Robust rendezvous for mobile autonomous agents via proximity graphs in arbitrary dimensions

Jorge Cortés, Sonia Martínez, Francesco Bullo

Abstract— **This paper presents coordination algorithms for networks of mobile autonomous agents. The objective of the proposed algorithms is to achieve rendezvous, that is, agreement over the location of the agents in the network. We provide analysis and design results for multi-agent networks in arbitrary dimensions under weak requirements on the switching and failing communication topology. The novel correctness proof relies on proximity graphs and their properties and on a general LaSalle Invariance Principle for nondeterministic discrete-time dynamical systems.**

I. Introduction

This work is a contribution to the emerging discipline of motion coordination for ad-hoc networks of mobile autonomous agents. With this loose terminology we refer to groups of robotic agents with limited mobility and communication capabilities. In the not too distant future these groups of coordinated devices will perform a variety of challenging tasks including, for example, search and recovery operations, surveillance, exploration and environmental monitoring. The potential advantages of employing arrays of agents have recently motivated vast interest in this topic. For example, from a control viewpoint, a group of agents inherently provides robustness to failures of single agents or of communication links.

The motion coordination problem for groups of autonomous agents is a control problem in the presence of communication constraints. Typically, each agents makes decisions based only on partial information about the state of the entire network that is obtained via communication with its immediate neighbors. One important difficulty is that the topology of the communication network depends on the agents' locations and, therefore, changes with the evolution of the network. A fundamental system-theoretical problem in the motion coordination of ad-hoc networks is the synthesis of control laws whose communication requirements scale nicely with the number of agents in the network.

The "multi-agent rendezvous" problem and a first "circumcenter algorithm" have been introduced by Ando and coworkers in [1]. The algorithm proposed in [1] has been extended to various synchronous and asynchronous stop-and-go strategies in [2], [3]. A related algorithm, in which connectivity constraints are not imposed, is proposed in [4]. A preliminary study on rendezvous under communication quantization is presented in [5]. These motion coordination schemes are memoryless (static feedback), anonymous (all agents are indistinguishable), and spatially distributed (only local information is required). An incomplete list of recent works on motion coordination algorithms includes [6], [7] on pattern formation, [8] on flocking, [9] on self-assembly, [10] on foraging, [11] on gradient climbing, and [12] on deployment. Consensus and control theoretical problems on dynamic graphs are discussed in [13] and in [14], respectively.

In this paper we provide novel analysis and design results on a class of rendezvous algorithms. First, we define and analyze a class of "circumcenter algorithms" defined over switching communication topologies. We classify communication topologies for our algorithms via the notion of "proximity graphs," see [15] and [12]. Admissible communication topologies for our algorithms are proximity graphs with the following properties: they are "spatially distributed" over the disk graph (i.e., they can be computed with only the local information encoded in the disk graph) and their connected components have the same vertices as the disk graph. This is a more general class of communication topologies than the one adopted in most works on motion coordination including for example [1], [2], [3], [4]. The ability to rely on general communication topologies is advantageous in the design of wireless communication strategies and is referred to as "topology control," see for example [16] and references therein. For the proximity graphs of interest in this paper, we prove some novel technical facts regarding connectivity.

Second, we consider networks of agents whose state space is $\mathbb{R}^d$, where d is an arbitrary number not restricted to $\{1, 2\}$. We prove that our proposed class of circumcenter algorithms is indeed correct in arbitrary dimensions and include simulations in two and three dimensions. As a natural outcome of this analysis, we prove that the original circumcenter algorithm in [1] can be adapted to work in higher dimensions, and that it is guaranteed to converge in finite time.

Third, we establish a general theorem on the robustness of the proposed class of circumcenter algorithms with respect to communication link failures. Rendezvous is guaranteed even if each agent experiences different link failures, provided the resulting directed communication graph is strongly connected at least once every finite number of time instants. Our results provide the first contribution to the theoretical explanation of the robustness properties of the circumcenter algorithm observed in computer simulations in [1].

Fourth, we develop an innovative method of proof based

Submitted on July 19, 2004; This version: April 5, 2005. Jorge Cortés is with the Department of Applied Mathematics and Statistics, University of California at Santa Cruz, Santa Cruz, California 95064, jcortes@ucsc.edu . Sonia Martínez and Francesco Bullo are with the Department of Mechanical and Environmental Engineering, University of California at Santa Barbara, Santa Barbara, California 93106, {smartine,bullo }@engineering.ucsb.edu

on a recently-developed LaSalle Invariance Principle for nondeterministic discrete-time dynamical systems, see [12]. This version of the invariance principle helps us establish robust convergence as follows. At each configuration of the network, we consider all the possible evolutions of the agents under all the possible choices of strongly connected communication topologies. In this way, the evolution of the proposed class of circumcenter algorithms is embedded into the (larger) set of evolutions of a non-deterministic discrete-time dynamical system. In turn, this system is analyzed via our novel version of the invariance principle.

This paper and our previous work in [12] use the same tools (generalized invariance principle and proximity graphs) to study two different and complementary motion coordination problems (deployment and rendezvous, respectively). We envision that these theoretical tools will play an important role in the emerging discipline of scalable motion coordination.

We emphasize that rendezvous problems are also important because of their relevance in network consensus problems [17], [13]. In a network consensus problem, the objective is to achieve agreement over the value of some logical variables. Our rendezvous algorithms can be applied to tackle consensus problems over dynamically changing and failing topologies, where the agents communicate and adjust the values of the agreement variables instead of their location. The proposed rendezvous algorithms are therefore comparable with feedback consensus algorithms for networks with failures.

The paper is organized as follows. In Section II we provide the necessary tools from stability theory and from geometry; these include the LaSalle Invariance Principle for nondeterministic discrete-time dynamical systems and the notion and properties of proximity graphs. Section III contains (i) a model of robotic network, (ii) the statement of the rendezvous problem, (iii) the statement of the so-called Circumcenter Algorithm over a proximity graph, and (iv) the theorems on asymptotic convergence and robustness to link failures. Section IV contains all the proofs, and Section V contains some instructive simulations in two and three dimensions. Finally, we provide a summary and future directions of research in Section VI.

II. Preliminary developments

Here we collect some known and some novel concepts that will be required in the later sections. First, we present a recently-developed version of the LaSalle Invariance Principle. Next, we review some geometric concepts related to proximity graphs. Finally, we provide a formally accurate notion of spatially distributed maps and obtain some fundamental properties associated with it.

A. LaSalle Invariance Principle for nondeterministic discrete-time dynamical systems

We review some concepts regarding the stability of discrete-time dynamical systems and set-valued maps following [18], [12]. For $d \in \mathbb{N}$, an *algorithm on* $\mathbb{R}^d$ is a set-valued map $T : \mathbb{R}^d \to 2^{(\mathbb{R}^d)}$ with the property that $T(p) \neq \emptyset$ for all $p \in \mathbb{R}^d$. Note that a map from $\mathbb{R}^d$ to $\mathbb{R}^d$ can be interpreted as a singleton-valued map. A *trajectory* of an algorithm T is a sequence $\{p_m\}_{m\in\mathbb{N}\cup\{0\}} \subset \mathbb{R}^d$ with the property that

$$p_{m+1} \in T(p_m)\,, \quad m \in \mathbb{N}\cup\{0\}\,.$$

In other words, given any initial $p_0 \in \mathbb{R}^d$, a trajectory of T is computed by recursively setting p_{m+1} equal to an arbitrary element in $T(p_m)$. An algorithm is therefore a nondeterministic discrete-time dynamical system.

An algorithm T is *closed at* $p \in \mathbb{R}^d$ if for all pairs of convergent sequences $p_k \to p$ and $p'_k \to p'$ such that $p'_k \in T(p_k)$, one has that $p' \in T(p)$. An algorithm is *closed on* $W \subset \mathbb{R}^d$ if it is closed at p, for all $p \in W$. In particular, every continuous map $T : \mathbb{R}^d \to \mathbb{R}^d$ is closed on $\mathbb{R}^d$. A set C is *weakly positively invariant with respect to* T if, for any $p_0 \in C$, there exists $p \in T(p_0)$ such that $p \in C$. A point p_0 is said to be a *fixed point of* T if $p_0 \in T(p_0)$. The function $V : \mathbb{R}^d \to \mathbb{R}$ is *non-increasing along* T *on* $W \subset \mathbb{R}^d$ if $V(p') \leq V(p)$ for all $p \in W$ and $p' \in T(p)$. We are ready to state the following result, whose proof is provided in [12].

Theorem 2.1: (LaSalle Invariance Principle for closed algorithms) Let T be a closed algorithm on $W \subset \mathbb{R}^d$ and let $V : \mathbb{R}^d \to \mathbb{R}$ be a continuous function non-increasing along T on W. Assume the trajectory $\{p_m\}_{m\in\mathbb{N}\cup\{0\}}$ of T takes values in W and is bounded. Then there exists $c \in \mathbb{R}$ such that

$$p_m \longrightarrow M \cap V^{-1}(c)\,,$$

where M is the largest weakly positively invariant set contained in

$$\{p \in \overline{W} \mid \exists p' \in T(p) \text{ such that } V(p') = V(p)\}.$$

Remark 2.2: If W is closed, then T is closed on W if and only if the graph of T restricted to W, $\mathrm{Graph}(T)_{|W} = \{(p,p') \mid p \in W, p' \in T(p)\}$ is a closed set. From [19, Lemma 14], if T is bounded on a neighborhood of W, then $\mathrm{Graph}(T)_{|W}$ being closed is equivalent to T being upper semi-continuous on W. □

B. Basic geometric notions and the circumcenter of a set

We review some notation for standard geometric objects; for additional information we refer the reader to [20] and references therein. For a bounded set $S \subset \mathbb{R}^d$, $d \in \mathbb{N}$, we let $\mathrm{co}(S)$ denote the convex hull of S. For $p, q \in \mathbb{R}^d$, we let $]p, q[= \{\lambda p + (1-\lambda)q \mid \lambda \in]0,1[\}$ and $[p, q] = \mathrm{co}(\{p, q\})$ denote the *open* and *closed segment* with extreme points p and q, respectively. For a bounded set $S \subset \mathbb{R}^d$, we let $\mathrm{CC}(S)$ and $\mathrm{CR}(S)$ denote the *circumcenter* and *circumradius* of S, respectively, that is, the center and radius of the smallest-radius d-sphere enclosing S. Note that the computation of the circumcenter and circumradius of a bounded set is a strictly convex problem and in particular a quadratically constrained linear program. For $p \in \mathbb{R}^d$, we let $B(p, r)$ and $\overline{B}(p, r)$ denote the *open* and *closed ball* of radius $r \in \mathbb{R}_+$ centered at p, respectively. Here, we let $\mathbb{R}_+$ and $\overline{\mathbb{R}}_+$ denote the positive and the nonnegative real numbers, respectively. A *polytope* is the convex hull of a finite point set.[1] We let $\mathrm{Ve}(Q)$ denote the set

[1] Note that with this definition polytopes are automatically convex.

of vertices of a polytope Q, and we emphasize that any vertex of Q is strictly convex, i.e., $v \in \mathrm{Ve}(Q)$ if and only if there exists $u \in \mathbb{R}^d$ such that $(s - v) \cdot u > 0$ for all $s \in Q \setminus \{v\}$.

Proposition 2.3: Let S be a bounded and finite set in $\mathbb{R}^d$. The following statements hold:

(i) $\mathrm{CC}(S) \in \mathrm{co}(S) \setminus \mathrm{Ve}(\mathrm{co}(S))$;
(ii) if $p \in S \setminus \mathrm{CC}(S)$ and $r \in \mathbb{R}_+$ satisfy $S \subset \overline{B}(p, r)$, then $]p, \mathrm{CC}(S)[$ has nonempty intersection with $\overline{B}(\frac{p+q}{2}, \frac{r}{2})$ for all $q \in S$.

Proof: The first statement follows directly from the definition of circumcenter and of vertex of a polytope. Let us provide a proof for the second statement. Since $p \neq \mathrm{CC}(S)$ and $S \subset \overline{B}(p, r)$, we deduce that $\mathrm{CR}(S) < r$. Let $q \in S$. If $\|p - q\| < r$, then $p \in B(\frac{p+q}{2}, \frac{r}{2})$, and therefore $]p, \mathrm{CC}(S)[$ has nonempty intersection with $\overline{B}(\frac{p+q}{2}, \frac{r}{2})$. Consider the case when $\|p - q\| = r$. Since $p, q \in \overline{B}(\mathrm{CC}(S), \mathrm{CR}(S))$ and $\mathrm{CR}(S) < r$, it follows that $\mathrm{CC}(S) \in B(p, r) \cap B(q, r)$. From this, we deduce that $]p, \mathrm{CC}(S)[$ has nonempty intersection with $\overline{B}(\frac{p+q}{2}, \frac{r}{2})$, as claimed. ■

C. Proximity graphs and their properties

We introduce some concepts regarding proximity graphs for point sets in $\mathbb{R}^d$. We assume the reader is familiar with the standard notions of graph theory as defined in [21, Chapter 1]. We begin with some notation. Given a vector space $\mathbb{V}$, let $\mathbb{F}(\mathbb{V})$ be the collection of finite subsets of $\mathbb{V}$. Accordingly, $\mathbb{F}(\mathbb{R}^d)$ is the collection of finite point sets in $\mathbb{R}^d$; we shall denote an element of $\mathbb{F}(\mathbb{R}^d)$ by $\mathcal{P} = \{p_1, \dots, p_n\} \subset \mathbb{R}^d$, where $p_1, \dots, p_n$ are distinct points in $\mathbb{R}^d$. Let $\mathbb{G}(\mathbb{R}^d)$ be the set of undirected graphs whose vertex set is an element of $\mathbb{F}(\mathbb{R}^d)$.

A *proximity graph function* $\mathcal{G} : \mathbb{F}(\mathbb{R}^d) \to \mathbb{G}(\mathbb{R}^d)$ associates to a point set $\mathcal{P}$ an undirected graph with vertex set $\mathcal{P}$ and edge set $\mathcal{E}_{\mathcal{G}}(\mathcal{P})$, where $\mathcal{E}_{\mathcal{G}} : \mathbb{F}(\mathbb{R}^d) \to \mathbb{F}(\mathbb{R}^d \times \mathbb{R}^d)$ has the property that $\mathcal{E}_{\mathcal{G}}(\mathcal{P}) \subseteq \mathcal{P} \times \mathcal{P} \setminus \mathrm{diag}(\mathcal{P} \times \mathcal{P})$ for any $\mathcal{P}$. Here, $\mathrm{diag}(\mathcal{P} \times \mathcal{P}) = \{(p, p) \in \mathcal{P} \times \mathcal{P} \mid p \in \mathcal{P}\}$. In other words, the edge set of a proximity graph depends on the location of its vertices. General properties of proximity graphs and the following examples are defined in [20], [15], [12]:

(i) the r-disk graph $\mathcal{G}_{\mathrm{disk}}(r)$, for $r \in \mathbb{R}_+$, with $(p_i, p_j) \in \mathcal{E}_{\mathcal{G}_{\mathrm{disk}}(r)}(\mathcal{P})$ if $\|p_i - p_j\| \leq r$;
(ii) the Delaunay graph $\mathcal{G}_{\mathrm{D}}$, with $(p_i, p_j) \in \mathcal{E}_{\mathcal{G}_{\mathrm{D}}}(\mathcal{P})$ if the Voronoi regions of p_i and p_j have non-empty intersection;
(iii) the Relative Neighborhood graph $\mathcal{G}_{\mathrm{RN}}$, with $(p_i, p_j) \in \mathcal{E}_{\mathcal{G}_{\mathrm{RN}}}(\mathcal{P})$ if, for all $p_k \in \mathcal{P} \setminus \{p_i, p_j\}$, $p_k \notin B(p_i, \|p_i - p_j\|) \cap B(p_j, \|p_i - p_j\|)$;
(iv) the Gabriel graph $\mathcal{G}_{\mathrm{G}}$, with $(p_i, p_j) \in \mathcal{E}_{\mathcal{G}_{\mathrm{G}}}(\mathcal{P})$ if, for all $p_k \in \mathcal{P} \setminus \{p_i, p_j\}$, $p_k \notin B(\frac{p_i + p_j}{2}, \frac{\|p_i - p_j\|}{2})$;
(v) the Euclidean Minimum Spanning Tree $\mathcal{G}_{\mathrm{EMST}}$, which for each $\mathcal{P}$, is a minimum-weight spanning tree of the complete graph $(\mathcal{P}, \mathcal{P} \times \mathcal{P} \setminus \mathrm{diag}(\mathcal{P} \times \mathcal{P}))$ whose edge (p_i, p_j) has weight $\|p_i - p_j\|$.

If needed, we shall write $\mathcal{G}_{\mathrm{disk}}(\mathcal{P}, r)$ to denote $\mathcal{G}_{\mathrm{disk}}(r)$ at $\mathcal{P}$. In what follows, we will consider the proximity graphs $\mathcal{G}_{\mathrm{RN}\,\cap\mathrm{disk}}(r)$ and $\mathcal{G}_{\mathrm{G}\,\cap\mathrm{disk}}(r)$ defined by the intersection of $\mathcal{G}_{\mathrm{RN}}$ and $\mathcal{G}_{\mathrm{G}}$ with $\mathcal{G}_{\mathrm{disk}}(r)$, $r \in \mathbb{R}_+$, respectively. A different proximity graph related to, but different from, the intersection $\mathcal{G}_{\mathrm{D}\,\cap\mathrm{disk}}(r)$ of $\mathcal{G}_{\mathrm{D}}$ with $\mathcal{G}_{\mathrm{disk}}(r)$ is the r-limited Delaunay graph $\mathcal{G}_{\mathrm{LD}}(r)$, as defined in [12].

To each proximity graph function $\mathcal{G}$, one can associate the *set of neighbors map* $\mathcal{N}_{\mathcal{G}} : \mathbb{R}^d \times \mathbb{F}(\mathbb{R}^d) \to \mathbb{F}(\mathbb{R}^d)$, defined by

$$\mathcal{N}_{\mathcal{G}}(p, \mathcal{P}) = \{q \in \mathcal{P} \mid (p, q) \in \mathcal{E}_{\mathcal{G}}(\mathcal{P} \cup \{p\})\}.$$

Typically, p is a point in $\mathcal{P}$, but the definition is well-posed for any $p \in \mathbb{R}^d$. Given $p \in \mathbb{R}^d$, it is convenient to define the map $\mathcal{N}_{\mathcal{G},p} : \mathbb{F}(\mathbb{R}^d) \to \mathbb{F}(\mathbb{R}^d)$ by $\mathcal{N}_{\mathcal{G},p}(\mathcal{P}) = \mathcal{N}_{\mathcal{G}}(p, \mathcal{P})$.

Let $\mathcal{G}_1$ and $\mathcal{G}_2$ be two proximity graph functions. We say that $\mathcal{G}_1$ is *spatially distributed over* $\mathcal{G}_2$ if, for all $p \in \mathcal{P}$,

$$\mathcal{N}_{\mathcal{G}_1,p}(\mathcal{P}) = \mathcal{N}_{\mathcal{G}_1,p}\big(\mathcal{N}_{\mathcal{G}_2,p}(\mathcal{P})\big).$$

It is straightforward to deduce that if $\mathcal{G}_1$ is spatially distributed over $\mathcal{G}_2$, then $\mathcal{G}_1$ is a subgraph of $\mathcal{G}_2$, that is, $\mathcal{G}_1(\mathcal{P}) \subset \mathcal{G}_2(\mathcal{P})$ for all $\mathcal{P} \in \mathbb{F}(\mathbb{R}^d)$. The converse is in general not true (for instance, the graph $\mathcal{G}_{\mathrm{D}\,\cap\mathrm{disk}}$ is a subgraph of $\mathcal{G}_{\mathrm{disk}}$, but it is not spatially distributed over it, see [12]).

We say that two proximity graph functions $\mathcal{G}_1$ and $\mathcal{G}_2$ *have the same connected components* if, for all point sets $\mathcal{P}$, the graphs $\mathcal{G}_1(\mathcal{P})$ and $\mathcal{G}_2(\mathcal{P})$ have the same number of connected components consisting of the same vertices.

Theorem 2.4: For $r \in \mathbb{R}_+$, the following statements hold:

(i) $\mathcal{G}_{\mathrm{EMST}} \subset \mathcal{G}_{\mathrm{RN}} \subset \mathcal{G}_{\mathrm{G}}$ and $\mathcal{G}_{\mathrm{G}\cap\mathrm{disk}}(r) \subset \mathcal{G}_{\mathrm{LD}}(r)$;
(ii) $\mathcal{G}_{\mathrm{disk}}(r)$ is connected if and only if $\mathcal{G}_{\mathrm{EMST}} \subset \mathcal{G}_{\mathrm{disk}}(r)$;
(iii) $\mathcal{G}_{\mathrm{RN}\,\cap\mathrm{disk}}(r)$, $\mathcal{G}_{\mathrm{G}\,\cap\mathrm{disk}}(r)$, and $\mathcal{G}_{\mathrm{LD}}(r)$ are spatially distributed over $\mathcal{G}_{\mathrm{disk}}(r)$;
(iv) $\mathcal{G}_{\mathrm{EMST}\,\cap\mathrm{disk}}(r)$, $\mathcal{G}_{\mathrm{RN}\,\cap\mathrm{disk}}(r)$, $\mathcal{G}_{\mathrm{G}\,\cap\mathrm{disk}}(r)$ and $\mathcal{G}_{\mathrm{LD}}(r)$ have the same connected components as $\mathcal{G}_{\mathrm{disk}}(r)$.

Proof: Fact (i) is mostly taken from [15] and [20]. Facts (ii) and (iii) are taken from [12]. Here we prove fact (iv). For $r \in \mathbb{R}_+$, it is enough to show that $\mathcal{G}_{\mathrm{EMST}\,\cap\mathrm{disk}}(r)$ has the same connected components as $\mathcal{G}_{\mathrm{disk}}(r)$, since this implies that the same result holds for $\mathcal{G}_{\mathrm{RN}\,\cap\mathrm{disk}}(r)$, $\mathcal{G}_{\mathrm{G}\,\cap\mathrm{disk}}(r)$ and $\mathcal{G}_{\mathrm{LD}}(r)$. Let $\mathcal{P} \in \mathbb{F}(\mathbb{R}^d)$. Since $\mathcal{G}_{\mathrm{EMST}\,\cap\mathrm{disk}}(r)$ is a subgraph of $\mathcal{G}_{\mathrm{disk}}(r)$, it is clear that vertices belonging to the same connected component of $\mathcal{G}_{\mathrm{EMST}\,\cap\mathrm{disk}}(\mathcal{P}, r)$ must also belong to the same connected component of $\mathcal{G}_{\mathrm{disk}}(\mathcal{P}, r)$. To prove the converse assume p_i and p_j in $\mathcal{P}$ verify $\|p_i - p_j\| \leq r$. Let $\mathcal{C}$ be the connected component of $\mathcal{G}_{\mathrm{disk}}(\mathcal{P}, r)$ to which they belong, with vertices $V(\mathcal{C})$. Since $\mathcal{C}$ is connected, then $\mathcal{G}_{\mathrm{EMST}}(V(\mathcal{C})) \subset \mathcal{C}$ by (ii). Now, using the definition of the Euclidean Minimum Spanning Tree and the fact that $\mathcal{C}$ is a connected component of $\mathcal{G}_{\mathrm{disk}}(\mathcal{P}, r)$, one can show that $\mathcal{G}_{\mathrm{EMST}}(V(\mathcal{C})) = \mathcal{G}_{\mathrm{EMST}}(\mathcal{P})[\mathcal{C}]$, where the latter denotes the subgraph of $\mathcal{G}_{\mathrm{EMST}}(\mathcal{P})$ induced by $\mathcal{C}$ (see [21] for the notion of induced graph). From this, we deduce that $\mathcal{G}_{\mathrm{EMST}}(V(\mathcal{C})) \subset \mathcal{G}_{\mathrm{EMST}\,\cap\mathrm{disk}}(\mathcal{P}, r)$, and therefore p_i and p_j belong to the same component of $\mathcal{G}_{\mathrm{EMST}\,\cap\mathrm{disk}}(\mathcal{P}, r)$. This implies the result. ■

We conclude this section with some examples of proximity graphs in $\mathbb{R}^2$ and $\mathbb{R}^3$; see Figures 1 and 2.

D. Proximity graphs over arrays of possibly coincident points and spatially distributed maps

The notion of proximity graph is defined for sets of distinct points $\mathcal{P} = \{p_1, \dots, p_n\}$. However, we will often consider

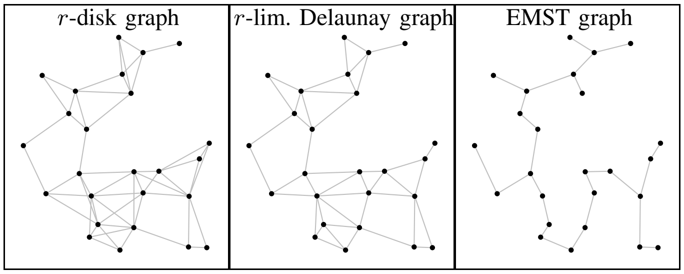

Fig. 1. From left to right, r-disk, r-limited Delaunay, and Euclidean Minimum Spanning Tree graphs in $\mathbb{R}^2$ for a configuration of 25 agents with coordinates uniformly randomly generated within the square $[-7,7]\times[-7,7]$. The parameter r is taken equal to 4.

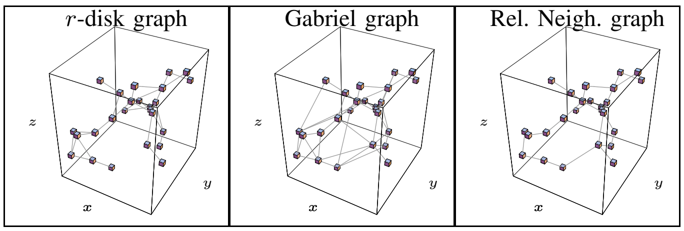

Fig. 2. From left to right, r-disk, Gabriel, and Relative Neighborhood graphs in $\mathbb{R}^3$ for a configuration of 25 agents with coordinates uniformly randomly generated within the square $[-7,7]\times[-7,7]\times[-7,7]$. The parameter r is taken equal to 4.

tuples of elements of $\mathbb{R}^d$ of the form $P=(p_1,\dots,p_n)$, i.e., ordered sets of possibly coincident points. Let $i_{\mathbb{F}}:(\mathbb{R}^d)^n\to\mathbb{F}(\mathbb{R}^d)$ be the natural immersion, i.e., $i_{\mathbb{F}}(P)$ is the point set that contains only the distinct points in $P=(p_1,\dots,p_n)$. Note that $i_{\mathbb{F}}$ is invariant under permutations of its arguments and that the cardinality of $i_{\mathbb{F}}(p_1,\dots,p_n)$ is in general less than or equal to n. In what follows, $\mathcal{P}=i_{\mathbb{F}}(P)$ will always denote the point set associated to $P\in(\mathbb{R}^d)^n$.

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

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

Additionally, we define the set of neighbors map $\mathcal{N}_{\mathcal{G}}:(\mathbb{R}^d)^n\to(\mathbb{F}(\mathbb{R}^d))^n$ as the function whose jth component is

$$\mathcal{N}_{\mathcal{G},j}(p_1,\dots,p_n)=\mathcal{N}_{\mathcal{G}}(p_j,i_{\mathbb{F}}(p_1,\dots,p_n)).$$

Note that coincident points in the tuple $(p_1,\dots,p_n)$ will have the same set of neighbors.

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

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

where T_j denotes the jth-component of T. In other words, the jth component of a spatially distributed map at $(p_1,\dots,p_n)$ can be computed with only the knowledge of the vertex p_j and the neighboring vertices in the undirected graph $\mathcal{G}(\{p_1,\dots,p_n\})$.

Remark 2.5: With this definition of spatially distributed map, one can see that the proximity graph function $\mathcal{G}_1$ is spatially distributed over the proximity graph function $\mathcal{G}_2$ if and only if the set of neighbors map $\mathcal{N}_{\mathcal{G}_1}$ is spatially distributed over $\mathcal{G}_2$. □

III. Rendezvous via proximity graphs

In this section we state the model, the control objective, the motion coordination algorithm, and the properties of the resulting closed-loop system.

A. Modeling a network of robotic agents

We begin by introducing the notions of *robotic agent* and of *network of robotic agents*. Let n be the number of agents in the network. Each agent has the following sensing, computation, communication, and motion control capabilities. The ith agent has a processor with the ability of allocating continuous and discrete states and performing operations on them. The ith agent occupies a location $p_i\in\mathbb{R}^d$, $d\in\mathbb{N}$, and it is capable of moving at any time $m\in\mathbb{N}$, for any unit period of time, according to the discrete-time control system

$$p_i(m+1)=p_i(m)+u_i. \tag{1}$$

Here, the control u_i takes values in a bounded subset of $\mathbb{R}^d$. We assume that there is a maximum step size $s_{\max}\in\mathbb{R}_+$ common to all agents, that is, $\|u_i\|\le s_{\max}$, for all $i\in\{1,\dots,n\}$. The sensing and communication model is the following. The processor of each agent has access to its location, and transmits this information to any other agent within a closed disk of radius $r\in\mathbb{R}_+$. Note that we are assuming the communication radius is the same for all agents.

Remarks 3.1:
- Equivalently, we shall consider groups of robotic agents without communication capabilities, but instead capable of measuring the relative position of each other agent within a closed disk of radius $r\in\mathbb{R}_+$.
- At first we assume that all communication between agents and all sensing of agents locations are accurate. We shall later analyze the robustness of our algorithms with respect to communication link failures. We will instead not address in this paper the correctness of our algorithms in the presence of measurement errors or communication quantization.
- Our network model is synchronous. Regarding asynchronous network models in rendezvous problems, we refer to [1] for early numerical results and to [3] for a thorough theoretical analysis. □

B. The rendezvous motion coordination problem

We now state the control design problem for the network of robotic agents. The *rendezvous objective* is to achieve agreement over the location of the agents in the network, that is, to steer each agent to a common location. This objective is to be achieved with the limited information flow described in the model above.

Typically, it will be impossible to solve the rendezvous problem if the agents are placed in such a way that they do not form a connected communication graph. Arguably, a good property of any algorithm for rendezvous is that of maintaining some form of connectivity between agents.

C. The Circumcenter Algorithm

Here is an informal description of what we shall refer to as the *Circumcenter Algorithm over a proximity graph* $\mathcal{G}$:

> Each agent performs the following tasks: (i) it detects its neighbors according to $\mathcal{G}$; (ii) it computes the circumcenter of the point set comprised of its neighbors and of itself, and (iii) it moves toward this circumcenter while maintaining connectivity with its neighbors.

This algorithm is an extension of the one introduced in [1]. Let us clarify which proximity graphs are allowable and how connectivity is maintained. Firstly, we are allowed to design motion coordination algorithms that are spatially distributed over the r-disk graph $\mathcal{G}_{\text{disk}}(r)$, or more generally, over any proximity graph $\mathcal{G}$ that is spatially distributed over $\mathcal{G}_{\text{disk}}(r)$. This is a direct consequence of our modeling assumption that each agent can acquire the location of each other agent within distance less than or equal to $r \in \mathbb{R}_+$. Secondly, we maintain connectivity by restricting the allowable motion of each agent. In particular, we will show that it suffices to restrict the motion of each agent as follows. If agents p_i and p_j are neighbors in the proximity graph $\mathcal{G}$, then their subsequent positions are required to belong to $\overline{B}(\frac{p_i+p_j}{2}, \frac{r}{2})$. If an agent p_i has its neighbors at locations $\{q_1, \dots, q_l\}$, then its *constraint set* $C_{p_i,r}(\{q_1, \dots, q_l\})$ is

$$C_{p_i,r}(\{q_1, \dots, q_l\}) = \bigcap_{q \in \{q_1, \dots, q_l\}} \overline{B}\Big(\frac{p_i + q}{2}, \frac{r}{2}\Big).$$

Before stating the algorithm in a more formal fashion, let us introduce one final concept. For q_0 and q_1 in $\mathbb{R}^d$, and for a convex closed set $Q \subset \mathbb{R}^d$ with $q_0 \in Q$, let $\lambda(q_0, q_1, Q)$ denote the solution of the strictly convex problem:

$$\begin{aligned} &\text{maximize } \lambda \\ &\text{subject to } \lambda \leq 1,\ (1-\lambda)q_0 + \lambda q_1 \in Q. \end{aligned} \tag{2}$$

Note that this convex optimization problem has the following interpretation: move along the segment from q_0 to q_1 the maximum possible distance while remaining in Q. Under the stated assumptions the solution exists and is unique. We are now ready to formally describe the algorithm.

Name:	Circumcenter Algorithm over $\mathcal{G}$
Goal:	Solve the rendezvous problem
Assumes:	(i) $s_{\max} \in \mathbb{R}_+$ is maximum step size (ii) $r \in \mathbb{R}_+$ is communication radius (iii) $\mathcal{G}$ is spatially distributed proximity graph over $\mathcal{G}_{\text{disk}}(r)$

For $i \in \{1, \dots, n\}$, agent i executes at each time instant in $\mathbb{N}$:

1: acquire $\{q_1, \dots, q_k\} := \mathcal{N}_{\mathcal{G}_{\text{disk}}(r), p_i}(\mathcal{P})$
2: compute $\mathcal{M}_i := \mathcal{N}_{\mathcal{G},p_i}(\{q_1, \dots, q_k\}) \cup \{p_i\}$
3: compute $Q_i := C_{p_i,r}(\mathcal{M}_i \setminus \{p_i\}) \cap \overline{B}(p_i, s_{\max})$
4: compute $\lambda_i^* := \lambda(p_i, \text{CC}(\mathcal{M}_i), Q_i)$
5: set $u_i := \lambda_i^*(\text{CC}(\mathcal{M}_i) - p_i)$, i.e.,
move from p_i to $(1-\lambda_i^*)p_i + \lambda_i^* \,\text{CC}(\mathcal{M}_i)$

In what follows we shall refer to the Circumcenter Algorithm over the proximity graph $\mathcal{G}$ as the map $T_{\mathcal{G}} : (\mathbb{R}^d)^n \to (\mathbb{R}^d)^n$.

D. Asymptotic correctness of the Circumcenter Algorithm

We are now ready to state the main convergence result, whose proof is postponed to the following section.

Theorem 3.2: Let $p_1, \dots, p_n$ be a network of robotic agents in $\mathbb{R}^d$, for $d \in \mathbb{N}$, with maximum step size $s_{\max} \in \mathbb{R}_+$ and communication radius $r \in \mathbb{R}_+$. Let the proximity graph $\mathcal{G}$ be spatially distributed over $\mathcal{G}_{\text{disk}}(r)$ and have the same connected components as $\mathcal{G}_{\text{disk}}(r)$. Any trajectory $\{P_m\}_{m \in \mathbb{N}\cup\{0\}}$ of $T_{\mathcal{G}}$ has the following properties:

(i) if the locations of two agents belong to the same connected component of $\mathcal{G}_{\text{disk}}(P_k, r)$ for some $k \in \mathbb{N} \cup \{0\}$, then they remain in the same connected component of $\mathcal{G}_{\text{disk}}(P_m, r)$ for all $m \geq k$;
(ii) there exists $P^* = (p_1^*, \dots, p_n^*) \in (\mathbb{R}^d)^n$ with the following properties: $P_m \to P^*$ as $m \to +\infty$, and $p_i^* = p_j^*$ or $\|p_i^* - p_j^*\| > r$ for each $i, j \in \{1, \dots, n\}$;
(iii) if $\mathcal{G} = \mathcal{G}_{\text{disk}}(r)$, then there exists $k \in \mathbb{N}$ such that $P_m = P^*$ for all $m \geq k$, that is, convergence is achieved in finite time.

Remarks 3.3:
- A consequence of Theorem 3.2(i) and (ii) is that, if the locations of two agents belong to the same connected component of $\mathcal{G}$ at some time, then they converge to the same point in $\mathbb{R}^d$.
- The statements Theorem 3.2(i) and (ii) were originally proved in [1] for the Circumcenter Algorithm over $\mathcal{G}_{\text{disk}}$ and for $d = 2$. This result was extended to other control policies by [2], [3] (still on the plane and with $\mathcal{G}_{\text{disk}}$ communication topology).
- It is instructive to consider two alternative strategies. With the same notation as in the Circumcenter Algorithm, they can be described as follows:
  (i) each agent moves to the orthogonal projection of the circumcenter $\text{CC}(\mathcal{M}_i)$ onto the convex set $Q_i \cap \text{co}(\mathcal{M}_i)$;
  (ii) each agent moves to the point in $Q_i \cap \text{co}(\mathcal{M}_i)$ that minimizes the maximum distance to each point in $\mathcal{M}_i$.

These algorithms are also the solutions to convex optimization problems. However, at this time, it is not clear what, if any, advantages they possess in comparison with the Circumcenter Algorithm. We conjecture that their correctness can be established along similar lines as the ones provided in the next section for Theorem 3.2. □

E. Robustness properties of the Circumcenter Algorithm

Here we characterize the robustness of the Circumcenter Algorithm with respect to link failures. We provide no physical model to motivate the occurrence for link failures; rather we analyze the resulting closed-loop network.

Definition 3.4: A *link failure in* $\mathcal{G}_{\text{disk}}(r)$ *at* $P \in (\mathbb{R}^d)^n$ is said to occur at agent p_i if (p_i, p_j) is an edge in $\mathcal{G}_{\text{disk}}(P, r)$ and the agent p_i does not detect agent p_j. For $\mathcal{P} = i_{\mathbb{F}}(P)$, we denote this link failure by the directed edge $(p_i, p_j) \in \mathcal{P} \times \mathcal{P}$. □

Remark 3.5: Consider an application of the Circumcenter Algorithm over a proximity graph $\mathcal{G}$ as described in the steps `1`-`5` above. If the link failure (p_i, p_j) takes place at step `1`, then the following two events will ensue:

(i) if p_j is a neighbor of p_i according to $\mathcal{G}$, then p_i looses the neighbor p_j at step `2`,
(ii) if p_k is not a neighbor of p_i according to $\mathcal{G}$ because of the presence of p_j, then p_i gains the neighbor p_k at step `2`.

Note that, after steps `1` and `2`, the collection of neighbors has been computed inaccurately. Nevertheless the execution of steps `3` through `5` can continue. □

Definition 3.6: For $P \in (\mathbb{R}^d)^n$, let $\mathcal{P} = i_{\mathbb{F}}(P)$. Let $\mathcal{G}$ be a proximity graph function that is spatially distributed over $\mathcal{G}_{\text{disk}}(r)$ and let $F \subset \mathcal{P} \times \mathcal{P}$ be a set of link failures. Let

(i) $\mathcal{G}_{\text{disk}}(\mathcal{P}, r) \nleftarrow F$ be the directed graph with vertex set $\mathcal{P}$ and with edge set $\mathcal{E}_{\text{disk}}(\mathcal{P}, r) \setminus F$;
(ii) $\mathcal{G}(\mathcal{P}) \nleftarrow F$ be the directed graph with vertex set $\mathcal{P}$ and with edges determined as follows; the neighbors of $p \in \mathcal{P}$ are
$$\mathcal{N}_{\mathcal{G},p}\big(\{q \mid (p,q) \in \mathcal{E}_{\text{disk}}(\mathcal{P}, r) \setminus F\}\big),$$
that is, the edges of $\mathcal{G}(\mathcal{P}) \nleftarrow F$ arise from the computation of $\mathcal{G}(\mathcal{P})$ with the link failures F, as described in Remark 3.5;
(iii) $T_{\mathcal{G} \nleftarrow F}(P)$ is the configuration obtained from applying the Circumcenter Algorithm over $\mathcal{G}$ (steps `1`-`5`) at configuration P with the link failures F at step `1`. □

Note that only a finite number of possible link failures can occur at any configuration. Consequently, the set of possible directed graphs arising from link failures is finite. We are now ready to state the main robust convergence result, whose proof is postponed to the following section.

Theorem 3.7: Let the network $p_1, \dots, p_n$ and the proximity graph $\mathcal{G}$ have the same properties as in Theorem 3.2. Given $P_0 \in (\mathbb{R}^d)^n$, consider the two sequences $\{P_m\}_{m \in \mathbb{N} \cup \{0\}}$ and $\{F_m\}_{m \in \mathbb{N} \cup \{0\}}$ defined recursively by

(i) F_m is a set of link failures in $\mathcal{G}_{\text{disk}}(r)$ at P_m, and
(ii) $P_{m+1} = T_{\mathcal{G} \nleftarrow F_m}(P_m)$.

If there exists $\ell \in \mathbb{N}$ such that at least one graph of any ℓ consecutive elements of $\{\mathcal{G}(P_m) \nleftarrow F_m\}_{m \in \mathbb{N} \cup \{0\}}$ is strongly connected, then there exists $p^* \in \mathbb{R}^d$ such that $P_m \to P^* = (p^*, \dots, p^*)$ as $m \to +\infty$.

Remarks 3.8:
- One could also state a version of this result for each connected component of the network, in a similar way to Theorem 3.2. We leave this to the reader.
- Theorem 3.7 provides the first theoretical explanation for the robustness behavior against sensor and control errors of the Circumcenter Algorithm over $\mathcal{G}_{\text{disk}}(r)$ observed in [1]. □

Corollary 3.9: With the same notation as in Theorem 3.7, if at each step $m \in \mathbb{N}$, the proximity graph $\mathcal{G}(P_m)$ is k_m-edge connected[2] and if F_m contains at most $k_m - 1$ link failures, then there exists $p^* \in \mathbb{R}^d$ such that $P_m \to P^* = (p^*, \dots, p^*)$ as $m \to +\infty$.

Next, we analyze the performance of the Circumcenter Algorithm when each agent of the mobile network at each time step is allowed to use a different proximity graph to compute its neighbors. The following definition formalizes this idea.

Definition 3.10: Let $\mathcal{S}$ be a set of proximity graph functions that are spatially distributed over $\mathcal{G}_{\text{disk}}(r)$. The *Circumcenter Algorithm over* $\mathcal{S}$ is the Circumcenter Algorithm where step `2` is replaced by

`2`(a): choose any $\mathcal{G} \in \mathcal{S}$
`2`(b): compute $\mathcal{M}_i := \mathcal{N}_{\mathcal{G},p_i}(\{q_1, \dots, q_k\}) \cup \{p_i\}$.

The selection algorithm for each agent at each execution of step `2`(a) is left unspecified. □

The following result guarantees that, under suitable conditions on the set $\mathcal{S}$, rendezvous is still attained by the mobile network executing the Circumcenter Algorithm over $\mathcal{S}$.

Corollary 3.11: Let the network $p_1, \dots, p_n$ be as in Theorem 3.2. Let $\mathcal{S}$ be a set of proximity graph functions that are spatially distributed over $\mathcal{G}_{\text{disk}}(r)$. Assume there exists a proximity graph $\mathcal{F}$ with the same connected components as $\mathcal{G}_{\text{disk}}(r)$ such that $\mathcal{F} \subset \mathcal{G}$, for all $\mathcal{G} \in \mathcal{S}$. Then any trajectory $\{P_m\}_{m \in \mathbb{N} \cup \{0\}}$ of the Circumcenter Algorithm over $\mathcal{S}$ has properties (i) and (ii) in Theorem 3.2.

We postpone the proof of this result to the following section. Note that, for $r \in \mathbb{R}_+$, the proximity graphs introduced in Section II-C, $\mathcal{G}_{\text{RN} \cap \text{disk}}(r)$, $\mathcal{G}_{\text{G} \cap \text{disk}}(r)$ and $\mathcal{G}_{\text{LD}}(r)$ are spatially distributed over $\mathcal{G}_{\text{disk}}(r)$ and contain $\mathcal{G}_{\text{EMST} \cap \text{disk}}(r)$, which has the same connected components as $\mathcal{G}_{\text{disk}}(r)$ (cf. Theorem 2.4). Therefore, any set $\mathcal{S} \subset \{\mathcal{G}_{\text{RN} \cap \text{disk}}(r), \mathcal{G}_{\text{G} \cap \text{disk}}(r), \mathcal{G}_{\text{LD}}(r)\}$ satisfies the hypothesis of Corollary 3.11.

IV. Convergence analysis

This section presents the proof of the main results of the paper. Before going into the details, let us introduce some useful notation. Let G be a directed graph with vertex set $\{1, \dots, n\}$ and edge set $E \subset \{1, \dots, n\} \times \{1, \dots, n\}$. Let $\mathcal{N}_G(i) = \{j \in \{1, \dots, n\} \mid (i,j) \in E\}$. Given $P \in (\mathbb{R}^d)^n$,

[2]An undirected graph is k-edge connected if it remains connected after any $k - 1$ edges have been removed, see [21].

let $P(\mathcal{N}_G(i)) = \{p_j \in \mathbb{R}^d \mid (i,j) \in E\}$. To a proximity graph function $\mathcal{G}$ that is spatially distributed over $\mathcal{G}_{\text{disk}}(r)$, a configuration $P \in (\mathbb{R}^d)^n$, and a set of link failures $F \subset \mathcal{P} \times \mathcal{P}$ (where $\mathcal{P} = i_{\mathbb{F}}(P)$), one may associate a directed graph $G_{\mathcal{G}(P)\nleftarrow F} = (\{1,\dots,n\}, E)$ by defining $(i,j) \in E$ if (p_i,p_j) is an edge of $\mathcal{G}(P) \nleftarrow F$. Note that if F is empty, then $(i,j) \in E$ if and only if $(p_i,p_j) \in \mathcal{E}_{\mathcal{G}}(P)$. Clearly, for each $P \in (\mathbb{R}^d)^n$, $P(\mathcal{N}_{G_{\mathcal{G}(P)\nleftarrow F}}(i))$ is equal to the set of neighbors of p_i with respect to the directed graph $\mathcal{G}(P) \nleftarrow F$.

Given a directed graph $G = (\{1,\dots,n\}, E)$ and $r \in \mathbb{R}_+$, define the *Circumcenter Algorithm at Fixed Topology* $T_{G,r} : (\mathbb{R}^d)^n \to (\mathbb{R}^d)^n$ whose ith component is

$$(T_{G,r})_i(p_1,\dots,p_n) = (1-\mu_i^*)p_i + \mu_i^* \operatorname{CC}\left(\{p_i\} \cup P(\mathcal{N}_G(i))\right),$$

where the coefficient of the convex combination is

$$\mu_i^* = \lambda(p_i, \operatorname{CC}\left(\{p_i\} \cup P(\mathcal{N}_G(i))\right), \widetilde{Q}_i),$$

and the constraint set is defined by

$$\widetilde{Q}_i = C_{p_i,r_i(P)}(P(\mathcal{N}_G(i))) \cap \overline{B}(p_i, s_{\max}),$$
$$r_i(P) = \max\{r, \max\{\|p_i - p_j\| \mid (i,j) \in E\}\}.$$

Note that if $\|p_i - p_j\| \le r$ for all $j \in \mathcal{N}_G(i)$, then $r_i(P) = r$.

There are two differences between $T_{G,r}$ and the algorithm $T_{\mathcal{G}}$ defined in Section III-C: (1) the topology of the network is fixed in $T_{G,r}$ and changing in $T_{\mathcal{G}}$, and (2) the constraint sets are, in general, bigger in $T_{G,r}$ than in $T_{\mathcal{G}}$. The reason for the latter difference is purely technical and will become clear in the proof of Theorem 4.6 below.

Lemma 4.1: Let $P \in (\mathbb{R}^d)^n$ and $r \in \mathbb{R}_+$. Let $\mathcal{G}$ be a proximity graph function that is spatially distributed over $\mathcal{G}_{\text{disk}}(r)$ and let $F \subset \mathcal{P} \times \mathcal{P}$ be a set of link failures. Then $T_{G_{\mathcal{G}(P)\nleftarrow F},r}(p_1,\dots,p_n) = T_{\mathcal{G}\nleftarrow F}(p_1,\dots,p_n)$. In particular, $T_{G_{\mathcal{G}(P)\nleftarrow \emptyset},r}(p_1,\dots,p_n) = T_{\mathcal{G}}(p_1,\dots,p_n)$.

Proof: The result follows from the definition of the directed graph $G_{\mathcal{G}(P)\nleftarrow F}$. ∎

With a slight abuse of notation, we introduce the convex hull function $\operatorname{co} : (\mathbb{R}^d)^n \to 2^{(\mathbb{R}^d)}$ by $\operatorname{co}(P) = \operatorname{co}(i_{\mathbb{F}}(P))$, where we implicitly represent a polytope in $\mathbb{R}^d$ by its set of vertexes.

Lemma 4.2: For $G = (\{1,\dots,n\}, E)$ and $r \in \mathbb{R}_+$, the map $T_{G,r} : (\mathbb{R}^d)^n \to (\mathbb{R}^d)^n$ has the following properties:

(i) $T_{G,r}$ is continuous;
(ii) $\operatorname{co}(T_{G,r}(P)) \subset \operatorname{co}(P)$, for $P \in (\mathbb{R}^d)^n$;

Proof: Statement (i) is a consequence of the following two facts: the circumcenter of a point set depends continuously on their location, and the solutions μ_i^*, $i \in \{1,\dots,n\}$, of the convex optimization problem (2) depend continuously on the data. From Proposition 2.3(i), we deduce $(T_{G,r})_i(p_1,\dots,p_n) \in \operatorname{co}(P)$ for all $i \in \{1,\dots,n\}$, which implies statement (ii). ∎

Given $r \in \mathbb{R}_+$, define the set-valued map $T : (\mathbb{R}^d)^n \to 2^{((\mathbb{R}^d)^n)}$ by

$$T_r(P) = \{T_{G,r}(P) \in (\mathbb{R}^d)^n \mid G = (\{1,\dots,n\}, E) \text{ is strongly connected}\}.$$

We shall refer to T_r as to the *Circumcenter Algorithm at All Strongly Connected Topologies*. Because there are a finite number of strongly connected directed graphs with n vertices, the set $T_r(P)$ is finite.

Proposition 4.3: For $r \in \mathbb{R}_+$, the map $T_r : (\mathbb{R}^d)^n \to 2^{((\mathbb{R}^d)^n)}$ has the following properties:

(i) $\operatorname{co}(P') \subset \operatorname{co}(P)$ for all $P' \in T_r(P)$ and $P \in (\mathbb{R}^d)^n$;
(ii) T_r is closed on $(\mathbb{R}^d)^n$.

Proof: Fact (i) is a consequence of Lemma 4.2(ii). Next, we prove fact (ii). Take $P_* \in (\mathbb{R}^d)^n$ and let us prove that T_r is closed at P. Consider two convergent sequences $P_m \to P_*$ and $P'_m \to P'_*$ with $P'_m \in T_r(P_m)$ for all $m \in \mathbb{N}$. We have to prove that $P'_* \in T_r(P_*)$. In order to do so, we reason by contradiction. Assume $P'_* \notin T_r(P_*)$, i.e, $P'_* \neq T_{G,r}(P_*)$ for any strongly connected directed graph $G = (\{1,\dots,n\}, E)$. Let $\varepsilon = \min\{\|P'_* - T_{G,r}(P_*)\| \mid G = (\{1,\dots,n\}, E) \text{ is strongly connected}\} > 0$. On the other hand, since for each directed graph G, the map $T_{G,r}$ is continuous at P_*, there exists $\delta_G > 0$ such that if $\|P - P_*\| \le \delta_G$, then $\|T_{G,r}(P) - T_{G,r}(P_*)\| \le \varepsilon/2$. Take $\delta = \min\{\delta_G \mid G = (\{1,\dots,n\}, E) \text{ is strongly connected}\} > 0$. Using the fact that the sequence $\{P_m\}_{m\in\mathbb{N}}$ converges to P_*, we deduce that there exists m_0 such that $\|P_m - P_*\| \le \delta$ for all $m \ge m_0$. Therefore, for all $m \ge m_0$, one has $\|T_{G,r}(P_m) - T_{G,r}(P_*)\| \le \varepsilon/2$ for any strongly connected directed graph G. From $P'_m \in T_r(P_m)$ for each $m \in \mathbb{N}$, we deduce that there exists a strongly connected directed graph G_m such that $P'_m = T_{G_m,r}(P_m)$. In particular, note that for all $m \ge m_0$, we have that $\|T_{G_m,r}(P_m) - T_{G_m,r}(P_*)\| \le \varepsilon/2$. Using these facts, we deduce the following chain of inequalities,

$$\|P'_* - P'_m\| = \|P'_* - T_{G_m,r}(P_m)\| \ge$$
$$\Big|\|P'_* - T_{G_m,r}(P_*)\| - \|T_{G_m,r}(P_*) - T_{G_m,r}(P_m)\|\Big| \ge \frac{\varepsilon}{2},$$

for all $m \ge m_0$, which contradicts $P'_m \to P'_*$. ∎

Next, let us study some properties of the diameter of a set. The *diameter* function $\operatorname{diam} : 2^{(\mathbb{R}^d)} \to \overline{\mathbb{R}}_+ \cup \{+\infty\}$ is defined by

$$\operatorname{diam}(S) = \sup\{\|p - q\| \mid p, q \in S\}.$$

Lemma 4.4: The function diam has the following properties:

(i) $\operatorname{diam}(S) = 0$ if and only if S is a singleton;
(ii) if $S \subset R \subset \mathbb{R}^d$, then $\operatorname{diam}(S) \le \operatorname{diam}(R)$;
(iii) $\operatorname{diam}(S) = \operatorname{diam}(\operatorname{co}(S))$ for all $S \subset \mathbb{R}^d$;
(iv) if $S \subset \mathbb{R}^d$ and Q a polytope in $\mathbb{R}^d$ satisfy $S \subset Q \setminus \operatorname{Ve}(Q)$, then $\operatorname{diam}(S) < \operatorname{diam}(Q)$.

Proof: The proof of these statements is straightforward and we do not include it here in the interest of space. ∎

It is now possible to define the function $V_{\text{diam}} = \operatorname{diam} \circ \operatorname{co} : (\mathbb{R}^d)^n \to \overline{\mathbb{R}}_+$, by

$$V_{\text{diam}}(P) = \operatorname{diam}(\operatorname{co}(P))$$
$$= \max\{\|p_i - p_j\| \mid i,j \in \{1,\dots,n\}\}.$$

Let $\operatorname{diag}((\mathbb{R}^d)^n) = \{(p,\dots,p) \in (\mathbb{R}^d)^n \mid p \in \mathbb{R}^d\}$.

Lemma 4.5: The function $V_{\text{diam}} = \operatorname{diam} \circ \operatorname{co} : (\mathbb{R}^d)^n \to \overline{\mathbb{R}}_+$ has the following properties:

(i) V_{diam} is continuous and invariant under permutations of its arguments;
(ii) $V_{\text{diam}}(P) = 0$ if and only if $P \in \operatorname{diag}((\mathbb{R}^d)^n)$;
(iii) V_{diam} is non-increasing along T_r, $r \in \mathbb{R}_+$, on $(\mathbb{R}^d)^n$.

Proof: Fact (i) is a straightforward consequence of the definition of V_{diam}. Fact (ii) is a consequence of Lemma 4.4(i). Proposition 4.3(i) implies fact (iii). ∎

We are now ready to analyze the asymptotic convergence properties of the algorithm T_r, for $r \in \mathbb{R}_+$.

Theorem 4.6: (Rendezvous via switching strongly connected graphs and suitable constraints): For $r \in \mathbb{R}_+$ and $P_0 \in (\mathbb{R}^d)^n$, any sequence $\{P_m\}_{m\in\mathbb{N}\cup\{0\}}$, defined by $P_{m+1} \in T_r(P_m)$, converges to a point of the form $(p, \dots, p) \in (\mathbb{R}^d)^n$.

Proof: From Lemma 4.5, we know that $V_{\text{diam}} : (\mathbb{R}^d)^n \to \overline{\mathbb{R}}_+$ is non-increasing along T_r on $(\mathbb{R}^d)^n$. Proposition 4.3(i) implies that the evolution of the sequence $\{P_m\}_{m\in\mathbb{N}\cup\{0\}}$ is contained in the compact set $\operatorname{co}(P_0)$. Since T_r is closed (cf. Proposition 4.3), we can resort to the LaSalle Invariance Principle for closed algorithms (cf. Theorem 2.1) to deduce that $P_m \to M$, where M is the largest weakly positively invariant set contained in

$$\{P \in (\mathbb{R}^d)^n \mid \exists P' \in T_r(P) \text{ such that } \operatorname{diam}(P') = \operatorname{diam}(P)\}\,.$$

Let us show that $M = \operatorname{diag}((\mathbb{R}^d)^n)$. Clearly, $\operatorname{diag}((\mathbb{R}^d)^n) \subset M$. To prove the other inclusion, we reason by contradiction. Assume $P \in M \setminus \operatorname{diag}((\mathbb{R}^d)^n)$, and therefore $\operatorname{diam}(P) > 0$. Let G be a strongly connected directed graph and consider $T_{G,r}(P)$. Clearly, by Proposition 2.3(i), for all $p_j \notin \operatorname{Ve}(\operatorname{co}(P))$, we have that $(T_{G,r})_j(P) \in \operatorname{co}(P) \setminus \operatorname{Ve}(\operatorname{co}(P))$. Let p_i be a (strictly convex) vertex of the polytope $\operatorname{co}(P)$. In general, there might exist more than one agent located at the same position p_i. Let us see that the application of $T_{G,r}$ will strictly decrease the number of agents $N_{p_i}(P)$ located at p_i. Since the directed graph G is strongly connected, there must exist i_* with $p_{i_*} = p_i$ such that there exist $j \in \mathcal{N}_G(i_*)$ with $p_j \neq p_{i_*}$. By Proposition 2.3(i), $\operatorname{CC}(\{p_{i_*}\} \cup \mathcal{P}(\mathcal{N}_G(i_*))) \in \operatorname{co}(P) \setminus \operatorname{Ve}(\operatorname{co}(P))$, which in particular implies that $\operatorname{CC}(\{p_{i_*}\} \cup P(\mathcal{N}_G(i_*))) \neq p_{i_*}$. Using this fact, together with $\{p_{i_*}\} \cup P(\mathcal{N}_G(i_*)) \subset \overline{B}(p_{i_*}, r_{i_*}(P))$, we deduce (cf. Proposition 2.3(iii)) that $]p_{i_*}, \operatorname{CC}(\{p_{i_*}\} \cup P(\mathcal{N}_G(i_*)))[$ has nonempty intersection with $\overline{B}(\frac{p_{i_*}+q}{2}, \frac{r_{i_*}(P)}{2})$ for all $q \in P(\mathcal{N}_G(i_*))$. Therefore, the solution $\mu_{i_*}^*$ of the convex optimization problem (2) is strictly positive. As a consequence, we have that $(T_{G,r})_{i_*}(P) \in \operatorname{co}(P) \setminus \operatorname{Ve}(\operatorname{co}(P))$. Therefore, $N_{p_i}(T_{G,r}(P)) < N_{p_i}(P)$.

Next, let us show that, after a finite number of steps, no agents will remain at the location p_i. Define $N = \max\{N_{p_i}(P) \mid p_i \in \operatorname{Ve}(\operatorname{co}(P))\} < n-1$. Then all agents in the configuration $T_{G_1,r}(T_{G_2,r}(\dots T_{G_N,r}(P)))$ are contained in $\operatorname{co}(P) \setminus \operatorname{Ve}(\operatorname{co}(P))$, for any strongly connected directed graphs $G_1, \dots, G_N$. Therefore, by Proposition 2.3(ii), $\operatorname{diam}(T_{G_1,r}(T_{G_2,r}(\dots T_{G_N,r}(P)))) < \operatorname{diam}(P)$, which contradicts the fact that M is weakly invariant.

Therefore, we have proved that for any initial condition $P_0 \in (\mathbb{R}^d)^n$, any sequence $\{P_m\}_{m\in\mathbb{N}\cup\{0\}}$, defined by $P_{m+1} \in T_r(P_m)$, converges to the set $\operatorname{diag}((\mathbb{R}^d)^n)$. To finish the proof, let us show that indeed $\{P_m\}_{m\in\mathbb{N}\cup\{0\}}$ must converge to a point that belongs to $\operatorname{diag}((\mathbb{R}^d)^n)$. Since the sequence is contained in the compact set $\operatorname{co}(P_0)$, there must exist a convergent subsequence $\{P_{m_k}\}_{k\in\mathbb{N}\cup\{0\}}$, $P_{m_k} \to (p_*, \dots, p_*)$ when $k \to +\infty$. Therefore, for any $\varepsilon > 0$, there exists k_0 such that for $k \geq k_0$ one has $\|(p_i)_{m_k} - p_*\| \leq \varepsilon/\sqrt{n}$, or equivalently, $\operatorname{co}(P_{m_k}) \subset B(p_*, \varepsilon/\sqrt{n})$. From Proposition 4.3(i) we deduce that $\operatorname{co}(P_m) \subset B(p_*, \varepsilon/\sqrt{n})$ for all $m \geq m_{k_0}$, which in turn implies that $\|P_m - (p_*, \dots, p_*)\| \leq \varepsilon$ for all $m \geq m_{k_0}$, as desired. ∎

Finally, we are ready to present the proof of Theorem 3.2.

Proof of Theorem 3.2: We start by proving fact (i). Let $k \in \mathbb{N}\cup\{0\}$ and take $\mathcal{C}$ a connected component of $\mathcal{G}_{\text{disk}}(P_k, r)$. By assumption, $\mathcal{G}$ and $\mathcal{G}_{\text{disk}}(r)$ have the same connected components, and therefore $\mathcal{C}$ is also a connected component of $\mathcal{G}(P_k)$. By definition of $T_{\mathcal{G}}$, if agents i and j are neighbors according to the graph $\mathcal{G}(P_k)$, then $(p_i)_{k+1}, (p_j)_{k+1} \in \overline{B}(\frac{(p_i)_k+(p_j)_k}{2}, \frac{r}{2})$, which in particular implies that $\|(p_i)_{k+1} - (p_j)_{k+1}\| \leq r$. Therefore, the agents in $\mathcal{C}$ remain connected in the r-disk graph at step $k+1$, i.e., the agents in $\mathcal{C}$ are contained in the same connected component of $\mathcal{G}_{\text{disk}}(P_{k+1}, r)$.

Now, let us prove fact (ii). From (i), we deduce that the number of vertices in each of the connected components of $\mathcal{G}(P_m)$ is non-decreasing. Since there is a finite number of agents, there must exist m_0 such that the identity of the agents in each connected component is fixed for all $m \geq m_0$ (i.e., no more agents are added to the connected component afterwards). Let $\mathcal{C} = \{p_{i_1}, \dots, p_{i_K}\}$ be any of these connected components. As a consequence of Theorem 4.6, we deduce that all the agents in $\mathcal{C}$ asymptotically converge to the same location in $\mathbb{R}^d$ (since their evolution under $T_{\mathcal{G}}$ is one of the many possible evolutions under the algorithm T, see Lemma 4.1).

Finally, we prove fact (iii). It suffices to prove that the agents in $\mathcal{C}$ will rendezvous in finite time. Let $a = \min\{s_{\text{max}}, \frac{r}{2}\} \in \mathbb{R}_+$. By the previous discussion, there exists $k \in \mathbb{N}$ such that the location of the agents in $\mathcal{C}$ belongs to a closed ball of radius $\sqrt{2}a/2$. In such a case, we deduce that (1) $\mathcal{G}_{\text{disk}}(r)$ at $\mathcal{C}$ is the complete graph, and therefore all agents in $\mathcal{C}$ compute the same circumcenter point CC, and (2) the corresponding circumradius can be seen to be less than or equal to a using a simple geometric argument. From the latter, we deduce that $\operatorname{CC} \in \overline{B}(p_i, s_{\text{max}})$ and $\operatorname{CC} \in C_{p_i,r}(\mathcal{P} \setminus \{p_i\})$, i.e., the circumcenter belongs to Q_i, for all $i \in \{1, \dots, n\}$. As a consequence, all mobile agents in $\mathcal{C}$ rendezvous at the same location CC at step $k+1$. ∎

Proof of Theorem 3.7: The proof of this result goes along the same lines as the one of Theorem 4.6. Given $r \in \mathbb{R}_+$, define the set-valued map $P \in (\mathbb{R}^d)^n \mapsto \widetilde{T}_r(P) = \{T_{G,r}(P) \in (\mathbb{R}^d)^n \mid G = (\{1, \dots, n\}, E) \text{ directed graph}\}$. Reasoning as in the proof of Proposition 4.3, one can show that $\widetilde{T}_r$ is closed. Given two set-valued maps $T_1, T_2 : (\mathbb{R}^d)^n \to 2^{((\mathbb{R}^d)^n)}$, define its composition as the set-valued map $T_1 \circ T_2 : (\mathbb{R}^d)^n \to 2^{((\mathbb{R}^d)^n)}$ given by $(T_1 \circ T_2)(P) = \{P'' \in (\mathbb{R}^d)^n \mid \exists P' \in (\mathbb{R}^d)^n \text{ such that } P'' \in T_1(P') \text{ and } P' \in T_2(P)\}$. For $k \in \mathbb{N}$, we denote by $\widetilde{T}_r^k$ the composition of k instances of $\widetilde{T}_r$. Now, let us define the set-valued map $P \in (\mathbb{R}^d)^n \mapsto \mathcal{T}_{r,\ell}(P) = \{P' \in (\mathbb{R}^d)^n \mid \exists k \in \{0, \dots, \ell-1\} \text{ such that } P' \in \widetilde{T}_r^k(T_r(P))\}$. Using Lemma 4.2(ii), together with the fact that

$\widetilde{T}_r$ and T_r are closed, we deduce that $\mathcal{T}_{r,\ell}$ is also closed. Reasoning as in the proof of Theorem 4.6, one can show that any sequence defined by $\mathcal{T}_{r,\ell}$ converges to a point that belongs to $\mathrm{diag}((\mathbb{R}^d)^n)$. This concludes the result, since the hypotheses of the statement of the theorem imply that the evolution of the network, $\{P_m\}_{m\in\mathbb{N}\cup\{0\}}$, is one of the many possible evolutions under $\mathcal{T}_{r,\ell}$, see Lemma 4.1. ■

Proof of Corollary 3.11: The proof of fact (i) is parallel to that of Theorem 3.2(i) invoking now that $\mathcal{F} \subset \mathcal{G}$, for all $\mathcal{G} \in \mathcal{S}$, and that $\mathcal{F}$ and $\mathcal{G}_{\text{disk}}(r)$ have the same connected components. Fact (ii) is a consequence of Theorem 3.7 since any execution of the Circumcenter Algorithm over $(\mathcal{S}_1, \dots, \mathcal{S}_n)$ can be seen as an instance of the Circumcenter Algorithm over $\mathcal{G}_{\text{disk}}(r)$ with appropriately selected link failures at each step. ■

V. Simulations

In order to illustrate the performance of our rendezvous algorithms, we developed a library of basic geometric routines. The resulting Mathematica® packages `PlanGeom.m` (containing the 2-dimensional routines) and `SpatialGeom.m` (containing the 3-dimensional routines) are freely available at `http://motion.mee.ucsb.edu`.

We implemented the Circumcenter Algorithm in the plane, $d = 2$, over the r-limited Delaunay proximity graph with link failures. The simulation run is illustrated in Figure 3. The 25 vehicles have a maximum step size $s_{\max} = .15$, and a communication radius $r = 4$. The initial configuration of the network is as in Figure 1 over the square $[-7, 7] \times [-7, 7]$.

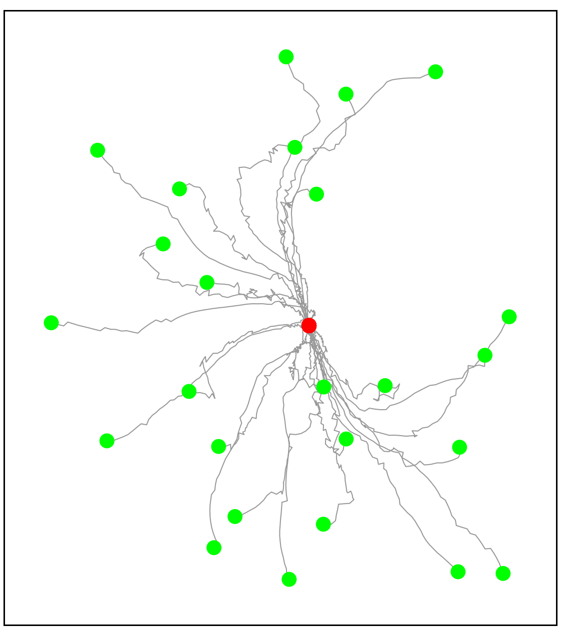

Fig. 3. Evolution (in light gray) of the Circumcenter Algorithm over the r-limited Delaunay graph $\mathcal{G}_{\text{LD}}(r)$ with link failures. The initial configuration of the network is as in Figure 1.

At each time step, a set consisting of 18 numbers between 1 and 25 is randomly selected, corresponding to the identities of the agents where link failures occur. For each of them, a randomly selected link failure in $\mathcal{G}_{\text{disk}}(r)$ is chosen. Note that, the identity of an agent might appear more than once in the random set, and therefore, more than one link failure may occur at the same agent. Nevertheless, rendezvous is asymptotically achieved according to Theorem 3.7 (indeed, in the various simulations that we ran, usually after 80 steps).

We also implemented the Circumcenter Algorithm in space, $d = 3$, over the set of proximity graphs $\{\mathcal{G}_{\text{disk}}(r), \mathcal{G}_{\text{G}}(r) \cap \mathcal{G}_{\text{disk}}(r), \mathcal{G}_{\text{RN}}(r) \cap \mathcal{G}_{\text{disk}}(r)\}$. The simulation run is illustrated in Figure 4. The 25 vehicles have, as before, a maximum step size $s_{\max} = .15$, and a communication radius $r = 4$. The initial configuration of the network is as in Figure 2 over the square $[-7, 7] \times [-7, 7] \times [-7, 7]$. At each time step, each agent randomly selects one of the proximity graphs in $\{\mathcal{G}_{\text{disk}}(r), \mathcal{G}_{\text{RN}\cap\text{disk}}(r), \mathcal{G}_{\text{G}\cap\text{disk}}(r)\}$ and computes its corresponding set of neighbors according to it. Then, it executes steps 3 through 5 of the Circumcenter Algorithm. Rendezvous is achieved in a finite number of steps (in the various simulations that we ran, usually after 100 steps).

VI. Conclusions

We have designed and analyzed a class of circumcenter algorithms over proximity graphs for multi-agent rendezvous. Additionally, we have provided a set of novel tools that we believe are important in the design and analysis of general motion coordination algorithms. Future directions of research in motion coordination include the study of increasingly realistic communication settings (asynchronicity, quantization, media access and power control issues), the analysis of the performance and complexity of the algorithms, and the formal design of other spatially distributed coordination primitives.

Acknowledgments

It is our pleasure to thank Stephen Morse and Brian Anderson for introducing us to the rendezvous problem and for being available for helpful discussions. This material is based upon work supported in part by ONR YIP Award N00014-03-1-0512 and NSF SENSORS Award IIS-0330008. Sonia Martínez's work was supported in part by a Fulbright Post-Doctoral Fellowship from the Spanish Ministry of Education and Science.

References

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

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

[3] J. Lin, A. S. Morse, and B. D. O. Anderson, "The multi-agent rendezvous problem: an extended summary," in *Proceedings of the 2003 Block Island Workshop on Cooperative Control* (V. Kumar, N. E. Leonard, and A. S. Morse, eds.), vol. 309 of *Lecture Notes in Control and Information Sciences*, pp. 257–282, New York: Springer Verlag, 2004.

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

[5] F. Fagnani, K. H. Johansson, A. Speranzon, and S. Zampieri, "On multi-vehicle rendezvous under quantized communication," in *Mathematical Theory of Networks and Systems*, (Leuven, Belgium), July 2004. Electronic proceedings.

[6] I. Suzuki and M. Yamashita, "Distributed anonymous mobile robots: Formation of geometric patterns," *SIAM Journal on Computing*, vol. 28, no. 4, pp. 1347–1363, 1999.

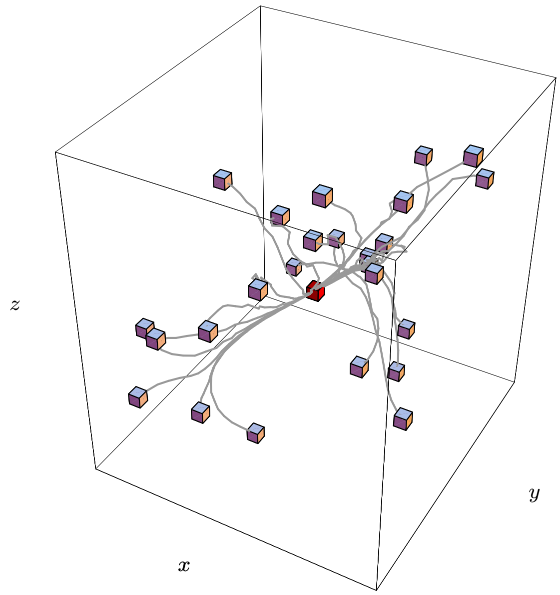

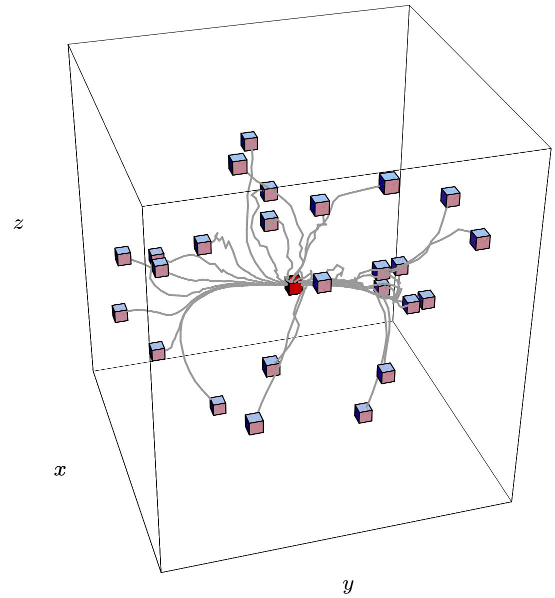

Fig. 4. Evolution (in light gray) of the Circumcenter Algorithm over $\{\mathcal{G}_{\text{disk}}(r), \mathcal{G}_{\text{G}}(r) \cap \mathcal{G}_{\text{disk}}(r), \mathcal{G}_{\text{RN}}(r) \cap \mathcal{G}_{\text{disk}}(r)\}$. The initial configuration of the network is as in Figure 2. The right figure is a rotated view of the left figure by 45 degrees.

[7] E. W. Justh and P. S. Krishnaprasad, "Equilibria and steering laws for planar formations," *Systems & Control Letters*, vol. 52, no. 1, pp. 25–38, 2004.

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

[9] E. Klavins, "Toward the control of self-assembling systems," in *Control Problems in Robotics* (A. Bicchi, H. Christensen, and D. Prattichizzo, eds.), vol. 4 of *STAR, Springer Tracts in Advanced Robotics*, pp. 153–168, Berlin Heidelberg: Springer Verlag, 2003.

[10] Y. Liu and K. M. Passino, "Stable social foraging swarms in a noisy environment," *IEEE Transactions on Automatic Control*, vol. 49, no. 1, pp. 30–44, 2004.

[11] P. Ögren, E. Fiorelli, and N. E. Leonard, "Cooperative control of mobile sensor networks: adaptive gradient climbing in a distributed environment," *IEEE Transactions on Automatic Control*, vol. 49, no. 8, pp. 1292–1302, 2004.

[12] J. Cortés, S. Martínez, and F. Bullo, "Spatially-distributed coverage optimization and control with limited-range interactions," *ESAIM. Control, Optimisation & Calculus of Variations*, Jan. 2004. To appear. Electronic version available at `http://www.me.ucsb.edu/bullo`.

[13] R. Olfati-Saber and R. M. Murray, "Consensus problems in networks of agents with switching topology and time-delays," *IEEE Transactions on Automatic Control*, vol. 49, no. 9, pp. 1520–1533, 2004.

[14] M. Mesbahi, "State-dependent graphs," in *IEEE Conf. on Decision and Control*, (Maui, Hawaii), pp. 3058–3063, Dec. 2003.

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

[16] X.-Y. Li, "Algorithmic, geometric and graphs issues in wireless networks," *Wireless Communications and Mobile Computing*, vol. 3, no. 2, pp. 119–140, 2003.

[17] N. A. Lynch, *Distributed Algorithms*. San Mateo, CA: Morgan Kaufmann Publishers, 1997.

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

[19] A. F. Filippov, *Differential Equations with Discontinuous Righthand Sides*, vol. 18 of *Mathematics and Its Applications*. Dordrecht, The Netherlands: Kluwer Academic Publishers, 1988.

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

[21] R. Diestel, *Graph Theory*, vol. 173 of *Graduate Texts in Mathematics*. New York: Springer Verlag, second ed., 2000.