Kuramoto models on spheres: Using hyperbolic geometry to explain their low-dimensional dynamics

Max Lipton,[1, a)] Renato Mirollo,[2, b)] and Steven H. Strogatz[1, c)]
[1)] *Cornell University*
[2)] *Boston College*

(Dated: 6 May 2021)

The aim of this paper is to explain and unify some work by Tanaka[1], Lohe[2] and Chandra et al.[3,4] on a generalization of Kuramoto oscillator networks to the case of higher dimensional "oscillators." Instead of oscillators represented by points on the unit circle S^1 in $\mathbb{R}^2$, the individual units in the network are represented by points on a higher dimensional unit sphere S^{d-1} in $\mathbb{R}^d$. Tanaka demonstrates in his 2014 paper that the dynamics of such a system can be reduced using Möbius transformations fixing the unit ball, similar[5] to the classic case when $d = 2$. Tanaka also presents a generalization of the famous Ott-Antonsen reduction[6] for the complex version of the system[7]. Lohe derives a similar reduction using Möbius transformations for the finite-N model, whereas Chandra et al. concentrate on the infinite-N or continuum limit system, and derive a dynamical reduction for a special class of probability densities on S^{d-1}, generalizing the Poisson densities used in the Ott-Antonsen reduction.

Previously studied oscillator systems[1–4] are intimately related to the natural hyperbolic geometry on the unit ball B^d in $\mathbb{R}^d$; as we shall show, once this connection is realized, the reduced dynamics, evolution by Möbius transformations and the form of the special densities[3,4] all follow naturally. This framework also allows one to see the seamless connection between the finite and infinite-N cases. In addition, we shall show that special cases of these networks have gradient dynamics with respect to the hyperbolic metric, and so their dynamics are especially easy to describe.

In 1975, Kuramoto introduced a model of collective synchronization in a large population of coupled oscillators with randomly distributed natural frequencies. Kuramoto's model displayed many remarkable features: It was exactly solvable (at least in some sense, and in the limit of infinitely many oscillators), despite being nonlinear and high-dimensional. Its solution shed analytical light on a phase transition to mutual synchronization that Winfree had previously discovered in a similar but less convenient system of oscillators. Since then, the Kuramoto model has been an object of fascination for nonlinear dynamicists, as well as a simplified model for many real-world instances of synchronization in physics, biology, chemistry, and engineering.

From a mathematical standpoint, one of the most intriguing problems has been to explain the tractability of the Kuramoto model. What symmetry or other hidden structure accounts for its solvability?

The first clues came from work on an adjacent topic: the dynamics of series arrays of N identical overdamped Josephson junctions. The governing equations for these superconducting oscillators are closely related to the those of the Kuramoto model, and themselves displayed remarkable dynamical features. These included ubiquitous neutral stability of splay states, invariant low-dimensional tori, and evidence of constants of motion, despite the presence of damping and driving in the governing equations. These features were explained in 1993 by the discovery of a certain change of variables, now called the Watanabe-Strogatz transformation, which showed that for the governing equations have $N-3$ constants of motion, for all $N \geq 3$. Goebel then simplified and rationalized this transformation by showing that it could be viewed as a time-dependent version of a linear fractional transformation, a standard tool in complex analysis. For more than a decade, however, these results did not attract much attention, perhaps because they were assumed to be restricted to problems about Josephson junctions, and within that specialized setting, even further restricted to junctions that were strictly identical.

A breakthrough occurred in 2008 with the work of Ott and Antonsen. They found an astonishing way to capture the macroscopic dynamics of the infinite-N Kuramoto model, even when the oscillators' frequencies were non-identical and randomly distributed. First, they wrote down an ansatz – seemingly pulled out of thin air – for the density function $\rho(\theta,\omega,t)$ of oscillators having phase θ and intrinsic frequency ω at time t.

Their ansatz had the form of a time-dependent Poisson density (a density better known for its role in the study of partial differential equations, specifically for the solution of Laplace's equation on a disk, given the values of the unknown function on the bounding circle). By making this ansatz of a Poisson density, Ott and Antonsen reduced the infinite-N Kuramoto model, an integro-partial differential equation, to an infinite set of coupled ordinary differential equations. Then, by further assuming the intrinsic frequencies were randomly distributed according to a Lorentzian (aka Cauchy) distribution, Ott and Antonsen showed that the order parameter dynamics of the Kuramoto model could be reduced tremendously, all the way down to an ordinary differential equation for a single scalar variable, the amplitude of the order parameter. With this discovery, the floodgates were now open. Almost immediately the Ott-Antonsen ansatz was used to solve many longstanding problems about the Kuramoto model and its variants, as well as to generate and solve many new prob-

a)Electronic mail: ml2437@cornell.edu
b)Electronic mail: mirollo@bc.edu
c)Electronic mail: shs7@cornell.edu

lems.

Still, a lot of old questions hung in the air. Both the Watanabe-Strogatz transformation and the Ott-Antonsen ansatz appeared somewhat unmotivated and almost miraculous. Where did they come from, really, and why did they work? It also wasn't clear whether they were connected, or perhaps even equivalent; the Watanabe-Strogatz transformation could be used on any system of globally, sinusoidally coupled, phase oscillators for all finite $N \geq 3$, but seemed restricted to identical oscillators, whereas the Ott-Antonsen ansatz allowed for non-identical oscillators but seemed restricted to the continuum limit of infinite N. Also, why were ideas like linear fractional transformations and Poisson densities – tools from other branches of mathematics – popping up in these questions about dynamical systems?

Later work made sense of all of this. The Josephson arrays and the Kuramoto model both turned out to have deep mathematical ties to group theory, hyperbolic geometry, and projective geometry, and both the Watanabe-Strogatz transformation and the Ott-Antonsen ansatz were tapping into these structures. For the Josephson arrays, the governing equations turned out to be generated by a group action, specifically the action of the Möbius group of linear fractional transformations of the unit disk to itself. Seen in this light, the constants of motion for the Josephson arrays were cross-ratios, and the invariant tori were group orbits. The same group-theoretic structure was found to underlie the Kuramoto model (in the special case where all the oscillator frequencies are identical) as well as other sinusoidally coupled systems of identical phase oscillators.

In the past few years, several researchers wondered how far this story could be pushed. Are there higher-dimensional or quantum extensions of the Kuramoto model that might show similar reducibility? A number of results along these lines have now been found. In particular, several researchers have explored a generalization of the Kuramoto model in which the oscillators move on spheres instead of the unit circle. These spheres could be either the ordinary two-dimensional sphere or higher-dimensional spheres. A counterpart of the Ott Antonsen ansatz has been discovered for the continuum version of the Kuramoto model on the d-dimensional sphere and used to reduce its infinite-dimensional dynamics to a lower dimensional set of ODEs. But as before, some of the results appear disconnected and a bit miraculous.

Our goal in this paper is to show that hyperbolic geometry and group theory can unify and clarify our understanding of the Kuramoto model on spheres, and make all the latest results seem natural, just as they did before for the traditional Kuramoto model. Our approach explains the new model's reducibility for any finite number of oscillators, as well as for the continuum limit, and it reveals why Poisson densities arise again in this setting. There is a close connection to Laplace's equation and harmonic analysis, as we'll see in Section V below. We also find that complex analysis is not really essential – which is just as well, since it does not generalize to the higher-dimensional spheres being considered here. Instead, the proper mathematical setting is harmonic analysis and hyperbolic geometry on higher -dimensional balls. Our work also allows us to do more than merely unify existing results. For instance, by establishing that that systems of identical Kuramoto oscillators on spheres have a "hyperbolic gradient" structure, we can prove new global stability results about convergence to the synchronized state, as described in Section VII.

I. PRELIMINARIES

A. The Kuramoto Model

A natural generalization of Kuramoto oscillator networks to higher dimensions is an "oscillator" system governed by equations of the form

$$\dot{x}_i = A_i x_i + Z - \langle Z, x_i \rangle x_i, \quad i = 1, \dots, N, \tag{1}$$

where x_i is a point on the unit sphere $S^{d-1} \subset \mathbb{R}^d$, each A_i is an antisymmetric $d \times d$ matrix, and $Z \in \mathbb{R}^d$ such that A_i and Z are functions of the population $(x_1, \dots, x_N)$. Note Z does not depend on i. A straightforward computation shows $\langle x_i, \dot{x}_i \rangle = 0$, which proves each higher oscillator is indeed confined to a sphere. The state space for this system is the N-fold product $X = (S^{d-1})^N$, which has dimension $N(d-1)$. (Later we will also consider the natural infinite-N analogue of (1), where a state is a probability measure on S^{d-1}). When $d = 2$ these are just Kuramoto networks given by equivalent equations

$$\dot{\theta}_i = \omega_i + B \cos \theta_i + C \sin \theta_i, \quad i = 1, \dots, N.$$

We think of the function Z as the system order parameter; it can be any smooth function on the space X, though in examples we usually restrict to fairly simple functions, like a linear combination of the form

$$Z = \sum_{i=1}^{N} a_i x_i.$$

B. Hyperbolic geometry and Möbius transformations

In this paper, a "Möbius transformation" is a composition of Euclidean isometries and spherical inversions of $\mathbb{R}^d$ mapping the unit ball homeomorphically to itself and preserving orientation. This is a more restrictive definition than the commonly defined Möbius transformations which in general do not need to preserve the unit ball.

As in the case $d = 2$, flows of the form (1) are intimately related to the natural hyperbolic geometry on the unit ball B^d with boundary S^{d-1}. This geometry has metric

$$ds = \frac{2|dx|}{1 - |x|^2},$$

where $|dx|$ is the ordinary Euclidean metric. Isometries are assumed to be with respect to this geometry, unless otherwise qualified as Euclidean. The metric ds has constant (sectional) curvature -1, and we can describe its isometries, which generalize the Möbius transformations preserving the unit disc for $d=2$. For $d=2$, let $w \in B^2$ and consider the Möbius transformation

$$M_w(x) = \frac{x-w}{1-\overline{w}x};$$

M_w preserves the unit disc B^2 and its boundary S^1. To generalize this to higher dimensions, we need to express $M_w(x)$ without reference to complex arithmetic operations or conjugation. Using the identity $2\langle w,x\rangle = \overline{w}x + w\overline{x}$, we see that

$$\begin{aligned}\frac{(x-w)(1-w\overline{x})}{(1-\overline{w}x)(1-w\overline{x})} &= \frac{x-w-w|x|^2+w^2\overline{x}}{1-2\langle w,x\rangle+|w|^2|x|^2}\\ &= \frac{x-w-w|x|^2+w(2\langle w,x\rangle-\overline{w}x)}{1-2\langle w,x\rangle+|w|^2|x|^2}\\ &= \frac{(1-|w|^2)x-(1-2\langle w,x\rangle+|x|^2)w}{1-2\langle w,x\rangle+|w|^2|x|^2}.\end{aligned}$$

This form of M_w generalizes to higher dimensions: Let $w \in B^d$ and define

$$\begin{aligned}M_w(x) &= \frac{(1-|w|^2)x-(1-2\langle w,x\rangle+|x|^2)w}{1-2\langle w,x\rangle+|w|^2|x|^2}\\ &= \frac{(1-|w|^2)(x-|x|^2w)}{1-2\langle w,x\rangle+|w|^2|x|^2}-w,\end{aligned}$$

where $x \in B^d$ or S^{d-1}. We call M_w a boost transformation. If $|x|=1$ this formula simplifies to

$$M_w(x) = \frac{(1-|w|^2)(x-w)}{|x-w|^2}-w.$$

Now see that

$$\begin{aligned}M_w(w) &= \frac{(1-|w|^2)w-(1-2\langle w,w\rangle+|w|^2)w}{1-2\langle w,w\rangle+|w|^2|w|^2}\\ &= \frac{(1-|w|^2)w-(1-|w|^2)w}{1-2\langle w,w\rangle+|w|^2|w|^2} = 0.\end{aligned}$$

Alternatively, we can use the second form to show

$$M_w(w) = \frac{(1-|w|^2)(w-|w|^2w)}{1-2\langle w,w\rangle+|w|^2|w|^2}-w = \frac{(1-|w|^2)^2w}{(1-|w|^2)^2}-w = 0.$$

Similar computations show that M_0 is the identity, $M_w^{-1} = M_{-w}$, and $M_w(0) = -w$.

Any orientation-preserving isometry of B^d can be expressed uniquely in the form

$$g(x) = \zeta M_w(x)$$

and also uniquely in the form

$$g(x) = M_{-z}(\xi x),$$

for some $w,z \in B^d$ and $\zeta,\, \xi \in SO(d)$, the group of orientation-preserving orthogonal linear transformations on $\mathbb{R}^d$. Computing the linearization with these two formulas at $x=0$ (a computation involving matrix calculus identities, which we leave to the reader) gives

$$g(x) \approx \zeta(-w+(1-|w|^2)x) \approx z+(1-|z|^2)\xi x,$$

which implies $z=-\zeta w$ (hence $|z|=|w|$) and $\xi=\zeta$.

The group G of all such isometries is isomorphic to the linear group $SO^+(d,1)$, which has dimension $d(d+1)/2$. The infinitesimal transformations are given by flows on B^d of the form

$$\dot{y} = Ay-\langle Z,y\rangle y+\frac{1}{2}(1+|y|^2)Z, \tag{2}$$

with A antisymmetric $d\times d$ and $Z\in\mathbb{R}^d$. Note that this flow extends to a flow on S^{d-1} of the form in (1). To derive these infinitesimal transformations, we can work separately with the boost and rotation components. Replace w by tw, expand to first order in t to see

$$M_{tw}(x) \approx \frac{x-|x|^2tw}{1-2t\langle w,x\rangle}-tw \approx x+t\left(2\langle w,x\rangle x-(1+|x|^2)w\right).$$

Now take the linear t coefficient to deduce the infinitesimal generator is an "infinitesimal boost" of the form (2) with $Z=-2w$ and $A=0$. The infinitesimal generators corresponding to the rotation components are flows of the form $\dot{x}=Ax$ with A antisymmetric; together with the infinitesimal boosts we get all flows of the form (2). The group G acts on the space X in the natural way (component by component) and the infinitesimal generators of this group action on X are flows of the form (1) with all A_i identical. Therefore the evolution of any initial point $p\in X$ under the system (1) with all $A_i=A$ lies in the group orbit Gp.

II. REDUCED EQUATIONS

The given Kuramoto system has Nd degrees of freedom, for some large N. However, since the flow of the system is determined via an action of the $d(d+1)/2$ dimensional Lie group G, we can alternatively study the auxiliary dynamical system on G, which we call the reduced equations. By ignoring rotations, we can further restrict our attention to a system on the d-dimensional quotient $G/SO(d)\cong B^d$. The dimensional reduction not only makes the reduced equations easier to analyze than the original Kuramoto system, but the reduced equations require fewer computational resources to numerically integrate.

Now suppose all the terms A_i in (1) are equal. Fix a base point $p=(p_1,\ldots,p_N)\in X$. Then if the points p_i are in sufficiently general position, every element in the G-orbit of p can be expressed uniquely as gp for some $g\in G$, with parameters w,z and ζ. We wish to derive the corresponding evolution equations for w,z and ζ. Let $(x_i(t))$ be any solution to (1) in the group orbit Gp; we do not require that the initial point

$(x_i(0)) = p$. Then $(x_i(t)) = g_t p$ for a unique $g_t \in G$, which determines the parameters w, z, ζ as functions of t. Now consider the equation (2), with coefficients A and Z evaluated at $(x_i(t))$. This is a non-autonomous ODE on $\overline{B^d}$, and its time-t flow must be given by some $\tilde{g}_t \in G$. This ODE has solutions $(x_i(t)) = g_t p = g_t g_0^{-1}(x_i(0))$, which implies that $\tilde{g}_t = g_t g_0^{-1}$.

So for any $y_0 \in \overline{B^d}$,

$$y(t) = g_t g_0^{-1}(g_0(y_0)) = g_t(y_0) = \zeta M_w(y_0) = M_{-z}(\zeta y_0)$$

must satisfy the ODE (2) with A and Z evaluated at $(x_i(t))$ at time t. In particular, if we let $y_0 = 0$, then $y(t) = -\zeta w = z$, so z satisfies the ODE (2).

Now expand $y = \zeta M_w(y_0) = M_{-z}(\zeta y_0)$ to first order in y_0, using the variables z and ζ:

$$y \approx z + (1-|z|^2)\zeta y_0,$$

so

$$\dot{y} \approx \dot{z} - 2\langle \dot{z}, z\rangle \zeta y_0 + (1-|z|^2)\dot{\zeta} y_0.$$

On the other hand, (2) gives

$$\begin{aligned}\dot{y} &= Ay + \frac{1}{2}\left(1+|y|^2\right)Z - \langle Z, y\rangle y \\ &\approx Az + \frac{1}{2}(1+|z|^2)Z - \langle Z, z\rangle z \\ &+ (1-|z|^2)\Big(A\zeta y_0 + \langle z, \zeta y_0\rangle Z - \langle Z, z\rangle \zeta y_0 - \langle Z, \zeta y_0\rangle z\Big).\end{aligned}$$

Setting $y_0 = 0$ gives the $\dot{z}$ equation

$$\dot{z} = Az + \frac{1}{2}(1+|z|^2)Z - \langle Z, z\rangle z \tag{3}$$

as expected, and this in turn implies that

$$\langle \dot{z}, z\rangle = \frac{1}{2}(1-|z|^2)\langle Z, z\rangle$$

(use $\langle Az, z\rangle = 0$). Equating the y_0 terms, factoring out $1-|z|^2$ and canceling the common term $\langle Z, z\rangle \zeta y_0$ gives

$$\dot{\zeta} y_0 = A\zeta x_0 + \langle z, \zeta y_0\rangle Z - \langle Z, \zeta y_0\rangle z.$$

Together, the last two terms above define a special type of antisymmetric operator of ζy_0: Given any $y_1, y_2 \in \mathbb{R}^d$, define the antisymmetric operator α as

$$\alpha(y_1, y_2)y = \langle y_1, y\rangle y_2 - \langle y_2, y\rangle y_1;$$

this operator has range = span(y_1, y_2) providing y_1 and y_2 are linearly independent; otherwise $\alpha(y_1, y_2) = 0$. Then for all $y_0 \in \mathbb{R}^d$,

$$\dot{\zeta} y_0 = A\zeta y_0 + \alpha(z, Z)\zeta y_0$$

and therefore

$$\dot{\zeta} = (A + \alpha(z, Z))\zeta.$$

Differentiating $z = -\zeta w$ gives

$$Az + \frac{1}{2}(1+|z|^2)Z - \langle Z, z\rangle z = -\zeta \dot{w} - \dot{\zeta} w$$

so

$$\begin{aligned}\zeta \dot{w} &= (A + \alpha(z, Z))z - Az - \frac{1}{2}(1+|z|^2)Z + \langle Z, z\rangle z \\ &= Az + |z|^2 Z - \langle Z, z\rangle z - Az - \frac{1}{2}(1+|z|^2)Z + \langle Z, z\rangle z \\ &= -\frac{1}{2}(1-|z|^2)Z;\end{aligned}$$

hence

$$\dot{w} = -\frac{1}{2}(1-|w|^2)\zeta^{-1}Z.$$

Summing up, the evolution equations for the (z, ζ) coordinate system on Gp are

$$\dot{z} = Az + \frac{1}{2}(1+|z|^2)Z - \langle Z, z\rangle z \tag{4a}$$

$$(1+|z|^2)Z - \langle Z, z\rangle z\dot{\zeta} = (A + \alpha(z, Z))\zeta \tag{4b}$$

with Z evaluated at $M_{-z}(\zeta p)$, and for the (w, ζ) coordinate system on Gp are

$$\dot{w} = -\frac{1}{2}(1-|w|^2)\zeta^{-1}Z \tag{5a}$$

$$\dot{\zeta} = (A - \alpha(\zeta w, Z))\zeta, \tag{5b}$$

with Z evaluated at $\zeta M_w(p)$. Note that these equations generalize the evolution equations for the parameters w and ζ given in Chen et. al.[8] for the classic case $d = 2$.

III. COMPARISON OF Z VERSUS W COORDINATES

The $\dot{z}$ equation is an extension of the system equation on S^{d-1}. However, for finite N, the $\dot{z}$ equation does not uncouple from ζ, since Z is evaluated at $M_{-z}(\zeta p)$. The exception to this is in the infinite-N limit: if the base point p is now the uniform density on S^{d-1}, then $\zeta p = p$ (the uniform density is invariant under rotations) and the density $M_{-z}(p)$ is a hyperbolic Poisson density on S^{d-1} whose centroid a function of z. In the case $d = 2$, this Poisson density has centroid z. Unfortunately this is false for $d \geq 3$ (we will give more details on this in the next section).

The advantage of the $\dot{w}$ equation is that for an order parameter function of the form

$$Z = \sum_{i=1}^{N} a_i x_i,$$

with $a_i \in \mathbb{R}$, ζ drops out of the $\dot{w}$ equation and we get the reduced equation

$$\dot{w} = -\frac{1}{2}(1-|w|^2)Z(M_w(p)).$$

The parameter w essentially defines the "phase relations" among the x_i; two configurations have the same w if and only if they are related by a rotation. So w is the key parameter that determines whether the system is approaching synchrony or incoherence.

The w variable also has a nice invariance under change of base points. Suppose $p' = M(p) \in Gp$; then we have coordinates w', ζ' associated to the base point p'. Any $q \in Gp$ has two expressions

$$q = \zeta M_w(p) = \zeta' M_{w'} p' = \zeta' M_{w'}(M(p)).$$

Assuming the coordinates of p are in sufficiently general position, this implies $\zeta M_w = \zeta' M_{w'} \circ M$, and hence

$$0 = \zeta M_w(w) = \zeta' M_{w'}(M(w)).$$

But the unique solution to $M_{w'}(y) = 0$ is w', and hence $w' = M(w)$. In other words, the coordinates w and w' transform exactly as the base points p and p'.

IV. CONTINUUM LIMIT

Next, we consider the dynamics of the network (1) in the limit $N \to \infty$. Let us assume first that the rotation terms $A_i = A$ are constant across the population; later we will consider the case where A varies depending on some distribution. Let us also assume that

$$Z = \frac{K}{N}\sum_{i=1}^{N} x_i$$

is proportional to the centroid of the population. In the continuum limit, a state of the system is a probability measure ρ on S^{d-1}, and

$$Z = K\int_{S^{d-1}} x\, d\rho(x).$$

The measure ρ evolves according to the continuity equation (AKA noiseless Fokker-Planck) associated to the flow in (1). Naturally, this flow must preserve group orbits under the action of G. Recall that if $M \in G$, then the measure $M_*\rho$ is defined by the adjunction formula

$$\int_{S^{d-1}} f(x)\, d(M_*\rho)(x) = \int_{S^{d-1}} f(M(x))\, d\rho(x).$$

In particular, we can consider the G-orbit of the uniform probability measure σ on S^{d-1}. This orbit is special; whereas a typical group orbit $G\rho$ has dimension equal to the dimension of G, namely $d(d+1)/2$, the orbit $G\sigma$ has dimension only d. This is because the stabilizer of σ is $SO(d)$; any rotation fixes σ, whereas the boosts deform σ. Hence the orbit $G\sigma$ has dimension d. Any element in $G\sigma$ can be written as $(M_{-z})_*\sigma$, with $z \in B^d$. The evolution equation for z is (3), with

$$Z(z) = K\int_{S^{d-1}} x\, d(M_{-z})_*\sigma(x) = K\int_{S^{d-1}} M_{-z}(x)\, d\sigma(x). \quad (6)$$

In the case $d = 2$ with $x = \zeta \in S^1$, we have

$$d\sigma(\zeta) = \frac{1}{2\pi i}\frac{d\zeta}{\zeta},$$

so the integral

$$Z(z) = \frac{K}{2\pi i}\int_{S^1} \frac{\zeta + z}{1+\bar{z}\zeta}\cdot\frac{d\zeta}{\zeta} = K\frac{\zeta+z}{1+\bar{z}\zeta}\bigg|_{\zeta=0} = Kz$$

by the Cauchy integral formula. Therefore (3) simplifies to the equation

$$\dot{z} = i\omega z + \frac{K}{2}(1-|z|^2)z$$

when $d = 2$. Unfortunately, the formula $Z(z) = Kz$ is not correct for $d \geq 3$; though as we shall see later, this formula *is* correct in higher dimensions for the complex hyperbolic model in even dimensions, which we discuss in the next section. For $d = 2$ the two geometries agree, which explains the coincidence for $d = 2$.

Any Riemannian manifold X has a Laplace-Beltrami operator Δ associated to its metric; functions f on X satisfying the equation $\Delta f = 0$ are called harmonic. For functions on the ball B^d with the hyperbolic metric, the operator is

$$\Delta_{hyp} = (1-|x|^2)^2\Delta_{euc} + 2(d-2)(1-|x|^2)\sum_{i=1}^{d} x_i\frac{\partial}{\partial x_i},$$

where

$$\Delta_{euc} = \sum_{i=1}^{d}\frac{\partial^2}{\partial x_i^2}$$

is the standard Laplace operator (see Stoll[9], Chapter 3). We will call solutions to the equation $\Delta_{hyp} f = 0$ *hyperbolic harmonic functions*; for $d = 2$ these coincide with ordinary (Euclidean) harmonic functions. We can consider the analogue of the Dirichlet problem: given a continuous function f on S^{d-1}, extend f to a hyperbolic harmonic function $\tilde{f}$ on B^d. Assuming this problem has a unique solution, then for any rotation $\zeta \in SO(d)$ we must have $\widetilde{f\circ\zeta} = \tilde{f}\circ\zeta$, since rotations preserve the hyperbolic metric. If we average $f\circ\zeta$ on S^{d-1} over all rotations $\zeta \in SO(d)$ we get the constant function

$$f_{ave} = \int_{S^{d-1}} f(x)\, d\sigma(x)$$

on S^{d-1}, and any constant is hyperbolic harmonic on B^d. Therefore the average on B^d of $\widetilde{f\circ\zeta} = \tilde{f}\circ\zeta$ over all $\zeta \in SO(d)$ must be the constant f_{ave}. But $\tilde{f}(\zeta(0)) = \tilde{f}(0)$ for all ζ, so we must have

$$\tilde{f}(0) = \int_{S^{d-1}} f(x)\, d\sigma(x).$$

Now let $z \in B^d$; since M_{-z} preserves the hyperbolic metric, we must have $\widetilde{f \circ M_{-z}} = \tilde{f} \circ M_{-z}$, which implies

$$\begin{aligned}\tilde{f}(z) &= \widetilde{f \circ M_{-z}}(0) \\ &= \int_{S^{d-1}} f(M_{-z}(x))\, d\sigma(x) \\ &= \int_{S^{d-1}} f(x)\, d(M_{-z})_* \sigma(x).\end{aligned}$$

As shown in Chapter 5 in Stoll, we can express the measure

$$d(M_{-z})_* \sigma(x) = P_{hyp}(z,x)\, d\sigma(x),$$

where the hyperbolic Poisson kernel function is

$$P_{hyp}(z,x) = \left(\frac{1-|z|^2}{|z-x|^2} \right)^{d-1}. \tag{7}$$

Thus the solution to the hyperbolic Dirichlet problem with boundary function f on S^{d-1} is given by the hyperbolic Poisson integral

$$\tilde{f}(z) = \int_{S^{d-1}} P_{hyp}(z,x) f(x)\, d\sigma(x), \quad z \in B^d.$$

The orbit $G\sigma$ consists of all *hyperbolic* Poisson measures $P(z,x)\, d\sigma(x)$, parametrized by $z \in B^d$. By contrast, the Euclidean Poisson kernel function is

$$P_{euc}(z,x) = \frac{1-|z|^2}{|z-x|^d},$$

so the hyperbolic Poisson measures agree with the Euclidean Poisson measures only if $d = 2$.

Now we can calculate the expression $Z(z)$ in the general case $d \geq 2$. We see from (6) that $Z(z)$ is the hyperbolic Poisson integral of the function Kx on S^{d-1}. The function Kx is (Euclidean) harmonic and homogeneous of degree 1 on $\mathbb{R}^d$; following the recipe in Chapter 5 in Stoll, we see that its extension from S^{d-1} to a hyperbolic harmonic function on B^d is given by

$$Z(z) = K \frac{F(1, 1-d/2; 1+d/2; |z|^2)}{F(1, 1-d/2; 1+d/2; 1)} z, \tag{8}$$

where F is the hypergeometric function

$$F(a,b;c;t) = \sum_{k=0}^{\infty} \frac{(a)_k (b)_k}{(c)_k} \frac{t^k}{k!},$$

with $(a)_0 = 1$ and $(a)_k = a(a+1)\cdots(a+k-1)$ for $k \geq 1$. Notice that if a or $b = 0$, then $F(a,b;c;t) = 1$; this gives $Z(z) = Kz$ for $d = 2$, as expected.

V. COMPLEX CASE

There is an alternative generalization of Kuramoto networks to higher-dimensional oscillators when $d = 2m$ is even. Then $\mathbb{R}^d = \mathbb{C}^m$, and we can study systems of the form

$$\dot{x}_j = A_j x + Z - \langle x_j, Z \rangle x_j, \quad i = 1, \ldots, N, \tag{9}$$

where now x_i is a point on the unit sphere $S^{2m-1} \subset \mathbb{C}^m$, A_i is an anti-Hermitian $m \times m$ complex matrix, $Z \in \mathbb{C}^m$ and $\langle , \rangle$ denotes the complex-valued Hermitian inner product. These systems are the same as the real case when $d = 2, m = 1$ but are different for $m \geq 2$. To see this, suppose

$$Ax + Y - \langle x, Y \rangle_{\mathbb{R}} x = Bx + Z - \langle x, Z \rangle_{\mathbb{C}} x$$

for all $x \in S^{2m-1} \subset \mathbb{C}^m = \mathbb{R}^d$, where A is antisymmetric, B is anti-Hermitian, $Y, Z \in \mathbb{C}^m$ and we use the subscripts $\mathbb{R}$ and $\mathbb{C}$ to distinguish the real and complex inner products. Then

$$(A-B)x = Z - Y + \Big(\langle x, Y \rangle_{\mathbb{R}} - \langle x, Z \rangle_{\mathbb{C}} \Big) x$$

and so $(A-B)(-x) = (A-B)x$ for all $x \in S^{2m-1}$, which implies $A = B$. This implies

$$Y - Z = \Big(\langle x, Y \rangle_{\mathbb{R}} - \langle x, Z \rangle_{\mathbb{C}} \Big) x$$

for all $x \in S^{2m-1}$, hence $Y - Z \in \mathrm{span}_{\mathbb{C}}(x)$ for all $x \in \mathbb{C}^m$; if $m \geq 2$, this implies $Y = Z$. But then we have

$$\langle x, Y \rangle_{\mathbb{R}} = \langle x, Y \rangle_{\mathbb{C}}$$

for all $x \in \mathbb{C}^m$, which can only hold if $Y = 0$. Hence for $m \geq 2$, the only flows simultaneously of the form (1) and (9) have $Z = 0$ and A anti-Hermitian.

Flows of the form (9) are related to the complex hyperbolic geometry on the complex unit ball B^m with the Bergman metric (see Rudin[7], Chapter 1). The orientation-preserving isometries of this metric are generated by unitary transformations $\zeta \in U(m)$ and "boost" transformations of the form

$$\begin{aligned}M_w(x) &= \frac{\sqrt{1-|w|^2}\, x + \left(\frac{\langle x,w \rangle}{1+\sqrt{1-|w|^2}} - 1 \right) w}{1 - \langle x, w \rangle} \\ &= \frac{x - w + \frac{\langle x,w \rangle w - |w|^2 x}{1+\sqrt{1-|w|^2}}}{1 - \langle x, w \rangle}.\end{aligned}$$

Notice that when $m = 1$, this reduces to the standard complex Möbius map M_w. As in the real case M_0 is the identity, $M_w^{-1} = M_{-w}$, $M_w(w) = 0$ and $M_w(0) = -w$. Any orientation-preserving isometry of B^d can be expressed uniquely in the form

$$g(x) = \zeta M_w(x) = M_{-z}(\xi x),$$

where $w, z \in B^m$ but now $\zeta, \xi \in U(m)$, the complex unitary group. Linearizing at $x = 0$ gives

$$\begin{aligned}g(x) &\approx \zeta \left(-w - \langle x, w \rangle w + \sqrt{1-|w|^2}\, x + \frac{\langle x, w \rangle w}{1 + \sqrt{1-|w|^2}} \right) \\ &\approx \zeta \left(-w + \sqrt{1-|w|^2}\, x - \frac{\sqrt{1-|w|^2} \langle x, w \rangle w}{1 + \sqrt{1-|w|^2}} \right) \\ &\approx z + \sqrt{1-|z|^2}\, \xi x - \frac{\sqrt{1-|z|^2} \langle \xi x, z \rangle z}{1 + \sqrt{1-|z|^2}}\end{aligned}$$

FIG. 1. A first order linear Kuramoto system on S^2 with equal weights, and randomly chosen initial conditions. The states shown are at $t=0, t=10$, and $t=40$ respectively. This simulation was written in Python and visualized with Plotly.

which implies $z=-\zeta w$ (hence $|z|=|w|$) and $\xi=\zeta$, as before. The corresponding infinitesimal transformations are given by flows on the complex unit ball B^m of the form

$$\dot{y}=Ay+Z-\langle y,Z\rangle y, \tag{10}$$

with A anti-Hermitian $m\times m$ and $Z\in\mathbb{C}^m$. This flow extends to a flow on S^{2m-1} of the form in (9). Note the absence of the quadratic term $|y|^2Z$ here. To derive these infinitesimal transformations, we again work separately with the boost and rotation components. Replace w by tw and expand to first order in t, using that to first order in t, $(1-|w|^2)^{1/2}=1$:

$$M_{tw}(x)\approx\frac{x-tw}{1-t\langle x,w\rangle}\approx x+t\left(\langle x,w\rangle x-w\right),$$

so the infinitesimal generator is an "infinitesimal boost" of the form (10) with $Z=-w$ and $A=0$. The infinitesimal generators corresponding to the rotation components are flows of the form $\dot{x}=Ax$ with A anti-Hermitian; together with the infinitesimal boosts we get all flows of the form (10).

VI. RELATIONS TO WORKS OF TANAKA, LOHE, CHANDRA-GIRVAN-OTT

Many of the results above can be found in some form in the papers of these authors: Tanaka is studying the same system as (1) (see his equation (9)). He writes his Möbius transformations differently, but he's using the same group of transformations as we use above, and he gets reduced equations for his Möbius parameters. Tanaka's equation (10b) looks similar to the $\dot{z}$ equation, except without the $|z|^2$ term, which is puzzling. He doesn't mention the reduction down to dimension d in the finite-N case that we get with the $\dot{w}$ equation. He also understands that the complex case when $d=2m$ is different, and generalizes the OA residue calculation to this case, which is the highlight of his paper. In the real case, Tanaka's equation (15) is similar to our equation (8), though we were not able to show that the two expressions are equivalent.

Lohe is also looking at the same system as (1) (see his equation (22)). His transformation (30) on S^{d-1} is our M_w (with $v=w$) and his equation (31) is the same as our $\dot{z}$ equation. He also has something that looks like the $\dot{w}$ equation (42), which he says is independent of the rest of the reduced system for (in our notation) an order parameter function of the form

$$Z=\frac{1}{N}\sum_{i=1}^{N}\lambda_iQ_ix_i,$$

where $Q_i\in O(d)$ and $\lambda_i\in\mathbb{R}$. But such a Z does not satisfy the identity $\zeta Z(p)=Z(\zeta p)$ for all rotations ζ, unless $Q_i=\pm I$, so we don't see how the ζ term cancels.

Lohe's map M in equation (55) (ignoring the R factor) agrees with our map M_{-v} on the sphere S^{d-1}, but not on the ball B^d. So it's not a Möbius transformation of the type we're using. For example, $M(-v)=v$ whereas $M_{-v}(-v)=0$. We're not sure why Lohe prefers these maps over the boosts; he claims that M preserves cross-ratios, but we don't see why this is advantageous. His map F in equation (63) (again ignoring the R factor) is exactly our M_{-v}.

Chandra, Girvan and Ott[4] proceed directly to the infinite-N version of (1). They make a very clever guess (their equation (7)) of the form of the special densities that generalize the Poisson densities for $d=2$, and then calculate the exponent in the denominator of their expression, getting exactly the hyperbolic Poisson kernel densities in (7) above. Their equation (15) is exactly the same as our $\dot{z}$ equation (4) in the infinite-N limit. The integral in their equation (19) can be evaluated, as shown above in (8).

VII. AN EXAMPLE: FIRST-ORDER LINEAR ORDER PARAMETER GIVES GRADIENT SYSTEM

We conclude with an analysis of the system (1) with order parameter function

$$Z=\sum_{i=1}^{N}a_ix_i, \tag{11}$$

where the a_i are real constants.

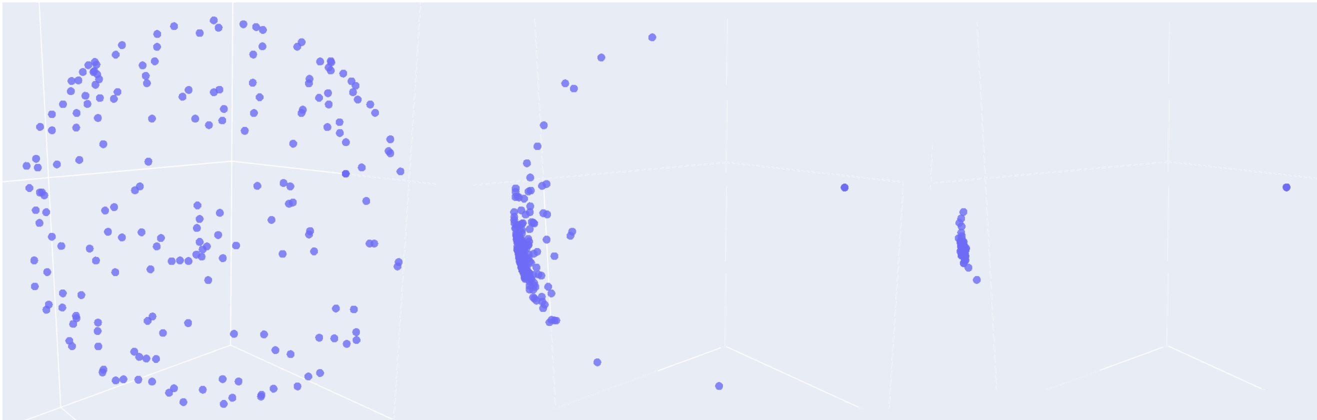

FIG. 2. A first order linear Kuramoto system on S^2 with a majority cluster, where one body is chosen to have a weight which exceeds the combined weights of all other bodies. The states shown are at $t = 0, t = -10$, and $t = -40$ respectively. In this simulation, one body was chosen to have a weight of 0.6, and the remaining 400 bodies were chosen to have equal weights of 1/1000.

A. A Computer Visualization

We implemented the Runge-Kutta algorithm for solving differential equations in the case where $d = 2$, $N = 1000$, $A = 0$, and $a_i = \frac{1}{N}$. See Fig. 1. Setting $A = 0$ is tantamount to ignoring the rotational influence, or equivalently, rotating the frame of reference along with the system as it evolves. Randomly chosen points on the sphere were used as initial conditions. As time increases, one can see that the bodies coalesce to a limit point, mimicking the synchronization that is well observed in the $d = 1$ case. These bodies never actually collide, since this would violate uniqueness of solutions to differential equations with distinct initial conditions. Later in this section, we will prove this synchronization behavior is generic in the space of Kuramoto systems of this form, provided the weighting satisfies an upper bound.

When time runs backwards, almost all configurations of the bodies will tend towards an equally spaced position where their centroid is at the origin. The exception is when we have a majority cluster, where one body has a weight which exceeds the weight of all other bodies. When this occurs, it is impossible to arrange the bodies so their weighted centroid is at the origin, so the backwards time limit will tend towards an antipodal configuration, where all bodies not in the majority cluster will coalesce around the antipode of the cluster. See Fig. 2.

B. Existence of Hyperbolic Gradient

As mentioned above, the $\dot{w}$ equation in (5) reduces to

$$\dot{w} = -\frac{1}{2}(1-|w|^2)Z(M_w(p)), \tag{12}$$

independent of the parameter ζ. We will show that this is a gradient flow on the unit ball B^d with respect to the hyperbolic metric. In the presence of a Riemannian metric we can associate a 1-form to any vector field, and the vector field is gradient if and only if the associated 1-form is closed (which is equivalent to exactness because the unit ball is simply connected). For the Euclidean metric on B^d (or any open subset of $\mathbb{R}^d$) and standard coordinates $w_1, \dots, w_d$, the 1-form associated to the vector field with components $f_1, \dots, f_d$ is

$$\omega = f_1\, dw_1 + \cdots + f_d\, dw_d.$$

If we scale the Euclidean metric by a positive smooth function ϕ, then the associated 1-form for the metric $ds = \phi|dw|$ is now

$$\omega = \phi^2(f_1\, dw_1 + \cdots + f_d\, dw_d).$$

Now let's consider the vector field defined by (12). By linearity, it suffices to treat the case $Z = x_i$, and we can take $i = 1$ without loss of generality. We have $\phi(w) = 2(1-|w|^2)^{-1}$ for the hyperbolic metric, so to prove that the flow in (12) is gradient we must prove that the 1-form

$$\begin{aligned}\omega = {} & \frac{4}{(1-|w|^2)^2}\left(-\frac{1}{2}(1-|w|^2)\right) \\ & \cdot \sum_{j=1}^{d}\left(\frac{(1-|w|^2)(p_{1,j}-w_j)}{|p_1-w|^2} - w_j\right) dw_j \\ = {} & -2\sum_{j=1}^{d}\left(\frac{p_{1,j}-w_j}{|p_1-w|^2} - \frac{w_j}{1-|w|^2}\right) dw_j\end{aligned}$$

is closed, where $p_{1,j}$ denotes the jth component of the point $p_1 \in S^{d-1}$. Let E_j denote the coefficient of dw_j in parentheses above; then

$$d\omega = -2\sum_{j,k=1}^{d} \frac{\partial E_j}{\partial w_k}\, dw_k \wedge dw_j.$$

Applying the chain and quotient rules gives

$$\frac{\partial E_j}{\partial w_k} = \frac{2(p_{1,j}-w_j)(p_{1,k}-w_k)}{|p_1-w|^2} + \frac{2w_j w_k}{(1-|w|^2)^2}$$

for $j \neq k$, which is symmetric in j and k; hence the sum above for $d\omega$ simplifies to $d\omega = 0$. Thus ω is closed and we see that the flow (12) is gradient for any order parameter function of the form (11).

Next, we show that the hyperbolic potential to W, up to an additive constant, is $\Phi(w) = \frac{1}{4N}\sum_{l=1}^{N} a_l \log \frac{1-|w|^2}{|w-p_l|^2}$. Note that this potential doesn't involve Poisson kernels in higher dimensions because the exponent in the denominator is 2, not d. We will make several uses of the identity $\nabla_E |w|^2 = 2w$, for $w \in \mathbb{R}^d$. Observe that

$$\begin{aligned}\nabla_E \Phi(w) &= \frac{1}{4N}\sum_{l=1}^{N} a_l \frac{|w-p_l|^2}{1-|w|^2} \\ &\quad \cdot \frac{\left(-2|w-p_l|^2 w - 2(1-|w|^2)(w-p_l)\right)}{|w-p_l|^4} \\ &= \frac{1}{4N}\sum_{l=1}^{N} a_l \frac{-2|w-p_l|^2 w - 2(1-|w|^2)(w-p_l)}{(1-|w|^2)|w-p_l|^2} \\ &= \frac{1}{4N}\sum_{l=1}^{N} a_l \left(\frac{w}{1-|w|^2} + \frac{w-p_l}{|w-p_l|^2}\right).\end{aligned}$$

Hence, we can see that $\nabla_H \Phi = (1-|w|^2)^2 \nabla_E \Phi = W$, as desired. The interpretation that the potential is the hyperbolic logarithmic average of the bodies still holds in higher dimensions.

C. Analysis of Dynamics

We can use the existence of the potential $\Phi(w)$ for the flow on B^d to prove a global synchrony result for the system (1) when the coefficients a_i in the order parameter Z are all positive. Specifically, we assume that $0 < a_i < 1/2$ for all i, and $\sum_{i=1}^N a_i = 1$. We also assume $N \geq 3$. Under these conditions, almost all trajectories for (1) converge to the $(d-1)$-dimensional diagonal manifold $\Delta \subset X$ as $t \to \infty$. As $t \to -\infty$, almost all trajectories for (1) converge to the codimension-d subspace $\Sigma \subset X$ consisting of states with $Z(p) = 0$.

The proof is modeled after Theorem 1 in Tanaka, and will be based on two preliminary lemmas. In each of these lemmas we assume the conditions on the a_i above, and that the base point $p = (p_i)$ for the flow (12) has all distinct coordinates.

We begin with a general observation about gradient flows in the ball B^d: if $w_0 \in B^d$ is any initial condition and $w^* \in B^d$ is in the forward limit set $\Omega_+(w_0)$, then w^* is a fixed point for the flow. To see this, let Φ be a potential for the flow, and suppose $w(t_n) \to w^* \in B^d$ for some sequence $t_n \to \infty$. Since the potential decreases along trajectories,

$$\lim_{t\to\infty} \Phi(w(t)) = \lim_{n\to\infty} \Phi(w(t_n)) = \Phi(w^*).$$

Let F_t denote the time-t flow map. If w^* is not a fixed point, then for any $s > 0$,

$$\begin{aligned}\lim_{t\to\infty} \Phi(w(t)) &= \lim_{n\to\infty} \Phi(w(t_n+s)) \\ &= \lim_{n\to\infty} \Phi(F_s(w(t_n))) \\ &= \Phi(F_s(w^*)) \\ &< \Phi(w^*),\end{aligned}$$

which is a contradiction, so w^* must be a fixed point. (Compact limit sets are connected, so $\Omega_+(w_0)$ cannot consist of two or more but finitely many fixed points; however it is possible that forward or backward limits sets for gradient flows consist of a continuum of fixed points. We will see that this is not the case for our system on B^d.)

Lemma 1. *Any fixed point for the flow (12) in B^d is repelling.*

Proof. Suppose $w^* \in B^d$ is a fixed point for (12). As discussed above, an advantage of using the w-parameter is the equivariance with respect to change of base point p. Consequently we can assume $w^* = 0$ without loss of generality, so $Z(p) = \sum_{i=1}^{N} a_i p_i = 0$. To first order in w,

$$\begin{aligned}M_w(p_i) &= \frac{p_i - w}{1 - 2\langle w, p_i\rangle} - w \\ &= (p_i - w)\left(1 + 2\langle w, p_i\rangle\right) - w \\ &= p_i - 2w + 2\langle w, p_i\rangle p_i.\end{aligned}$$

The linearization of (12) at the fixed point $w^* = 0$ is

$$\begin{aligned}\dot{w} &= -\frac{1}{2}\sum_{i=1}^{N} a_i \Big(p_i - 2w + 2\langle w, p_i\rangle p_i\Big) \\ &= w - \sum_{i=1}^{N} a_i \langle w, p_i\rangle p_i.\end{aligned}$$

We claim that the linear map

$$Tw = \sum_{i=1}^{N} a_i \langle w, p_i\rangle p_i$$

has $||T|| < 1$; to see this, suppose $|w| = 1$. Then $|\langle w, p_i\rangle p_i| \leq 1$ and Tw is a convex combination of the vectors $\langle w, p_i\rangle p_i$. We can only obtain $|Tw| = 1$ if all terms $\langle w, p_i\rangle p_i = u$ with $|u| = 1$, which implies all $p_i = \pm u$, and this cannot happen if at least three of the p_i are distinct. Hence $||T|| < 1$ and so the eigenvalues μ_i of T satisfy $|\mu_i| < 1$. The eigenvalues for the $\dot{w}$ linearization are $\lambda_i = 1 - \mu_i$, so wee see that $\operatorname{Re}\lambda_i > 0$ for all i, establishing that the fixed point w^* is repelling. □

Lemma 2. $\lim_{|w|\to 1} \Phi(w) = -\infty$.

Proof. The potential function, up to a positive multiplicative constant, is given by

$$\Phi(w) = \sum_{i=1}^{N} a_i \log \frac{1-|w|^2}{|w-p_i|^2}.$$

To prove Lemma 2, it suffices to show that

$$\lim_{n\to\infty}\Phi(w_n) = -\infty$$

for any sequence $w_n \in B^d$ with $w_n \to x \in S^{d-1}$. The result is clear if $x \neq p_i$, since as $n \to \infty$ the terms $|w_n - p_i|$ will be bounded away from 0, and $1-|w_n|^2 \to 0$. So let's say that $w_n \to p_1$. We rewrite $\Phi(w_n)$ as

$$\begin{aligned}\Phi(w_n) &= \log(1-|w_n|^2) - 2a_1\log|w_n - p_1| \\ &\quad - 2\sum_{i=2}^{N} a_i \log|w_n - p_i| \\ &= \log(1-|w_n|) - 2a_1\log|w_n-p_1| \\ &\quad + \log(1+|w_n|) - 2\sum_{i=2}^{N} a_i\log|w_n - p_i|.\end{aligned}$$

The latter two terms above have finite limit as $n \to \infty$, so we focus on the first two terms. We have $1-|w_n| \le |p_1 - w_n|$, so

$$\log(1-|w_n|) - 2a_1\log|w_n - p_1| \le (1-2a_1)\log|w_n - p_1| \to -\infty$$

as $n \to \infty$, which proves our result. Notice that we need the assumption $a_i < 1/2$ for this argument. □

Theorem 3. *Under the conditions above, almost all trajectories for (1) converge to Δ as $t \to \infty$ and to Σ as $t \to -\infty$.*

Proof. Let $p = (p_1, \dots, p_N) \in X$ be any point with all distinct coordinates. The points on Gp are parametrized by $w \in B^d$ and $\zeta \in SO(d)$, and the dynamics for these parameters are given by (5). We begin with the dynamics as $t \to -\infty$. Let $w(t)$ be a trajectory for (12) with initial condition $w_0 \in B^d$, and consider the backward time limit set $\Omega_-(w_0)$; this is a nonempty, compact, connected subset of $\overline{B^d}$. The potential Φ is decreasing along all trajectories $w(t)$, hence bounded below as $t \to -\infty$, so Lemma 2 implies that the limit set $\Omega_-(w_0)$ must be contained in the interior B^d. We know that any $w^* \in \Omega_-(w_0)$ is a fixed point for the flow. By Lemma 1, w^* is repelling and so any trajectory $w(t)$ which comes sufficiently close to w^* must have $w(t) \to w^*$ as $t \to -\infty$; therefore $\Omega_-(w_0) = \{w^*\}$. This proves the existence of fixed points for (12), and that every trajectory $w(t)$ converges to a fixed point as $t \to -\infty$. If the flow had multiple fixed points, we would obtain a partition of B^d into the disjoint open basins of repulsion of the fixed points, violating connectedness of the ball. Therefore (12) has a unique fixed point w^*, and $w(t) \to w^*$ as $t \to -\infty$ for all trajectories. The fixed point w^* has $Z(M_{w^*}(p)) = 0$, so all trajectories in Gp converge to Σ as $t \to -\infty$.

In forward time, the limit set $\Omega_+(w_0)$ for any $w_0 \neq w^*$ must be completely contained in the boundary S^{d-1}, since the unique fixed point $w^* \in B^d$ is repelling. Suppose we remove the factor $(1/2)(1-|w|^2)$ in the flow (12); the scaled vector field on B^d given by

$$\dot{w} = -\sum_{i=1}^{N} a_i M_w(p_i) = w - \sum_{i=1}^{N} a_i \left(\frac{(1-|w|^2)(p_i - w)}{|p_i - w|^2}\right) \quad (13)$$

has the same trajectories as the original flow, just with different time parametrizations. Observe that this scaled vector field extends smoothly to $\mathbb{R}^d - \{p_i\}$, and reduces to the radial vector field x at any $x \in S^{d-1}$ with $x \neq p_i$. Therefore there is a unique trajectory passing through each point $x \in S^{d-1}$, flowing from the interior to the exterior of the sphere, as long as $x \neq p_i$. Consequently the original flow (12) has a unique trajectory $w(t)$ in B^d with $w(t) \to x$ as $t \to \infty$, as long as $x \neq p_i$. This also shows that there is a neighborhood U of $S^{d-1} - \{p_i\}$ such that if $w(t_0) \in U$ for some t_0, then $w(t) \to x \neq p_i$ for some $x \in S^{d-1}$. So if $\Omega_+(w_0)$ contains some $x \neq p_i$, then the trajectory $w(t)$ of w_0 must enter the neighborhood U, and therefore $w(t) \to x \in S^{d-1}$ as $t \to \infty$.

Since limit sets are connected, the only other possibility is $\Omega_+(w_0) = \{p_i\}$ for some i; equivalently, $w(t) \to p_i$. We will show that there is a unique trajectory with this behavior for each p_i. Assuming this, we see that with $N+1$ exceptions, any trajectory $w(t)$ converges to a point $x \in S^{d-1}$ with $x \neq p_i$ (the exceptions are the N trajectories converging to the base point coordinates p_i, and the fixed point trajectory w^*). The corresponding trajectory in Gp has coordinates

$$\zeta(t)M_{w(t)}(p_i) = \zeta(t)\left(\frac{(1-|w(t)|^2)(p_i - w(t))}{|p_i - w(t)|^2} - w(t)\right).$$

We have $|w(t)| \to 1$ and $|p_i - w(t)|$ is bounded away from 0 as $t \to \infty$, so $M_{w(t)}(p_i) \to -x$ and the trajectory $\zeta(t)M_{w(t)}(p)$ in Gp converges to Δ as $t \to \infty$.

This analysis breaks down at $x = p_i$ because the scaled vector field above does not have a unique limit as $w \to p_i$; rather, its limit depends on the direction of the approach. To see this, write $w = p_1 - ru$, where $0 < r < 1$ and $|u| = 1$ (with this convention, $u = p_1$ corresponds to w approaching p_1 radially). Then $|p_1 - w| = r$ and

$$|w|^2 = 1 - 2r\langle p_1, u\rangle + r^2$$

so

$$\frac{(1-|w|^2)(p_1 - w)}{|p_1 - w|^2} = \frac{(2r\langle p_1,u\rangle - r^2)ru}{r^2} = (2\langle p_1, u\rangle - r)\,u.$$

As $r \to 0$, the magnitude of this term is $2\langle p_1, u\rangle$, which depends on the angle of approach given by u (which is positive because $-u$ points inwards at p_1).

To complete the proof, we will examine the scaled system (13) using the polar representation (r,u), and show that the polar system has the unique fixed point $r^* = 0, u^* = p_1$, which has a unique attracting trajectory because it's a saddle with a $(d-1)$-dimensional unstable manifold.

We see that the scaled system has

$$\begin{aligned}\dot{w} &= p_1 - ru - a_1\left(2\langle p_1, u\rangle - r\right)u + O(r) \\ &= p_1 - 2a_1\langle p_1, u\rangle u + O(r),\end{aligned}$$

where the $O(r)$ term is a smooth function of r and u for $|r| < \varepsilon = \min|p_i - p_1|$, $i \ge 2$. This condition insures that $|p_i - w| \ge |p_i - p_1| - |r| > 0$, so the $i \ge 2$ terms in the scaled $\dot{w}$ equation are all smooth functions of r and u. And we can allow $r < 0$

here, even though it's not relevant to the $\dot{w}$ system. Now $r^2 = |w-p_1|^2$, so

$$r\dot{r} = \langle w-p_1, \dot{w}\rangle = -r\langle u, \dot{w}\rangle,$$

which gives

$$\dot{r} = -(1-2a_1)\langle p_1, u\rangle + O(r).$$

Differentiating $ru = p_1 - w$ gives

$$\begin{aligned} r\dot{u} &= -\dot{r}u - \dot{w} \\ &= (1-2a_1)\langle p_1,u\rangle u - \Big(p_1 - 2a_1\langle p_1,u\rangle u\Big) + O(r) \\ &= \langle p_1,u\rangle u - p_1 + O(r). \end{aligned}$$

Hence the scaled system in polar form can be written

$$\begin{aligned} r\dot{r} &= -(1-2a_1)r\langle p_1,u\rangle + O(r^2), \\ r\dot{u} &= \langle p_1,u\rangle u - p_1 + O(r). \end{aligned}$$

We emphasize that the $O(r)$ and $O(r^2)$ terms are smooth functions of r,u as long as $|r| < \varepsilon$. We consider the "semi-scaled" polar system

$$\dot{r} = -(1-2a_1)r\langle p_1,u\rangle + O(r^2), \tag{14a}$$

$$\dot{u} = \langle p_1,u\rangle u - p_1 + O(r), \tag{14b}$$

which has the same trajectories as the original system, just with different time parametrizations. The advantage of this modified system is that the equations are smooth on $(-\varepsilon,\varepsilon)\times S^{d-1}$.

Observe that (14) has $\{0\}\times S^{d-1}$ invariant, and has fixed point $(r^*,u^*) = (0,p_1)$. The fixed point p_1 is repelling on the invariant manifold $\{0\}\times S^{d-1}$; to see this, observe that

$$\langle p_1,u\rangle^{\cdot} = \langle p_1,u\rangle^2 - 1$$

when $r=0$. In fact, if we assign the coordinate θ on any great circle joining p_1 and $-p_1$ on S^{d-1} so that $u = e^{i\theta}$ and $p_1 = 1$, then the system reduces to $\dot{\theta} = \sin\theta$. We also see that the $\dot{r}$ equation linearized at $(0,p_1)$ is $\dot{r} = -(1-2a_1)r$, so the linearization of (14) has the single negative eigenvalue $-(1-2a_1)$ and $d-1$ positive eigenvalues $+1$. Therefore $(0,p_1)$ is a saddle with a one-dimensional stable manifold, and hence has a unique trajectory $(r(t),u(t)) \to (0,p_1)$ with $r(t) > 0$.

Now suppose we have a trajectory $w(t) \to p_1$ in our original system (12). The corresponding trajectory for (14) will have $r(t)\to 0$; we cannot achieve $r(t) = 0$ in finite time because the manifold $\{r=0\}$ is invariant for (14). We must prove that $u(t)\to p_1$, so that this trajectory is in fact the saddle stable manifold. Observe that

$$\langle p_1,u\rangle^{\cdot} = \langle p_1,u\rangle^2 - 1 + O(r).$$

Also note that $\langle p_1,u(t)\rangle > 0$ since $|w(t)| < 1$. Let $0 < c < 1$; then for some $T \geq 0$, $t \geq T$ implies $O(r(t)) \leq (1-c^2)/2$. If for some $t_0 \geq T$ we have $0 < \langle p_1,u(t_0)\rangle \leq c$, then

$$\langle p_1,u(t)\rangle^{\cdot} \leq c^2 - 1 + \frac{1}{2}(1-c^2) \leq -\frac{1}{2}(1-c^2)$$

for all $t \geq t_0$. But then eventually $\langle p_1,u(t)\rangle < 0$, which is a contradiction. Hence we must have $c < \langle p_1,u(t)\rangle$ for sufficiently large t, which proves that $u(t)\to p_1$. □

VIII. SUMMARY

The natural hyperbolic geometry on the unit ball, with isometries consisting of the higher-dimensional Möbius group, is key to understanding the dynamics of networks of the form (1). Using this framework, we see that dynamical trajectories are constrained to lie on group orbits, and we can explicitly give the equations for the reduced dynamics on group orbits. For the special class of linear order parameters, the dynamics can be further reduced to holonomic constraints determining a flow on the unit ball B^d, which is gradient with respect to the hyperbolic metric.

ACKNOWLEDGMENTS

This article was supported in part by the NSF Research Training Group Grant: Dynamics, Probability, and PDEs in Pure and Applied Mathematics, DMS-1645643.

[1] T. Tanaka, "Solvable model of the collective motion of heterogeneous particles interacting on a sphere," New Journal of Physics **16** (2014).
[2] M. Lohe, "Higher-dimensional generalizations of the Watanabe-Strogatz transform for vector models of synchronization," Journal of Physics A: Mathematical and Theoretical **51** (2018).
[3] S. Chandra, M. Girvan, and E. Ott, "Continuous versus discontinuous transitions in the D-dimensional generalized Kuramoto model: Odd D is different," Phys. Rev. X **9** (2019).
[4] S. Chandra, M. Girvan, and E. Ott, "Complexity reduction ansatz for systems of interacting orientable agents: Beyond the Kuramoto model," Chaos **29** (2019).
[5] S. Marvel, R. Mirollo, and S. Strogatz, "Identical phase oscillators with global sinusoidal coupling evolve by Möbius group action," Chaos **19** (2009).
[6] E. Ott and T. Antonsen, "Low dimensional behavior of large systems of globally coupled oscillators," Chaos **18** (2008).
[7] W. Rudin, *Function Theory in the Unit Ball of* $\mathbb{C}^n$ (Springer, 1980).
[8] B. Chen, J. Engelbrecht, and R. Mirollo, "Dynamics of the Kuramoto-Sakaguchi oscillator network with asymmetric order parameter," Chaos **29** (2019).
[9] M. Stoll, "Harmonic function theory on real hyperbolic space," `https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.561.4447&rep=rep1&type=pdf`.