	Law 2
$\mathbb{E}[\epsilon]$	$1.0 \cdot 10^{-4}$	$8.0 \cdot 10^{-5}$
$\sigma^2(\epsilon)$	$2.7 \cdot 10^{-4}$	$9.0 \cdot 10^{-5}$
$\max \epsilon$	$1.7 \cdot 10^{-3}$	$4.1 \cdot 10^{-4}$

approximately hexagonal [16]. Simulations illustrate how this algorithm does indeed provide "fat" equitable partitions, even though a proof is not yet available.

VII. SIMULATION

In this section, we compare the performance of the different control laws proposed in the previous sections. For conciseness, we will refer to the Gradient Descent law (17) as "Law 1" and to the Lloyd Descent law (18) as "Law 2". In all simulations we assume that 10 agents provide service in the unit square A and that the measure λ is uniform over A : $\lambda \equiv 1$. Then, agents should converge to a configuration such that all regions of dominance have the same measure equal to $\bar{a} = 0.1$. For each law, we run 50 simulations. Agents' initial positions are independently and uniformly distributed over A ; the initial position of each virtual generator coincides with the initial position of the corresponding agent, and all weights are initialized to zero. Time is discretized with a step $dt = 0.01$, and each simulation run consists of 800 iterations (thus, the final time is $T = 8$). Define the error ϵ as the difference, at $T = 8$, between the measure of the region of dominance with maximum measure and the measure of the region of dominance with minimum measure. Table I summarizes simulation results. Expectation, variance and worst case error are with respect to 50 runs.

Recalling that the desired measure of each region of dominance is 0.1, we argue that both control laws 1 and 2 are effective in providing a convex equipartition. Notice from the third row of Table I that the worst case error ($\max \epsilon$), at $T = 8$, is within 2% from \bar{a} . Therefore, convergence to a convex equipartition seems to be robust for both policies. In particular, in all 50 simulations of Law 1, the solution of gradient flow (17) never reached Γ_{coinc} . Figure 4 shows how Law 2 provides "more regular" equipartitions.



(a) Typical equipartition of A with Law 1

(b) Typical equipartition of A with Law 2

Fig. 4. Comparison between the typical equipartitions achieved by using, respectively, Law 1 and Law 2.

VIII. APPLICATIONS

In this section we discuss some applications of the control policies proposed in the previous sections.

A. *Distributed Policies for the DTRP Problem*

The first application that we consider is the Dynamic Traveling Repairman Problem (DTRP). In the DTRP, m agents operating in a workspace A must service demands whose time of arrival, location and on-site service are stochastic; the objective is to find a 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. Any distribution system which receives orders in real time and makes deliveries based on these orders (e.g., courier services) is a clear candidate. Equitable partitioning policies (with respect to a suitable measure λ related to the probability distribution of demand locations) are, indeed, optimal for the DTRP (see [1], [17], [18]). Therefore, combining the optimal equitable partitioning policies in [17] with the algorithms presented in this paper, we immediately obtain optimal distributed policies for the DTRP problem. Notice that, since each agent is required to travel inside its own region of dominance, this scheme is inherently safe against collisions.

B. *Hybrid Networks*

An 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 relief efforts, etc., wireless ad hoc networks become valuable alternatives to fixed infrastructure networks for nodes to communicate with each other. To improve throughput capacity, a sparse network of more sophisticated nodes (supernodes) is placed within an ad hoc network. Supernodes provide long-distance communication. Assuming that normal nodes are independently and uniformly located in the workspace, supernodes should divide the area according to a hexagonal tessellation [3], where all hexagon cells have the same area. If we let supernodes play the role of agents in our framework, and we set $\lambda \equiv 1$, then the Lloyd Descent algorithm is a candidate distributed control law to allow a hybrid network to reach an efficient configuration.

IX. CONCLUSION

We have presented distributed control policies to allow a team of agents to achieve a convex equipartition configuration, using the Power Diagram partitioning policy. Our algorithms could find applications in many problems, including dynamic vehicle routing, and wireless networks. This paper leaves numerous important extensions open for further research. First, we plan to analytically characterize the set Ω . Second, we would like to extend our algorithms to guarantee control on the shape of the cells. Third, we envision considering the setting of structured environments (ranging from simple nonconvex polygons to more realistic ground environments). 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

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. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the supporting organizations.

REFERENCES

- [1] D. J. Bertsimas and G. J. van Ryzin. Stochastic and dynamic vehicle routing in the Euclidean plane with multiple capacitated vehicles. *Advances in Applied Probability*, 25(4):947–978, 1993.
- [2] O. Baron, O. Berman, D. Krass, and Q. Wang. The equitable location problem on the plane. *European Journal of Operational Research*, 183(2):578–590, 2007.
- [3] B. Liu, Z. Liu, and D. Towsley. On the capacity of hybrid wireless networks. In *IEEE INFOCOM 2003*, pages 1543–1552, San Francisco, CA, April 2003.
- [4] J. Carlsson, D. Ge, A. Subramaniam, A. Wu, and Y. Ye. Solving min-max multi-depot vehicle routing problem. *Report*, 2007.
- [5] S. P. Lloyd. Least-squares quantization in PCM. *IEEE Trans. Information Theory*, 28(2):129–137, 1982.
- [6] A. Kwok and S. Martinez. Energy-balancing cooperative strategies for sensor deployment. In *Proc. IEEE Conf. on Decision and Control*, pages 6136–6141, New Orleans, LA, December 2007.
- [7] J.G. Carlsson, B. Armbruster, and Y. Ye. Finding equitable convex partitions of points in a polygon efficiently. *To appear in The ACM Transactions on Algorithms*, 2008.
- [8] I. Chavel. *Eigenvalues in Riemannian Geometry*. Academic Press, New York, NY, 1984.
- [9] A. Okabe, B. Boots, K. Sugihara, and S. N. Chiu. *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*. John Wiley & Sons, New York, NY, 2000.
- [10] H. Imai, M. Iri, and K. Murota. Voronoi diagram in the Laguerre geometry and its applications. *SIAM Journal on Computing*, 14(1):93–105, 1985.
- [11] F. Aurenhammer. Power diagrams: properties, algorithms and applications. *SIAM Journal on Computing*, 16(1):78–96, 1987.
- [12] M. Cao and C. N. Hadjicostis. Distributed algorithms for Voronoi diagrams and applications in ad-hoc networks. Technical Report UILU-ENG-03-2222, UIUC Coordinated Science Laboratory, 2003.
- [13] 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*, pages 4869–4874, New Orleans, LA, December 2007.
- [14] A. Bressan. *Personal Communication*, 2008.
- [15] J. Cortés, S. Martínez, T. Karatas, and F. Bullo. Coverage control for mobile sensing networks. *IEEE Transactions On Robotics and Automation*, 20(2):243–255, 2004.
- [16] Q. Du, V. Faber, and M. Gunzburger. Centroidal voronoi tessellations: Applications and algorithms. *SIAM Review*, 41(4):637–676, 1999.
- [17] D. J. Bertsimas and G. J. van Ryzin. Stochastic and dynamic vehicle routing with general interarrival and service time distributions. *Advances in Applied Probability*, 25:947–978, 1993.
- [18] H. Xu. *Optimal policies for stochastic and dynamic vehicle routing problems*. Dept. of Civil and Environmental Engineering, Massachusetts Institute of Technology, Cambridge, MA., 1995.