Proceedings of the 17th World Congress
The International Federation of Automatic Control
Seoul, Korea, July 6-11, 2008

Global synchronization on the circle ⋆

Alain Sarlette [*] S. Emre Tuna [**] Vincent D. Blondel [***]
Rodolphe Sepulchre [*]

[*] *Dept. of Electrical Engineering and Computer Science, Université de Liège, B-4000 Liège, Belgium (e-mail: alain.sarlette@ulg.ac.be, r.sepulchre@ulg.ac.be)*
[**] *Dept. of Electrical and Electronics Engineering, Middle East Technical University, 06531 Ankara, Turkey (e-mail: tuna@eee.metu.edu.tr)*
[***] *Dept. of Mathematical Engineering, Université catholique de Louvain, B-1348 Louvain-la-Neuve, Belgium (e-mail: blondel@inma.ucl.ac.be)*

Abstract: The convexity arguments used in the consensus literature to prove synchronization in vector spaces can be applied to the circle only when all agents are initially located on a semicircle. Existing strategies for (almost-)global synchronization on the circle are either restricted to specific interconnection topologies or use auxiliary variables. The present paper first illustrates this problem by showing that weighted, directed interconnection topologies can be designed to make any reasonably chosen configuration of the agents on the circle a stable equilibrium of a basic continuous-time consensus algorithm. Then it proposes a so-called "gossip algorithm", which achieves global asymptotic synchronization on the circle with probability 1 for a large class of interconnections, without using auxiliary variables, thanks to the introduction of randomness in the system.

Keywords: consensus on circle; multi-agent system; time-variant communication topology; gossip algorithm.

1. INTRODUCTION

Multi-agent systems have drawn considerable attention in the systems and control community for different reasons: they appear in natural phenomena that are not fully understood, in existing interconnected (mostly communication) systems, and their potential capabilities indicate that they may be increasingly used in future applications — see Belta and Kumar (2003); Beugnon et al. (2005); Nair and Leonard (2006); Tsitsiklis and Athans (1984); Tsitsiklis and Bertsekas (1986); Strogatz (2000); Hopfield (1982); Jadbabaie et al. (2003); Vicsek et al. (1995); Blondel et al. (2005); Leonard et al. (2007); Izzo and Pettazzi (2005); Reynolds (1987); Strogatz (2003); Justh and Krishnaprasad (2002); McInnes (1996); Cortes et al. (2004) to name a few.

A basic task for autonomous control of multi-agent systems is to make all the agents *agree on some quantity*. For instance a group of agents may have to decide on a meeting point and time, or on the value of a parameter to use for their individual operations, say an amount of data, a power level,... ; these examples usually involve *synchronization* (i.e. reaching a common value) in vector spaces. But the consensus state space may also be nonlinear. For example, consider a group of agents moving on the plane which want to agree on a common direction of motion in order to stay together, or a group of oscillators that must synchronize their phases; this involves synchronization on the circle. Other applications may involve synchronization on the sphere (e.g. equal direction of motion in 3D), on the space of rotation matrices (e.g. equal satellite attitudes), or on more abstract spaces like the Grassmannian manifolds (e.g. in computational applications involving averaging on those spaces).

The design of individual control laws for a set of interconnected agents such that they asymptotically agree on a common point in a *vector space* is a well-studied problem, widely known as the *consensus problem*. It has been solved in many settings — see among others Tsitsiklis and Athans (1984); Blondel et al. (2005); Moreau (2005, 2004), and Olfati-Saber et al. (2007) for a review. The convergence analysis rests on a convexity argument. For the sake of illustration, consider the real line, with each agent moving towards a weighted average of the values of connected agents; then the minimal value in the set can only increase and the maximal value only decrease, until (under connectedness assumptions) they become equal.

On nonlinear spaces (manifolds), this convexity argument cannot be used globally: for a set of agents distributed on the circle, there is no 'minimal value' or 'maximal value'

⋆ This paper presents research results of the Belgian Network DYSCO (Dynamical Systems, Control, and Optimization), funded by the Interuniversity Attraction Poles Programme, initiated by the Belgian State, Science Policy Office. The scientific responsibility rests with its authors. The first author is supported as an FNRS fellow (Belgian Fund for Scientific Research).

and in general, every agent is attracted by some agents to move in one direction and by others to move in the opposite direction. Direct convexity arguments can be used only when all agents are initially located within a semi-circle (Jadbabaie et al. (2003); Moreau (2005)). For arbitrarily distributed agents, synchronization can be ensured for particular interconnection topologies, like trees and all-to-all interconnection with equal weights (Sepulchre et al. (2007); Scardovi et al. (2007); Sarlette and Sepulchre (200?)); for other interconnections, the traditional Kuramoto-like consensus algorithm admits locally stable configurations different from synchronization (Jeanna et al. (2005); Sarlette and Sepulchre (200?)). To recover global synchronization properties with weak assumptions on the interconnection structure, Sepulchre et al. (to appear); Scardovi et al. (2007); Sarlette and Sepulchre (200?) propose an algorithm that "cheats" the manifold structure by using auxiliary variables in the embedding space (e.g. evolving in the whole plane $\mathbb{R}^2$ for synchronization on the circle). However, this assumes that agents can memorize and communicate such auxiliary variables, which may not be an appropriate model in several practical situations.

The present paper proposes an algorithm that achieves global synchronization on the circle with weak assumptions on the interconnection structure and without using auxiliary variables. Instead, it introduces randomness to build a "gossip algorithm" as in Boyd et al. (2006): at each iteration, each agent randomly selects at most one of its neighbors and moves towards it. This algorithm is shown to ensure global synchronization with probability one asymptotically. Before studying the gossip algorithm, the problems encountered by traditional consensus algorithms on the circle are illustrated by showing that positively weighted directed interconnection structures can be designed to locally asymptotically stabilize arbitrary configurations that are sufficiently "spread".

The paper is organized as follows. Section II presents a classical consensus algorithm on the circle and shows how it can stabilize arbitrary configurations. Section III presents the gossip algorithm and shows how it achieves global synchronization with probability 1, for a very general class of interconnection topologies; the theoretical statements are confronted with simulation.

2. CONSENSUS ON THE CIRCLE

Consider a set of N agents evolving on the circle S^1; the position of agent k on the circle is denoted by the angular variable $\theta_k \in S^1$. The interconnection (or communication) among agents is represented with a directed graph G: each of the N graph nodes represents an agent; an edge from node j to node k means that the corresponding agents are interconnected or linked, in the sense that agent k gets information from agent j, also denoted $j \rightsquigarrow k$. Agent j is called a neighbor of k. A nonvanishing weight $a_{jk} \geq a_m$ can be associated to the link from j to k, for some fixed $a_m > 0$, and $a_{jk} = 0$ for $j \not\rightsquigarrow k$.

A classical (continuous-time) consensus algorithm in the vector space $\mathbb{R}$ reads

$$\tfrac{d}{dt} x_k = \sum_{j \rightsquigarrow k} a_{jk}\,(x_j - x_k) \quad , \; x_k \in \mathbb{R},\; k = 1...N. \qquad (1)$$

The convergence analysis rests on the fact that the convex hull $[\min x_k,\ \max x_k]$ cannot expand along the solutions. Convergence is ensured if G is *root-connected*, that is, G contains an agent k from which any other agent can be reached by following a directed path (equivalently, G contains a spanning tree). If G is time-varying, root-connectedness over a uniform horizon is still sufficient to ensure exponential convergence (Moreau (2004); Blondel et al. (2005)).

A natural adaptation of (1) on the circle is

$$\tfrac{d}{dt} \theta_k = \sum_{j \rightsquigarrow k} a_{jk}\, \sin(\theta_j - \theta_k) \quad , \; k = 1...N. \qquad (2)$$

This is in fact the celebrated Kuramoto model with equal natural frequencies (Kuramoto (1975)). Defining $z_k = e^{i\theta_k}$, (2) is equivalent to

$$\tfrac{d}{dt} z_k = \mathrm{Proj}_{z_k}\left(\sum_{j \rightsquigarrow k} a_{jk}\,(z_j - z_k) \right) \qquad (3)$$

where $\mathrm{Proj}_{z_k}(r_k)$ denotes the orthogonal projection of $r_k \in \mathbb{C}$ onto the direction tangent to the unit circle at $z_k = e^{i\theta_k}$, that is $\mathrm{Proj}_{z_k}(r_k) = i z_k \langle i z_k, r_k \rangle$. The geometric interpretation is that (3) defines a consensus update similar to (1) but constrained to the manifold where $\|z_k\| = 1$. Studies of model (2) as such in the synchronization context (Sepulchre et al. (2007); Jeanna et al. (2005); Scardovi et al. (2007); Sarlette and Sepulchre (200?)) have shown that it (almost) globally converges towards synchronization for tree graphs and the equally weighted complete graph. For other graphs, (2) may fail to converge to fixed positions (e.g. "cyclic pursuit" problem for a directed cycle graph) or may locally converge to a stable configuration that is different from synchronization (e.g. agents uniformly distributed around the circle for an undirected cycle graph); the latter are called *consensus configurations* in Sarlette and Sepulchre (200?) — on the circle and other nonlinear manifolds, there are thus graph-dependent consensus configurations different from synchronization.

The following result shows that stable consensus configurations different from synchronization are not exceptional: in fact, any configuration sufficiently "spread" on the circle is a stable consensus for a well-chosen directed graph.

Proposition 1. Consider N agents distributed on the circle in a configuration $\{\theta_k\}$ such that for every k, there is at least one agent located in $(\theta_k, \theta_k + \pi/2)$ and one agent located in $(\theta_k - \pi/2, \theta_k)$; such a configuration requires $N \geq 5$. Then there exists a positively weighted, directed and root-connected interconnection graph making this configuration locally exponentially stable under (2).

Proof. From (3), a necessary and sufficient condition for a configuration to be an equilibrium is that $r_k = \sum_{j \rightsquigarrow k} a_{jk}(z_j - z_k)$ must be aligned with z_k. Since the positive linear combinations of two vectors in the plane span the cone between them, it is easy to assign non-negative weights to edges $j \rightsquigarrow k$ with $\theta_j \in (\theta_k - \pi/2, \theta_k + \pi/2)$, and zero weight to all other edges, such that r_k and z_k are aligned, for each k. The requirement $a_{jk} \in \{0\} \cup$

$[a_m, +\infty)$ is obtained by multiplying all weights by a constant. This ensures that the configuration is an equilibrium. Linearization of (2) around the equilibrium yields a linear consensus algorithm in $\mathbb{R}$ with weights $a_{jk}\cos(\theta_j - \theta_k) \in \{0\} \cup (a_m\alpha, +\infty)$, where $\alpha > 0$ is the minimal value of $\cos(\theta_j - \theta_k)$ among edges with $a_{jk} > 0$. Moreover, it is easy to verify that the graph associated to positive weights is root-connected. This ensures exponential stability of the equilibrium configuration (in the shape space, that is, modulo a rigid rotation of all agents). ■

Proposition 1 identifies how specific configurations can be made locally exponentially stable by choosing appropriate weights for the directed graph edges. For any of these choices, synchronization is also exponentially stable — but thus only locally; in particular, agents initially located within a semicircle always converge towards synchronization (Jadbabaie et al. (2003); Moreau (2005)).

Proposition 1 emphasizes that global convergence analysis of a consensus algorithm is trickier on S^1 than in a vector space. Deciding whether a given graph is S^1-synchronizing, that is, admits no other stable configuration than synchronization, seems to be a difficult question, even though simulations suggest that spurious stable consensus configurations are rare and possess a small basin of attraction.

3. GOSSIP ALGORITHM

The discussion in the previous section raises the question of designing an alternative algorithm that ensures (almost) global convergence towards synchronization for a larger class of interconnection graphs. A recent solution in Scardovi et al. (2007); Sepulchre et al. (to appear); Sarlette and Sepulchre (200?) uses auxiliary variables $x_k \in \mathbb{C}$, $k = 1...N$: the agents actually run a consensus algorithm on these auxiliary variables and then move towards the projection of the consensus value in $\mathbb{C}$ on the circle. This allows to recover convergence properties similar to vector spaces, but it requires an auxiliary dynamical system, to memorize, update and communicate the x_k. Although this can be justified for some engineering applications, it is perhaps a questionable model to describe natural phenomena.

This motivates the introduction of *randomness* instead of auxiliary variables to overcome spurious "consensus" configurations on the circle. In particular, the present paper discusses a so-called "gossip algorithm" (Boyd et al. (2006)), described in discrete-time for ease of formulation: at each time instant, each agent randomly selects at most one of its neighbors in G, according to a fixed probability distribution, to update its phase value; a natural probability distribution for link selection would follow the weights of the links.

Gossip algorithm (directed graph). At each update t,

1. each agent k randomly selects an agent $j \rightsquigarrow k$ with probability $a_{jk}/(\beta + \sum_{l \rightsquigarrow k} a_{lk})$, where $\beta > 0$ is the weight for choosing no agent [1];
2. $\theta_k(t+1) = \theta_j(t)$ if agent k chooses neighbor j at time t, and $\theta_k(t+1) = \theta_k(t)$ if it chooses no neighbor.

[1] Note that the links chosen at time $t+1$ are thus independent of the links chosen at time t.

Gossip algorithm (undirected graph: $j \rightsquigarrow k \Leftrightarrow k \rightsquigarrow j$ and $a_{jk} = a_{kj}$). At each update t,

1. each agent k randomly selects one neighbor or none, as in the directed case;
2. if k chooses j AND j chooses k at time t, then k and j move towards the midpoint of the shortest arc between them, i.e. $\theta_k(t+1) = \theta_j(t+1) = (\theta_k(t) + \theta_j(t))/2 + \mathrm{s_a}\pi$ where $\mathrm{s_a} \in \{0, 1\}$ ensures that the average is located on the *shortest* arc between $\theta_j(t)$ and $\theta_k(t)$. If k chooses no neighbor or a neighbor j which does not choose k, $\theta_k(t+1) = \theta_k(t)$.

Boyd et al. (2006) perform a detailed analysis of a gossip algorithm for synchronization in vector space. As convergence towards synchronization is not an issue on vector spaces, the problem is to quantify the convergence *rate* as a function of the interconnection graph and probability (i.e. weights) distribution. On the circle, convergence towards synchronization is not obvious a priori and the goal of the present analysis is to clarify this issue.

3.1 Convergence analysis

Definition 1. In the present stochastic setting, the agents are said to *asymptotically converge towards synchronization with probability 1* if for any initial condition, for any $\varepsilon > 0$ and $\delta \in (0, 1)$, there exists a time T after which the maximal distance $|\theta_k(T) - \theta_j(T)|$ between any pair of agents is smaller than ε with probability larger than δ.

The strategy to prove asymptotic convergence towards synchronization with probability 1 in different settings is based on the following facts, valid for *time-varying and directed* interconnection graphs.

Lemma 2. If all agents are located within an open semicircle, asymptotic synchronization with probability 1 is ensured with the Gossip Algorithms if there exists a finite horizon T and a probability $p_0 > 0$ such that $\forall t$, the graph formed by the links selected during $[t, t+T]$ is root-connected with probability p_0 at least.

Proof. The proof uses the exponential synchronization result for time-varying, directed graphs on *vector spaces* (Tsitsiklis and Athans (1984); Moreau (2005); Blondel et al. (2005)): a sufficient condition for exponential synchronization is that the graph formed by the union of all links appearing during a finite time span $[t, t+T_0)$ is root-connected for $t = t_0, t_0 + T_0, t_0 + 2T_0, ...$ Consensus on the semicircle can be mapped to consensus on the real line (Moreau (2005)). Moreover, one run of a Gossip Algorithm can be viewed as a deterministic algorithm for a particular time-varying graph (featuring at most one link per agent at a time). Therefore exponential synchronization is ensured if the root-connectedness assumption is satisfied.

Denote by $N_\varepsilon T_0$ the maximal time for the time-varying algorithm to ensure that all agents are in an interval smaller than ε, over all possible graphs with $a_{jk} \in \{0\} \cup [a_m, +\infty)$ that are root-connected on T_0-time spans; by time-scale symmetry, N_ε is independent of T_0. If a graph is root-connected on T with probability p_0, it is root-connected on $T_0 = mT$ with probability at least $p_1 = (1 - (1-p_0)^{\lfloor m \rfloor})$ which can be made arbitrarily close to 1 with a sufficiently large m. The graph is then root-connected on N_ε consecutive T_0-time spans with probability at least $p_2 = p_1^{N_\varepsilon}$; it suffices to select m such that $p_2 > \delta$ to prove the property. ■

The study of synchronization on the circle can thus be reduced to the study of "bringing all agents within a semicircle". Indeed, write $\delta = \delta_1\delta_2$ with $\delta_1, \delta_2 < 1$. If it can be ensured that for any $\delta_1 \in (0,1)$, all agents are within a semicircle after a finite time T_{δ_1} with probability δ_1, then applying Lemma 2 with δ replaced by δ_2 ensures asymptotic synchronization with probability 1.

Lemma 3. Consider a sequence σ of link selections over a finite time span T_σ, whose probability to appear at least once in a time span $[t, t + T_s)$ is at least $p_{\sigma s} > 0$ $\forall t$ and for a finite $T_s \geq T_\sigma$. If applying a Gossip Algorithm with sequence σ ensures that all agents end up in a semicircle, *for any initial condition*, then for any $\delta \in (0, 1)$ there exists a finite time T_h after which the Gossip Algorithm (with random sequence) has driven all agents within a semicircle with probability $p > \delta$.

Proof. If σ appears with probability $p_{\sigma s}$ in any time span of length T_s, it appears at least once in any time span of length $T_h = mT_s$ with probability at least $(1 - (1 - p_{\sigma s})^{\lfloor m \rfloor})$, which can be made $> \delta$ by taking m sufficiently large. Thus for m sufficiently large, there is probability $p_3 > \delta$ that a link sequence appears during T_h such that all agents are within a semicircle at the end of that link sequence; after the sequence, convexity arguments ensure that the agents never leave this semicircle. ■

A finite sequence σ of link selections driving all agents within a semicircle, regardless of the initial condition, is called a *synchronizing sequence* in this paper. Thanks to Lemma 2 and Lemma 3, the study of asymptotic synchronization with probability 1 is reduced to the search for synchronizing sequences that appear with probability $p \geq p_s > 0$ in every time span of some bounded length T_s. The following results identify such sequences for both types of Gossip Algorithms.

Proposition 4. Consider a set of N agents interconnected according to a time-varying directed graph G. Assume that $\forall t$, the weights of G belong to $\{0\} \cup (a_m, a_M)$ for some fixed $a_M \geq a_m > 0$. Also assume that there exists $T_c > 0$ such that the union of all links appearing in G over $[t, t + T_c)$ forms a graph which contains a rooted tree (see Figure 1) $\forall t$. If the agents apply the Gossip Algorithm for directed graphs, with a fixed finite $\beta > 0$, they asymptotically synchronize with probability 1.

Proof. The proof is based on the construction of a synchronizing sequence appearing with probability at least p_s in some time span T_s.

The links appearing in $[t_0 + nT_c, t_0 + (n + 1)T_c)$ contain a rooted tree for $n = 0, 1, 2...$. Since there are N possible roots, there exists an agent r serving as a root for at least $N - 1$ trees $tr_1, tr_2, ...tr_{N-1}$ not ovelapping in time over $[t_0, t_0 + N(N - 1)T_c)$. Consider the following link sequence σ. Over time intervals not corresponding to tr_i, $i = 1...N - 1$, no agent chooses a neighbor (no updates). Over the time interval of tr_1, a child c of r (i.e. an agent c such that $r \rightsquigarrow c$ in the r-rooted tree) chooses its link with r, such that $\theta_c(t + 1) = \theta_r(t)$, and all other agents choose no neighbor [2] . Over the time interval of further trees tr_i, $i = 2...N - 1$, choose just one update $\theta_j(t + 1) = \theta_k(t)$, where k is either r or a previously updated agent and j is any of k's childs in tr_i that was not updated yet; no other agents move. One easily checks that this is always possible until all agents have been updated and are thus located at the initial position of agent r. It remains to show that this link sequence appears with finite probability p_s in a time span T_s.

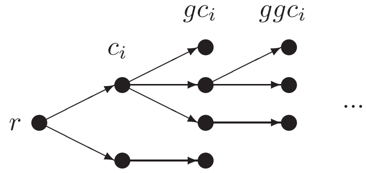

Fig.1: A directed rooted tree. Replacing the arrows by simple lines yields an undirected tree.

At any time t, the probability for agent k to choose a particular link is at least $\frac{a_m}{(N-1)a_M+\beta} =: \eta_1 > 0$, while the probability of k choosing no link is at least $\frac{\beta}{(N-1)a_M+\beta} =: \eta_2 > 0$. Thus the probability to choose a particular link *and no other link in the graph* is at least $\eta_1\,\eta_2^{N-1}$ and the probability to choose no link at all is η_2^N. Therefore, the probability to choose the particular sequence σ in time span $T_s = N(N - 1)T_c$ is at least $p_s = (\eta_1\,\eta_2^{N-1})^{N-1}\,(\eta_2^N)^{N(N-1)T_c-(N-1)} > 0$. ■

Proposition 5. Consider a set of N agents interconnected according to a time-varying undirected graph G. Assume that $\forall t$, the weights of G belong to $\{0\} \cup (a_m, a_M)$ for some fixed $a_M \geq a_m > 0$. Also assume that there exists $T_c > 0$ such that the union of all links appearing in G over $[t, t + T_c)$ forms a root-connected graph $\forall t$. If the agents apply the Gossip Algorithm for undirected graphs, with a fixed finite $\beta > 0$, they asymptotically synchronize with probability 1.

Proof. The proof is based on the construction of a synchronizing sequence appearing with probability at least p_s in some time span T_s.

First assume that G is *time-invariant*. Then G contains an undirected spanning tree (see Figure 1). The latter can be built by starting with 2 interconnected agents and adding agents one at a time. The link sequence pulls agents within a semicircle in the following way.
Suppose that a link sequence σ_q is known to bring $q < N$ agents in the partly constructed tree S_q within $q\frac{\alpha}{N}$, with $\alpha \in (0, \pi/2)$. Denote by k the new agent to add and by j the agent to which it will be connected in order to build tree S_{q+1}. The link sequence $\sigma_{q+1} = \{j \leftrightsquigarrow k\}$; σ_q; $\{j \leftrightsquigarrow k\}$; σ_q; ... (repeat $x > \log_2(\pi N/\alpha)$ times), brings the $q+1$ agents within $(q + 1)\frac{\alpha}{N}$ degrees. Indeed, denote by γ the distance from k to the furthest point of the arc containing $\{\theta_i : i \in S_q\}$, along the shortest arc between k and j; γ_0 is

[2] Although it could be possible to choose further links for tr_1, this is the minimum achievable for any graph, namely when r has only one child c in tr_1 and the link $r \rightsquigarrow c$ appears at the end of the time interval of tr_1.

the initial value of γ. In other words, the $q+1$ agents are initially "at worst within γ_0", with $\gamma_0 \leq \pi + q\frac{\alpha}{N}$. After one iteration of $\{j \leftrightsquigarrow k\}$; σ_q, the agents of S_q are within $q\frac{\alpha}{N}$ again, but $\gamma - q\frac{\alpha}{N} \leq (\gamma_0 - q\frac{\alpha}{N})/2$. Thus after x iterations, $\gamma - q\frac{\alpha}{N} \leq (\gamma_0 - q\frac{\alpha}{N})/(2^x) \leq \pi/(2^x) < \frac{\alpha}{N}$ such that the $q+1$ agents are within $(q+1)\frac{\alpha}{N}$.
Now build the sequence $\sigma = \sigma_2; \sigma_3; ...\sigma_N$, with σ_{q+1} built iteratively on σ_q as just explained; σ_2 contains a single link, such that two agents average their values and are thus within $0 < \frac{2\alpha}{N}$. The sequence σ is synchronizing, as it brings all agents within $\frac{N\alpha}{N} = \alpha < \pi/2$ for any initial condition. Moreover, it has finite length T_σ (although T_σ grows exponentially with N). It contains a single link at each time t. The probability to choose a particular link at time t is at least $\eta_3 := \eta_1^2 \eta_2^{N-2} > 0$ (the two agents to be connected choose each other mutually, the others choose no neighbor), where $\eta_1 = \frac{a_m}{(N-1)a_M + \beta}$ and $\eta_2 = \frac{\beta}{(N-1)a_M + \beta}$. Thus sequence σ is chosen at least with finite probability $(\eta_3)^{T_\sigma}$ in time span T_σ.

If G is time-varying, one checks that, for an arbitrarily large B, choosing a long enough time span $T_B \geq BT_c$ ensures that there is an undirected tree tr_1 appearing on B non-overlapping time spans of length T_c during T_B. Then the synchronizing sequence σ (built above) of the fixed tr_1 can be implemented by choosing the next link in the sequence if it is available, and no link else; if $B \geq T_\sigma$, the whole sequence can be applied, leaving the agents in the same configuration as with invariant G. The probability of this sequence is then at least $(\eta_3)^{T_\sigma}(\eta_2^N)^{T_B - T_\sigma}$. ■

The sequences proposed in the proofs are just examples leading to easy discussion and probability bounds. Although absent from the proofs, synchronizing sequenc involving simultaneous moves of several agents do exi For instance, if G is directed and *time-invariant* and has rooted tree tr_1 of depth $d < N-1$, selecting an arbitra number of links *only chosen in* tr_1 at each update t is possible way towards synchronization. It is even possib to choose all links of tr_1; doing this d consecutive tim yields synchronization.

Although this paper considers the particular case of t circle, synchronization problems can be formulated many other manifolds (Sarlette and Sepulchre (200? Proposition 4 is valid for any manifold, as its proof do not involve the structure of the space on which the agents evolve. Proposition 5, associated to Lemma 2, involves the structure of the circle. Lemma 2 can be generalized to ensure asymptotic synchronization when the agents are initially located within a *convex set* of the manifold. Thus to generalize Proposition 5, a link sequence of finite probability must be found that drives all the agents within a convex set of the manifold, for any initial condition.

3.2 Simulation results

The theoretical analysis proves that a set of agents applying the Gossip Algorithm globally asympotically synchronize with probability 1, but it says little about the convergence rate. Bounds on the convergence rate can be obtained from the synchronizing sequences constructed in the proofs, but they would be very conservative.

It is expected (and confirmed in simulations) that some randomness in the selection of neighbors favors global synchronization and destroys the stability of spurious consensus configurations, synchronization being the only common consensus configuration of all the partial graphs. Nevertheless, convergence can be slow if the initial condition is a spurious stable consensus configuration for the deterministic algorithm. As an illustration, consider a set of N agents connected by an undirected cycle graph (each agent connected to two others such that the whole forms a "closed cycle"). For this graph, the consensus algorithm possesses a locally stable configuration with the agents uniformly distributed around the circle, separated by $2\pi/N$ (called *splay state* in Sepulchre et al. (2007)). For the Gossip Algorithm, this graph does not look too bad a priori: suppressing one link makes it a tree, for which synchronization is the only stable configuration, so if one link is not selected during a longer time span, the agents can be driven close to synchronization.

Figure 2 represents two simulations starting close to the splay state with $N = 9$, $\beta = 1$ and $a_{jk} = 1\ \forall j \rightsquigarrow k$. The left simulation has made no progress towards synchronization after 800 iterations; repeating simulations, this situation appears roughly 3 times out of 4, which means that in Definition 1, $T > 800$ for $\delta = 0.25$ and any small ε for this particular initial condition. The right simulation shows a case where synchronization is achieved: at one point in time, the agents end up within a semicircle and from this point on convergence is much faster. The synchronization strategy of Scardovi et al. (2007); Sarlette and Sepulchre (200?); Sepulchre et al. (to appear), with auxiliary variables, obtains such fast convergence from $t = 0$ on.

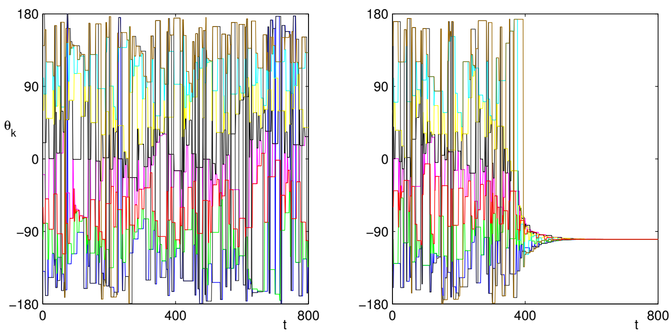

Fig.2: Evolution of the θ_k (in degrees) from an initial *splay state*, for the Gossip Algorithm on an undirected cycle graph with $N = 9$ agents and probability distribution $\beta = 1$, $a_{jk} = 1\ \forall j \rightsquigarrow k$. About 75% of trials exhibit no synchronization after 800 iterations (left figure). Convergence in the remaining trials is fast once agents are in a semicircle (right figure).

4. CONCLUSION

The present paper investigates the global synchronization properties of consensus algorithms defined on the circle. This problem is fundamentally different from the traditional consensus problem defined in vector spaces, because convexity arguments no longer hold.

In the first part, the issue of global synchronization on the circle is illustrated by showing that, for $N \geq 5$ agents, a positively weighted directed interconnection graph can be designed to make any reasonably chosen configuration of the agents on the circle a stable equilibrium of a basic continuous-time consensus algorithm proposed in Scardovi et al. (2007); Sarlette and Sepulchre (200?). Thus there exist many stable equilibria of the consensus algorithm different from synchronization (Jeanna et al. (2005); Sarlette and Sepulchre (200?)).

In the second part, a new approach is proposed to drive a set of agents with arbitrary initial condition on the circle towards a common point. The basic feature of the resulting discrete-time algorithm is to introduce randomness in the update process: at each time instant, each agent randomly chooses one of its neighbors at most to update its own position on the circle; this feature and the associated term "gossip algorithm" are inspired by Boyd et al. (2006).

The gossip algorithm on the circle ensures global asymptotic convergence towards synchronization with probability 1 for any interconnection graph satisfying the connectedness assumptions for synchronization on vector spaces. Moreover, it does so while avoiding to introduce auxiliary variables as in Scardovi et al. (2007); Sarlette and Sepulchre (200?); Sepulchre et al. (to appear). This looks more natural at least to describe natural phenomena. However, the slow convergence observed in simulations when starting from a spurious stable configuration of the deterministic algorithm indicates that it is not an optimal strategy for engineering applications.

For future work on this gossip algorithm, it could be interesting to quantify the convergence rate, in the sense of finding estimates / better bounds of the probability to obtain a synchronizing sequence in a given time span. After this, a study similar to Boyd et al. (2006) could examine optimal probability distributions for fast convergence.

REFERENCES

C. Belta and V. Kumar. Motion generation for groups of robots. *ASME Journal of Mechanical Design*, 2003.

C. Beugnon, E. Buvat, M. Kersten, and S. Boulade. GNC design for the DARWIN spaceborne interferometer. *Proc. 6th ESA Conf. on Guidance, Navigation and Control Systems*, 2005.

V.D. Blondel, J.M. Hendrickx, A. Olshevsky, and J.N. Tsitsiklis. Convergence in multiagent coordination, consensus and flocking. *Proc. 44th IEEE Conf. on Decision and Control*, 2005.

S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. Randomized gossip algorithms. *IEEE Transactions on Information Theory (Special issue)*, 52(6):2508–2530, 2006.

J. Cortes, S. Martinez, and F. Bullo. Coordinated deployment of mobile sensing networks with limited-range interactions. *Proc. 43rd IEEE Conf. on Decision and Control*, pages 1944–1949, 2004.

J.J. Hopfield. Neural networks and physical systems with emergent collective computational capabilities. *Proc. National Academy of Sciences*, 79:2554–2558, 1982.

D. Izzo and L. Pettazzi. Equilibrium shaping: distributed motion planning for satellite swarm. *Proc. 8th Intern. Symp. on Artificial Intelligence, Robotics and Automation in space*, 2005.

A. Jadbabaie, J. Lin, and A.S. Morse. Coordination of groups of mobile autonomous agents using nearest neighbor rules. *IEEE Transactions on Automatic Control*, 48(6):988–1001, 2003.

J. Jeanna, D. Paley, and N.E. Leonard. On the stable phase relationships of loop-coupled planar particles. *Proc. 44th IEEE Conf. on Decision and Control*, 2005.

E. Justh and P. Krishnaprasad. A simple control law for UAV formation flying. Technical report, TR 2002-38, ISR, University of Maryland, 2002.

Y. Kuramoto. In *International Symposium on Mathematical Problems in Theoretical Physics*, volume 39 of *Lecture Notes in Physics*, page 420. Springer, 1975.

N.E. Leonard, D. Paley, F. Lekien, R. Sepulchre, D. Frantantoni, and R. Davis. Collective motion, sensor networks and ocean sampling. *Proceedings of the IEEE*, 95 (1):48–74, January 2007.

C.R. McInnes. Distributed control for on-orbit assembly. *Advances in the Astronautical Sciences*, 90:2079–2092, 1996.

L. Moreau. Stability of continuous-time distributed consensus algorithms. *Proc. 43rd IEEE Conf. on Decision and Control*, pages 3998–4003, 2004.

L. Moreau. Stability of multi-agent systems with time-dependent communication links. *IEEE Transactions on Automatic Control*, 50(2):169–182, 2005.

S. Nair and N.E. Leonard (advisor). Stabilization and synchronization of networked mechanical systems. *PhD Thesis, Princeton University*, 2006.

R. Olfati-Saber, J.A. Fax, and R.M. Murray. Consensus and cooperation in networked multi-agent systems. *Proceedings of the IEEE*, 95(1):215–233, 2007.

C.W. Reynolds. Flocks, herds, and schools: a distributed behavioral model. *Computer Graphics (ACM SIGGRAPH Conf. Proceedings)*, 21:25–34, July 1987.

A. Sarlette and R. Sepulchre. Consensus optimization on manifolds. *Submitted to SIAM J. Control and Optimization.*

L. Scardovi, A. Sarlette, and R. Sepulchre. Synchronization and balancing on the N-torus. *Systems and Control Letters*, 56(5):335–341, 2007.

R. Sepulchre, D. Paley, and N.E. Leonard. Stabilization of planar collective motion with all-to-all communication. *IEEE Transactions on Automatic Control*, 52(5):811–824, 2007.

R. Sepulchre, D. Paley, and N.E. Leonard. Stabilization of planar collective motion with limited communication. *IEEE Transactions on Automatic Control*, to appear.

S.H. Strogatz. From Kuramoto to Crawford: exploring the onset of synchronization in populations of coupled nonlinear oscillators. *Physica D*, 143:1–20, 2000.

S.H. Strogatz. *Sync: The emerging science of spontaneous order.* Hyperion, 2003.

J.N. Tsitsiklis and M. Athans (advisor). Problems in decentralized decision making and computation. *PhD Thesis, MIT*, 1984.

J.N. Tsitsiklis and D.P. Bertsekas. Distributed asynchronous optimal routing in data networks. *IEEE Transactions on Automatic Control*, 31(4):325–332, 1986.

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(6): 1226–1229, 1995.