Characterizing robust coordination algorithms via proximity graphs and set-valued maps

Jorge Cortés

***Abstract*— This note studies correctness and robustness properties of motion coordination algorithms with respect to link failures in the network topology. The technical approach relies on computational geometric tools such as proximity graphs, nondeterministic systems defined via set-valued maps and Lyapunov stability analysis. The manuscript provides two general results to help characterize the asymptotic behavior of spatially distributed coordination algorithms. These results are illustrated in rendezvous and flocking coordination algorithms.**

I. INTRODUCTION

Problem motivation: In the context of multi-agent networks, there are at least three relevant scenarios that concern robustness: (i) errors in sensor and/or communication data, (ii) agents arrivals and departures, and (iii) failures in the links of the network topology. Ideally, one would like to develop motion coordination algorithms that are robust in all three situations. Our main objective in this paper is to develop tools to analyze and carefully characterize the robustness properties of coordination algorithms with respect to link failures in the network topology. A second objective is to get insight into the features that make algorithms robust, and hopefully be able to provide tools that will guide researchers in the design of distributed algorithms that are more robust to link failures. In the future, we plan to develop tools that help study other notions of robustness.

Literature review: Recently, many works on cooperative control and multi-agent systems have incorporated various novel techniques to analyze the behavior of coordination algorithms. An incomplete list include ergodic, stochastic and circulant matrices [1], [2] from linear algebra, graph Laplacians and algebraic connectivity [1], [3] from algebraic graph theory, symmetries of differential equations [4], and LaSalle Invariance Principles for nondeterministic systems [5], [6]. Recent research efforts at developing widely applicable analysis tools for coordination algorithms include [7], [8] on set-valued Lyapunov theory, [9] on formal models for networks of robotic agents, and [10] on the role of network topology in multi-agent formation. The stability analysis for nondeterministic systems employed here is related to the approach for discrete-time systems in [11].

Statement of contributions: This paper generalizes in a systematic way the technical approach taken in [6] to analyze the circumcenter algorithm for the rendezvous problem, and further develops it to provide broadly applicable analysis tools for arbitrary coordination algorithms executed in discrete time. The network topology is induced by the communication or sensing capabilities of the agents. In general, coordination algorithms are not amenable to standard Lyapunov stability techniques because they are defined by maps which are discontinuous in the agents' states. This discontinuity is due to the sudden changes in neighboring relationships resulting from the agents' mobility. Here, we build on computational geometric notions (proximity graphs) and analysis tools (continuous and closed maps with power sets in their domain and codomain) to model the synchronous execution of (possibly discontinuous) coordination algorithms. The idea is to subsume the network trajectories into a larger set of trajectories corresponding to suitably-chosen set-valued maps. Our first contribution is a systematic procedure to define these set-valued maps, and a characterization of their continuity properties (Propositions 3.2 and 4.1).

This approach makes possible to apply stability techniques for nondeterministic dynamical systems in the analysis of the asymptotic behavior of arbitrary cooperative strategies. The generality provided by nondeterministic systems allows to consider also executions subject to failures in the links of the network topology. Under appropriate assumptions on the network evolution, we provide two convergence results that characterize the correctness and robustness properties of general coordination algorithms (Propositions 4.2 and 4.3). From a practical viewpoint, this procedure has no effect on the implementation of the coordination algorithm itself and, therefore, does not impose any additional complexity or computational load on the actual network execution. To illustrate the soundness of the approach, we apply the results obtained to rendezvous and flocking coordination tasks. In the rendezvous case, we study the robustness properties of the circumcenter algorithm, and obtain a characterization (cf. Theorem 5.2) stronger than the one in [6, Theorem 3.7]. In the flocking case, we recover the results obtained in [1], [7] for the average-heading algorithm (cf. Theorem 5.4). For reasons of space, we omit the proof of all results.

Organization: Section II introduces tools from graph theory and computational geometry. Section III discusses the continuity and closedness properties of maps with power sets in their domain and/or codomain. Section IV blends the previous two sections to synthesize analysis tools for correctness and robustness of coordination algorithms. Section V uses these tools in analyzing rendezvous and flocking algorithms. Finally, we present our conclusions in Section VI.

Notation: Given a set X, $\mathfrak{P}(X)$ (resp. $\mathbb{F}(X)$) denotes the collection of all subsets (resp. finite subsets) of X. Accordingly, elements of $\mathbb{F}(\mathbb{R}^d)$ are of the form $\{p_1, \dots, p_m\} \subset \mathbb{R}^d$,

Jorge Cortés is with the Department of Applied Mathematics and Statistics, Baskin School of Engineering, University of California, Santa Cruz, CA 95064, USA, http://www.ams.ucsc.edu/~jcortes, jcortes@ucsc.edu

where $p_1, \dots, p_m$ are distinct points in $\mathbb{R}^d$. Let $i_{\mathbb{F}} : X^n \to \mathbb{F}(X)$ be the natural immersion, i.e., $i_{\mathbb{F}}(P)$ is the point set that contains only the distinct points in $P = (p_1, \dots, p_n) \in X^n$. We denote $\operatorname{diag}(X^n) = \{(x, \dots, x) \in X^n \mid x \in X\}$.

II. Graph-theoretic tools

In this section, we review some standard notions of graph theory and computational geometry, and introduce some basic constructions that will be useful later.

A (*directed*) graph G is a pair $G = (V, E)$, with $E \subset V \times V$. An *undirected* graph G' is a pair $G' = (V', E')$, with $E' \subset V' \times V'$ and the property that $(i, j) \in E'$ implies $(j, i) \in E'$. Under this definition, any undirected graph is trivially a directed graph. The set of neighbors of $v \in V$ is defined as $\mathcal{N}_G(v) = \{w \in V \mid (w, v) \in E\}$.

A graph $G_1 = (V_1, E_1)$ is a *subgraph* of another graph $G_2 = (V_2, E_2)$ if $V_1 \subseteq V_2$ and $E_1 \subseteq E_2$. Alternatively, G_2 is said to be a *supergraph* of G_1. Formally, we set $G_1 \subseteq G_2$. We denote by $\mathfrak{P}(G)$ the set of all subgraphs of a graph G. Note that any undirected graph has directed graphs as subgraphs. Given a collection of graphs $\{G_1, \dots, G_s\}$, with a slight abuse of notation, we denote

$$\mathfrak{P}(\{G_1, \dots, G_s\}) = \bigcup_{k=1}^{s} \mathfrak{P}(G_k).$$

A. Proximity graphs

The notion of proximity graph will allow us later to rigorously characterize the distributed character of a motion coordination algorithm. Here, we introduce some concepts regarding proximity graphs for point sets in a d-dimensional space X. Let $\mathbb{G}(X)$ be the set of undirected graphs whose vertex set is an element of $\mathbb{F}(X)$. A *proximity graph* $\mathcal{G} : X^n \to \mathbb{G}(X)$ *on* X associates to a tuple $P \in X^n$ an undirected graph with vertex set $i_{\mathbb{F}}(P)$ and edge set $\mathcal{E}(P)$, where $\mathcal{E} : X^n \to \mathbb{F}(X \times X)$. In other words, the edge set of a proximity graph depends on the location of its vertices. To each proximity graph $\mathcal{G}$, one associates the *set of neighbors map* $\mathcal{N}_{\mathcal{G}} : X^n \to \mathbb{F}(X)^n$, defined by

$$\mathcal{N}_{\mathcal{G},i}(P) = \{p_j \in i_{\mathbb{F}}(P) \mid j \neq i \text{ and } (p_i, p_j) \in \mathcal{E}(P)\}.$$

Examples of proximity graphs on $\mathbb{R}^d$ include the complete graph, the r-disk graph, the Euclidean Minimum Spanning Tree, the Delaunay graph, etc. see [12], [13], [6]. For instance, the r-disk graph $\mathcal{G}_{\text{disk}}(r)$ is defined by $(p_i, p_j) \in \mathcal{E}_{\mathcal{G}_{\text{disk}}(r)}(P)$ if and only if $\|p_i - p_j\| \leq r$.

B. Fixed-topology graphs associated with a proximity graph

Let $\mathcal{G}$ be a proximity graph on X. For each $P \in X^n$, consider the undirected graph $\mathcal{G}(P) = (i_{\mathbb{F}}(P), \mathcal{E}(P))$. Let us define the undirected graph $G_P = (\{1, \dots, n\}, E_P)$, where $(i, j) \in E_P$ if and only if $(p_i, p_j) \in \mathcal{E}(P)$. A point $P \in X^n$ is *regular with respect to* $\mathcal{G}$ if there exists a neighborhood $\mathcal{U}$ of P in X^n such that $G_{P'} = G_P$ for all $P' \in \mathcal{U}$. Otherwise, the point P is called *singular with respect to* $\mathcal{G}$. Roughly speaking, the singular points correspond to the configurations where topology changes occur in the proximity graph.

Let $FT_{\mathcal{G}}(P) = \{G \text{ undirected} \mid \exists P'_m \to P \text{ with } G_{P_m} = G\}$ denote the set of undirected graphs associated with $\mathcal{G}$ at P. Note that P is regular if and only if $FT_{\mathcal{G}}(P) = \{G_P\}$. The set of *fixed-topology (undirected) graphs* associated with $\mathcal{G}$ is then

$$FT_{\mathcal{G}} = \bigcup_{P \in X^n} FT_{\mathcal{G}}(P) = \{G_P \mid P \in X^n\}.$$

Because there is a finite number of graphs with vertex set $\{1, \dots, n\}$, the set $FT_{\mathcal{G}}$ is finite.

In the presence of link failures, an agent located at p_i might receive information from (or might detect) an agent at p_j, and instead, the agent at p_j might not receive information from (or might not detect) the agent at p_i. The effective interaction topology is hence best described as a directed graph. This is the reason why we define $\mathfrak{P}(FT_{\mathcal{G}})$ as the set of *fixed-topology (directed) graphs with link failures* associated with $\mathcal{G}$. Note that

$$\mathfrak{P}(FT_{\mathcal{G}}) = \bigcup_{P \in X^n} \mathfrak{P}(G_P).$$

Clearly $FT_{\mathcal{G}} \subset \mathfrak{P}(FT_{\mathcal{G}})$. Moreover, since $\mathfrak{P}(G)$ is finite for each $G \in FT_{\mathcal{G}}$, then $\mathfrak{P}(FT_{\mathcal{G}})$ is also finite.

III. Maps with power sets in their domain and codomain

This section deals with the continuity and closedness properties of maps with power sets in their domain and codomain. The need to consider maps which take sets as both arguments and values arises from the fact that each agent has in general a set of other agents as neighbors, and this set changes discontinuously with the agents' states. Along the section, we state various definitions and results that link together the treatment here with various notions related to proximity graphs.

A. Continuity properties

Consider a map of the form $F : X \times \mathfrak{P}(Y) \to \mathfrak{P}(Z)$, where X, Y and Z are Hausdorff topological spaces. We say that F is *upper semicontinuous at* $(x, A) \in X \times \mathfrak{P}(Y)$ if for all $\epsilon > 0$, there exists $\delta > 0$ such that, for $x' \in B(x, \delta)$ and $A' \subset A + B(0, \delta)$, one has

$$F(x', A') \subset F(x, A) + B(0, \epsilon).$$

Similarly, F is *lower semicontinuous at* $(x, A) \in X \times \mathfrak{P}(Y)$ if for all $\epsilon > 0$, there exists $\delta > 0$ such that, for $x' \in B(x, \delta)$ and $A' \subset A + B(0, \delta)$, one has

$$F(x, A) \subset F(x', A') + B(0, \epsilon).$$

The map F is *continuous at* (x, A) if it is both upper and lower semicontinuous at (x, A). Note that for single-valued maps, upper semicontinuity is equivalent to lower semicontinuity, which in turn is also equivalent to continuity. The same notions can also be defined for maps of the form $F : X \times \mathfrak{P}(Y) \to U \times \mathfrak{P}(Z)$, simply by requiring the desired property on each component $F_1 : X \times \mathfrak{P}(Y) \to U$, $F_2 : X \times \mathfrak{P}(Y) \to \mathfrak{P}(Z)$. It is not difficult to verify that (i) the composition of upper semicontinuous maps is upper semicontinuous, and (ii) the composition of an upper semicontinuous map with a lower semicontinuous map is also lower semicontinuous. As a consequence, the composition of continuous maps is also continuous.

Fixed-topology graphs give naturally rise to set of neighbors maps which are continuous, as stated next.

Lemma 3.1: For a directed graph $G = (\{1, \ldots, n\}, E)$ and a set X, the map $\mathrm{Ev}_G : X^n \to \mathbb{F}(X)^n$ whose ith component is $P \mapsto P(\mathcal{N}_G(i)) = \{p_j \in i_{\mathbb{F}}(P) \mid p_j \in \mathcal{N}_G(i)\}$, for $i \in \{1, \ldots, n\}$, is continuous.

B. Closedness properties

We say that $T : X \to \mathfrak{P}(Y)$ is *closed at* $x \in X$ if for all pairs of convergent sequences $x_m \to x$ and $y_m \to y$ such that $y_m \in T(x_m)$, one has that $y \in T(x)$. In particular, every continuous map $T : X \to Y$ at $x \in X$ is closed at $x \in X$. The map T is closed at $A \subset X$ if it is closed at x, for all $x \in A$. The notions of upper semicontinuity and closedness are in general different. However, if $A \subset X$ is closed and T is bounded on a neighborhood of A, then T is closed at A if and only if T is upper semicontinuous at A (see [14]).

Proposition 3.2: Given two Hausdorff topological spaces X and Y, let $T : X \to \mathfrak{P}(Y)$ be a set-valued map and $x \in X$. Assume there exists a neighborhood U of x in X and functions $f_1, \ldots, f_s : U \subset X \to Y$ continuous at x such that $T(x') = \{f_1(x'), \ldots, f_s(x')\}$ for all $x' \in U$. Then T is closed at x.

The following result readily follows from Proposition 3.2.

Corollary 3.3: Given two topological spaces X and Y, let $f_1, \ldots, f_s : X \to Y$ be continuous functions. Then the set-valued map $T : X \to \mathfrak{P}(Y)$ defined by $T(x) = \{f_1(x), \ldots, f_s(x)\}$ is closed.

C. Spatially distributed maps

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

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

where T_j denotes the jth-component of T. We refer to $\widetilde{T}$ as the map *associated to* T *for* $\mathcal{G}$. In other words, the jth component of a spatially distributed map at $(p_1, \ldots, p_n)$ can be computed with only the knowledge of the vertex p_j and the neighboring vertices in the graph $\mathcal{G}(p_1, \ldots, p_n)$.

D. Trajectories of set-valued maps

Let $T : X \to \mathfrak{P}(X)$ have the property that $T(x) \neq \emptyset$ for all $x \in X$. A *trajectory* of T is a sequence $\{x_m\}_{m \in \mathbb{N}_0} \subset X$ with the property that

$$x_{m+1} \in T(x_m), \quad m \in \mathbb{N}_0.$$

In other words, given any initial $x_0 \in X$, a trajectory of T is computed by recursively setting x_{m+1} equal to an arbitrary element in $X(p_m)$. A set $C \subset X$ is *weakly positively invariant with respect to* T if, for any $x_0 \in C$, there exists $x \in T(x_0)$ such that $x \in C$. A point x_0 is said to be a *fixed point of* T if $x_0 \in T(x_0)$.

Lemma 3.4: Let $T : X \to \mathfrak{P}(X)$ be a closed map with non-empty values. The ω-limit set of any bounded trajectory $\{x_m\}_{m \in \mathbb{N}_0}$ of T is positively invariant with respect to T.

A function $V : X \to \mathbb{R}$ is *non-increasing along* $T : X \to \mathfrak{P}(X)$ *on* $A \subset X$ if $V(x') \leq V(x)$ for all $x \in A$ and $x' \in T(x)$. We are ready to state the following result (see [6]).

Theorem 3.5: (LaSalle Invariance Principle for closed set-valued maps): Let $T : X \to \mathfrak{P}(X)$ be a closed set-valued map on $W \subset X$ and let $V : X \to \mathbb{R}$ be a continuous function non-increasing along T on W. Assume the trajectory $\{x_m\}_{m \in \mathbb{N}_0}$ of T takes values in W and is bounded. Then there exists $c \in \mathbb{R}$ such that $x_m \longrightarrow M \cap V^{-1}(c)$, where M denotes the largest weakly positively invariant set contained in $\{x \in \overline{W} \mid \exists x' \in T(x) \text{ such that } V(x') = V(x)\}$.

In general, the constant c in Theorem 3.5 depends on the specific trajectory of T.

IV. Synchronous executions of coordination algorithms

In this section, we begin by introducing the notion of network of robotic agents. Next, we define what a synchronous execution of a motion coordination algorithm by the network of robotic agents is. Then we associate to a spatially distributed coordination algorithm various set-valued maps that serve to model its possible executions. Finally, we build on the tools presented in the previous sections to provide correctness and robustness results for general motion coordination algorithms.

Let us 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's state p_i belongs to a d-dimensional topological space X, $d \in \mathbb{N}$, i.e., $p_i \in X$. Typically, X corresponds to the Euclidean space. The agent can change its state at any time $m \in \mathbb{N}$, for any unit period of time, according to the discrete-time control system

$$p_i(m+1) = \mathrm{dyn}(p_i(m), u_i), \tag{1}$$

where $\mathrm{dyn} : X \times U \to X$ and the control u_i takes values in a bounded subset U of $\mathbb{R}^d$. The sensing and communication model is determined by a proximity graph $\mathcal{G}$ over X. The processor of each agent has access to its state, and transmits this information to any other neighboring agent in the graph $\mathcal{G}$. Equivalently, we shall consider groups of robotic agents without communication capabilities, but instead capable of measuring the relative state of each other neighboring agent in the graph $\mathcal{G}$.

A *coordination algorithm* for a network of robotic agents is a map $f : X^n \to X^n$. We denote by $f_i : X^n \to X$ the ith component of f. We say that f is *executable* by the network of robotic agents if for each $p \in X$, there exist $u_1, \ldots, u_n \in U$ such that $\mathrm{dyn}(p, u_i) = f_i(p)$, $i \in \{1, \ldots, n\}$. We distinguish between the absence and presence, respectively, of link failures in the network topology. A *synchronous execution* of the coordination algorithm f by the network of robotic agents is a trajectory of f, i.e.,

$$p_i(m+1) = f_i(p_1(m), \ldots, p_n(m)), \quad i \in \{1, \ldots, n\}. \tag{2}$$

A convenient short-hand notation is given by the sequence $\{P_m = f^m(P_0)\}_{m\in\mathbb{N}_0}$, with $P_0 = (p_1(0), \dots, p_n(0))$. Provided the map f is spatially distributed over the proximity graph $\mathcal{G}$ –which we assume in the following– the coordination algorithm is implementable over the robotic network.

Let us now introduce the notion of synchronous execution with link failures of f. Let $\widetilde{f}$ be the map associated to f for $\mathcal{G}$. For each $G \in \mathfrak{P}(FT_{\mathcal{G}})$, consider the map $f_G : X^n \to X^n$ whose jth component is given by

$$f_{G,j}(p_1, \dots, p_n) = \widetilde{f}(p_j, P(\mathcal{N}_G(j))).$$

In particular, note that $f(P) = f_{G_P}(P)$ for all $P \in X^n$. A *synchronous execution* of the coordination algorithm f *with link failures* by the network of robotic agents is a sequence $\{P_m\}_{m\in\mathbb{N}_0}$ with the property that for each $m \in \mathbb{N}_0$, there exists $G_m \in \mathfrak{P}(G_{P_m})$ such that

$$P_{m+1} = f_{G_m}(P_m). \tag{3}$$

We refer to $\{G_m\}_{m\in\mathbb{N}_0}$ as the *graphs with link failures* associated to the execution $\{P_m\}_{m\in\mathbb{N}_0}$. Note that a synchronous execution of f as defined in (2) is simply a synchronous execution of f with (no) link failures as defined in (3) with $G_m = G_{P_m}$ for all $m \in \mathbb{N}_0$.

A. Closed set-valued maps associated with a spatially distributed coordination algorithm

Set-valued maps provide a flexible tool to model executions of coordination algorithms and help establish their correctness and robustness properties. A relevant instance of this assertion is given by the following situation: often times, a motion coordination algorithm is not continuous in the agents' position because of the discontinuous changes in neighboring relationships of the network topology. If this is the case, one can invoke various closed set-valued maps in the stability analysis of the executions of the coordination algorithm. This is what we establish next.

Given a coordination algorithm $f : X^n \to X^n$, spatially distributed over a proximity graph $\mathcal{G}$, consider

- the set-valued map $T_f : X^n \to \mathfrak{P}(X^n)$ defined by
$$T_f(P) = \{f_G(P) \mid G \in FT_{\mathcal{G}}(P)\},$$
- for $\ell \in \mathbb{N}$ and $\mathcal{S} \subset \mathfrak{P}(FT_{\mathcal{G}})^\ell$, the set-valued map $T_{f,\mathcal{S}} : X^n \to \mathfrak{P}(X^n)$
$$T_{f,\mathcal{S}}(P) = \{f_{G_\ell} \circ \cdots \circ f_{G_1}(P) \mid (G_1, \dots, G_\ell) \in \mathcal{S}\}.$$

The difference between these maps lies in the fact that the number of fixed-topology graphs can change with P in the definition of T_f, whereas it remains fixed for all $P \in X^n$ in the definition of $T_{f,\mathcal{S}}$. We will show in Section IV-B that, depending on the properties of the coordination algorithm, each of these set-valued maps will play a key role in characterizing the correctness and robustness of its executions.

Proposition 4.1: Let $f : X^n \to X^n$ be spatially distributed over a proximity graph $\mathcal{G}$. Assume the map associated to f for $\mathcal{G}$ is continuous. Then

(i) T_f is closed;
(ii) for any $\ell \in \mathbb{N}$ and $\mathcal{S} \subset \mathfrak{P}(FT_{\mathcal{G}})^\ell$, $T_{f,\mathcal{S}}$ is closed.

B. Correctness and robustness analysis of coordination algorithms

Here, we leverage on the set-valued map approach presented in Section IV-A. The main idea is that, by associating suitable set-valued maps (i.e., T_f or $T_{f,\mathcal{S}}$, for some $\mathcal{S} \subset \mathfrak{P}(FT_{\mathcal{G}})^\ell$) to the coordination algorithm, we will be able to provide theoretical guarantees of its correctness and robustness properties. From a practical point of view, this "proof technique" has no effect on the actual implementation of the coordination algorithm itself.

We start by stating a useful correctness result.

Proposition 4.2 (Correctness): Let $f : X^n \to X^n$ be spatially distributed over a proximity graph $\mathcal{G}$. Assume the map associated to f for $\mathcal{G}$ is continuous, and let $V : X^n \to \mathbb{R}$ be continuous and non-increasing along f. Let $P_0 \in X^n$ and assume the sequence $\{P_m = f^m(P_0)\}_{m\in\mathbb{N}_0}$ is bounded. Then there exists $c \in \mathbb{R}$ such that

(i) $P_m \longrightarrow M_1 \cap V^{-1}(c)$, where M_1 denotes the largest weakly positively invariant set with respect to T_f contained in $\{P \in X^n \mid \exists P' \in T_f(P)$ such that $V(P') = V(P)\}$.
(ii) for $\mathcal{S} \subset \mathfrak{P}(FT_{\mathcal{G}})$ such that $\{P_m\}_{m\in\mathbb{N}_0}$ is a trajectory of $T_{f,\mathcal{S}}$, then $P_m \longrightarrow M_2 \cap V^{-1}(c)$, where M_2 denotes the largest weakly positively invariant set with respect to $T_{f,\mathcal{S}}$ contained in $\{P \in X^n \mid \exists P' \in T_{f,\mathcal{S}}(P)$ such that $V(P') = V(P)\}$.

It is important to realize that the choice of $\mathcal{S} = FT_{\mathcal{G}}$ is always possible in the statement of Proposition 4.2(ii) (given that any trajectory of f is in particular a trajectory of the set-valued map $T_{f,FT_{\mathcal{G}}}$). In general, however, the smaller the set $\mathcal{S}$, the smaller the corresponding invariant set M, and the tighter the resulting convergence results on f.

Note also that to establish the result in Proposition 4.2, it is enough to require that the function $V : X^n \to \mathbb{R}$ is non-increasing with respect to f, i.e., it is not necessary that V is non-increasing with respect to T_f or $T_{f,\mathcal{S}}$. Under stronger assumptions on the evolution of V with respect to the fixed-topology graphs associated with the proximity graph $\mathcal{G}$, one can establish an important robustness result.

Proposition 4.3 (Robustness): Let $f : X^n \to X^n$ be spatially distributed over a proximity graph $\mathcal{G}$. Assume the map associated to f for $\mathcal{G}$ is continuous. Let $\ell \in \mathbb{N}$ and $\mathcal{S} \subset \mathfrak{P}(FT_{\mathcal{G}})^\ell$. For $P_0 \in X^n$, let $\{P_m\}_{m\in\mathbb{N}_0}$ be a bounded synchronous execution of f with link failures, i.e., $P_{m+1} = f_{G_m}(P_m)$ for $m \in \mathbb{N}_0$ and $G_m \in \mathfrak{P}(G_{P_m})$, with the property that any ℓ consecutive graphs with link failures belong to $\mathcal{S}$, i.e., $(G_m, \dots, G_{m+\ell-1}) \in \mathcal{S}$ for all $m \in \mathbb{N}_0$. Assume $V : X^n \to \mathbb{R}$ is a continuous function non-increasing along $T_{f,\mathcal{S}}$. Then there exists $c_0, \dots, c_{\ell-1} \in \mathbb{R}$ such that

$$P_m \longrightarrow M \cap (V^{-1}(c_0) \cup \cdots \cup V^{-1}(c_{\ell-1})),$$

where M denotes the largest weakly positively invariant set with respect to $T_{f,\mathcal{S}}$ contained in $\{P \in X^n \mid \exists P' \in T_{f,\mathcal{S}}(P)$ such that $V(P') = V(P)\}$.

V. Applications to sample problems

In this section we illustrate the suitability of the tools developed in the previous sections to characterize the behavior of mobile networks executing coordination algorithms.

A. Rendezvous

Here, we consider the rendezvous coordination task originally introduced in [15], and later studied in [16], [17], [6]. Roughly speaking, rendezvous consists of achieving agreement over the location of the agents. Consider a network of robotic agents evolving on $X = \mathbb{R}^d$, $d \in \mathbb{N}$, with dynamics $p_i(m+1) = p_i(m) + u_i$, for $i \in \{1, \dots, n\}$. Here, we first introduce the circumcenter coordination algorithm $f_{\rm CC}$, we study some of its properties, and then we characterize its asymptotic convergence features.

1) Circumcenter coordination algorithm: In order to define the circumcenter algorithm, we need to introduce some preliminary notation. 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:

$$\begin{aligned} &\text{maximize } \lambda \\ &\text{subject to } \lambda \leq 1, \ (1-\lambda)q_0 + \lambda q_1 \in Q. \end{aligned} \tag{4}$$

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. The solution exists and is unique. Moreover, it depends continuously on the data q_0, q_1 and Q. Another piece of useful notation is the following. For $p \in \mathbb{R}^d$, $A \in \mathbb{F}(\mathbb{R}^d)$ and $r \in \mathbb{R}_+$, define

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

We are now ready to define the circumcenter coordination algorithm. Let $r \in \mathbb{R}_+$, and let $\mathcal{G}$ be a proximity graph spatially distributed over $\mathcal{G}_{\rm disk}(r)$. Define the map $f_{\rm CC} : (\mathbb{R}^d)^n \to (\mathbb{R}^d)^n$ with ith component given by

$$f_{{\rm CC},i}(p_1, \dots, p_n) = (1 - \lambda_i^*)\, p_i + \lambda_i^* \ {\rm CC}(\mathcal{M}_i),$$

where $\lambda_i^* = \lambda(p_i, {\rm CC}(\mathcal{M}_i), C_{p_i,r}(\mathcal{N}_{\mathcal{G},p_i}(\mathcal{P})))$ and $\mathcal{M}_i = \mathcal{N}_{\mathcal{G},p_i}(\mathcal{P}) \cup \{p_i\}$. Here ${\rm CC}(S)$ denotes the circumcenter of S (the center of the smallest-radius sphere that encloses S).

2) Properties of the coordination algorithm: By definition of $f_{\rm CC}$, if $\|p_i - p_j\| \leq r$, then $f_{{\rm CC},i}(P), f_{{\rm CC},j}(P) \in \overline{B}(\frac{p_i+p_j}{2}, \frac{r}{2})$, which in particular implies that $\|f_{{\rm CC},i}(P) - f_{{\rm CC},j}(P)\| \leq r$. Therefore, $f_{\rm CC}$ preserves neighboring relationships in the graph $\mathcal{G}_{\rm disk}(r)(P)$.

Another important property of $f_{\rm CC}$ is that it is spatially distributed over the proximity graph $\mathcal{G}$. The associated map $\widetilde{f_{\rm CC}} : \mathbb{R}^d \times \mathbb{F}(\mathbb{R}^d) \to \mathbb{R}^d$ is given by

$$\widetilde{f_{\rm CC}}(p, A) = (1 - \lambda^*)\, p + \lambda^* \ {\rm CC}(A \cup \{p\}),$$

where $\lambda^* = \lambda(p, {\rm CC}(A \cup \{p\}), C_{p,r(p,A)}(A))$ and $r(p, A) = \max\{r, \|p - q\| \mid q \in A\}$. Given that the circumcenter of a finite set of points depends continuously on their location, one can show that the function $\widetilde{f_{\rm CC}}$ is continuous.

A final property of interest in the discussion is the fact that the map $f_{{\rm CC},G}$ associated with each $G \in \mathfrak{P}(FT_{\mathcal{G}})$ verifies

$${\rm co}(f_{{\rm CC},G}(P)) \subset {\rm co}(P), \quad P \in (\mathbb{R}^d)^n. \tag{5}$$

Here, ${\rm co}(S)$ is the convex hull of S. Equation (5) is a consequence of the fact that for $S \subset \mathbb{R}^d$ finite, the circumcenter satisfies ${\rm CC}(S) \in {\rm co}(S) \setminus V({\rm co}(S))$, where $V(Q)$ denotes the set of strictly convex vertices of the polytope $Q \subset \mathbb{R}^d$.

3) Correctness and robustness characterization: Consider the function $V_{\rm diam} : (\mathbb{R}^d)^n \to \mathbb{R}$ defined by

$$V_{\rm diam}(p_1, \dots, p_n) = \max\{\|p_i - p_j\| \mid i, j \in \{1, \dots, n\}\}.$$

This function is continuous. Moreover, from the discussion in Section V-A.2, we deduce that $V_{\rm diam}$ is non-increasing along $T_{f_{\rm CC}, \mathfrak{P}(FT_{\mathcal{G}})}$. The following result characterizes the correctness properties of the circumcenter coordination algorithm (see [15], [16], [6]).

Proposition 5.1: Let $r \in \mathbb{R}_+$, and consider the proximity graph $\mathcal{G}_{\rm disk}(r)$. Then, for any synchronous execution $\{P_m\}_{m \in \mathbb{N}_0}$ of $f_{\rm CC}$ such that $\mathcal{G}_{\rm disk}(r)(P_{m_0})$ is connected for some $m_0 \in \mathbb{N}_0$, there exists $(p^*, \dots, p^*) \in {\rm diag}((\mathbb{R}^d)^n)$ such that $P_m \to (p^*, \dots, p^*)$ as $m \to +\infty$.

Proposition 5.1 is a corollary of the next, stronger result. This theorem builds on the tools introduced in Section IV to establish the robustness properties of the circumcenter coordination algorithm.

Theorem 5.2: Let $r \in \mathbb{R}_+$, and consider the proximity graph $\mathcal{G}_{\rm disk}(r)$. Let $\{P_m\}_{m \in \mathbb{N}_0}$ be a synchronous execution with link failures of $f_{\rm CC}$ with the following property: there exits $\ell \in \mathbb{N}$ such that the union of any ℓ consecutive graphs with link failures of $\{P_m\}_{m \in \mathbb{N}_0}$ is strongly connected. Then, there exists $(p^*, \dots, p^*) \in {\rm diag}((\mathbb{R}^d)^n)$ such that $P_m \to (p^*, \dots, p^*)$ as $m \to +\infty$.

Fig. 1 illustrates the result in Theorem 5.2.

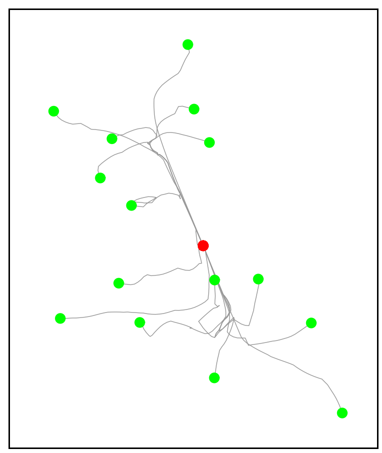

Fig. 1. Synchronous execution of $f_{\rm CC}$ with link failures over $\mathcal{G}_{\rm disk}(r)$. The 15 vehicles have a communication radius $r = 4$ and are initially deployed over the square $[-7, 7] \times [-7, 7]$. At each time step, two randomly selected agents fail and are not able to detect any of its neighbors. Nevertheless, rendezvous is asymptotically achieved (cf. Theorem 5.2).

Remark 5.3: For a synchronous execution with no link failures, note that if $\mathcal{G}_{\rm disk}(r)(P_{m_0})$ is connected, then $\mathcal{G}_{\rm disk}(r)(P_m)$ is connected for $m \geq m_0$ (since $f_{\rm CC}$ preserves neighboring relationships, cf. Section V-A.2). Therefore, Theorem 5.2 with $\ell = 1$ implies Proposition 5.1. Both results actually are also valid for the circumcenter algorithm executed over any proximity graph which is spatially distributed over $\mathcal{G}_{\rm disk}(r)$ and has the same connected components, but here we only consider $\mathcal{G}_{\rm disk}(r)$ for simplicity. •

B. Flocking

Here, we consider the average-heading coordination algorithm first analyzed in [1] and later in [7], [8] (see also [18]). Under some assumptions, this algorithm achieves flocking, i.e., agreement upon the direction of motion of the agents. We show how the application of the tools introduced above yields similar convergence results. Consider a network of robotic agents evolving on $X = \mathbb{R}^2 \times \mathbb{S}^1$ (position in $\mathbb{R}^2$ and orientation in $\mathbb{S}^1$, respectively) with dynamics $(p_i(m+1), \theta_i(m+1)) = (p_i(m), \theta_i(m)) + u_i$, for $i \in \{1, \dots, n\}$. Here, we first define the average-heading coordination algorithm f_{Ave}, we study some of its properties, and then we characterize its asymptotic convergence features.

1) Average-heading coordination algorithm: Let $\mathcal{G}$ be a proximity graph on $\mathbb{R}^2 \times \mathbb{S}^1$. Define the map $f_{\text{Ave}} : (\mathbb{R}^2 \times \mathbb{S}^1)^n \to (\mathbb{R}^2 \times \mathbb{S}^1)^n$ with ith component given by

$$f_{\text{Ave},i}((p_1, \theta_1), \dots, (p_n, \theta_n)) = \big(p_i + (\cos\theta_i, \sin\theta_i), \\ \text{Average}\big(\theta_i \cup \{\theta_j \mid (p_j, \theta_j) \in \mathcal{N}_{\mathcal{G}}(p_i, \theta_i)\}\big)\big).$$

For simplicity, we will use the short-hand notation $(P, \Theta) = ((p_1, \theta_1), \dots, (p_n, \theta_n)) \in (\mathbb{R}^2 \times \mathbb{S}^1)^n$.

2) Properties of the coordination algorithm: An important property of f_{Ave} is that it is spatially distributed over the proximity graph $\mathcal{G}$. The associated map $\widetilde{f_{\text{Ave}}} : (\mathbb{R}^2 \times \mathbb{S}^1) \times \mathbb{F}(\mathbb{R}^2 \times \mathbb{S}^1) \to \mathbb{R}^2 \times \mathbb{S}^1$ is given by

$$\widetilde{f_{\text{Ave}}}((p, \theta), A) = \big(p + (\cos\theta, \sin\theta), \\ \text{Average}\big(\theta \cup \{\bar{\theta} \mid (\bar{p}, \bar{\theta}) \in A\}\big)\big).$$

Clearly, the function $\widetilde{f_{\text{Ave}}}$ is continuous.

3) Correctness and robustness characterization: Consider the function $V_{\text{max-min}} : (\mathbb{S}^1)^n \to \mathbb{R}$ defined by

$$V_{\text{max-min}}(\theta_1, \dots, \theta_n) = \\ \max\{\theta_i \mid i \in \{1, \dots, n\}\} - \min\{\theta_i \mid i \in \{1, \dots, n\}\}.$$

This function is continuous, and so is the composition $V_{\text{max-min}} \circ \pi_{(\mathbb{S}^1)^n}$, with $\pi_{(\mathbb{S}^1)^n} : (\mathbb{R}^2 \times \mathbb{S}^1)^n \to (\mathbb{S}^1)^n$ the natural projection. From the definition of f_{Ave}, we deduce that $V_{\text{max-min}} \circ \pi_{(\mathbb{S}^1)^n}$ is non-increasing along $T_{f_{\text{Ave}}, \mathfrak{P}(FT_{\mathcal{G}})}$. With these ingredients and the tools presented above, one is ready to characterize the correctness and robustness properties of the average-heading coordination algorithm as in [1], [7], [8]. We include here the proof for reference.

Theorem 5.4: Let $\mathcal{G}$ be a proximity graph on $\mathbb{R}^2 \times \mathbb{S}^1$. Let $\{(P_m, \Theta_m)\}_{m \in \mathbb{N}_0}$ be a synchronous execution with link failures of f_{Ave} with the following property: there exits $\ell \in \mathbb{N}$ such that the union of any ℓ consecutive graphs with link failures of $\{(P_m, \Theta_m)\}_{m \in \mathbb{N}_0}$ is strongly connected. Then, there exists $(\theta^*, \dots, \theta^*) \in \text{diag}((\mathbb{R}^d)^n)$ such that $\Theta_m \to (\theta^*, \dots, \theta^*)$ as $m \to +\infty$.

VI. Conclusions

We have developed some promising analysis tools for the characterization of the correctness and robustness properties of motion coordination algorithms. The technical approach has built on computational geometric tools (proximity graphs), nondeterministic dynamical systems (set-valued maps) and Lyapunov stability analysis (LaSalle Invariance Principles). The results obtained here have been applied to coordination algorithms achieving rendezvous and flocking. Future work will be devoted to the establishment of similar convergence properties of other cooperative strategies, the application of the results to the design of robust coordination algorithms for deployment and coverage, the exploration of the applicability of this approach to asynchronous executions, and the development of similar tools to characterize robustness to agents departures and arrivals.

Acknowledgments

This research was partially supported by NSF CAREER award ECS-0546871 and NASA University Aligned Research Program Award TO.030.MM.D.

References

[1] 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.

[2] J. A. Marshall, M. E. Broucke, and B. A. Francis, "Formations of vehicles in cyclic pursuit," *IEEE Transactions on Automatic Control*, vol. 49, no. 11, pp. 1963– 1974, 2004.

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

[4] 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.

[5] J. Cortés and F. Bullo, "Coordination and geometric optimization via distributed dynamical systems," *SIAM Journal on Control and Optimization*, vol. 44, no. 5, pp. 1543–1574, 2005.

[6] J. Cortés, S. Martínez, and F. Bullo, "Robust rendezvous for mobile networks via proximity graphs in arbitrary dimensions," *IEEE Transactions on Automatic Control*, vol. 51, no. 7, 2006. To appear.

[7] L. Moreau, "Stability of multiagent systems with time-dependent communication links," *IEEE Transactions on Automatic Control*, vol. 50, no. 2, pp. 169–182, 2005.

[8] D. Angeli and P.-A. Bliman, "Stability of leaderless multi-agent systems. Extension of a result by Moreau," Nov. 2004. http://arxiv.org/abs/math/0411338.

[9] S. Martínez, F. Bullo, J. Cortés, and E. Frazzoli, "On synchronous robotic networks – Part I: models, tasks and complexity notions," *IEEE Transactions on Automatic Control*, Mar. 2005. Submitted.

[10] J. A. Marshall and M. E. Broucke, "On invariance of cyclic group symmetries in multiagent formations," in *IEEE Conf. on Decision and Control*, (Seville, Spain), pp. 746–751, Dec. 2005.

[11] A. R. Teel, "Nonlinear systems: discrete-time stability analysis." Lecture Notes, University of California at Santa Barbara, 2004.

[12] M. de Berg, M. van Kreveld, M. Overmars, and O. Schwarzkopf, *Computational Geometry: Algorithms and Applications*. New York: Springer Verlag, 2 ed., 2000.

[13] 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.

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

[15] 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.

[16] 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.

[17] 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.

[18] H. Tanner, A. Jadbabaie, and G. J. Pappas, "Stable flocking of mobile agents, Part I: Fixed topology," in *IEEE Conf. on Decision and Control*, (Maui, HI), pp. 2010–2015, Dec. 2003.