

Distributed Consensus on Manifolds using the Riemannian Center of Mass

Spencer Kraisler, *Student Member, IEEE*, Shahriar Talebi, *Student Member, IEEE*,
and Mehran Mesbahi, *Fellow, IEEE*

Abstract—In this paper, we develop a distributed consensus algorithm for agents whose states evolve on a manifold. This algorithm is complementary to traditional consensus, predominantly developed for systems with dynamics on vector spaces. We provide theoretical convergence guarantees for the proposed manifold consensus provided that agents are initialized within a geodesically convex (g-convex) set. This required condition on initialization is not restrictive as g-convex sets may be comparatively “large” for relevant Riemannian manifolds. Our approach to manifold consensus builds upon the notion of Riemannian Center of Mass (RCM) and the intrinsic structure of the manifold to avoid projections in the ambient space. We first show that on a g-convex ball, all states coincide if and only if each agent’s state is the RCM of its neighbors’ states. This observation facilitates our convergence guarantee to the consensus submanifold. Finally, we provide simulation results that exemplify the linear convergence rate of the proposed algorithm and illustrates its statistical properties over randomly generated problem instances.

Index Terms—Consensus, Riemannian Manifolds, Riemannian Center of Mass

I. INTRODUCTION

Consensus techniques are a ubiquitous class of algorithms that arise in fields such as distributed computation, optimization, and control of multi-agent systems. At a foundational level, consensus algorithms steer a set of dynamical agents towards a single point in a distributed fashion [1]. Such a setup is essential in multi-agent systems that lack all-to-all communication. Applications of such a setup include constellations around planetary bodies, an intelligent system of cameras tracking moving targets, and resource allocation for an electric grid [2]. Consensus type algorithms also appear in nature, e.g., synchronization in neuronal networks and flashing of fireflies in unison [3].

However, consensus algorithms have been predominantly studied in the context of discrete and Euclidean spaces. There is a strong motivation to generalize this fundamental niche of multi-agent control to particular regimes of Riemannian manifolds. The state space of many systems in robotics and aerospace are manifolds. This includes orientation of satellites, dome cameras, and robot arms [4], [5]. As a result for many applications, it is more natural to utilize the

manifold structure of the underlying state space, rather than linearizing the dynamics in a local neighborhood.

Consensus algorithms for manifolds have been studied previously in the literature. One of the earliest papers that studies consensus on manifolds and provides numerical schemes is [6], which introduces a discrete distributed consensus algorithm on $SO(3)$ and $SE(3)$. The proof of convergence of this algorithm appeared in [7]. Furthermore, the works [8], [9] provide a globally convergent algorithm for a particular class of homogeneous manifolds. This algorithm hinges on the projection from the ambient space, and hence is *not intrinsic*. Finally, another relevant work is [10] that proposes a discrete distributed consensus algorithm with almost-global convergence on $SO(3)$.

In this paper, we propose a discrete-time distributed consensus algorithm for Riemannian manifolds. The algorithm is intrinsic to the manifold and does not require any embedding. Also, the radius of initialization set for guaranteed convergence is reasonably large which is also shown to be forward invariance. The rest of the paper is organized as follows. §II introduces the problem statement and provides a brief overview of the algorithm. §III provides a short background on essential constructs for manifolds and consensus algorithms. In §IV, we explicate the proposed distributed consensus algorithm. Subsequently, we provide a proof of convergence and elaborate on properties of our algorithm in §V. §VI showcases a concrete application to $SO(n)$, including numerical results for several examples. Lastly, we conclude our paper and discuss future directions in §VII.

II. PROBLEM STATEMENT

Consider $\mathcal{M}$ as a complete Riemannian manifold. Let $x_1, \dots, x_N \in \mathcal{M}$ be the states of N dynamic agents and $\mathcal{G} = ([N], E)$ be a connected undirected communication graph. The problem considered is to develop a consensus algorithm with the following three properties: **(1)** The designed dynamics are discrete, memoryless, and *distributed*, i.e., the dynamics for the i th agent depends only on the information of itself and its neighbors $\mathcal{N}_i$. In other words, we want to design mappings $T_i : \mathbb{N} \times \mathcal{M} \times \mathcal{M}^{|\mathcal{N}_i|} \rightarrow \mathcal{M}$ for every $i \in [N]$ which can be implemented as

$$x_i(k+1) = T_i(k, x_i(k), x_j(k) : j \sim i),$$

where $j \sim i$ denotes any neighbor $j \in \mathcal{N}_i$ of the i th agent. **(2)** The algorithm is *intrinsic* to the manifold. In

The authors are with the William E. Boeing Department of Aeronautics and Astronautics, University of Washington, Seattle, WA, USA. Emails: {kraisler+shahriar+mesbahi}@uw.edu. The research of the authors has been supported by NSF grant ECCS-2149470 and AFOSR grant FA9550-20-1-0053.

particular, the way the manifold is embedded in the ambient Euclidean space has no effect on the algorithm. This implies the algorithm does not use projections. One benefit of this feature is that the choice of parameterization has no effect on the convergence properties of the algorithm, such as error rate, convergence point, etc. For instance, when $\mathcal{M} = SE(3)$ the algorithm should have the same performance and convergence guarantees if we use dual quaternions or (vector, rotation matrix) pairs. And lastly, (3) there is a proven reasonably *large radius* such that consensus is guaranteed if agents are initialized within a geodesic ball characterized by that bound. By “large” we mean a quantified lower bound on the radius of the domain on which consensus is guaranteed.

To the best of our knowledge, the only algorithm that meets all these goals is presented in [7]; this paper is a key resource for the topic of distributed consensus on manifolds. However, the paper shows that a small domain is forward invariant under the proposed dynamics. This domain depends on the manifold convexity radius and the diameter of the graph. In contrast, our proposed algorithm provides convergence guarantees for a larger domain is invariant of the graph structure. Our approach builds on an observation we refer to as the “Mean-Consensus Lemma.”

III. MANIFOLDS AND CONSENSUS

In this section, we provide a brief overview of geometric concepts used in this paper. We then summarize the intuition on consensus algorithms in the Euclidean setup, illustrating our motivation in developing the analogous algorithm for the Riemannian case.

A. Manifolds

Many geometric constructs—such as curvature, direction, and length on a smooth manifolds—can be characterized by a *Riemannian metric* [11], [12]. Specifically, it is a mapping from the manifold $\mathcal{M}$ to 2-tensor fields, that smoothly assigns to each point $x \in \mathcal{M}$ an *inner product* defined on the tangent space $T_x\mathcal{M}$ of that point:

$$\langle \cdot, \cdot \rangle_x : T_x\mathcal{M} \times T_x\mathcal{M} \rightarrow \mathbb{R};$$

this inner product naturally defines a norm $\|\cdot\|_x$ that depends on x . Then, a *Riemannian manifold* is just a smooth manifold equipped with a Riemannian metric.

Next, given a smooth curve $\gamma : J \rightarrow \mathcal{M}$ on $\mathcal{M}$ parameterized by an interval $J \subset \mathbb{R}$, we naturally define its length as

$$L(\gamma) := \int_J \|\dot{\gamma}(t)\|_{\gamma(t)} dt. \quad (1)$$

Also, for any $x, y \in \mathcal{M}$, let $\mathcal{C}_{x,y}$ be the family of constant-speed curves γ starting from x and ending at y . Then, restricting (1) to $\mathcal{C}_{x,y}$, any extremum curve¹ is called a *geodesic* and any globally minimizing curve is called a *minimizing geodesic*. Finally, the *geodesic distance* between x and y is defined as,

$$d(x, y) := \inf\{L(\gamma) : \gamma \in \mathcal{C}_{x,y}\}. \quad (2)$$

¹Over $\mathcal{C}_{x,y}$, (1) only admits minima and saddle points, and never maxima.

Suppose $\mathcal{M}$ is complete. The *manifold exponential* at $x \in \mathcal{M}$, denoted by $\exp_x : T_x\mathcal{M} \rightarrow \mathcal{M}$, is defined as $\exp_x(v) = \gamma(1)$, where $\gamma(\cdot)$ is the unique geodesic with $\gamma(0) = x$ and $\dot{\gamma}(0) = v$. The manifold exponential is invertible on a neighborhood $U \subset T_x\mathcal{M}$ of 0_x , and its inverse, called the *manifold logarithm*, is denoted by $\log_x : \exp_x(U) \subset \mathcal{M} \rightarrow T_x\mathcal{M}$.

B. Consensus Algorithms

Herein, we review the analogy between standard consensus algorithm in an Euclidean space and its “standard counterpart” on Riemannian manifolds. We then clarify our intuition behind the developed algorithm and its convergence properties.

1) *Consensus algorithm in the Euclidean case*: Let $x_1, \dots, x_N \in \mathbb{R}^n$ represent the states of discrete dynamical agents under the connected communication graph $\mathcal{G}$. The dynamics of the i th agent assumes the form,

$$x_i(k+1) = x_i(k) + \epsilon \sum_{j \sim i} (x_j(k) - x_i(k)), \quad (3)$$

for some $\epsilon > 0$. This is standard consensus algorithm that can compactly be represented as,

$$\mathbf{x}(k+1) = \mathbf{x}(k) - \epsilon(L \otimes I)\mathbf{x}(k),$$

where $\mathbf{x}(k) = [x_1(k), \dots, x_N(k)]^T$ and L is the graph Laplacian of $\mathcal{G}$. For sufficiently small $\epsilon > 0$, it is guaranteed for this algorithm that the states eventually reach consensus, i.e., $\mathbf{x}(k) \rightarrow x^* \mathbf{1}$ for some $x^* \in \mathbb{R}^n$; in fact, x^* will be the average of agents’ initial states.

2) *An analogous algorithm on manifolds*: Here, we present a standard consensus algorithm on Riemannian manifolds. The paper [7] is the first work providing rigorous convergence analysis for the Riemannian consensus. Like before, let $\mathbf{x} = (x_1, \dots, x_N) \in \mathcal{M}^N$ be agents’ states evolving on the product manifold $\mathcal{M}^N$ in discrete-time and under the connected communication graph $\mathcal{G}$. The dynamics of the i th agent is then designed as,

$$x_i(k+1) = \exp_{x_i(k)} \left(\epsilon \sum_{j \sim i} \log_{x_i(k)}(x_j(k)) \right), \quad (4)$$

for some $\epsilon > 0$. Define

$$\varphi(\mathbf{x}) := \frac{1}{2} \sum_{\{i,j\} \in E} d(x_i, x_j)^2 = \frac{1}{4} \sum_{i=1}^N \sum_{j \sim i} d(x_i, x_j)^2. \quad (5)$$

We refer to $\varphi(\mathbf{x})$ as the *consensus error with respect to $\mathcal{G}$* as $\varphi(\mathbf{x}) = 0$ if and only if $\mathbf{x}$ is at consensus. Notice that self-loops in $\mathcal{G}$ does not affect $\varphi(\mathbf{x})$. We also define the local consensus error $\varphi_i(\mathbf{x}) := \frac{1}{2} \sum_{j \sim i} d(x_i, x_j)^2$. Then, compactly, we can write the dynamics in (4) as,

$$\mathbf{x}(k+1) = \exp_{\mathbf{x}(k)}(-\epsilon \nabla \varphi(\mathbf{x}(k))).$$

We draw attention to the similarity between (4) and (3), justifying calling the former a *direct* generalization of the latter.

It was proven that if $\mathcal{M}$ is complete and has sectional curvature bounded above, then as long as agents are initialized such that $\varphi(\mathbf{x}(0)) < (r^*)^2/2D$, where $D = \text{diam}(\mathcal{G})$, consensus is guaranteed with a sublinear rate. Relaxing this constraint, if one assumes that $\mathbf{x}(0)$ are initialized within any geodesic ball B with radius $r < r^*$, and that the iterates remain in B indefinitely, then consensus is guaranteed. This assumption is not uncommon in the field of optimization on manifolds. Alternatively, if $\mathcal{M}$ is complete, simply connected, and has non-positive sectional curvature (making $\mathcal{M}$ a *Hadamard manifold*), consensus is guaranteed regardless of where the agents have been initialized. This observation follows from the fact that Hadamard manifolds are diffeomorphic to Euclidean space. Such manifolds include projective spaces and the manifold of positive-definite matrices under certain Riemannian metrics [13]. Variants of (4) have also been studied. [10] considers one such variant for $\mathcal{M} = SO(3)$ with almost-global consensus guarantee, where a continuous version is studied in [14].

IV. THE RCM AND OUR ALGORITHM

In this section, we present our algorithm that builds on the properties of Riemannian Center of Mass (RCM) as its cornerstone. As a generalization of the Euclidean mean, RCM retains many desirable properties of being a “mean” not reflected in other such generalizations [15].

Let $A \subset \mathcal{M}$. We say U is g -convex if for any $x, y \in A$, there exists a unique minimizing geodesic connecting x and y contained in A . Some authors refer to this as *strong* g -convexity [12]. Also, the *convexity radius* of $\mathcal{M}$ is defined as

$$r^* := \frac{1}{2} \min(\text{inj}(\mathcal{M}), \frac{\pi}{\sqrt{\Delta}}),$$

where $\text{inj}(\mathcal{M})$ is the injectivity radius and Δ is the upper-bound on the sectional curvature of $\mathcal{M}$ [16].

The RCM of states x_i ($i = 1, 2, \dots, N$) is defined as a point that globally minimizes the sum of squared distances, i.e.

$$\text{RCM}(\mathbf{x}) \in \arg \min_{y \in \mathcal{M}} \frac{1}{2} \sum_{i=1}^N d(y, x_i)^2.$$

The RCM exists and is unique if all x_i ’s are contained within some geodesic ball with radius $r < r^*$ [16]. Therefore, $\text{RCM} : \mathcal{C} \rightarrow \mathcal{M}$ is well-defined where,

$$\mathcal{C} := \{\mathbf{x} \in \mathcal{M}^N \mid \exists y \in \mathcal{M}, r < r^* : x_i \in B_y(r) \forall i\}$$

is the *convexity submanifold* of the product manifold $\mathcal{M}^N$ [7]. The RCM of $\mathbf{x}$ can be computed by fixing a tolerance $\tau > 0$ and step size $\epsilon > 0$, and performing gradient descent on the *sum of squared distances*

$$f_{\mathbf{x}}(y) := \frac{1}{2} \sum_{i=1}^m d(y, x_i)^2. \quad (6)$$

The details are laid out in Algorithm 1. Due to the strong g -convexity of sum of squared distances cost, gradient descent under a fixed step size enjoys a linear rate of convergence. See [17] for optimally chosen stepsizes.

Algorithm 1: RCM Subroutine

Input: $(x_1, \dots, x_m) \in \mathcal{C}$, stepsize $\epsilon > 0$, tolerance $\tau > 0$
Initialize: $\bar{x}(0) = x_1$
do
 Compute $\nabla f_{\mathbf{x}}(\bar{x}(k)) = -\sum_{i=1}^m \log_{\bar{x}(k)}(x_i)$
 Update $\bar{x}(k+1) = \exp_{\bar{x}(k)}(-\epsilon \nabla f_{\mathbf{x}}(\bar{x}(k)))$
while $\|\nabla f_{\mathbf{x}}(\bar{x}(k))\|_{\bar{x}(k)} > \tau$
return $\bar{x}(k+1)$

A. Brief Overview of the Algorithm

In our algorithm, each agent moves to the Riemannian Center of Mass (RCM) of its neighbors *one at a time*:

$$x_i(k+1) = \begin{cases} \text{RCM}(x_j(k) : j \sim i) & i-1 \equiv k \pmod{N} \\ x_i(k) & \text{else} \end{cases} \quad (7)$$

We emphasize that the RCM step above excludes self-loops as $j \neq i$. Yet, we observed that consensus is always achieved experimentally—in the case of $SO(3)$ —even if self-loops are included. However, our proof technique for arbitrarily Riemannian manifolds builds on the consensus error (5) that naturally excludes these self-loops.

Note this algorithm is memoryless and distributed since we are only taking the RCM of an agent’s neighbors at the current iteration. Also, this algorithm is intrinsic since RCM is defined in terms of the geodesic distance function—as we will subsequently discuss in this paper. A simple illustration of our algorithm for four agents evolving on a sphere is depicted in Figure 1. Here, we perform four iterations of the algorithm, moving each agent once. Notice how agents 2 and 4 both end up at the same position after one iteration due to the fact that agent 2 is the only neighbor of agent 4.

A similar asynchronous method where one agent moves at a time has been briefly mentioned in [9]. However, this reference was in regards to an extrinsic center known as the *induced arithmetic mean*. Furthermore, no convergence nor domain analysis were discussed. In the rest of this paper, we provide convergence analysis of our algorithm.

B. The Algorithm

Given a set of agents with initial states $\mathbf{x}(0) \in \mathcal{C} \subset \mathcal{M}^N$, a connected graph $\mathcal{G}$, and a fixed tolerance level $\tau > 0$, the proposed distributed algorithm proceeds as Algorithm 2.

Algorithm 2: Distributed RCM-based Consensus

Input: $(x_1(0), \dots, x_N(0)) \in \mathcal{C}$, tolerance $\tau > 0$
for $k = 0, 1, \dots$ **do**
 for $i = 1, \dots, N$ **do**
 if $\varphi_i(\mathbf{x}(k)) > 2\tau/N$ & $i-1 \equiv k \pmod{N}$
 then $x_i(k+1) \leftarrow \text{RCM}(x_j(k) : j \sim i)$ **else**
 $x_i(k+1) \leftarrow x_i(k)$




Fig. 1: (Left) illustrates four agents under the communication graph with states to be evolved on a sphere. (Right) illustrates how the algorithm is updating the states of i th agent on the sphere from i to i' , sequentially.

Note that the local stopping criterion

$$\varphi_i(\mathbf{x}(k)) := \frac{1}{2} \sum_{j \sim i} d(x_i(k), x_j(k))^2 > \frac{2\tau}{N}, \quad (8)$$

ensures that $\varphi(\mathbf{x}(k)) = \frac{1}{2} \sum_{i=1}^N \varphi_i(\mathbf{x}(k)) \leq \tau$ whenever the agents halt. We note that there is no distributed global stopping criterion in Algorithm 2. This can be achieved simply by performing Boolean consensus with (8) after each iteration. We emphasize that despite there being no distributed global stopping criterion for Algorithm 2, our proof of convergence ensures the agents will themselves eventually stop moving in finite iterations.

For the remainder of this paper, we will refer to k as an “iteration” of Algorithm 2 where a single agent moves per iteration. A “time step” of Algorithm 2 then denote N consecutive iterations within which every agent has moved exactly once. We subsequently show that our algorithms converges to a consensus point denoted as $x^* \in \mathcal{M}$.

C. Properties of Algorithm 2

In what follows, we discuss mean-like properties of Algorithm 2 in certain regimes of Riemannian manifolds that play an important role in different applications.

1) *Equivariance*: Let G be a group acting on sets X, Y . We say a map $f : X \rightarrow Y$ is G -equivariant if for any $g \in G$ and $x \in X$, we have $f(g \cdot x) = g \cdot f(x)$.

Now, suppose $\mathcal{M}$ is a Riemannian manifold acted on transitively and isometrically by a Lie group G . For example, if $\mathcal{M}$ is a sphere, then we may choose $G = SO(3)$. Then, $\text{RCM} : \mathcal{C} \rightarrow \mathcal{M}$ is G -equivariant, and so is Algorithm 2. In particular, if $(x_1, \dots, x_N) \in \mathcal{C}$, $g \in \mathcal{M}$, and x^* was the point of consensus via Algorithm 2, then the point of consensus when the input is $(g \cdot x_1, \dots, g \cdot x_N)$ will become $g \cdot x^*$ —similar to the Euclidean mean that is translation-equivariant.

2) *Forward Invariance and Contractability*: Like the Euclidean mean, the consensus point is always inside the g -convex hull of the initial points. This is due to the property of the RCM that it is always contained in the interior of the g -convex hull of the points [16]. In fact, at every iteration of the algorithm, the agents move inside the current g -convex hull. This property induces a behavior such that the initial positions contract to “their center” as the consensus point.

V. CONVERGENCE ANALYSIS

Define the *consensus submanifold* of $\mathcal{M}^N$ as

$$\mathcal{A} := \{(x_1, \dots, x_N) \in \mathcal{M}^N : x_1 = \dots = x_N\}.$$

To prove convergence, we first introduce an intuitive lemma. We also provide a proof since, to the best of our knowledge, is missing from literature.

Lemma 1. (Mean-Consensus Lemma) *Let $\mathcal{M}$ be a Riemannian manifold and let $\mathcal{C}$ be the convexity submanifold of $\mathcal{M}^N$. Suppose $\mathbf{x} = (x_1, \dots, x_N) \in \mathcal{C}$ and let $\mathcal{G} = ([N], E)$ be a connected graph. Then $\mathbf{x} \in \mathcal{A}$ if and only if $x_i = \text{RCM}(x_j : j \sim i)$ for each i . That is, each point is at the RCM of its neighbors.*

Proof. If $\mathbf{x} \in \mathcal{A}$ then the claim is obvious, so to show the converse and for the sake of contradiction, suppose these points are not in consensus. Since $x_i = \text{RCM}(x_j : j \sim i)$, then we have for each i

$$\nabla f_{(x_j : j \sim i)}(x_i) = \sum_{j \sim i} \log_{x_i}(x_j) = 0, \quad (9)$$

where $f_{(x_j : j \sim i)}(\cdot)$ is defined in (6). Remark the gradient of φ at $\mathbf{x}$ is

$$\nabla \varphi(\mathbf{x}) = \left[\sum_{j \sim i} \log_{x_i}(x_j) \right]_i \in T_{\mathbf{x}} \mathcal{M}^N,$$

where $[\xi_i]_i \in T_{x_i} \mathcal{M}$ denotes the i th component of a tangent vector in the product topology to be ξ_i . Note that the coordinates of $\mathbf{x}$ fit within a g -convex geodesic ball, yet are not at consensus. Thus, by connectivity of the graph, $\mathbf{x}$ cannot be a critical point of φ [7]. Thus, we must have $\nabla \varphi(\mathbf{x}) \neq 0$. This contradicts (9). Therefore if each point is at the RCM of its neighbors, they are necessarily at consensus, which completes the proof. $\square$

Theorem 1. *Let $\mathcal{M}$ be a complete Riemannian manifold with curvature bounded above, and let $\mathcal{C}$ be the convexity submanifold of $\mathcal{M}^N$. Let $x_1, \dots, x_N \in \mathcal{M}$ be discrete dynamic agents under the connected communication graph $\mathcal{G} = ([N], E)$. Suppose $\mathbf{x}(0) = (x_1(0), \dots, x_N(0)) \in \mathcal{C}$, then under Algorithm 2, the agents eventually reach consensus, i.e. $\mathbf{x}(k) \rightarrow \mathcal{A}$ as $k \rightarrow \infty$.*

Proof. In our proof, we use the cost function (5) and show that it strictly decreases over each iteration unless the agent that is supposed to move is already at the RCM of its neighbors. We then show convergence of the cost implies convergence of the agents to a consensus configuration.

Suppose agent i moved at iteration k . Let $E(i) \subset E$ be all edges adjacent to i . Then, by decomposing $\varphi(\mathbf{x}(k))$ into edge distances of $E(i)$ and $E(i)^c$, we arrive at

$$\begin{aligned} \varphi(\mathbf{x}(k)) &= \frac{1}{2} \sum_{\{i,j\} \in E(i)} d(x_i(k), x_j(k-1))^2 \\ &\quad + \frac{1}{2} \sum_{\{l,m\} \in E(i)^c} d(x_l(k-1), x_m(k-1))^2 \end{aligned}$$

Since $x_i(k)$ is the *unique* point that minimizes the function $f_{(x_j(k-1):j \sim i)}(\cdot)$, it follows

$$\begin{aligned} \frac{1}{2} \sum_{\{i,j\} \in E(i)} d(x_i(k), x_j(k-1))^2 \\ \leq \frac{1}{2} \sum_{\{i,j\} \in E(i)} d(x_i(k-1), x_j(k-1))^2 \end{aligned}$$

with equality if and only if

$$x_i(k-1) = \text{RCM}(x_j(k-1) : j \sim i). \quad (10)$$

Therefore

$$\begin{aligned} \varphi(\mathbf{x}(k)) &\leq \frac{1}{2} \sum_{\{i,j\} \in E(i)} d(x_i(k-1), x_j(k-1))^2 \\ &+ \frac{1}{2} \sum_{\{l,m\} \in E(i)^c} d(x_l(k-1), x_m(k-1))^2 \\ &= \varphi(\mathbf{x}(k-1)), \end{aligned}$$

for all k , with equality if and only if (10). Since $\{\varphi(\mathbf{x}(k))\}$ is bounded below by 0, the sequence converges.

For $1 \leq i \leq N$, define $T_i : \mathcal{C} \rightarrow \mathcal{C}$ as

$$\begin{aligned} T_i(x_1, \dots, x_N) \\ = (x_1, \dots, x_{i-1}, \text{RCM}(x_j : j \sim i), x_{i+1}, \dots, x_N). \end{aligned}$$

Next, define $T : \mathcal{C} \rightarrow \mathcal{C}$ as the composition

$$T := T_N \circ T_{N-1} \circ \dots \circ T_1. \quad (11)$$

In other words, (11) represents one time-step of (7), effectively moving each agent exactly once. Remark T is time-invariant and $\varphi(T(\mathbf{x})) - \varphi(\mathbf{x}) \leq 0$ for all $\mathbf{x} \in \mathcal{C}$. Define

$$\mathcal{I} := \{\mathbf{x} \in \mathcal{C} : \varphi(T(\mathbf{x})) = \varphi(\mathbf{x})\}.$$

Note $\mathbf{x} \in \mathcal{I}$ if and only if (10) holds for each $1 \leq i \leq N$. It was then showed in Lemma 1 that coordinate of $\mathbf{x}$ is at the RCM of its neighbors iff $\mathbf{x} \in \mathcal{A}$. Therefore $\mathcal{I} = \mathcal{A}$.

Next, by Whitney Embedding Theorem, every smooth n -dimensional manifold admits a proper embedding ι into $\mathbb{R}^{2n+1}$ ([11, Corollary 6.16]). But such an embedding is a closed mapping [11, Theorem A.57]. That is, there exists a smooth embedding (i.e., a smooth mapping with injective differential) $\iota : \mathcal{M}^N \rightarrow \mathbb{R}^{2nN+1}$ such that $\iota(C)$ is closed in $\mathbb{R}^{2nN+1}$ for any $C \subset \mathcal{M}^N$ closed in $\mathcal{M}^N$ under the product manifold topology. For the remainder of this proof, we will identify $\mathcal{M}^N$ with $\iota(\mathcal{M}^N) \subset \mathbb{R}^{2nN+1}$.

Let $\Omega(\mathbf{x}(0))$ be the forward limit set of $\mathbf{x}(0)$ under T . That is, $\Omega(\mathbf{x}(0))$ is the set of limit points of all convergent subsequences of $\{\mathbf{x}(k)\}$. Let $B \subset \mathcal{M}$ be a closed g -convex geodesic ball with radius $r < r^*$ with each $x_1(0), \dots, x_N(0) \in B$. Since ι is a closed smooth embedding (in particular continuous), the N -product set B^N is compact in $\mathbb{R}^{2nN+1}$.

Note that the RCM is always contained inside the closed g -convex hull of the points [16]. Thus, for every iteration of the algorithm, the agents always move inside the current closed g -convex hull. As such, $\{\mathbf{x}(k)\} \subset B^N$ for all k . But

then we must have $\Omega(\mathbf{x}(0)) \subset B^N$ since B^N is closed in $\mathbb{R}^{2nN+1}$. Since B^N is compact and $\{\mathbf{x}(k)\} \subset B^N$, then by the discrete version of LaSalle invariance principle [18, Theorem 2], we have $\mathbf{x}(k) \rightarrow \mathcal{I} = \mathcal{A}$. $\square$

Remark 1. Note that the proposed algorithm and its convergence guarantees provided here are more general. In particular, our RCM-based algorithm can be generalized by replacing the RCM with any other “mean-type construct” arising from locally minimizing “a distributed consensus error”. The convergence guarantees then directly follows by the same arguments in this section.

VI. APPLICATION: SYNCHRONIZATION ON LIE GROUPS

Herein, we let $\mathcal{M} \subset \mathbb{R}^{n \times n}$ to be a matrix Lie group equipped with the following left-invariant Riemannian metric defined at $x \in \mathcal{M}$ as

$$\langle \xi, \eta \rangle_x := \frac{1}{2} \text{tr}[(x^{-1}\xi)^T x^{-1}\eta], \quad \forall \xi, \eta \in T_x \mathcal{M}. \quad (12)$$

This metric is commonly equipped to matrix Lie groups due to its nice geometric properties. First, all connected matrix Lie groups equipped with (12) are geodesically complete since they are homogeneous [19]. Second, the induced norm on the corresponding Lie algebra $\mathfrak{g}$ coincides with half the Frobenius norm, i.e.,

$$\|\xi\|_f^2 := \frac{1}{2} \text{tr}(\xi^T \xi), \quad \forall \xi \in \mathfrak{g}.$$

Let $\text{Exp}(\cdot)$ and $\text{Log}(\cdot)$ be the matrix exponential and logarithm, respectively. In the case that $\mathcal{M} = SO(n) \subset \mathbb{R}^{n \times n}$, the squared geodesic distance under (12) is

$$d(x, y)^2 = -\frac{1}{2} \text{tr}[\text{Log}(x^T y)^2]. \quad (13)$$

For matrix Lie groups equipped with any bi-invariant metrics, such as the case when $\mathcal{M} = SO(n)$ equipped with (12), the unique geodesic starting at $x \in \mathcal{M}$ with initial velocity $v \in T_x \mathcal{M}$ is always²

$$\gamma(t) = x \text{Exp}(tx^{-1}v).$$

Under this setup, the geodesic connecting $x, y \in \mathcal{M}$ attains the form

$$\gamma(t) = x \text{Exp}[t \text{Log}(x^{-1}y)]. \quad (14)$$

Note that while (2) is defined everywhere, (13) and (14) only hold whenever the argument of $\text{Log}(\cdot)$ is contained in its domain of definition.

A. Simulations and Numerical Results

In this section, we exemplify the performance of our algorithm through several simulation scenarios. We also provide a comparison with the algorithm presented in [6]. We illustrates these numerical simulations for the case of $SO(3)$ equipped with (12). The MATLAB code for these simulations are available online in [20].

²Note that this does not necessarily hold when the metric is merely left-invariant.




Fig. 2: Statistical visualization of the normalized consensus error at each time step, for (top) the proposed Algorithm 2, and (bottom) the algorithm in [7] over the same 40 random problem instances.

We generate 40 different problem instances with $N = 10$ agents randomly initialized in $\mathcal{C} \subset SO(3)^{10}$. We run Algorithm 2 on each instance for 30 time steps. We illustrate the result in Figure 2 showing the statistics of the consensus error $\varphi(x)$ at each time step of Algorithm 2, confirming a linear convergence rate. Although this observation is not formally proved in this work, several numerical experiments with different communication graphs and initial states have confirmed this hypothesis.

Next, for comparison, we also ran the algorithm presented in [7] on the same problem instances for 30 time steps. The user-tuned step size for this algorithm is set to 0.1 as suggested therein. The resulting consensus error are illustrated in Figure 2. In order to compare our algorithm to this one, note that each time step in Figure 2 corresponds to N iterations of Algorithm 2, within which every agent moves exactly once. While both algorithms converge on all problem instances, the proposed algorithm in this paper tends to have a faster convergence (seemingly linear) with less variance.

VII. CONCLUDING REMARKS

In this paper, we presented a novel discrete-time distributed consensus algorithm for Riemannian manifolds with bounded curvature. We proved that consensus is guaranteed if agents are initialized within a geodesic ball whose radius is less than the convexity radius of the manifold. A key observation that enabled our result is established in Lemma 1 that we believe is of independent interest and can be used to analyze other classes of consensus algorithms. An immediate future direction of this work is to justify the linear rate of convergence which is apparent in the numerical simulations. Also, several variants of the proposed algorithm can be considered. This includes updating the agents randomly, rather than in an

order which would remove the need for a global agent index. The case in which agent are moving simultaneously is yet another interesting variant to study.

VIII. ACKNOWLEDGMENTS

The authors would like to acknowledge the insightful discussions with Saptarshi Bandyopadhyay and Amir Rahmani at the Jet Propulsion Laboratory.

REFERENCES

- [1] R. Olfati-Saber and R. M. Murray, "Consensus problems in networks of agents with switching topology and time-delays," *IEEE Transactions on automatic control*, vol. 49, no. 9, pp. 1520–1533, 2004.
- [2] S. S. Kia, B. Van Scoy, J. Cortes, R. A. Freeman, K. M. Lynch, and S. Martinez, "Tutorial on dynamic average consensus: The problem, its applications, and the algorithms," *IEEE Control Systems Magazine*, vol. 39, no. 3, pp. 40–72, 2019.
- [3] S. H. Strogatz, *Sync: How order emerges from chaos in the universe, nature, and daily life*. Hachette UK, 2012.
- [4] A. Sarlette, S. Bonnabel, and R. Sepulchre, "Coordinated motion design on Lie groups," *IEEE Transactions on Automatic Control*, vol. 55, no. 5, pp. 1047–1058, 2010.
- [5] J. M. Selig, "Lie groups and Lie algebras in robotics," in *Computational Noncommutative Algebra and Applications*, pp. 101–125, Springer, 2004.
- [6] R. Tron, R. Vidal, and A. Terzis, "Distributed pose averaging in camera networks via consensus on $SE(3)$," in *2008 Second ACM/IEEE International Conference on Distributed Smart Cameras*, pp. 1–10, IEEE, 2008.
- [7] R. Tron, B. Afsari, and R. Vidal, "Riemannian consensus for manifolds with bounded curvature," *IEEE Transactions on Automatic Control*, vol. 58, no. 4, pp. 921–934, 2012.
- [8] A. Sarlette and R. Sepulchre, "Consensus optimization on manifolds," *SIAM journal on Control and Optimization*, vol. 48, no. 1, pp. 56–76, 2009.
- [9] R. Sepulchre, "Consensus on nonlinear spaces," *IFAC Proceedings Volumes*, vol. 43, no. 14, pp. 1029–1039, 2010.
- [10] R. Tron, B. Afsari, and R. Vidal, "Intrinsic consensus on $SO(3)$ with almost-global convergence," in *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*, pp. 2052–2058, IEEE, 2012.
- [11] J. M. Lee, *Introduction to Smooth Manifolds*. Springer, 2nd ed., 2013.
- [12] N. Boumal, "An introduction to optimization on smooth manifolds," *Available online*, May, vol. 3, p. 4, 2020.
- [13] M. Moakher and M. Zéraï, "The riemannian geometry of the space of positive-definite matrices and its application to the regularization of positive-definite matrix-valued data," *Journal of Mathematical Imaging and Vision*, vol. 40, no. 2, pp. 171–187, 2011.
- [14] C. Mostajeran, *Invariant Differential Positivity*. PhD thesis, University of Cambridge, 2018.
- [15] N. Miolane, "Defining a mean on Lie group," Master's thesis, Imperial College London, Sept. 2013.
- [16] B. Afsari, "Riemannian L_p center of mass: existence, uniqueness, and convexity," *Proceedings of the American Mathematical Society*, vol. 139, no. 2, pp. 655–673, 2011.
- [17] B. Afsari, R. Tron, and R. Vidal, "On the convergence of gradient descent for finding the riemannian center of mass," *SIAM Journal on Control and Optimization*, vol. 51, no. 3, pp. 2230–2260, 2013.
- [18] W. Mei and F. Bullo, "LaSalle invariance principle for discrete-time dynamical systems: A concise and self-contained tutorial," *arXiv preprint arXiv:1710.03710*, 2017.
- [19] J. M. Lee, *Introduction to Riemannian Manifolds*. Cham, Switzerland: Springer Nature, 2nd ed., 2018.
- [20] S. Kraislner, S. Talebi, and M. Mesbahi, "Distributed RCM-Based Consensus," *GitHub repository*, 2023. Available online at <https://github.com/spencerkraislner/Distributed-RCM-Based-Consensus>.