Joint 48th IEEE Conference on Decision and Control and
28th Chinese Control Conference
Shanghai, P.R. China, December 16-18, 2009

Consensus on homogeneous manifolds

Alain Sarlette   Rodolphe Sepulchre

Abstract— **The present paper considers distributed consensus algorithms for agents evolving on a connected compact homogeneous (CCH) manifold. The agents track no external reference and communicate their relative state according to an interconnection graph. The paper first formalizes the consensus problem for synchronization (i.e. maximizing the consensus) and balancing (i.e. minimizing the consensus); it thereby introduces the induced arithmetic mean, an easily computable mean position on CCH manifolds. Then it proposes and analyzes various consensus algorithms on manifolds: natural gradient algorithms which reach local consensus equilibria; an adaptation using auxiliary variables for almost-global synchronization or balancing; and a stochastic gossip setting for global synchronization. It closes by investigating the dependence of synchronization properties on the attraction function between interacting agents on the circle. The theory is also illustrated on $SO(n)$ and on the Grassmann manifolds.**

I. Introduction

The distributed computation of means/averages of datasets (in an algorithmic setting) and the synchronization or spreading of a set of agents (in a control setting) have attracted growing interest in the literature, with applications like swarms/formations (e.g. [16], [26]), distributed decision making (e.g. [20], [31]), networks (e.g. [30]), optimal coding or covering (e.g. [8], [9]), etc. The modeling and understanding of swarm behavior in nature has also led to many studies (e.g. [14], [29], [32]).

Recent results have contributed to a good understanding of synchronization of interacting agents in Euclidean space, based on the linear consensus algorithm (e.g. [18], [31], [20])

$$\tfrac{d}{dt} y_k = \textstyle\sum_j a_{jk} \left(y_j - y_k\right) \qquad (1)$$

where $y_k \in \mathbb{R}^n$, $k = 1, 2, ..., N$ are the agents' states and the $a_{jk} \geq 0$ characterize how strongly they are attracted towards each other ($a_{jk} = 0$ implying no interconnection, i.e. no attraction of agent k towards j). Global exponential synchronization is ensured even with varying a_{jk}, as long as the agents are *uniformly connected* (see below).

However, many interesting applications involve manifolds that are not homeomorphic to an Euclidean space, like the circle S^1 for (e.g. oscillator) phase variables or the group of rotations $SO(n)$ for rigid body orientations.

The goal of the present paper is to extend the framework of consensus algorithm (1) to connected compact homogeneous manifolds (which include S^1, $SO(n)$, Grassmann manifolds $Grass(p, n)$ and spheres S^{n-1}) and to propose algorithms for global synchronization on these manifolds. Indeed, unlike for Euclidean spaces, agents attracted towards each other on manifolds do not always reach synchronization. We therefore first define particular "(anti-)consensus" configurations on manifolds (Section II). A cost function is then built and gradient algorithms are derived which drive an interacting swarm to (anti-)consensus configurations (Section III). The whole framework can be obtained from an easily computable "mean position" on manifolds. The convenient idea behind it is to embed the manifold in an Euclidean space $\mathbb{R}^m$ and measure distances between agents in $\mathbb{R}^m$. We then propose two modifications to reach global synchronization with weak conditions on agent interconnections: adding auxiliary "estimator variables" with values in $\mathbb{R}^m$, which the agents update and communicate (Section IV); and letting each agent at each time interact with at most one other agent, which is stochastically chosen (Section V). In both cases, the resulting algorithms generically ensure global convergence to synchronization if the swarm of agents is uniformly connected; the estimator algorithm also has a variant which distributes the agents on the compact manifold in a configuration which we call "balanced". Finally, we show with two examples how the convergence properties of a consensus algorithm on the circle depend on the attraction profile between agents as a function of distance (Section VI). The concepts are illustrated on S^1, $SO(n)$ and $Grass(p, n)$.

In the literature so far, the study of global synchronization or balancing properties in non-Euclidean manifolds is not widely covered. The circle is often addressed: oscillator synchronization studies mostly derive from the Kuramoto model (see [28] for a review); recently, we addressed consensus on S^1 from a control perspective [26], [23], [25], [27]. Manifold $SO(3)$ has attracted attention in recent years in the context of satellite attitudes: reference- or leader-dependent synchronization is studied e.g. in [15], [3], local synchronization studies with a geometric approach are found in [6], [19]. The computation of means on manifolds has triggered some research, including classical but computationally heavy definitions like [13], [7], as well as developments for particular spaces which are covered by our "induced arithmetic mean" approach (see [17] on $SO(3)$ and [1], [11], [8] on $Grass(p, n)$). The topic of optimization-based algorithm design on manifolds has considerably developed over the last decades (see e.g. [5], [10], [12], [2]).

The present paper is based on [21] and associated publications. Sections II, III, IV are based on [22]. Section V is based on [24]. The reader is invited to consult these references for more mathematical background, discussion and proofs, as well as a deeper treatment of examples.

This paper is part of the SIAM session; it covers results from [22], [24] among others. A. Sarlette and R. Sepulchre are with the Department of Electrical Engineering and Computer Science, University of Liege, 4000 Liege Sart-Tilman, Belgium {alain.sarlette,r.sepulchre}@ulg.ac.be

978-1-4244-3872-3/09/$25.00 ©2009 IEEE

Preliminaries

Interconnections among agents are represented by *a graph* G, whose N vertices are the N agents, and containing edge (j,k) iff agent j sends information to agent k, which is denoted $j \rightsquigarrow k$ or $(j,k) \in E$, the set of edges. A weight a_{jk} is associated to each ordered pair of agents, such that $a_{jk} > 0$ iff $j \rightsquigarrow k$, $a_{jk} = 0$ else. By convention, $a_{kk} = 0\ \forall k$ is assumed. Matrix A containing the a_{jk} is called the adjacency matrix of G. The in-Laplacian of G is $L^{(i)} = D^{(i)} - A$ where diagonal matrix $D^{(i)}$ contains the in-degrees $d^{(i)}_{kk} := \sum_{j=1}^N a_{jk}$. By construction, $L^{(i)}$ has zero column sums. G is undirected if $A = A^T$. G is balanced if $d^{(i)}_{kk} = \sum_{j=1}^N a_{kj}$. G is strongly connected if it contains a directed path from any vertex j to any vertex k; G is weakly connected if such paths exist in the associated undirected graph, with adjacency matrix $A + A^T$. Time-varying interconnections are represented by time-varying edges. We always assume that the elements of $A(t)$ are bounded and satisfy some threshold $a_{jk}(t) \geq \delta > 0\ \forall (j,k) \in E(t)$ and $\forall t$, i.e. $G(t)$ is a δ-digraph. In a δ-digraph $G(t)$, vertex j is said to be connected to vertex k across $[t_1, t_2]$ if there is a directed path from j to k in the digraph defined by adjacency matrix $\bar{A}$ where

$$\bar{a}_{jk} = \begin{cases} \int_{t_1}^{t_2} a_{jk}(t)dt & \text{if } \int_{t_1}^{t_2} a_{jk}(t)dt \geq \delta \\ 0 & \text{if } \int_{t_1}^{t_2} a_{jk}(t)dt < \delta \end{cases} .$$

A δ-digraph $G(t)$ is uniformly connected if there exist a vertex k and a time horizon $T > 0$ such that $\forall t$, k is connected to all other vertices across $[t, t+T]$.

A homogeneous manifold $\mathcal{M}$ is a manifold with a transitive group action by a Lie group $\mathcal{G}$: it is isomorphic to the quotient manifold $\mathcal{G}/\mathcal{H}$ of a group $\mathcal{G}$ by one of its subgroups $\mathcal{H}$. Informally, it can be seen as a manifold on which "all points are equivalent". The present paper considers connected compact homogeneous (CCH) manifolds satisfying the following embedding property.

Assumption 1: $\mathcal{M}$ is a CCH manifold smoothly embedded in $\mathbb{R}^m$ with the Euclidean norm $\|y\| = r_{\mathcal{M}}$ constant over $y \in \mathcal{M}$. The Lie group $\mathcal{G}$ acts as a subgroup of the orthogonal group on $\mathbb{R}^m$.

It is sometimes preferred to represent $y \in \mathcal{M}$ by a matrix instead of a vector. Componentwise identification $\mathbb{R}^{n_1 \times n_2} \cong \mathbb{R}^m$ is then assumed; the corresponding norm is the Frobenius norm $\|B\| = \sqrt{\mathrm{trace}(B^T B)}$.

The *special orthogonal group* $SO(n)$ is the set of rotation matrices in $\mathbb{R}^n$. A point of $SO(n)$ is characterized by a real $n \times n$ orthogonal matrix Q, i.e. $Q^T = Q^{-1}$, with determinant $+1$. $SO(n)$ is a homogeneous (as any Lie group), compact and connected manifold. It has dimension $n(n-1)/2$.

Each point on the *Grassmann manifold* $Grass(p,n)$ denotes a p-dimensional subspace $\mathcal{Y}$ of $\mathbb{R}^n$. The dimension of $Grass(p,n)$ is $p(n-p)$. Since $Grass(n-p,n)$ is isomorphic to $Grass(p,n)$ by identifying orthogonally complementary subspaces, we assume w.l.o.g. that $p \leq \frac{n}{2}$. $Grass(p,n)$ is connected, compact and homogeneous as the quotient manifold of $O(n)$ by $O(p) \times O(n-p)$. In order to embed $Grass(p,n)$ in $\mathbb{R}^m \cong \mathbb{R}^{n \times n}$, we use the "projector representation" $\Pi = Y Y^T$, where $Y \in \mathbb{R}^{n \times p}$ is any orthonormal basis of subspace $\mathcal{Y}$ in $\mathbb{R}^n$.

II. Mean and Consensus on Manifolds

A. The induced arithmetic mean

Consider a set of N agents on a manifold $\mathcal{M}$ satisfying Assumption 1. The position of agent k is denoted by y_k and its weight by $w_k > 0$.

Definition 1: The *induced arithmetic mean IAM* $\subseteq \mathcal{M}$ is the set of points in $\mathcal{M}$ that globally minimize the weighted sum of squared Euclidean distances in $\mathbb{R}^m$ to each y_k:

$$\begin{aligned} IAM &= \underset{c \in \mathcal{M}}{\operatorname{argmin}} \sum_{k=1}^N w_k\, d^2_{\mathbb{R}^m}(y_k, c) & (2) \\ &= \underset{c \in \mathcal{M}}{\operatorname{argmin}} \sum_{k=1}^N w_k\, (y_k - c)^T (y_k - c)\ . & (3) \end{aligned}$$

The *anti-[induced arithmetic mean] AIAM* $\subseteq \mathcal{M}$ is the set of points in $\mathcal{M}$ that globally maximize this weighted sum.

The point in Definition 1 is that distances are measured *in the embedding space* $\mathbb{R}^m$. It thereby differs from the canonical *Karcher mean* [13], which uses the geodesic distance on $\mathcal{M}$. The *IAM* satisfies several properties of a mean, see [22]. It does not always reduce to a single point, but this seems unavoidable (imagine e.g. points uniformly distributed on a circle). The main advantage of the *IAM* over the Karcher mean is computational: defining the *centroid* $C_e \in \mathbb{R}^m$ by

$$C_e = \frac{1}{W} \sum_{k=1}^N w_k\, y_k\ , \qquad \text{where } W = \sum_{k=1}^N w_k$$

it holds

$$IAM = \underset{c \in \mathcal{M}}{\operatorname{argmax}} (c^T C_e)\ . \qquad (4)$$

Thus computing the *IAM* just requires to maximize a linear function of $\mathbb{R}^m$ in a very regular search space $\mathcal{M}$. For $SO(n)$, $Grass(p,n)$ and the n-dimensional spheres $S^{n-1} \subset \mathbb{R}^n$, the linear function has no local minima, so local optimization is sufficient.

Assumption 2: The local maxima of any linear function $f(c) = c^T b$ over $c \in \mathcal{M}$, with b fixed in $\mathbb{R}^m$, are all global maxima.

Example 1: Tthe circle S^1 embedded in $\mathbb{R}^2$ with its center at the origin satisfies Assumptions 1 and 2. The *IAM* is the central projection of C_e onto S^1. It reduces to a single point if $C_e \neq 0$, else it contains the whole circle. The *IAM* uses the chordal distance between points, while the Karcher mean would use arclength distance.

$SO(n)$, embedded as orthogonal matrices $Q \in \mathbb{R}^{n \times n}$, $\det(Q) > 0$, satisfies Assumptions 1 and 2. The *IAM* is the orthogonal component of the polar decomposition of C_e if $\det(C_e) > 0$; if $\det(C_e) \leq 0$ it is given by a related formula [22].

$Grass(p,n)$ is represented as the set of orthonormal p-rank projectors Π_k, embedded in the symmetric positive semidefinite cone of $\mathbb{R}^{n \times n}$, to satisfy Assumptions 1 and 2. The *IAM* is the dominant p-dimensional eigenspace of C_e.

B. Consensus on manifolds

Consider that the N agents are interconnected according to a fixed digraph G of adjacency matrix $A = [a_{jk}]$. For simplicity we take $w_k = 1\ \forall k$.

Definition 2: Synchronization is the configuration where $y_j = y_k\ \forall j, k$. A *consensus configuration* for G satisfies

$$y_k \in \underset{c\in\mathcal{M}}{\operatorname{argmax}} \left(c^T \sum_{j=1}^{N} a_{jk}\, y_j \right) \qquad \forall k \qquad (5)$$

i.e. each agent is located at a point of the *IAM* of its neighbors $j \rightsquigarrow k$. An *anti-consensus configuration* is similarly defined with *IAM* replaced *AIAM*. The agents are *balanced* if their *IAM* contains all $\mathcal{M}$.

Note that (anti-)consensus is defined as a Nash equilibrium: each agent minimizes its cost function *assuming the others fixed*. Consensus and anti-consensus are graph-dependent notions. Synchronization and balancing are graph-independent and can be seen as situations of "complete" consensus and anti-consensus respectively.

Proposition 1: If G is the equally-weighted complete graph, then synchronization is the only consensus configuration. All balanced configurations are anti-consensus configurations for the equally-weighted complete graph.

The second part of Proposition 1 does not establish a necessary and sufficient condition; anti-consensus configurations for the equally-weighted complete graph that are not balanced, though exceptional, do exist. Balancing implies some spreading of the agents on the manifold. A full characterization of balanced configurations seems complicated.

Example 2: We limit ourselves to the circle S^1.

Consider the equally-weighted *undirected ring graph*, in which each agent is connected to two neighbors such that the graph forms a single closed undirected path. Regular consensus configurations correspond to situations with consecutive agents in the path always separated by the same angle $0 \le \chi \le \pi/2$; regular anti-consensus configurations have $\pi/2 \le \chi \le \pi$. In addition, for $N \ge 4$, irregular consensus and anti-consensus configurations exist where non-consecutive angles of the regular configurations are replaced by $(\pi - \chi)$. The reader is encouraged to discuss implications of this example (also see [22]); for instance, there is no common anti-consensus state for all ring graphs.

Anti-consensus configurations for the equally-weighted complete graph are fully characterized in [26]: the only anti-consensus configurations that are not balanced correspond to $(N+1)/2$ agents at some θ^* and $(N-1)/2$ agents at $\theta^* + \pi$, for N odd. Balanced configurations are unique for $N = 2$ and $N = 3$ and form a continuum for $N > 3$.

III. Gradient consensus algorithms

For a graph G with adjacency matrix $A = [a_{jk}]$ and associated Laplacian $L^{(i)} = [l_{jk}^{(i)}]$, associated to $y = (y_1, ..., y_N) \in \mathcal{M}^N$, define

$$\begin{aligned} P_L(y) &= \tfrac{1}{2N^2} \textstyle\sum_{j,k} a_{jk}\, y_j^T y_k \;=\; \xi_1 - \tfrac{1}{2N^2} \sum_{j,k} l_{jk}^{(i)}\, y_j^T y_k \\ &= \xi_2 - \tfrac{1}{4N^2} \textstyle\sum_{j,k} a_{jk}\, \|y_j - y_k\|^2 \end{aligned} \qquad (6)$$

where ξ_1, ξ_2 are constants. In [23], [27], P_L is studied on S^1 for undirected equally-weighted G. For the unit-weighted complete graph, $P := P_L + \frac{r_{\mathcal{M}}^2}{2N} = \frac{1}{2}\|C_e\|^2$, the squared norm of centroid C_e; this is a classical measure of oscillator synchrony in the literature, e.g. in the context of the Kuramoto model [14], [28].

Proposition 2: Synchronization of the N agents on $\mathcal{M}$ is the unique global maximum of P_L whenever G is weakly connected. Further, if $\mathcal{M}$ satisfies Assumptions 1 and 2, then given an undirected graph G, a local maximum (resp. minimum) of the associated $P_L(y)$ necessarily corresponds to a consensus (resp. anti-consensus) configuration for G.

In [26], P is used on S^1 to derive gradient algorithms for synchronizing (by maximizing P) or balancing (by minimizing P) headings of particles in planar motion. We extend this to CCH manifolds and to general consensus configurations. For simplicity, we limit ourselves to continuous-time gradient algorithms, where the gradient is defined with the canonical metric induced by the embedding of $\mathcal{M}$ in $\mathbb{R}^m$.

A gradient algorithm for P_L yields, for $k = 1...N$,

$$\begin{aligned} \tfrac{d}{dt} y_k(t) &= 2N^2 \alpha\, \mathrm{grad}_{k,\mathcal{M}}(P_L) && (7) \\ &= \alpha\, \mathrm{Proj}_{T\mathcal{M},k} \left(\textstyle\sum_j (a_{jk} + a_{kj}) y_j \right) && (8) \\ &= \alpha\, \mathrm{Proj}_{T\mathcal{M},k} \left(\textstyle\sum_j (a_{jk} + a_{kj})(y_j - y_k) \right) && (9) \end{aligned}$$

where $\alpha > 0$ (resp. $\alpha < 0$) for consensus (resp. anti-consensus), $\mathrm{grad}_{k,\mathcal{M}}(f)$ denotes the gradient of f with respect to y_k along $\mathcal{M}$, and $\mathrm{Proj}_{T\mathcal{M},k}$ is the orthogonal projection onto the tangent space to $\mathcal{M}$ at y_k. Algorithm (9) requires each agent k to know the relative position with respect to itself of all agents j for which $j \rightsquigarrow k$ or $k \rightsquigarrow j$. Since information flow is restricted to $j \rightsquigarrow k$, (9) can only be implemented for undirected G, for which it becomes

$$\tfrac{d}{dt} y_k(t) = 2\alpha\, \mathrm{Proj}_{T\mathcal{M},k} \left(\textstyle\sum_j a_{jk}(y_j - y_k) \right) \,. \qquad (10)$$

In the special case of a complete unit-weighted graph,

$$\tfrac{d}{dt} y_k(t) = 2\alpha N\, \mathrm{Proj}_{T\mathcal{M},k} \left(C_e(t) - y_k \right) \,. \qquad (11)$$

Proposition 3: A swarm of N agents moving according to (10) on a manifold $\mathcal{M}$ satisfying Assumptions 1 and 2, with fixed undirected G, always converges to a set of equilibrium points. If $\alpha < 0$, all asymptotically stable equilibria are anti-consensus configurations for G. If $\alpha > 0$, all asymptotically stable equilibria are consensus configurations for G (in particular, for the equally-weighted complete graph, the only asymptotically stable configuration is synchronization).

Note that in Proposition 2, optimizing P_L is a *sufficient* condition to reach (anti-)consensus configurations. Therefore, all stable equilibria of the gradient algorithm are (anti-) consensus configurations, but there may also be (anti-) consensus configurations that are unstable. For instance, for a tree, maximization of P_L always leads to synchronization, although other consensus configurations can exist.

Formally, algorithm (10) can be written for directed and even time-varying graphs, although the gradient property is lost. Nevertheless, using the argument of [18], it can be shown that synchronization is still a stable equilibrium (asymptotically stable if disconnected graph sequences are excluded). Its basin of attraction includes the configurations where all the agents are located in a convex set of $\mathcal{M}$. On the other hand, examples where algorithm (10) runs into a limit cycle, quasi-periodic behavior,... can be built with undirected varying G or with fixed directed G; see [21], [22] for simple examples on S^1. With directed *and* varying G, there are even more possibilities. However, these examples seem to be non-generic: performing simulations with random graph sequences and initial conditions on S^1, $SO(n)$ and $Grass(p,n)$, the swarm seems to always eventually converge to synchronization when $\alpha > 0$.

It can be noted that the discrete-time version of (10)

$$y_k(t+1) \ \in \ IAM\left(\{y_j(t)|j \rightsquigarrow k \text{ in } G(t)\} \cup \{y_k(t)\}\right) \quad (12)$$

exactly corresponds to Vicsek's phase update law on S^1 (see [32]), and readily generalizes it to manifolds.

Example 3: Denoting angular positions on S^1 by θ_k, the specific form of (10) is

$$\tfrac{d}{dt}\theta_k = \alpha' \ \textstyle\sum_j a_{jk} \sin(\theta_k - \theta_j) \ , \qquad k = 1...N \ . \quad (13)$$

For the equally-weighted complete graph, this is strictly equivalent to the Kuramoto model [14] with identical (zero) natural frequencies. Algorithm (13) can e.g. run into a limit cycle when the graph switches between two different undirected rings.

On $SO(n)$ and $Grass(p,n)$, P_L with matrix forms $y_k \in \mathbb{R}^{n \times n}$ becomes

$$P_L(y) \ = \ \tfrac{1}{2N^2} \textstyle\sum_{j,k} a_{jk} \, \mathrm{trace}(y_j^T y_k) \ . \quad (14)$$

On $SO(n)$, $Q_j^T Q_k \in SO(n)$ is the unique rotation matrix translating Q_j to Q_k by matrix (group) multiplication on the right. Previous work [6], [19] already use $\mathrm{trace}(Q_k^T Q_j)$ as a measure of disagreement on $SO(3)$. The explicit form of (10) is

$$Q_k^T \tfrac{d}{dt} Q_k = \alpha \ \textstyle\sum_j a_{jk} \ \left(Q_k^T Q_j - Q_j^T Q_k\right) \ . \quad (15)$$

On $Grass(p,n)$, (14) can be rewritten as

$$P_L(\mathcal{Y}) = \tfrac{1}{2N^2} \textstyle\sum_{j,k} a_{jk} \ \left(\sum_{i=1}^{p} \cos^2(\phi_{jk}^i)\right)$$

with ϕ_{jk}^i the i^{th} principal angle between subspaces $\mathcal{Y}_j$ and $\mathcal{Y}_k$. This formulation has previously appeared in e.g. [8]. Algorithm (10) writes

$$\tfrac{d}{dt}\Pi_k = 2\alpha \ \textstyle\sum_j a_{jk} \left(\Pi_k \Pi_j \Pi_{\perp k} + \Pi_{\perp k} \Pi_j \Pi_k\right) \quad (16)$$

where $\Pi_\perp = I - \Pi$, with I the identity matrix.

IV. ALGORITHMS WITH ESTIMATOR VARIABLES

The (anti-)consensus configurations reached with (10) are directly linked to the interconnection graph G. In many applications, G is just a restriction on communication possibilities, under which one actually wants to achieve a consensus for the equally-weighted complete graph, i.e. synchronization or balancing. This section presents algorithms achieving the same performance as the complete graph gradient algorithm, but under very weak conditions on the actual $G(t)$. The reduction of information channels is compensated by adding a *consensus estimator variable* $x_k \in \mathbb{R}^m$ to the state space and communication of each agent.

For synchronization, the agents run a linear consensus algorithm on their arbitrarily initialized estimator variables x_k in $\mathbb{R}^m$; agent k's position y_k on $\mathcal{M}$ independently tracks (the projection on $\mathcal{M}$ of) x_k. This yields

$$\begin{aligned} \tfrac{d}{dt}x_k &= \beta \textstyle\sum_j a_{jk} \left(x_j - x_k\right) && (17) \\ \tfrac{d}{dt}y_k &= \gamma_S \, \mathrm{grad}_{k,\mathcal{M}}(y_k^T x_k) = \gamma_S \, \mathrm{Proj}_{T\mathcal{M},k}(x_k) && (18) \end{aligned}$$

with $\beta, \gamma_S > 0$ for $k = 1...N$. Equation (17) is a classical consensus algorithm in $\mathbb{R}^m$, exactly equivalent to (1). According to e.g. [18], the x_k exponentially converge to a common value x_∞ if $G(t)$ is piecewise continuous in time and uniformly connected. This leads to the following convergence property, where IAM_g generalizes (4) when the points defining C_e are not on $\mathcal{M}$.

Proposition 4: Assume that $\mathcal{M}$ satisfies Assumptions 1 and 2, and $G(t)$ is piecewise continuous and uniformly connected. Then the only stable limit configuration of the y_k under (17),(18), with the x_k initialized arbitrarily but independently and such that they can take any value in an open subset of $\mathbb{R}^m$, is synchronization at $y_\infty = \mathrm{Proj}_{T\mathcal{M},k}(x_\infty)$; if $G(t)$ is balanced, then $y_\infty = IAM_g\{x_k(0),\, k = 1...N\}$.

For anti-consensus, by analogy, each y_k uses a gradient algorithm to maximize its distance to $x_k(t)$. Algorithm

$$\begin{aligned} \tfrac{d}{dt}x_k &= \beta \textstyle\sum_j a_{jk} \left(x_j - x_k\right) + \tfrac{d}{dt}y_k && (19) \\ \tfrac{d}{dt}y_k &= \gamma_B \, \mathrm{grad}_{k,\mathcal{M}}(y_k^T x_k) = \gamma_B \, \mathrm{Proj}_{T\mathcal{M},k}(x_k) && (20) \end{aligned}$$

with $\beta > 0, \gamma_B < 0$ for $k = 1...N$, ensures that all $x_k(t)$ asymptotically converge to $C_e(t)$ if $G(t)$ is balanced $\forall t$ and $x_k(0) = y_k(0)\ \forall k$; then the motion (20) of y_k asymptotically becomes equivalent to (11). Note that the variables x_k and y_k are fully coupled. This makes the convergence proof more involved, but the general result remains.

Proposition 5: Assume that $\mathcal{M}$ satisfies Assumptions 1 and 2, and $G(t)$ is piecewise continuous, uniformly connected and balanced. Then algorithm (19),(20) with initial conditions $x_k(0) = y_k(0)\ \forall k$ converges to an equilibrium configuration of (11) with $\alpha < 0$.

In simulations, a swarm applying (19),(20) with $x_k(0) = y_k(0)\ \forall k$ seems to generically converge to an anti-consensus configuration of the equally-weighted complete graph, that is a *stable* equilibrium configuration of (11) with $\alpha < 0$.

Example 4: Applying this strategy to the circle S^1 yields the results of [25]; the x_k are vectors of $\mathbb{R}^2$. On $SO(n)$ and $Grass(p,n)$, introducing estimator variables $X_k \in \mathbb{R}^{n \times n}$, (17) may be transcribed verbatim; (18) becomes respectively

$$\begin{aligned} Q_k^T \tfrac{d}{dt} Q_k &= \tfrac{\gamma_S}{2} \left(Q_k^T X_k - X_k^T Q_k\right) && (21) \\ \tfrac{d}{dt}\Pi_k &= \gamma_S \left(\Pi_k X_k \Pi_{\perp k} + \Pi_{\perp k} X_k \Pi_k\right) \ . && (22) \end{aligned}$$

In [25], the algorithms including estimator variables are expressed "relative to agent positions" on S^1. The algorithms for $SO(n)$ can similarly be expressed completely in the agents' body frames: defining $Z_k = Q_k^T X_k$, the agents only need to know relative positions $Q_k^T Q_j$ and communicate arrays of scalars Z_k to implement e.g. (17),(18) by

$$\begin{aligned} \tfrac{d}{dt} Z_k &= (Q_k^T \tfrac{d}{dt} Q_k)^T Z_k \qquad (23) \\ &\quad + \beta \; \textstyle\sum_j a_{jk} \left((Q_k^T Q_j) Z_j - Z_k \right) \\ Q_k^T \tfrac{d}{dt} Q_k &= \tfrac{\gamma_S}{2} \left(Z_k - Z_k^T \right) \; . \qquad (24) \end{aligned}$$

V. Gossip algorithm

The algorithms of the previous section use auxiliary variables x_k which agents must memorize, update and communicate. This is not always possible in applications, nor realistic to describe natural phenomena. Another possibility to achieve global synchronization is to use a so-called "gossip algorithm" [4] where at each time, each agent randomly selects at most one of its neighbors in $G(t)$ to update its own phase value. For simplicity we here use discrete-time dynamics; the convergence proof can be repeated with an appropriate continuous-time version where agents are ε-close to the discrete-time values at the end of a period. We consider two variants.

Directed gossip: at each update t,

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

Undirected gossip: at each update t,

1. same procedure as in the directed case;
2. if at time t, k chooses j AND j chooses k, then $y_k(t+1) = y_j(t+1) \in \mathit{IAM}(y_k(t), y_j(t))$. If k chooses no neighbor or a neighbor j which does not choose k, then $y_k(t+1) = y_k(t)$.

In the directed variant, agents move between a finite set of points fixed by their initial positions; the manifold structure and topology plays no role. The undirected variant was already proposed and analyzed on vector spaces (e.g. convergence speed optimization in [4]), where it maintains $\frac{1}{N} \sum_k y_k(t) = \frac{1}{N} \sum_k y_k(0) \; \forall t$. This does not carry over to manifolds, because the average on manifolds cannot be computed by consecutive pairwise averaging; nevertheless the more symmetric character of undirected gossip may sometimes be preferred.

Proposition 6: Assume that G is uniformly connected and $\beta > 0$ is fixed. Then N agents applying the *directed* gossip algorithm, on any set, asymptotically synchronize with probability 1. Also, N agents applying the *undirected* gossip algorithm, on $\mathcal{M}$ the circle S^1 or a sphere S^{n-1}, asymptotically synchronize with probability 1.

[1]The neighbor chosen at $t{+}1$ is thus independent of the one chosen at t.

It must be noted that with both variants, the convergence speed can be quite slow. The undirected variant speeds up once all agents are located within a convex set of $\mathcal{M}$.

VI. Sensitivity to attraction profile

The algorithms in the previous sections are based on an attraction between agents proportional to their distance in $\mathbb{R}^m$ (e.g. chordal distance for S^1). This can also be viewed as a particular dependence on the more classical geodesic distance (e.g. sinusoidal function of arclength distance on S^1). One could naturally imagine other possibilities for agent interactions, a.o. mimicking physical attraction laws. On vector spaces, all these "attraction profiles" always lead to synchronization, as can be seen e.g. by rewriting them as a linear consensus algorithm with varying weights a_{jk}. The following shows that synchronization properties on the circle S^1 are sensitive to the attraction profile.

First consider N agents on S^1 which instead of (13) apply, for some $b > 0$ (see Fig.1),

$$\tfrac{d}{dt}\theta_k = \alpha \; \textstyle\sum_j a_{jk} \, g(\theta_j - \theta_k) \; , \qquad k = 1 \ldots N \; , \quad (25)$$

$$g(\theta) = \begin{cases} \frac{-b}{N-1}(\pi + \theta) & \text{for } \theta \in [-\pi, -\frac{\pi}{N}] \\ b\,\theta & \text{for } \theta \in [-\frac{\pi}{N}, \frac{\pi}{N}] \\ \frac{b}{N-1}(\pi - \theta) & \text{for } \theta \in [\frac{\pi}{N}, \pi] \; . \end{cases} \quad (26)$$

Proposition 7: For any equally-weighted fixed undirected G, synchronization is the only asymptotically stable equilibrium for N agents applying (25),(26) on S^1.
(proof see [21])

Remember that in contrast, when applying (13) e.g. with an undirected ring graph, there are stable "consensus" configurations different from synchronization.

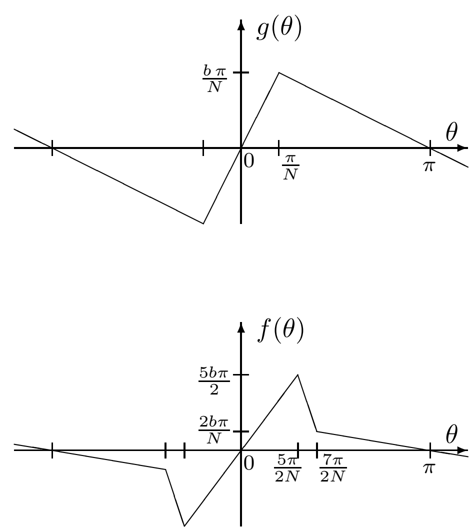

Fig. 1. Schematic representation of attraction profiles $g(\theta)$ and $f(\theta)$.

Next consider the interaction (see Fig.1)

$$\frac{d}{dt}\theta_k = \alpha \sum_j a_{jk} f(\theta_j - \theta_k) , \qquad k = 1...N , \qquad (27)$$

$$f(\theta) = \begin{cases} Nb\,\theta & \text{for } \theta \in [0, \frac{5\pi}{2N}] \\ \frac{4-5N}{2} b\,\theta + (8.75 - \frac{5b}{N})\pi & \text{for } \theta \in [\frac{5\pi}{2N}, \frac{7\pi}{2N}] \\ \frac{4b}{7-2N}(\theta - \pi) & \text{for } \theta \in [\frac{7\pi}{2N}, \pi] \\ -f(-\theta) & \text{for } \theta \in [-\pi, 0) . \end{cases} \qquad (28)$$

For an equally-weighted complete graph, synchronization is the only stable equilibrium under (13). One shows that, in contrast, the configuration with $N > 3$ agents uniformly distributed on the circle (i.e. separated by π/N) is a stable equilibrium when applying (27),(28) with $b > 0$.

Interactions that stabilize distributed configurations on S^1 for equally-weighted complete G are proposed in [26]. The goal here is to show that, locally, this can happen with a *nowhere repulsive* interaction $f(\theta)$ *close to the nicely synchronizing* $g(\theta)$. In conclusion, modifying the attraction profile w.r.t. (13) can both enhance or deteriorate convergence to synchronization. The proposed alternative attraction profiles derive from alternative distance measures among agents. They are not smooth, but our conclusions still hold with smoothed (e.g. finite Fourier series) approximations.

VII. Conclusion

This paper extends the consensus algorithm framework from vector spaces to connected compact homogeneous manifolds. It builds gradient algorithms which can be seen as the projection of linear consensus algorithms onto the embedded manifold. These algorithms can converge to several configurations, formalized as "consensensus configurations", depending on the communication graph and initial positions. It is shown that unlike for vector spaces, convergence properties on the circle depend on the attraction profile among connected agents. Further, means to generically achieve global synchronization are proposed, using estimator variables or a stochastic "gossip" setting. The (anti-)consensus algorithms can also be used to distribute points on compact manifolds, which may be useful for some applications.

VIII. Acknowledgements

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 currently a postdoc of the FNRS (Belgian Fund for Scientific Research).

References

[1] P. Absil, R. Mahony, and R. Sepulchre. Riemannian geometry of Grassmann manifolds with a view on algorithmic computation. *Acta Appl. Math.*, 80(2):199–220, 2004.

[2] P. Absil, R. Mahony, and R. Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, 2007.

[3] A. Bondhus, K. Pettersen, and J. Gravdahl. Leader/follower synchronization of satellite attitude without angular velocity measurements. *Proc. 44th IEEE Conf. Decision and Control*, pages 7270–7277, 2005.

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

[5] R. Brockett. Dynamical systems that sort lists, diagonalize matrices, and solve linear programming problems. *Lin. Alg. Appl.*, 146:79–91, 1991.

[6] F. Bullo and R. M. (advisor). Nonlinear control of mechanical systems: a Riemannian geometry approach. *PhD Thesis, CalTech*, 1998.

[7] S. Buss and J. Fillmore. Spherical averages and applications to spherical splines and interpolation. *ACM Trans. Graphics*, 20(2):95–126, 2001.

[8] J. Conway, R. Hardin, and N. Sloane. Packing lines, planes, etc.: packings in Grassmannian spaces. *Exper.Math.*, 5(2):139–159, 1996.

[9] J. Cortés, S. Martínez, and F. Bullo. Coordinated deployment of mobile sensing networks with limited-range interactions. *Proc. 43rd IEEE Conf. Decision and Control*, pages 1944–1949, 2004.

[10] A. Edelman, T. Arias, and S. Smith. The geometry of algorithms with orthogonality constraints. *SIAM J. Matrix Anal. Appl.*, 20(2):303–353, 1999.

[11] P. Gruber and F. Theis. Grassmann clustering. *Proc. 14th Eur. Signal Processing Conf.*, 2006.

[12] U. Helmke and J. Moore. *Optimization and Dynamical Systems*. Springer, 1994.

[13] H. Karcher. Riemannian center of mass and mollifier smoothing. *Comm. Pure and Appl. Math.*, 30:509–541, 1977.

[14] Y. Kuramoto. Self-entrainment of population of coupled nonlinear oscillators. In *Internat. Symp. on Math. Problems in Theoretical Physics*, volume 39 of *Lect. N. Physics*, page 420. Springer, 1975.

[15] J. Lawton and R. Beard. Synchronized multiple spacecraft rotations. *Automatica*, 38:1359–1364, 2002.

[16] N. Leonard, D. Paley, F. Lekien, R. Sepulchre, D. Frantantoni, and R. Davis. Collective motion, sensor networks and ocean sampling. *Proc. IEEE*, 95(1):48–74, 2007.

[17] M. Moakher. Means and averaging in the group of rotations. *SIAM J. Matrix Anal. Appl.*, 24(1):1–16, 2002.

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

[19] S. Nair and N. Leonard. Stabilization of a coordinated network of rotating rigid bodies. *Proc. 43rd IEEE Conf. Decision and Control*, pages 4690–4695, 2004.

[20] R. Olfati-Saber, J. Fax, and R. Murray. Consensus and cooperation in networked multi-agent systems. *Proc. IEEE*, 95(1):215–233, 2007.

[21] A. Sarlette and R. S. (advisor). Geometry and symmetries in coordination control. *PhD Thesis, University of Liège*, 2009.

[22] A. Sarlette and R. Sepulchre. Consensus optimization on manifolds. *SIAM J. Control and Optimization*, 48(1):56–76, 2009.

[23] A. Sarlette, R. Sepulchre, and N. Leonard. Discrete-time synchronization on the N-torus. *Proc. 17th Intern. Symp. Math. Theory of Networks and Systems*, pages 2408–2414, 2006.

[24] A. Sarlette, S. Tuna, V. Blondel, and R. Sepulchre. Global synchronization on the circle. *Proc. 17th IFAC World Congress*, 2008.

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

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

[27] R. Sepulchre, D. Paley, and N. Leonard. Stabilization of planar collective motion with limited communication. *IEEE Trans. Automatic Control*, 53(3):706–719, 2008.

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

[29] S. Strogatz. *Sync: The emerging science of spontaneous order*. Hyperion, 2003.

[30] J. Tsitsiklis and D. Bertsekas. Distributed asynchronous optimal routing in data networks. *IEEE Trans. Automatic Control*, 31(4):325–332, 1986.

[31] J. Tsitsiklis, D. Bertsekas, and M. Athans. Distributed asynchronous deterministic and stochastic gradient optimization algorithms. *IEEE Trans. Automatic Control*, 31(9):803–812, 1986.

[32] T. Vicsek, A. Czirók, E. Ben-Jacob, I. Cohen, and O. Shochet. Novel type of phase transition in a system of self-driven particles. *Phys. Rev. Lett.*, 75(6):1226–1229, 1995.