Agreeing Asynchronously: Announcement of Results*

M. Cao
Yale University
m.cao@yale.edu

A. S. Morse
Yale University
as.morse@yale.edu

B. D. O. Anderson
Australian National University
and National ICT Australia Ltd.
brian.anderson@anu.edu.au

***Abstract*— This paper formulates and solves a continuous-time version of the widely studied Vicsek consensus problem in which each agent independently updates its heading at times determined by its own clock. It is not assumed that the agents' clocks are synchronized or that the "event" times between which any one agent updates its heading are evenly spaced. Heading updates need not occur instantaneously. Using the concept of "analytic synchronization" together with several key results concerned with properties of "compositions" of directed graphs, it is shown that the conditions under which a consensus is achieved are essentially the same as those applicable in the synchronous discrete-time case provided the notion of an agent's neighbor between its event times is appropriately defined.**

I. Introduction

In a recent paper Vicsek and co-authors [1] consider a simple discrete-time model consisting of n autonomous agents or particles all moving in the plane with the same speed but with different headings. Each agent's heading is updated using a local rule based on the average of the headings of its "neighbors." In their paper, Vicsek *et al.* provide a variety of interesting simulation results which demonstrate that the nearest neighbor rule they are studying can cause all agents to eventually move in the same direction despite the absence of centralized coordination and despite the fact that each agent's set of nearest neighbors can change with time. Vicsek's problem is what in computer science is called a "consensus problem" [2] or an "agreement problem." Roughly speaking, one has a group of agents which are all trying to agree on a specific value of some quantity. Each agent initially has only limited information available. The agents then try to reach a consensus by communicating what they know to their neighbors either just once or repeatedly, depending on the specific problem of interest. For the Vicsek problem, each agent always knows only its own heading and the headings of its neighbors. One feature of the Vicsek problem which sharply distinguishes it from other consensus problems, is that each agent's neighbors change with time, because all agents are in motion. It has recently been explained why Vicsek's agents are able to reach a common heading [3], [4], [5], [6], [7].

* The research of Cao and Morse was supported in part, by grants from the Army Research Office and the National Science Foundation and by a gift from the Xerox Corporation. The research of Anderson was supported by National ICT Australia, which is funded by the Australian Government's Department of Communications, Information Technology and the Arts and the Australian Research Council through the Backing Australia's Ability initiative and the ICT Centre of Excellence Program.

In this paper we consider a continuous-time version of the Vicsek problem in which each agent independently updates its heading at times determined by its own clock. We do not assume that the agents' clocks are synchronized or that the "event times" between which any one agent updates its heading are evenly spaced. In contrast to prior work addressed to asynchronous consensus [8], [9], heading updates need not occur instantaneously. As a consequence, it is not so clear at the outset how to construct from the asynchronous update model we consider, the type of discrete-time state equation upon which the formulation of the problem addressed in [8] depends. For the problem considered in this paper, the deriving of conditions under which all agents eventually move with the same heading requires the analysis of the asymptotic behavior of an overall *asynchronous* continuous-time process which models the n-agent system. We carry out the analysis by first embedding this asynchronous process in a suitably defined *synchronous* discrete-time, dynamical system $\mathbb{S}$ using the concept of *analytic synchronization* outlined previously in [10], [11]. This enables us to bring to bear key results derived in [12] to characterize a rich class of system trajectories under which consensus is achieved. In particular, we prove that the conditions under which a consensus is achieved are essentially the same as those in the synchronous discrete-time case studied in [4], [5], [12] provided the notion of an agent's neighbor between its event times is appropriately defined.

II. Asynchronous System

The system to be studied consists of n autonomous agents, labelled 1 through n, all moving in the plane with the same speed but with different headings. Each agent's heading is updated using a simple local rule based on the average of its own heading plus the headings of its "neighbors." Agent i's *neighbors* at time t, are those agents, including itself, which are either in or on a closed disk of pre-specified radius r_i centered at agent i's current position. In the sequel $\mathcal{N}_i(t)$ denotes the set of labels of those agents which are neighbors of agent i at time t. In contrast to earlier work [3], [4], [5], [6], [7], this paper considers a version of the flocking problem in which each agent independently updates its heading at times determined by its own clock. We assume for $i \in \{1, 2, \ldots, n\}$ that agent i's *event times* $t_{i0}, t_{i1}, \ldots, t_{ik}, \ldots$ satisfy the constraints

$$\bar{T}_i \geq t_{i(k+1)} - t_{ik} \geq T_i, \quad k \geq 0 \tag{1}$$

where $t_{i0} = 0$ and $\bar{T}_i$ and T_i are positive numbers.

Updating of agent i's heading is done as follows. At its kth event time t_{ik}, agent i senses the headings $\theta_j(t_{ik}),\ \ j \in \mathcal{N}_i(t_{ik})$ of its current neighbors and from this data computes its kth way-point $w_i(t_{ik})$. We will consider way point rules based on averaging. In particular

$$w_i(t_{ik}) = \frac{1}{n_i(t_{ik})}\left(\sum_{j\in\mathcal{N}_i(t_{ik})} \theta_j(t_{ik})\right),\ i \in \{1, 2, \ldots, n\} \tag{2}$$

where $n_i(t_{ik})$ is the number of indices in $\mathcal{N}_i(t_{ik})$. Agent i then changes its heading from $\theta_i(t_{ik})$ to $w_i(t_{ik})$ on the interval $(t_{ik}, t_{i(k+1)}]$. Thus

$$\theta_i(t_{i(k+1)}) = w_i(t_{ik}), \quad i \in \{1, 2, \ldots, n\}, \quad k \geq 0 \tag{3}$$

Although we will not be concerned about the precise manner in which the value of each θ_i changes between way-points, we will assume that for each $i \in \{1, 2, \ldots, n\}$, there is a piece-wise continuous signal $\mu_i : [0, \infty) \to [0,\ 1]$ satisfying $\mu(t_{ik}) = 1$ and $\lim_{t\downarrow t_{ik}} \mu_i(t) = 0$ for all $k \geq 0$, such that

$$\begin{aligned}&\theta_i(t) = \theta_i(t_{ik}) + \mu_i(t)(w_i(t_{ik}) - \theta_i(t_{ik})),\\ &t \in (t_{ik}, t_{i(k+1)}],\ \ k \geq 0,\ i \in \{1, 2, \ldots, n\}\end{aligned} \tag{4}$$

For $i \in \{1, 2, \ldots, n\}$, let $\mathcal{M}_i$ denote the class of all piecewise continuous signals $\rho : [0, \infty) \to [0,\ 1]$ satisfying $\lim_{t\downarrow t_{ik}} \rho(t) = 0$ and $\rho(t_{ik}) = 1$ for all $k \geq 0$. The assumption that (4) holds for some $\mu_i \in \mathcal{M}_i$, is equivalent to assuming that θ_i is at least piecewise continuous and that

$$\begin{aligned}|\theta_i(t) - \theta_i(t_{ik})| \leq |w_i(t_{ik}) - \theta_i(t_{ik})|,&\\ t \in (t_{ik}, t_{i(k+1)}], \quad k \geq 0&\end{aligned} \tag{5}$$

Clearly (4) implies (5); on the other hand if (5) holds and we define $\mu_i : [0, \infty) \to [0, 1]$ on $(t_{ik}, t_{i(k+1)}]$ as

$$\mu_i(t) = \begin{cases} \frac{\theta_i(t)-\theta_i(t_{ik})}{w_i(t_{ik})-\theta_i(t_{ik})} & \text{if } w_i(t_{ik}) \neq \theta_i(t_{ik}) \\ 1 & \text{if } w_i(t_{ik}) = \theta_i(t_{ik}) \end{cases}$$

then μ_i will be in $\mathcal{M}_i$ and (4) will hold.

For μ_i to be in $\mathcal{M}_i$ means that μ_i could be constant at the value 1 on each interval $(t_{ik}, t_{i(k+1)})$; this would mean that just after t_{ik}, θ_i would jump discontinuously from its value at t_{ik} to $w_i(t_{ik})$ and remain constant at this value until just after $t_{i(k+1)}$ [9]. More realistically, μ_i might change continuously from 0 to 1 on $(t_{ik}, t_{i(k+1)})$ which would imply that θ_i is continuous on $[0, \infty)$. Under any conditions equations (2) and (4) completely describe the temporal evolution of the n agent asynchronous system of interest.

A. Extended Neighbor Graph

The explicit form of the update equations determined by (2) and (4) depends on the relationships between neighbors which exist at each agent's event times. It is possible to describe all neighbor relationships at any time t using a directed graph $\mathbb{N}(t)$ with vertex set $\mathcal{V} = \{1, 2, \ldots n\}$ and arc set $\mathcal{A}(\mathbb{N}) \subset \mathcal{V} \times \mathcal{V}$ which is defined in such a way so that (i, j) is an *arc* or directed edge from i to j just in case agent i is a neighbor of agent j at time t. Thus $\mathbb{N}(t)$ is a directed graph on n vertices with at most one arc between each ordered pair of vertices and with exactly one self-arc at each vertex. We write $\mathcal{G}_{sa}$ for the set of all such graphs. It is natural to call a vertex i a *neighbor* of vertex j in any graph $\mathbb{G}$ in $\mathcal{G}_{sa}$ if (i, j) is an arc in $\mathbb{G}$.

Although the neighbors of each agent i are well defined at event times of other agents, what's important for modelling agent i's updates are the headings of neighboring agents *only* at agent i's own event times. We deal with this matter by re-defining each agent's neighbor set at times between its own event times to consist of only itself. Our reason for doing this will become clear later when, for purposes of analysis, we use analytic synchronization to embed the n agent asynchronous model defined by (2) and (4) in a synchronous dynamical system.

To proceed, let $\mathcal{T}$ denote the set of all event times of all n agents. Relabel the elements of $\mathcal{T}$ as $t_0, t_1, t_2, \cdots$ in such a way so that $t_0 = 0$ and $t_\tau < t_{\tau+1},\ \tau \in \{0, 1, 2, \ldots\}$. For $i \in \{1, 2, \ldots, n\}$, let $\mathcal{T}_i$ denote the set of $t_\tau \in \mathcal{T}$ which are event times of agent i. For each $i \in \{1, 2, \ldots, n\}$ define

$$\bar{\mathcal{N}}_i(\tau) = \begin{cases} \mathcal{N}_i(t_\tau) & \text{if } t_\tau \in \mathcal{T}_i \\ i & \text{if } t_\tau \notin \mathcal{T}_i, \end{cases} \tag{6}$$

Thus $\bar{\mathcal{N}}_i(\tau)$ coincides with $\mathcal{N}_i(t_\tau)$ whenever t_τ is an event time of agent i, and is simply the single index i otherwise.

Much like $\mathbb{N}(t)$ which describes the original neighbor relations of system (2), (3) at time t, we describe all re-defined neighbor relationships at time $\tau \in \{0, 1, \ldots\}$ to be the directed graph $\bar{\mathbb{N}}(\tau)$ with vertex set $\mathcal{V}$ and arc set $\mathcal{A}(\bar{\mathbb{N}}(\tau)) \subset \mathcal{V} \times \mathcal{V}$ which is defined so that (i, j) is an arc from i to j just in case agent j is in the neighbor set $\bar{\mathcal{N}}_i(\tau)$. Thus like the neighbor graphs $\mathbb{N}(t)$, each $\bar{\mathbb{N}}(\tau)$ is a directed graph on n vertices with at most one arc between each ordered pair of vertices and with exactly one self-arc at each vertex. We call $\bar{\mathbb{N}}(\tau)$ the *extended neighbor graph* of the system (2) and (3) at time τ.

B. Objective

A complete description of the asynchronous system defined by (2) and (4) would have to include a model which explains how the $\mu_i(t)$ and $\mathcal{N}_i(t)$ change over time as functions of the positions of the n agents in the plane. While such a model is easy to derive and is essential for simulation purposes, it would be difficult to take into account in a convergence analysis. To avoid this difficulty, we shall adopt a more conservative approach which ignores how the $\mathcal{N}_i(t)$ and the $\mu_i(t)$ depend on the agent positions in the plane and assumes instead that each might be any function in some suitably defined set of interest. Our ultimate objective is to show for any initial set of agent headings, any set of $\mu_i \in \mathcal{M}_i,\ i \in \{1, 2, \ldots n\}$ and for a large class of functions $t \longmapsto \mathcal{N}_i(t)$, that the headings of all n agents will converge to the same steady state value θ_{ss}.

III. Main Results

To state our main result, we need a few ideas from [12]. We call a vertex i of a directed graph $\mathbb{G}$, a *root* of $\mathbb{G}$ if for each other vertex j of $\mathbb{G}$, there is a path from i to j. Thus i is a root of $\mathbb{G}$, if it is the root of a directed spanning tree of $\mathbb{G}$. We will say that $\mathbb{G}$ is *rooted at* i if i is in fact a root. Thus $\mathbb{G}$ is rooted at i just in case each other vertex of $\mathbb{G}$ is *reachable* from vertex i along a path within the graph. $\mathbb{G}$ is *strongly rooted at* i if each other vertex of $\mathbb{G}$ is reachable from vertex i along a path of length 1. Thus $\mathbb{G}$ is strongly rooted at i if i is a neighbor of every other vertex in the graph. By a *rooted graph* $\mathbb{G}$ is meant a graph which possesses at least one root. Finally, a *strongly rooted graph* is a graph which has at least one vertex at which it is strongly rooted.

By the *composition* of two directed graphs $\mathbb{G}_p, \mathbb{G}_q$ with the same vertex set we mean that graph $\mathbb{G}_q \circ \mathbb{G}_p$ with the same vertex set and arc set defined such that (i, j) is an arc of $\mathbb{G}_q \circ \mathbb{G}_p$ if for some vertex k, (i, k) is an arc of $\mathbb{G}_p$ and (k, j) is an arc of $\mathbb{G}_q$. Let us agree to say that a finite sequence of directed graphs $\mathbb{G}_{p_1},\ \mathbb{G}_{p_2}, \dots, \mathbb{G}_{p_m}$ with the same vertex set is *jointly rooted* if the composition $\mathbb{G}_{p_m} \circ \mathbb{G}_{p_{m-1}} \circ \cdots \circ \mathbb{G}_{p_1}$ is rooted. An infinite sequence of graphs $\mathbb{G}_{p_1}, \mathbb{G}_{p_2}, \dots,$ with the same vertex set is *repeatedly jointly rooted* if there is a positive integer m for which each finite sequence $\mathbb{G}_{p_{m(k+1)}},\ \dots, \mathbb{G}_{p_{mk+1}},\quad k \geq 0$, is jointly rooted.

Equations (2) and (4) can be combined. What results is a description of the evolution of θ_i on agent i's event time set.

$$\theta_i(t_{i(k+1)}) = \frac{1}{n_i(t_{ik})} \left(\sum_{j \in \mathcal{N}_i(t_{ik})} \theta_j(t_{ik}) \right),\ i \in \{1, 2, \dots, n\}$$

In the synchronous version of the problem treated in [3], [4], [5], [6], [7], for each $k \geq 0$, the kth event times $t_{1k}, t_{2k}, \dots, t_{nk}$ of all n agents are the same. Thus in this case each agent's heading update equation can be written as

$$\theta_i(t) = \frac{1}{n_i(t_k)} \left(\sum_{j \in \mathcal{N}_i(t_k)} \theta_j(t_k) \right), \quad t \in (t_k, t_{k+1}], \quad k \geq 0 \tag{7}$$

where $t_0 = 0$ and $t_k = t_{ik}$. The main result of [3] is as follows.

Theorem 1: Let the $\theta_i(0)$ be fixed. For any trajectory of the synchronous system determined by (7) along which the sequence of neighbor graphs $\mathbb{N}(0), \mathbb{N}(1), \dots$ is repeatedly jointly rooted, there is a constant θ_{ss} for which

$$\lim_{t \to \infty} \theta_i(t) = \theta_{ss} \tag{8}$$

where the limit is approached exponentially fast.

The aim of this paper is to prove that essentially the same result holds in the face of asynchronous updating.

Theorem 2: Let the $\theta_i(0)$, $w_i(0)$, and $\mu_i \in \mathcal{M}_i$ be fixed. For any trajectory of the asynchronous system determined by (2) and (4) along which the sequence of extended neighbor graphs $\bar{\mathbb{N}}(0), \bar{\mathbb{N}}(1), \dots$ is repeatedly jointly rooted, there is a constant θ_{ss} for which

$$\lim_{t \to \infty} \theta_i(t) = \theta_{ss} \tag{9}$$

where the limit is approached exponentially fast.

It is worth noting that the validity of this theorem depends critically on the fact that there are finite positive numbers, namely $T_{\max} = \max\{\bar{T}_1, \bar{T}_2, \dots, \bar{T}_n\}$ and $T_{\min} = \min\{T_1, T_2, \dots, T_n\}$, which uniformly bound from above and below respectively, the time between any two successive event times of any agent. This is a consequence of the assumption that inequality (1) holds.

As noted in the last section, for the asynchronous problem under consideration, the only vertices of $\bar{\mathbb{N}}(\tau)$ which can have more than one neighbor, are those corresponding to agents for whom t_τ is an event time. Thus in the most likely situation when distinct agents have only distinct event times, there will be at most one vertex in each graph $\bar{\mathbb{N}}(\tau)$ which has more than one neighbor. It is this situation we want to explore further. Toward this end, let $\mathcal{G}^*_{sa} \subset \mathcal{G}_{sa}$ denote the subclass of all graphs which have at most one vertex with more than one neighbor. Note that for $n > 2$, there is no rooted graph in $\mathcal{G}^*_{sa}$. Nonetheless, in the light of Theorem 2 it is clear that convergence to a common steady state heading will occur if the infinite sequence of graphs $\bar{\mathbb{N}}(0), \bar{\mathbb{N}}(1), \dots$ is repeatedly jointly rooted. This of course would require that there exist jointly rooted sequences of graphs from $\mathcal{G}^*_{sa}$. We will now explain why such sequences do in fact exist.

Let us agree to call a graph $\mathbb{G} \in \mathcal{G}_{sa}$ an *all neighbor graph centered at* v if every vertex of $\mathbb{G}$ is a neighbor of v. Note that all neighbor graphs are maximal in $\mathcal{G}^*_{sa}$ with respect to the partial ordering of $\mathcal{G}^*_{sa}$ by inclusion, where in this context $\mathbb{G}_p \in \mathcal{G}^*_{sa}$ is contained in $\mathbb{G}_q \in \mathcal{G}^*_{sa}$ if $\mathcal{A}(\mathbb{G}_p) \subset \mathcal{A}(\mathbb{G}_q)$. Note also the composition of any all neighbor graph with itself is itself. On the other hand, because the arcs of any two graphs in $\mathcal{G}_{sa}$ are arcs in their composition, the composition of n all neighbor graphs with distinct centers must clearly be a graph in which each vertex is a neighbor of every other; i.e., the complete graph. Thus the composition of n all neighbor graphs from $\mathcal{G}^*_{sa}$ with distinct centers is strongly rooted. In summary, the hypothesis of Theorem 2 is not at all vacuous for the asynchronous problem under consideration. When that hypothesis is satisfied, convergence to a common steady state heading will occur.

IV. Analytic Synchronization

To prove Theorem 2 requires the analysis of the asymptotic behavior of the n mutually unsynchronized processes $\mathbb{P}_1, \mathbb{P}_2, \dots \mathbb{P}_n$ which the n pairs of heading equations (2), (4) define. Despite the apparent complexity of the resulting *asynchronous* system which these n interacting processes determine, it is possible to capture its salient features using a suitably defined *synchronous* discrete-time, hybrid dynamical system $\mathbb{S}$. The sequence of steps involved in defining $\mathbb{S}$ has been discussed before and is called *analytic synchronization* [10], [11]. First, all n event time sequences are merged into a single ordered sequence of event times $\mathcal{T}$, as we've already done. This clever idea has been used before in [13] to study the convergence of totally asynchronous iterative algorithms. Second, between event times each agent's neighbor set is defined to have exactly one neighbor, namely itself; this we have also already done. Third, the "synchronized" state of $\mathbb{P}_i$ is then defined to be the original of $\mathbb{P}_i$ at $\mathbb{P}_i$'s event times $\{t_{i1}, t_{i2}, \ldots\}$ plus possibly some additional state variables; at values of $t \in \mathcal{T}$ between event times t_{ik} and $t_{i(k+1)}$, the synchronized state of $\mathbb{P}_i$ is taken to be the same at the value of its state at time t_{ik}. Although it is not always possible to carry out all of these steps, in this case it is. What ultimately results is a synchronous dynamical system $\mathbb{S}$ evolving on the index set of $\mathcal{T}$, with state composed of the synchronized states of the n individual processes under consideration. We now use these ideas to develop such a synchronous system $\mathbb{S}$ for the asynchronous process under consideration.

A. Definition of $\mathbb{S}$

For each such i and each $q \in \mathcal{T}_i$ define

$$\bar{\theta}_i(\tau) \;=\; \theta_i(t_q), \;\; q \leq \tau < q' \qquad (10)$$
$$\bar{w}_i(\tau) \;=\; w_i(t_q), \;\; q \leq \tau < q' \qquad (11)$$

where $t_{q'}$ is the first event time of agent i after t_q. Note that for any $t_q \in \mathcal{T}_i$ there is always such a q' because we've assumed via (1) that the time between any two successive event times of agent i is bounded above. In the full length version of this paper it is shown that for $i \in \{1, 2, \ldots, n\}$ and $\tau > 0$

$$\bar{\theta}_i(\tau) \;=\; \bar{w}_i(\tau - 1), \quad t_\tau \in \mathcal{T}_i \qquad (12)$$
$$\bar{\theta}_i(\tau) \;=\; \bar{\theta}_i(\tau - 1), \quad t_\tau \notin \mathcal{T}_i \qquad (13)$$
$$\bar{w}_i(\tau) \;=\; \frac{1}{\bar{n}_i(\tau)} \sum_{j \in \bar{\mathcal{N}}_i(\tau)} \{(1 - \bar{\mu}_j(\tau))\bar{\theta}_j(\tau - 1) + \bar{\mu}_j(\tau)(\bar{w}_j(\tau - 1)\}, \quad t_\tau \in \mathcal{T}_i \qquad (14)$$
$$\bar{w}_i(\tau) \;=\; \bar{w}_i(\tau - 1), \quad t_\tau \notin \mathcal{T}_i \qquad (15)$$

where for $\tau \in \{0, 1, \ldots\}$, $\bar{\mu}_j(\tau) = \mu_j(t_\tau)$ for $j \in \{1, 2, \ldots, n\}$, and $\bar{n}_i(\tau)$ is the number of indices in $\bar{\mathcal{N}}_i(\tau)$. This set of equations constitute the synchronous system $\mathbb{S}$ we intend to analyze.

B. State Space Model

The equations defining $\mathbb{S}$, namely (12) – (15), determine a state space system of the form

$$x(\tau + 1) = F(\tau)x(\tau), \quad \tau \in \{1, 2, \ldots\} \qquad (16)$$

where

$$x(\tau) = [\bar{\theta}_1(\tau - 1) \;\cdots\; \bar{\theta}_n(\tau - 1) \; \bar{w}_1(\tau - 1) \;\cdots\; \bar{w}_n(\tau - 1)]' \qquad (17)$$

Each $F(\tau)$ is a $2n \times 2n$ stochastic matrix which can be described as follows.

Let $\mathcal{R}$ denote the set of all lists of n numbers $\bar{\mu} = \{\bar{\mu}_1, \bar{\mu}_2, \ldots, \bar{\mu}_n\}$ with each $\bar{\mu}_i$ taking a value in the real closed interval $[0, \; 1]$. Let $\mathcal{B}$ denote the set of all lists of n integers $b = \{b_1, b_2, \ldots, b_n\}$ with each b_i taking a value in the binary integer set $\{0, 1\}$. Each such triple $(\bar{\mathbb{N}}, \bar{\mu}, b) \in \mathcal{G}_{sa} \times \mathcal{R} \times \mathcal{B}$ determines a $2n \times 2n$ stochastic matrix $\mathbf{F}(\bar{\mathbb{N}}, \bar{\mu}, b)$ whose entries for $i \in \{1, 2, \ldots, n\}$ are

$$f_{ij} = \delta_{(i+n)j}, \qquad \text{and}$$
$$f_{(i+n)j} = \begin{cases} \frac{1}{\bar{n}_i}(1 - \bar{\mu}_j) & j \in (\bar{\mathcal{N}}_i - i) \\ \frac{1}{\bar{n}_i}\bar{\mu}_j & j \in (\bar{\mathcal{N}}_i - i) + \{n\} \\ \frac{1}{\bar{n}_i}\delta_{(i+n)j} & j \notin (\bar{\mathcal{N}}_i - i) \cup ((\bar{\mathcal{N}}_i - i) + \{n\}) \end{cases}$$

if $b_i = 1$ and

$$f_{ij} = \delta_{ij} \qquad \text{and} \qquad f_{(i+n)j} = \delta_{(i+n)j}$$

if $b_i = 0$. Here $\bar{\mathcal{N}}_i$ is the set of neighbors of vertex i in $\bar{\mathbb{N}}$, $\bar{n}_i$ is the number of elements in $\bar{\mathcal{N}}_i$, $\bar{\mathcal{N}}_i - i$ is the complement of i in $\bar{\mathcal{N}}_i$, δ_{ij} is the Kronecker delta, and for any set of integers $\mathcal{I}$, $\mathcal{I} + \{n\}$ is the set $\mathcal{I} + \{n\} = \{i + n : i \in \mathcal{I}\}$. We call any such matrix F an *asynchronous flocking matrix*. Thus the image of $\mathbf{F}$ is the set of all possible asynchronous flocking matrices.

It is easy to verify that the matrix $F(\tau)$ in (16) is of the form $\mathbf{F}(\bar{\mathbb{N}}(\tau), \bar{\mu}(\tau), b(\tau))$ where $\bar{\mathbb{N}}(\tau)$ is that graph in $\mathcal{G}_{sa}$ with neighbor sets $\bar{\mathcal{N}}_1(\tau), \bar{\mathcal{N}}_2(\tau), \ldots, \bar{\mathcal{N}}_n(\tau)$, $\bar{\mu}(\tau)$ is that list in $\mathcal{R}$ whose ith element is $\bar{\mu}_i(\tau)$, and $b(\tau)$ is that list in $\mathcal{B}$ whose ith element is $b_i(\tau) = 1$ if $t_\tau \in \mathcal{T}_i$ or $b_i(\tau) = 0$ if $t_\tau \notin \mathcal{T}_i$.

Note that unlike the other flocking problems considered in the past where the $F(\tau)$ were flocking matrices from a finite set, the set of all asynchronous flocking matrices which arise here, namely image $\mathbf{F}$, is not a finite set because $\mathcal{R}$ is not a finite set. Nonetheless image $\mathbf{F}$ is a closed and therefore compact subset of the set of all $2n \times 2n$ stochastic matrices $\mathcal{S}$. To understand why this is so, note first that for each fixed $b \in \mathcal{B}$ and $\mathbb{N} \in \mathcal{G}_{sa}$, the mapping $\mathcal{R} \to \mathcal{S}, \; \mu \longmapsto \mathbf{F}(\mathbb{N}, \mu, b)$ is continuous on $\mathcal{R}$. Therefore its image must be compact because $\mathcal{R}$ is. Next note that $\mathcal{G}_{sa}$ and $\mathcal{B}$ are each finite sets. Since the union of a finite number of compact sets is compact, it must therefore be true that the image of $\mathbf{F}$ is compact as claimed.

V. Analysis

The ultimate aim of this section is to give a proof of Theorem 2. We begin with the notion of the graph of a stochastic matrix.

Any $2n \times 2n$ stochastic matrix S such as those in image $\mathbf{F}$, determines a directed graph $\gamma(S)$ with vertex set $\{1, 2, \ldots, n, n+1, n+2, \ldots, 2n\}$ and arc set defined is such a way so that (i, j) is an arc of $\gamma(S)$ from i to j just in case the jith entry of S is non-zero. It is easy to verify that for any two such matrices S_1 and S_2,

$$\gamma(S_2 S_1) = \gamma(S_2) \circ \gamma(S_1) \qquad (18)$$

We now define a set of directed graphs $\mathcal{G}$ on vertex set $\{1, 2, \ldots, n, n + 1, n + 2, \ldots, 2n\}$ which contains all $\gamma(F)$, $F \in \text{image}\, \mathbf{F}$, and which is large enough to be closed under composition. For this purpose it is convenient to adopt the notation $[v]$ for the subset $\{v, v + n\}$ whenever $v \in \mathcal{V}$, and to say that $([v], u)$ is an arc of a graph $\mathbb{G}$ in $\mathcal{G}$ if either

(v,u) or $(v+n,u)$ is. Similarly we say that $(v,[u])$ is an arc of $\mathbb{G}$ if either (v,u) or $(v,u+n)$ is and $([v],[u])$ is an arc of $\mathbb{G}$ if either $(v,[u])$ or $(v+n,[u])$ is.

We define $\mathcal{G}$ to be the set of all directed graphs with vertex set $\{1,2,\dots,2n\}$ whose graphs have the following properties. For each $\mathbb{G}\in\mathcal{G}$ and each pair of vertices $u\in\{1,2,\dots,2n\}$ and $v\in\mathcal{V}$:

- p1: $v+n$ has a self-arc in $\mathbb{G}$.
- p2: $([v],v)$ is an arc in $\mathbb{G}$.
- p3: If (u,v) is an arc in $\mathbb{G}$ and $u\neq v$, then $(u,v+n)$ is an arc in $\mathbb{G}$.
- p4: If $(u,[v])$ is an arc in $\mathbb{G}$ and $u\neq v$, then $(v+n,v)$ is an arc in $\mathbb{G}$.

It is straightforward to verify that for each $F\in\text{image }\mathbf{F}$, $\gamma(F)$ is a graph in $\mathcal{G}$. In view of the structure of the matrices in image $\mathbf{F}$ it is natural to call a graph $\mathbb{G}\in\mathcal{G}$ an *event graph of agent* $i\in\mathcal{V}$ if $(i+n,i)$ is the *only* incoming arc to vertex i. Note that the graph of every matrix $\mathbf{F}(\bar{\mathbb{N}},\bar{\mu},b)$ for which $b_i=1$ is an event graph of agent i. Thus $\gamma(F(\tau))$ is an event graph of agent i if t_τ is an event time of agent i. It is easy to see that there are graphs in $\mathcal{G}$ which are not the graphs of any matrix in image $\mathbf{F}$. Let us agree to say that $\mathbb{G}\in\mathcal{G}$ is *attached at* $i\in\mathcal{V}$ if vertex i has *at least* $(i+n,i)$ as an incoming arc. A graph $\mathbb{G}\in\mathcal{G}$ is *attached* if it is attached at every vertex in $\mathcal{V}$. Thus $\gamma(F(\tau))$ would be attached if and only if t_τ were an event time of every agent. Note that the definition $\mathcal{G}$ allows this set to contain graphs which are attached at i which are not event graphs of agent i. In other words, an event graph of agent i must be attached at i, but the converse is not necessarily so.

We begin our analysis with the following observation.

Proposition 1: The set of graphs $\mathcal{G}$ is closed under composition.

The following results from [12] are key to establishing this convergence.

Proposition 2: Let $\mathcal{S}_{sr}$ be any closed set of stochastic matrices which are all of the same size and whose graphs $\gamma(S)$, $S\in\mathcal{S}_{sr}$ are all strongly rooted. As $j\to\infty$, any product $S_j\cdots S_1$ of matrices from $\mathcal{S}_{sr}$ converges exponentially fast to a matrix of the form $\mathbf{1}c$ at a rate no slower than λ, where c is a non-negative row vector depending on the sequence and λ is a non-negative constant less than 1 depending only on $\mathcal{S}_{sr}$.

Proposition 3: Suppose $n>1$ and let $\mathbb{G}_{p_1},\mathbb{G}_{p_2},\dots,\mathbb{G}_{p_m}$ be a finite sequence of rooted graphs with the same vertex set. If each vertex of each graph has a self-arc and $m\geq(n-1)^2$, then $\mathbb{G}_{p_m}\circ\mathbb{G}_{p_{m-1}}\circ\cdots\circ\mathbb{G}_{p_1}$ is strongly rooted.

Unfortunately the graphs of importance in the asynchronous case, namely the $\gamma(F(\tau))$, do not have self arcs at all vertices. Thus Proposition 3 cannot be directly applied.

To describe the analog of Proposition 3 appropriate to the asynchronous problem at hand we need another concept. Note that each $\mathbb{G}\in\mathcal{G}$ determines a *quotient graph* $Q(\mathbb{G})\in\mathcal{G}_{sa}$ defined in such a way that $Q(\mathbb{G})$ has an arc from i to j just in case $\mathbb{G}$ has an arc from at least one vertex in the set $[i]$ to at least one vertex in the set $[j]$. Note that $Q(\gamma(\mathbf{F}(\bar{\mathbb{N}},\bar{\mu},b)))=\bar{\mathbb{N}}$. The following is the analog of Proposition 3.

Proposition 4: Let $\mathbb{G}_{p_1},\dots,\mathbb{G}_{p_{2m+1}}$ be a sequence of $2m+1$ attached graphs in $\mathcal{G}$ whose quotients are rooted. If $m\geq(n-1)^2$ then $\mathbb{G}_{p_{2m+1}}\circ\cdots\circ\mathbb{G}_{p_1}$ is strongly rooted.

A more in depth study of the graphs in $\mathcal{G}$ leads us to the following observation.

Proposition 5: Let $\mathbb{G}_{p_1},\dots,\mathbb{G}_{p_m}$ be a sequence of graphs from $\mathcal{G}$ which for each $i\in\mathcal{V}$, contains a graph which is attached at i. Then $\mathbb{G}_{p_m}\circ\cdots\circ\mathbb{G}_{p_1}$ is an attached graph.

Let h be the smallest positive integer such that $T_{\max}\leq hT_{\min}$, then there will be at least one event time of any one agent within a sequence of at most $h+1$ consecutive event times of any other agent. We are led to the following conclusion.

Lemma 1: In any sequence of $(n-1)h+1$ or more consecutive event times, there will be at least one event time of each of the n agents.

The following proposition shows that for any sequence of graphs $\mathbb{G}_{p_1},\dots,\mathbb{G}_{p_m}$ from $\mathcal{G}$ whose quotients constitute a jointly rooted sequence, the quotient of the composition of the sequence is rooted.

Proposition 6: Let $\mathbb{G}_{p_1},\dots,\mathbb{G}_{p_m}$ be a sequence of $m>1$ graphs from $\mathcal{G}$ for which $Q(\mathbb{G}_{p_m})\circ\cdots\circ Q(\mathbb{G}_{p_1})$ is a rooted graph. Then $Q(\mathbb{G}_{p_m}\circ\cdots\circ\mathbb{G}_{p_1})$ is also rooted at the same vertex as $Q(\mathbb{G}_{p_m})\circ\cdots\circ Q(\mathbb{G}_{p_1})$.

In proving Theorem 2, we will need to exploit the compactness of a particular subset of stochastic matrices in $\mathcal{S}$ which can be described as follows. Let $p\geq n$ be any given positive integer. Write $\mathcal{G}_{sa}^p$ for the subset of all sequences of p graphs in $\mathcal{G}_{sa}$ which are jointly rooted and $\mathcal{B}^p$ for the set of all lists of p binary vectors in $\mathcal{B}$ with the property that for each $i\in\{1,2,\dots,n\}$, each list $\{b_1,b_2,\dots,b_p\}$ contains at least one vector whose ith row is 1. Since $p\geq n$, $\mathcal{B}^p$ is nonempty. Let $\mathcal{R}^p$ be the Cartesian product of $\mathcal{R}$ with itself p times. We claim that the image of the mapping $\mathbf{F}^p:\mathcal{B}^p\times\mathcal{R}^p\times\mathcal{G}_{sa}^p\to\mathcal{S}$ defined by

$$(\{\mathbb{N}_1,\mathbb{N}_2,\dots,\mathbb{N}_p\},\{\mu_1,\mu_2,\dots,\mu_p\},\{b_1,b_2,\dots,b_p\})\longmapsto \mathbf{F}(\mathbb{N}_p,\mu_p,b_p)\cdots\mathbf{F}(\mathbb{N}_2,\mu_2,b_2)\mathbf{F}(\mathbb{N}_1,\mu_1,b_1)$$

is compact. The reason for this is essentially the same as the reason image $\mathbf{F}$ is compact. In particular, for any fixed $\{\mathbb{N}_1,\mathbb{N}_2,\dots,\mathbb{N}_p\}\in\mathcal{G}_{sa}^p$ and $\{b_1,b_2,\dots,b_p\}\in\mathcal{B}^p$, the restricted mapping $\{\mu_1,\mu_2,\dots,\mu_p\}\longmapsto\mathbf{F}^p(\{\mathbb{N}_1,\mathbb{N}_2,\dots,\mathbb{N}_p\},\{\mu_1,\mu_2,\dots,\mu_p\},\{b_1,b_2,\dots,b_p\})$ is continuous so its image must be compact. Since $\mathcal{B}^p$ and $\mathcal{G}_{sa}^p$ are finite sets, the image of $\mathbf{F}^p$ must therefore be compact as well.

Set $q=2(n-1)^2+1$ and let $\mathcal{F}^p(q)$ denote the set of all products of q matrices from image $\mathbf{F}^p$. Then $\mathcal{F}^p(q)$ is compact because image $\mathbf{F}^p$ is. More is true.

Proposition 7: The graph of each matrix in $\mathcal{F}^p(q)$ is strongly rooted.

We are now finally in a position to prove our main result.

Proof of Theorem 2: As already noted, it is sufficient to prove that the matrix product $F(\tau)\cdots F(1)$ converges exponentially fast to a matrix of the form $\mathbf{1}c$ as $\tau \to \infty$. Observe first that there is a vector binary vector $b(\tau) \in \mathcal{B}$ and a vector $\mu(\tau) \in \mathcal{R}$ such that

$$F(\tau) = \mathbf{F}(\bar{N}(\tau), \mu(\tau), b(\tau)),\ \tau \geq 0 \tag{19}$$

because each $F(\tau) \in \text{image } \mathbf{F}$.

By hypothesis, the sequence of extended neighbor graphs $\bar{\mathbb{N}}(0), \bar{\mathbb{N}}(1), \ldots,$ is repeatedly jointly rooted. This means that there is an integer m for which each of the sequences $\bar{\mathbb{N}}(km+1), \ldots, \bar{\mathbb{N}}((k+1)m),\ k \geq 0$, is jointly rooted. Let h be as is in Lemma 1 and define $p = rm$ where r is any positive integer large enough so that $p \geq (n-1)h + 1$. Set $q = 2(n-1)^2 + 1$ and let $\mathcal{G}^p_{sa},\ \mathcal{R}^p,\ \mathcal{B}^p,\ \mathbf{F}^p$, and $\mathcal{F}^p(q)$ be as defined just above Proposition 7.

Since each $\bar{\mathbb{N}}(km+1), \ldots, \bar{\mathbb{N}}((k+1)m),\ k \geq 0$, is jointly rooted, each of the compositions $\bar{\mathbb{N}}((k+1)m) \circ \cdots \circ \bar{\mathbb{N}}(km+1),\ k \geq 0$, is rooted. This implies that each graph $\bar{\mathbb{N}}((k+1)p) \circ \cdots \circ \bar{\mathbb{N}}(kp+1),\ k \geq 0$, is rooted because $p = rm$ and because the composition of r rooted graph is rooted. Therefore each sequence $\bar{\mathbb{N}}(kp+1), \ldots, \bar{\mathbb{N}}((k+1)p),\ k \geq 0$, is jointly rooted. It follows that

$$\{\bar{\mathbb{N}}(kp+1), \ldots, \bar{\mathbb{N}}((k+1)p\} \in \mathcal{G}^p_{sa},\ k \geq 0 \tag{20}$$

Note next that for each $i \in \{1, 2, \ldots, n\}$ and each $k \geq 0$, at least one of the graphs in the sequence $\gamma(F(kp+1)), \ldots, \gamma(F((k+1)p))$ must be attached at i because of Lemma 1 and the assumption that $p \geq (n-1)h + 1$. This implies that for each $i \in \{1, 2, \ldots, n\}$ there must be at least one vector in each list $\{b(kp+1), \ldots, b((k+1)p)\},\ k \geq 0$ whose ith row is 1. Therefore

$$\{b(kp+1), \ldots, b((k+1)p)\} \in \mathcal{B}^p,\ k \geq 0 \tag{21}$$

For $k \geq 0$, define

$$S(k) = F((k+1)p) \cdots F(kp+1) \tag{22}$$

In view of (19) - (21) and the definition of $\mathbf{F}^p$, it must be true that $S(k) \in \text{image } \mathbf{F}^p, k \geq 0$. Thus if we define

$$\bar{S}(k) = S((k+1)q - 1) \cdots S(kq),\ k \geq 0 \tag{23}$$

then each $\bar{S}(k)$ must be in $\mathcal{F}^p(q)$. Therefore by Proposition 7, the graph of each $\bar{S}(k)$ is strongly rooted. Therefore by Proposition 2, the matrix product $\bar{S}(k) \cdots \bar{S}(0)$ converges exponentially fast as $k \to \infty$ to a matrix of the form $\mathbf{1}c$ as $k \to \infty$.

The definitions of $S(\cdot)$ and $\bar{S}(\cdot)$ in (22) and (23) respectively imply that

$$\bar{S}(k) \cdots \bar{S}(0) = F((k+1)pq) \cdots F_1,\ \ k \geq 0$$

For $\tau \geq 0$, let $\kappa(\tau)$ and $\rho(\tau)$ denote respectively, the integer quotient and remainder of τ divided by pq. Then

$$F(\tau) \cdots F(1) = \widehat{S}(\tau)\bar{S}(k(\tau)) \cdots \bar{S}(0)$$

where $k(\tau) = \kappa(\tau) - 1$, and $\widehat{S}(\tau)$ is the bounded function

$$\widehat{S}(\tau) = \begin{cases} F(\tau) \cdots F((k(\tau)+1)pq + 1) & \text{if } \rho(\tau) \neq 0 \\ 1 & \text{if } \rho(\tau) = 0 \end{cases}$$

Since $k(\tau)$ is an unbounded monotone nondecreasing function and $\bar{S}(k) \cdots \bar{S}(0)$ converges exponentially fast as $k \to \infty$, it follows that $F(\tau) \cdots F(1)$ converges exponentially fast as $\tau \to \infty$ to a matrix of the form $\mathbf{1}c$. ■

VI. Concluding Remarks

The version of the asynchronous consensus considered here significantly generalizes our earlier work [9]. In particular, the present version of the problem can deal with continuous heading changes whereas the version of the problem solved in [9] cannot.

It is possible to formulate and solve a "continuous" version of Viscek's problem in which each agent's heading is adjusted by controlling its differential rate. Because of changing neighbors this leads to a differential equation model with a discontinuous vector field in which chattering may occur. To avoid this one can introduce "dwell times" as was done in [3] for the leader-follower version of the problem. As a result, the question of synchronization again arises, in this case with event times being the times at which each agent's dwell time periods begin. Thus although one might think that the question of synchronization is irrelevant in the continuous-time case, this appears to only be true if one is willing to accept generalized solutions to differential equations and the possibility of chattering.

References

[1] T. Vicsek, A. Czirók, E. Ben-Jacob, I. Cohen, and O. Shochet. Novel type of phase transition in a system of self-driven particles. *Physical Review Letters*, 75:1226–1229, 1995.

[2] M. J. Fischer, N. A. Lynch, and M. S. Paterson. Impossibility of distributed consensus with one faulty process. *Journal of the Association for Computing Machinery*, 32:347–382, 1985.

[3] A. Jababaie, J. Lin, and A. S. Morse. Coordination of groups of mobile autonomous agents using nearest neighbor rules. *IEEE Transactions on Automatic Control*, 48:988–1001, 2003. Also in Proc. 2002 IEEE CDC, pages 2953 - 2958.

[4] L. Moreau. Stability of multi-agent systems with time-dependent communication links. *IEEE Transactions on Automatic Control*, 50:169–182, 2005.

[5] W. Ren and R. Beard. Consensus seeking in multiagent systems under dynamically changing interaction topologies. *IEEE Transactions on Automatic Control*, 50:655–661, 2005.

[6] D. Angeli and P. A. Bliman. Extension of a result by Moreau on stability of leaderless multi-agent systems. In *Proc. of the 44th IEEE CDC*, pages 759–764, 2005.

[7] V. D. Blondel, J. M. Hendrichx, A. Olshevsky, and J. N. Tsitsiklis. Convergence in multiagent coordination, consensus, and flocking. In *Proc. of the 44th IEEE CDC*, pages 2996–3000, 2005.

[8] L. Fang, P. J. Antsaklis, and A. Tzimas. Asynchronous consensus protocols: preliminary results,simulations and open questions. In *Proc. of the 44nd IEEE CDC*, pages 2194–2199, 2005.

[9] M. Cao, A. S. Morse, and B. D. O. Anderson. Coordination of an asynchronous multi-agent system via averaging. In *Proc. of the IFAC world congress, Prague, Czech Republic*, 2005.

[10] J. Lin, A. S. Morse, and B. D. O. Anderson. The multi-agent rendezvous problem - the asynchronous case. In *Proc. of the 43rd IEEE CDC*, pages 1926–1931, 2004.

[11] J. Lin, A. S. Morse, and B. D. O. Anderson. The multi-agent rendezvous problem - part 2: The asynchronous case. 2004. Submitted to *SIAM Journal on Control and Optimization*.

[12] M. Cao, A. S. Morse, and B. D. O. Anderson. Reaching a consensus in a dynamically changing environment - A graphical approach. April 2006. Submitted to *SIAM Journal on Control and Optimization*.

[13] D. P. Bertsekas and J. N. Tsitsiklis. *Parallel and Distributed Computation*. Prentice Hall, 1989.