

Continuum Swarm Tracking Control: A Geometric Perspective in Wasserstein Space

Max Emerick and Bassam Bamieh

Abstract—We consider a setting in which one swarm of agents is to service or track a second swarm, and formulate an optimal control problem which trades off between the competing objectives of servicing and motion costs. We consider the continuum limit where large-scale swarms are modeled in terms of their time-varying densities, and where the Wasserstein distance between two densities captures the servicing cost. We show how this non-linear infinite-dimensional optimal control problem is intimately related to the geometry of Wasserstein space, and provide new results in the case of absolutely continuous densities and constant-in-time references. Specifically, we show that optimal swarm trajectories follow Wasserstein geodesics, while the optimal control tradeoff determines the time-schedule of travel along these geodesics. We briefly describe how this solution provides a basis for a model-predictive control scheme for tracking time-varying and real-time reference trajectories as well.

I. INTRODUCTION

Low-cost sensing, processing, and communication hardware is driving the use of autonomous robotic swarms in diverse settings, including emergency response, transportation, logistics, data collection, and defense. Large swarms in particular can have significant advantages in efficiency and robustness, but modeling these large swarms can become an issue. For sufficiently large swarms, modeling the swarm as a density distribution over the domain (i.e. as a continuum) provides a significant model reduction as well as improved insight into the global behavior of the swarm. Thus, the development of effective motion planning and control strategies for systems of distributions is a problem of interest.

One natural mathematical setting in which to study these continuum swarm problems is the *Wasserstein space* of optimal transport theory. This space equips the set of normalized distributions over a domain with the *Wasserstein distance* – a metric based on the cost of transport whose utility lies largely in the fact that it respects the topology of the underlying physical space¹. In recent years, numerous approaches have been taken to swarm control using the Wasserstein distance and other tools from optimal transport theory. These works fall into two main approaches: one based on distributed optimization, and one based on optimal control. This work falls into the latter category.

In the distributed optimization approach, each agent communicates and acts locally to steer the entire swarm towards a given target distribution. The most relevant works

here include [1]–[3]. This approach has the advantage of being decentralized, and thus balancing the computation and communication loads in a way that is desirable in a practical implementation. However, these approaches have also been limited in their treatment of the objective and constraints: they seek to converge to the target distribution while minimizing the transport distance. In a real-world setting, it is desirable to have a model that can accommodate more general behaviors, especially since convergence to a target distribution is not always possible.

On the other hand, in the optimal control approach, a central planner controls the entire swarm to minimize a given cost function. This approach has the advantages of being able to accommodate more general objectives and constraints and yielding greater insight into the nature of optimal swarm behavior generally. While the solutions here are centralized and thus impractical for a large-scale implementation, one expects that the results from this approach will lend tools to design better distributed swarm control algorithms.

There is a large recent literature on optimal control in Wasserstein spaces, with the most relevant works in the area of multi-agent systems including [4]–[12]. These works focus primarily on the analytic aspects of the problem such as existence and uniqueness of solutions and necessary conditions for optimality. They also focus on systems governed by various types of nonlocal PDEs which arise in connection with mean-field models for self-organizing swarms.

In this paper, we continue the investigation of a model for continuum swarm tracking control originally proposed in [13]. The approach we take is in line with the second group of approaches in that it is based on optimal control in Wasserstein space. However, it differs from these approaches in that it is primarily concerned with tracking as opposed to self-organization and that it uses more explicit models to obtain sharper results. In short, the focus of this work is to answer the question “what principles underlie optimal motion planning and control in swarm-based tracking scenarios?” In contrast to our previous work on this problem [13], [14] which focused on the special case of swarms in one spatial dimension, this work is more general in that it treats swarms in n dimensions (of course, $n = 1, 2,$ or 3 in a practical setting). In addition, this work takes a distinctively geometric approach, which not only provides more powerful tools for solving these problems but yields additional insight into the problem’s underlying structure. The main contributions of this paper are thus:

- The introduction of several geometric tools for solving swarm tracking control problems,

M. Emerick and B. Bamieh are with the Department of Mechanical Engineering, University of California, Santa Barbara, USA
memerick@ucsb.edu, bamieh@ucsb.edu

¹Compare, for example, the L^p distance, which can be arbitrarily sensitive to spatial perturbations.

- Analytic solutions to our model in the n -dimensional case where the swarm distribution is absolutely continuous and the reference distribution² is static.

We first introduce our problem formulation and present our main result. The rest of the paper is then devoted to developing the tools necessary to prove it.

II. PROBLEM FORMULATION AND MAIN RESULT

In previous work [13] we proposed a model for swarm tracking control. Since this proposed model forms the basis of this work, we briefly review the problem formulation here. For a full development and motivation, see [13].

The key feature of our proposed problem setting is that there are two distributions, which we refer to as *demand* and *resource* respectively. The demand can be mobile or stationary and represents the spatial distribution of entities which require attention (e.g. locations to monitor, facilities to supply, independent mobile robots to be serviced, etc.). The resource represents the spatial distribution of controlled mobile agents which are able to service the needs of the demand entities. Physical space is taken to be \mathbb{R}^n where n is a positive integer. Both the demand D and resource R are modeled as time-varying densities over the domain, taking the form³

$$D_t(x) = d(x, t) + \sum_{i=1}^N m_i(t) \delta(x - p_i(t)) \quad (1)$$

and similarly for R . Here, $x \in \mathbb{R}^n$ is the spatial location while $t \in [0, T]$ is the time. The function $d: \mathbb{R}^n \times [0, T] \rightarrow \mathbb{R}_{\geq 0}$ is the continuum portion of the density, which describes a continuum approximation of a large swarm. The summation represents the discrete portion of the density, which describes N discrete components having mass m_i at location p_i . In this work, we present results exclusively for continuum swarms, and so from here on we assume that $N = 0$. We also assume that both resource and demand distributions have been scaled to integrate to 1

$$\int_{\mathbb{R}^n} R_t dx = \int_{\mathbb{R}^n} D_t dx = 1. \quad (2)$$

We then say that the resource and demand distributions are *normalized*, and denote the set of normalized distributions over \mathbb{R}^n as $\mathfrak{D}(\mathbb{R}^n)$. We also make the additional technical assumption that the resource and demand distributions are supported within a compact subset $\Omega \subset \mathbb{R}^n$ for all time.

The resource agents service the needs of the demand entities through an assignment process where each resource particle is coupled to a set of demand particles and vice versa. This coupling is called an *assignment kernel* and is one of the decision variables in our optimal control problem. The assignment kernel incurs an *assignment cost* which is related to the efficiency of the assignment. In the problem setting discussed in [13], the resource provides communication services

to the demand, and the assignment cost is thus proportional to the squared distance between coupled resource and demand particles. Supposing that the communication loads must be balanced, each demand particle must be serviced, and the coupling is chosen optimally, we obtain the assignment cost

$$\begin{aligned} \inf_{\mathcal{K}} \int_{\mathbb{R}^n \times \mathbb{R}^n} \|x - y\|_2^2 \mathcal{K}(x, y, t) dx dy \\ \text{s.t. } \int \mathcal{K} dy = R \\ \int \mathcal{K} dx = D. \end{aligned} \quad (3)$$

It turns out that this assignment cost is exactly the square of the 2-Wasserstein distance \mathcal{W}_2 of optimal transport theory [15]. Endowed with \mathcal{W}_2 , the space of normalized distributions $\mathfrak{D}(\mathbb{R}^n)$ becomes a metric space, which we call *2-Wasserstein space* and denote \mathbb{W}_2 . This will be the setting that we work in in this paper.

It is important to note that while the Wasserstein distance is often motivated as a “transport cost” in optimal transport theory, that is not the way it appears in this problem setting. Here, the Wasserstein distance represents the cost of servicing, which is proportional to the distance between coupled particles, but is *unrelated to motion*. There is indeed a physical cost of motion in our problem setting as well, but this is quantified differently, as we describe next.

In order to track the demand, the resource swarm is controlled through a time-varying vector field V . The equations of motion of the system are given by the *transport equation*

$$\partial_t R_t(x) = -\nabla \cdot (V_t(x) R_t(x)). \quad (4)$$

The vector field V is the main decision variable in this optimal control problem. V can be used to steer the resource distribution closer to the demand distribution, lowering the assignment cost. However, this incurs a *motion cost* which is related to the total energy expended in motion. In our problem setting, the motion cost represents the total aerodynamic drag on the resource swarm, taking the form

$$\int_{\mathbb{R}^n} |V_t(x)|_2^2 R_t(x) dx =: \|V_t\|_{L^2(R_t)}^2. \quad (5)$$

Lastly, we have an *objective function* which defines optimal behavior for the resource swarm. The objective function is given by the integral of the assignment cost plus the motion cost over the time horizon $[0, T]$. We remark again that the assignment cost and the motion cost are competing objectives: if R_t and D_t are far apart, then the assignment cost will be large, but this cost can be reduced at the expense of motion (and vice versa). To capture the tradeoff between assignment and motion costs, we use a weighting parameter α to control the relative importance of the two costs. All in all, the proposed problem is to find the control which minimizes the total cost of a maneuver. The problem is stated formally as follows.

Problem 1 (Original Formulation). Given an initial resource distribution R_0 and demand trajectory D over $[0, T]$, solve

$$\begin{aligned} \inf_V \int_0^T \mathcal{W}_2^2(R_t, D_t) + \alpha \|V_t\|_{L^2(R_t)}^2 dt \\ \text{s.t. } \partial_t R_t(x) = -\nabla \cdot (V_t(x) R_t(x)). \end{aligned} \quad (6)$$

²In this paper we use the terms *density* and *distribution* synonymously.

³We conceptualize the demand and resource as parameterized curves in the space of distributions, thus we write D to refer to the entire curve and D_t to refer to the particular distribution at the parameter value t .

As stated, this problem is a nonlinear, infinite-dimensional optimal tracking problem: R plays the role of the state, D the reference trajectory, and V the control input. The assignment cost (given by \mathcal{W}_2^2) plays the role of the tracking error, while the motion cost (given by $\|V_t\|^2$) plays the role of the control energy. The transport equation is the dynamic constraint.

Notice that the decision variable \mathcal{K} does not appear explicitly in this problem formulation. This is because optimal solutions always have \mathcal{K} chosen optimally at each (static) instant in time. Assuming each of these static subproblems to be solved, we obtain the above formulation. For further discussion on this, see [13].

A cartoon depiction of this problem is shown below in Figure 1.




Fig. 1: The resource R is paired to the demand D by the assignment kernel \mathcal{K} and is transported by the vector field V . The cost is minimized over the whole maneuver.

A. Main Result

The main result of this paper is a complete characterization of solutions to Problem (1) when R_0 is absolutely continuous and D is constant in time. To state it, we first need to introduce some background.

Definition 2 (Pushforward). Let μ be a density defined on Ω_1 and f be a function from Ω_1 to Ω_2 . The *pushforward* of μ through f is a density $f_{\#}\mu$ defined on Ω_2 such that

$$\int_{\Omega_2} \psi(y) (f_{\#}\mu)(y) dy = \int_{\Omega_1} (\psi \circ f)(x) \mu(x) dx \quad (7)$$

for all test functions $\psi : \Omega_2 \rightarrow \mathbb{R}$.

The pushforward can be conceptualized as the density formed by “moving the mass in μ forward through f ”. The above definition in terms of integration against test functions is equivalent to other standard definitions, but will be more useful to us.

The notion of the pushforward allows us to define *transport maps* between densities.

Definition 3 (Transport Maps). Given two normalized densities μ, ρ on \mathbb{R}^n , a *transport map* from μ to ρ is a map $M : \mathbb{R}^n \rightarrow \mathbb{R}^n$ such that $\rho = M_{\#}\mu$. M is said to be *optimal* if it minimizes (over all such maps) the functional

$$\int_{\mathbb{R}^n} |M(x) - x|^2 \mu(x) dx =: \|M - \mathcal{I}\|_{L^2(\mu)}^2. \quad (8)$$

Lastly, we introduce the concept of a geodesic.

Definition 4 (Wasserstein Geodesic). Given two normalized densities μ, ρ on \mathbb{R}^n , a (constant-speed) *Wasserstein geodesic* from μ to ρ is a curve $\gamma : [0, 1] \rightarrow \mathbb{W}_2$ with $\gamma_0 = \mu$ and $\gamma_1 = \rho$ having the property

$$\mathcal{W}_2(\gamma_{t_1}, \gamma_{t_2}) = |t_1 - t_2| \mathcal{W}_2(\mu, \rho) \quad (9)$$

for all $t_1, t_2 \in [0, 1]$. We write $\Gamma(\mu, \rho)$ to denote the range of the geodesic (i.e. the set of points in \mathbb{W}_2).

Intuitively, a geodesic is a curve which achieves the minimum length between μ and ρ . That these curves always exist is a useful but nontrivial fact [15].

We can now state our main result.

Theorem 5. *When R_0 is absolutely continuous and D is constant in time, the solution to Problem 1 is given implicitly by the feedback controller*

$$V_t = -f(t) (\mathcal{I} - \bar{M}_t) / \alpha, \quad (10)$$

where \mathcal{I} is the identity map on \mathbb{R}^n , \bar{M}_t is the optimal transport map taking R_t to D , and

$$f(t) = \sqrt{\alpha} \tanh((T - t) / \sqrt{\alpha}). \quad (11)$$

The trajectory generated by this control law takes the form

$$R_t = [(1 - \sigma(t)) \mathcal{I} + \sigma(t) \bar{M}_0]_{\#} R_0 \quad (12)$$

where

$$\sigma(t) = 1 - \cosh((T - t) / \sqrt{\alpha}). \quad (13)$$

In particular, R follows the Wasserstein geodesic from R_0 to D . Furthermore, this solution attains the cost

$$\mathcal{J}(R_0, T; \alpha; D) = \frac{1}{2} \mathcal{W}_2^2(R_0, D) \sqrt{\alpha} \tanh(T / \sqrt{\alpha}). \quad (14)$$

In other words, the optimal control steers the state along the shortest path to the reference distribution at a rate which depends on the initial distance $\mathcal{W}_2(R_0, D)$, T , and α .

From a practical standpoint, we point out that even though this solution is based on a static reference, it provides a basis for a receding-horizon Model-Predictive Control (MPC) scheme for tracking time-varying, stochastic, and a priori unknown demand trajectories. This MPC scheme would be similar in spirit to that described in [16], where tracking of a time-varying reference is achieved by treating it as a piecewise constant signal. In this control scheme, the demand distribution would be updated at each timestep, with the optimal control (10) being recomputed and applied each time the demand is updated. Developing a control scheme of this sort will be the subject of future work.

From a theoretical standpoint, it is somewhat surprising that the problem (6) admits tractable solutions at all, let alone the highly structured analytic solutions we see here. At first glance, problem (6) is a nonlinear, infinite-dimensional, multi-level optimization problem, and we might expect that computationally-intensive numerical solutions would be the only line of approach. However, this is not the case. The reason that this problem turns out to be nice is that it has a rich geometry, which we can leverage for some very powerful

tools. It is worth giving this geometry attention then, as it seems that similar tools may be applied to solve other sorts of nonlinear optimization problems as well. Thus, the rest of this paper will focus on developing the geometric picture and the tools necessary to demonstrate our main result.

III. THE GEOMETRIC PICTURE

The central idea in this section is that there are two representations for our system, which we call *Eulerian* and *Lagrangian*. The Eulerian representation describes the system in terms of particle densities, while the Lagrangian representation describes the system in terms of particle coordinates. These two representations are equivalent, and we can learn a great deal about our problem by formalizing this. In this section we describe these two representations and their correspondence following [17]. In the next section we will use the tools developed here to reformulate and solve our problem.

A. The Eulerian Setting

In this setting we study the evolution of densities in the Wasserstein space \mathbb{W}_2 . It turns out that \mathbb{W}_2 forms an infinite-dimensional Riemannian manifold (with boundary) [18], as we explain next. Points in \mathbb{W}_2 represent normalized densities, which we denote in this section by ρ , μ , or χ_t (where $\chi : [0, T] \rightarrow \mathbb{W}_2$ is a curve in \mathbb{W}_2). Recall that the derivative of a curve χ at time t is formalized as the tangent vector χ'_t , and the set of all tangent vectors at a given point ρ forms the *tangent space* at ρ , denoted $\mathcal{T}_\rho \mathbb{W}_2$. Intrinsically, this tangent space is identified with the set of variations on the density ρ , roughly taking the form of signed distributions on \mathbb{R}^n having zero total mass [17]. However, since we frequently generate curves in \mathbb{W}_2 via the transport equation

$$\chi'_t = -\nabla \cdot (v_t \chi_t), \quad (15)$$

it is natural to ask what the relation is between the intrinsic tangent vectors (i.e. the variations on χ_t) and objects of the form $-\nabla \cdot (v_t \chi_t)$. It turns out that every sufficiently regular curve can be generated this way (as we show next) and thus we choose to identify the tangent space $\mathcal{T}_\rho \mathbb{W}_2$ with the set of objects $\{-\nabla \cdot (v \rho) : v : \mathbb{R}^n \rightarrow \mathbb{R}^n\}$.

Lemma 6 ([15, Theorem 5.14]). *A curve $\chi : [0, T] \rightarrow \mathbb{W}_2$ is absolutely continuous⁴ if and only if there exists a time-varying vector field v with*

$$\int_0^T \|v_t\|_{L^2(\chi_t)} dt < \infty \quad (16)$$

such that (χ, v) satisfy the transport equation (15). Furthermore, the speed of the curve $|\chi'_t|$ satisfies

$$|\chi'_t| = \min_{\chi'_t = -\nabla \cdot (v_t \chi_t)} \|v_t\|_{L^2(\chi_t)} \quad (17)$$

for almost all t .

⁴Note that absolute continuity is defined both for distributions and for curves. We will state which definition we are using in each context.

Observe that this lemma not only relates absolutely continuous curves in \mathbb{W}_2 to solutions of the transport equation, but by giving us a formula for the speed, identifies the metric tensor in \mathbb{W}_2 as well

$$\langle \tau_1, \tau_2 \rangle_\rho = \min_{\substack{\tau_1 = -\nabla \cdot (v_1 \rho) \\ \tau_2 = -\nabla \cdot (v_2 \rho)}} \langle v_1, v_2 \rangle_{L^2(\rho)}. \quad (18)$$

Recall that the metric tensor completely defines the geometry of the manifold, providing the notions of length, speed, distance, geodesics, and curvature.

B. The Lagrangian Setting

In this setting, we study the evolution of particle coordinates in the space of maps. Recall that the trajectory of a particle in a vector field v is described by an *integral curve* of the vector field, and that the collection of all integral curves gives us the *flow map* ϕ :

$$\begin{aligned} \partial_t \phi_t(x) &= v_t(\phi_t(x)) \\ \phi_0(x) &= x. \end{aligned} \quad (19)$$

Recall that $\phi_t(x)$ describes the position at time t of the particle that started from position x at time 0. Thus, at a fixed instant in time, ϕ_t describes the coordinates of all particles as a map from \mathbb{R}^n to \mathbb{R}^n . Thus the state of the system is represented as a point in the space of maps, denoted \mathbb{M} .

The space \mathbb{M} also forms an infinite-dimensional Riemannian manifold. (In fact, \mathbb{M} forms a vector space, which tells us a bit more.) Points in \mathbb{M} represent maps from \mathbb{R}^n to \mathbb{R}^n , and are denoted by M (or ϕ_t if the map represents the flow of a particular velocity field v). The tangent space at M is the set of variations u on the map M . Since \mathbb{M} is a vector space, this set is isomorphic to the space of maps itself, $\mathcal{T}_M \mathbb{M} \cong \mathbb{M}$. This allows us to endow this space with the inner product

$$\langle u_1, u_2 \rangle_{L^p(\mu)} := \int_{\mathbb{R}^n} u_1^T(x) u_2(x) \mu(x) dx, \quad (20)$$

where μ is some fixed reference distribution. Notice that an inner product gives a constant metric tensor, implying that \mathbb{M} has zero curvature, i.e. that \mathbb{M} is *flat*.

C. Correspondence Between Settings

The correspondence between the Eulerian and Lagrangian settings is identified by the following lemma.

Lemma 7 ([15, Theorem 4.4]). *Let $\chi_0 \in \mathbb{W}_2$ be an absolutely continuous initial distribution and $v : \mathbb{R}^n \times [0, T] \rightarrow \mathbb{R}^n$ be a sufficiently regular time-varying vector field. Then (χ, v) satisfy the transport equation (15) if and only if*

$$\chi_t = [\phi_t]_{\#} \chi_0, \quad (21)$$

where ϕ is the flow map of v .

There are several important points to make here. First, this correspondence is induced by the pushforward: given a fixed reference distribution μ , the pushforward defines a function Π which takes each map $M \in \mathbb{M}$ to a distribution $M_{\#} \mu \in \mathbb{W}_2$. Thus the pushforward induces a map between manifolds $\Pi : \mathbb{M} \rightarrow \mathbb{W}_2$. Second, Lemma 7 is a statement

about how dynamics transform under this map Π . Taking $\mu = \chi_0$ to be the reference distribution, the dynamics (15) and (19) are equivalent under Π in the sense that the curves generated by these dynamics agree under this transformation. Thus we write

$$\Pi : \phi_t \mapsto \chi_t = [\phi_t]_{\#} \chi_0. \quad (22)$$

Third, knowing how the dynamics transform tells us how tangent vectors transform as well. Recall that the map between tangent spaces $\mathcal{T}_M \mathbb{M}$ and $\mathcal{T}_{\Pi(M)} \mathbb{W}_2$ is the *differential* of the map Π at M , denoted $D\Pi(M)$. Lemma 7 tells us immediately that $D\Pi(\phi_t) : \phi'_t \mapsto \chi'_t$. We can manipulate this into an explicit form as follows. Using (19), we can write $v_t = \phi'_t \circ \phi_t^{-1}$. Substituting this into (15) along with (21) yields

$$\chi'_t = -\nabla \cdot \left((\phi'_t \circ \phi_t^{-1})([\phi_t]_{\#} \chi_0) \right). \quad (23)$$

From this, we deduce that the differential is

$$D\Pi(M) : u \mapsto -\nabla \cdot \left((u \circ M^{-1})(M_{\#} \mu) \right). \quad (24)$$

We take this moment to point out that we will use representations of tangent vectors in $\mathcal{T}_M \mathbb{M}$ in two different coordinate systems, related by $u = v \circ M$. For example, in these transformed coordinates, the differential takes the form

$$D\Pi(M) : v \mapsto -\nabla \cdot (v(M_{\#} \mu)). \quad (25)$$

We will always use the letters u and v for tangent vectors in each of these coordinate systems.

The next natural question to ask is what properties the map Π has (e.g. continuous, differentiable, 1-1, onto, etc.). We can see that Π is differentiable (and thus continuous) from (25). It turns out that Π is not 1-1, which we can see by observing that $M_{\#} \mu = [M \circ \mathcal{P}]_{\#} \mu$ for any measure-preserving transformation \mathcal{P} . However, Π is onto, provided that the reference distribution μ is chosen to be absolutely continuous. We can see this because we can define a right inverse for Π as

$$\Pi^{-R} : \rho \mapsto \operatorname{argmin}_{\Pi(M)=\rho} \|M - \mathcal{I}\|_{L^2(\mu)}. \quad (26)$$

To show that this is actually a right inverse, we need to show that this optimization problem attains a unique minimum for all ρ . To see this, observe that (26) is actually just the Monge problem of optimal transport theory: the $L^2(\mu)$ -norm of $(M - \mathcal{I})$ is the square root of the transport cost, the constraint $\Pi(M) = \rho$ expresses that $M_{\#} \mu = \rho$, i.e. that M is a transport map from μ to ρ , and thus the minimizing argument $\bar{M} = \Pi^{-R}(\rho)$ is the optimal transport map. Recall that when μ is absolutely continuous, a unique minimizer to the Monge problem is guaranteed to exist [15, Theorem 1.22], and thus Π^{-R} is indeed a right inverse. Therefore Π and Π^{-R} define a 1-1 correspondence between \mathbb{W}_2 and the set of optimal transport maps in \mathbb{M} , which we denote \mathbb{O}

$$\mathbb{W}_2 \xrightleftharpoons[\Pi^{-R}]{\Pi} \mathbb{O}. \quad (27)$$

By understanding the structure of \mathbb{O} we can therefore understand the structure of \mathbb{W}_2 . We start with the following characterization.

Proposition 8 ([15, Theorems 1.22, 1.48]). *A transport map is optimal if and only if it is equal to the gradient of a convex function.*

Thus we identify $\mathbb{O} = \{M = \nabla F : F \text{ convex}\}$. Remarkably, this implies that \mathbb{O} is *independent of the distribution* μ^5 . The picture is completed by identifying the kernel of the transformation Π . Recall our earlier counterexample that $M_{\#} \mu = [M \circ \mathcal{P}]_{\#} \mu$ for any measure-preserving transformation \mathcal{P} , thus any such \mathcal{P} is in the kernel. In fact, this is a complete characterization of the kernel.

Proposition 9 ([15, Theorem 1.53]). *Consider a map $M : \Omega \rightarrow \mathbb{R}^n$ where $\Omega \subset \mathbb{R}^n$ is convex, and μ an absolutely continuous distribution on Ω . Then there exists a convex function $F : \Omega \rightarrow \mathbb{R}^n$ and a map $\mathcal{P} : \Omega \rightarrow \Omega$ preserving μ such that $M = (\nabla F) \circ \mathcal{P}$. Furthermore, ∇F is unique up to μ -a.e. equivalence.*

To summarize, then: Π maps \mathbb{M} onto \mathbb{W}_2 . The kernel $\ker(\Pi) = \{\mathcal{P}\}$ forms a group. Two elements in \mathbb{M} are equivalent under Π if they are related by an element of $\{\mathcal{P}\}$. Each equivalence class $[M]$ has a unique representative \bar{M} in \mathbb{O} . Then \mathbb{O} can be identified with the quotient space $\mathbb{M}/\{\mathcal{P}\}$. Figure 2 below shows a cartoon depiction of this situation.




Fig. 2: Relationship between \mathbb{M} , \mathbb{O} , \mathbb{W}_2 . Π maps \mathbb{M} onto \mathbb{W}_2 . Each equivalence class $[M]$ is the preimage of a point ρ in \mathbb{W}_2 . Each equivalence class has a unique representative $\bar{M} \in \mathbb{O}$. \mathbb{O} is a quotient space of \mathbb{M} and is 1-1 with \mathbb{W}_2 .

In this way, \mathbb{O} is an embedding of \mathbb{W}_2 in \mathbb{M} . However, \mathbb{O} carries a different geometry than that of the latent space. The natural geometry for \mathbb{O} (and the one that makes it isometric to \mathbb{W}_2) is given by minimizing the metric tensor over all elements in the equivalence class of each point. Formally, let $\Psi = \Pi^{-R} \circ \Pi$ be the projection from \mathbb{M} to \mathbb{O} . Then the metric tensor in \mathbb{O} is defined to be

$$\langle w_1, w_2 \rangle_{\bar{M}} = \min_{\substack{D\Psi(M)[u_1]=w_1 \\ D\Psi(M)[u_2]=w_2}} \langle u_1, u_2 \rangle_M. \quad (28)$$

We can see that this coincides with the metric tensor in \mathbb{W}_2 (18) by applying our coordinate change $u = v \circ M$ and using

⁵Note, however, that the associations Π and Π^{-R} are *not* independent of μ , although this is suppressed in the notation for simplicity.

properties of the pushforward

$$\begin{aligned}\langle u_1, u_2 \rangle_M &= \langle u_1, u_2 \rangle_{L^2(\mu)} = \langle v_1 \circ M, v_2 \circ M \rangle_{L^2(\mu)} \\ &= \langle v_1, v_2 \rangle_{L^2(\mathcal{M}_{\#}\mu)} = \langle v_1, v_2 \rangle_{L^2(\rho)},\end{aligned}\quad (29)$$

so that by taking the minimum of both sides,

$$\begin{aligned}\langle w_1, w_2 \rangle_{\bar{M}} &:= \min \langle u_1, u_2 \rangle_M \\ &= \min \langle v_1, v_2 \rangle_{L^2(\rho)} =: \langle \tau_1, \tau_2 \rangle_{\rho}.\end{aligned}\quad (30)$$

In the language of differential geometry, the map Π is called a *Riemannian submersion*. This means that the differential $D\Pi(M)$ is an orthogonal projection onto each tangent space $\mathcal{T}_{\Pi(M)}\mathbb{W}_2$. The restriction of $D\Pi(M)$ to the orthogonal complement of its kernel is thus an isometry between $\ker(D\Pi(M))^\perp$ and $\mathcal{T}_{\Pi(M)}\mathbb{W}_2$. A tangent vector in $\ker(D\Pi(M))$ is called *vertical*, while a tangent vector in $\ker(D\Pi(M))^\perp$ is called *horizontal*. We have the following characterization of vertical and horizontal tangent vectors.

Proposition 10 ([17, Section 4.1]). *A tangent vector $v \in \mathcal{T}_M\mathbb{M}$ is vertical with respect to Π if and only if $\nabla \cdot (v\rho) = 0$, where $\rho = M_{\#}\mu$. A tangent vector $v \in \mathcal{T}_M\mathbb{M}$ is horizontal if and only if $v = \nabla H$ for some scalar function H .*

The machinery of Riemannian submersions now allows us to develop a number of useful results. First, we compute the differential of Π^{-R} . Due to the isometries between \mathbb{W}_2 and \mathbb{O} and their tangent spaces, we know that $D\Pi^{-R}(\rho)$ must preserve the norms of tangent vectors. Since $D\Pi(M)$ is an orthogonal projection onto each tangent space, there is exactly one tangent vector in the preimage of τ with the same norm (i.e. that which is horizontal). Thus we can write

$$D\Pi^{-R}(\rho) : \tau \mapsto \operatorname{argmin}_{D\Pi(M)[v]=\tau} \|v\|_{L^2(\rho)}. \quad (31)$$

The subsequent facts also follow directly from the properties of Π as a Riemannian submersion [17]:

- 1) If g is geodesic in \mathbb{M} and g'_t is horizontal for some t , then g'_t is horizontal for all t . In this case, g is called a *horizontal geodesic*.
- 2) The image of a horizontal geodesic g under a Riemannian submersion is a geodesic γ in the codomain.

Using these facts, we can now write down an expression for a geodesic in \mathbb{W}_2 . Take two distributions μ and ρ in \mathbb{W}_2 , with μ absolutely continuous. Take μ as our fixed reference distribution. Let \mathcal{I} be the identity map and $\bar{M} = \Pi^{-R}(\rho)$ be the optimal transport map taking μ to ρ . Since \mathbb{M} is flat, the geodesic from \mathcal{I} to \bar{M} in \mathbb{M} is given by

$$g_t = (1-t)\mathcal{I} + t\bar{M}. \quad (32)$$

Furthermore, we have $u_t := g'_t = \bar{M} - \mathcal{I}$, or $v_t = \bar{M} \circ g_t^{-1} - g_t^{-1}$ in our transformed coordinates. Observe that since $g_0 = \mathcal{I}$, $v_0 = \bar{M} - \mathcal{I}$. Since \bar{M} is an optimal transport map, $\bar{M} = \nabla F$ for some convex function F . Similarly, $\mathcal{I} = \nabla \frac{1}{2}|\cdot|^2$, and thus we can write

$$v_0 = \nabla \left(F - \frac{1}{2}|\cdot|^2 \right) = \nabla H. \quad (33)$$

Thus v_0 is horizontal, so v_t is horizontal for all t , so g is a horizontal geodesic in \mathbb{M} . Therefore the image of g under Π

$$\gamma_t = [(1-t)\mathcal{I} + t\bar{M}]_{\#}\mu \quad (34)$$

is a geodesic in \mathbb{W}_2 . This completes our development of the geometric picture.

IV. GEOMETRIC FORMULATION OF PROBLEM

The first step in approaching the higher-dimensional case is a geometric formulation of our original problem. In this formulation, instead of considering minimizing control inputs, we consider minimizing curves in Wasserstein space. To do this, we need to pose the problem explicitly in \mathbb{W}_2 . Examining Problem 1, we see that the assignment cost is already an object in \mathbb{W}_2 , but we still need to transform the motion cost and constraint. By Lemma (6), we see that the motion cost is equal to the square of the speed of the curve R and the constraint becomes simply that the curve R be absolutely continuous (A.C.). We call this the *geometric formulation*.

Problem 11 (*Geometric Formulation*). Given an initial resource distribution R_0 and demand trajectory D over $[0, T]$, solve

$$\begin{aligned}\inf_R \int_0^T \mathcal{W}_2^2(R_t, D_t) + \alpha |R'_t|^2 dt \\ \text{s.t. } R : [0, T] \rightarrow \mathbb{W}_2 \text{ A.C.}\end{aligned}\quad (35)$$

Notice that this formulation has no explicit dependence on V . However, since V is related to the tangent of R via $D\Pi^{-R}$, it is straightforward to recover one from the other. When the initial resource distribution R_0 is absolutely continuous, the two problems are equivalent.

Lemma 12. *Supposing R_0 is absolutely continuous, Problem 1 and Problem 11 are equivalent. The solutions to the two problems attain the same cost and are related by*

$$\bar{V}_t = \operatorname{argmin}_{\bar{R}'_t = -\nabla \cdot (V_t \bar{R}_t)} \|V_t\|_{L^2(\bar{R}_t)} \quad (36)$$

$$\bar{R}_t = [\bar{\phi}_t]_{\#} R_0, \quad (37)$$

where $\bar{\phi}$ is the flow generated by \bar{V} .

Sketch of Proof. Taking R_0 to be the fixed reference distribution, observe that $\bar{R}_t = \Pi(\bar{\phi}_t)$ and that $\bar{V}_t = D\Pi^{-R}(\bar{R}_t)[\bar{R}'_t]$. Then \bar{R} and $\bar{\phi}$ are corresponding curves in \mathbb{W}_2 and \mathbb{O} respectively, so by the isometry between these spaces, \bar{V} and \bar{R} attain the same cost. \square

V. CONTINUOUS RESOURCE / STATIC DEMAND

In the simple case where the resource distribution is absolutely continuous and the demand distribution is static, the solution has a particularly nice form. Specifically, the solution \bar{R} follows the geodesic from R_0 to D .

Lemma 13. *Consider Problem 11 and suppose that R_0 is absolutely continuous and D is constant in time. Then the*

solution \bar{R} is such that $\bar{R}_t \in \Gamma(R_0, D)$ for all $t \in [0, T]$, where Γ is the range of the Wasserstein geodesic.

Sketch of Proof. For any feasible solution R_1 , we can construct a solution R_2 with $|R'_2| = |R'_1|$ which flows along the geodesic. We find that $\text{cost}(R_2) \leq \text{cost}(R_1)$, so if optimal solutions exist, one must flow along the geodesic. \square

This result should be intuitive. In balancing the motion penalty with the assignment cost (i.e. the distance to D), the optimal solution will minimize distance using the least effort possible, and will thus move along the shortest path between R_0 and D . Since this fully defines the path that the solution takes, we can now parameterize the problem with a single variable representing the distance along this path. This allows us to reduce the problem to a familiar form.

Let $\sigma(t)$ define the fraction of the geodesic which has been traversed by time t . Then the following hold:

$$|R'_t| = \mathcal{W}_2(R_0, D) \sigma'(t) \quad (38)$$

$$\mathcal{W}_2(R_t, D) = \mathcal{W}_2(R_0, D) (1 - \sigma(t)). \quad (39)$$

Defining

$$\zeta := \mathcal{W}_2(R_0, D) \quad (40)$$

$$\eta(t) := \mathcal{W}_2(R_0, D) \sigma(t) \quad (41)$$

$$u(t) := \mathcal{W}_2(R_0, D) \sigma'(t), \quad (42)$$

and substituting into (35), we obtain

$$\inf_u \int_0^T (\zeta(t) - \eta(t))^2 + \alpha u^2(t) dt \quad (43)$$

$$\text{s.t. } \dot{\eta}(t) = u(t). \quad (44)$$

Observe that this is exactly the form of the scalar linear-quadratic tracking problem discussed in Appendix I. Then by Corollary 17 and equation (34) we obtain the following solution to Problem 11.

Lemma 14. *When R_0 is absolutely continuous and D is constant in time, the solution to Problem 11 takes the form*

$$\bar{R}_t = [(1 - \sigma(t))\mathcal{I} + \sigma(t)\bar{M}]_{\#} R_0 \quad (45)$$

where \mathcal{I} is the identity map on \mathbb{R}^n , \bar{M} is the optimal transport map from R_0 to D , and

$$\sigma(t) = 1 - \cosh((T - t)/\sqrt{\alpha}). \quad (46)$$

Furthermore, the solution attains the cost

$$\mathcal{J}(R_0, T; \alpha; D) = \frac{1}{2} \mathcal{W}_2^2(R_0, D) \sqrt{\alpha} \tanh(T/\sqrt{\alpha}). \quad (47)$$

Sketch of Proof. Follows directly as outlined above. \square

This gives us a portion of the main result stated in Theorem 5. However, it is not the complete picture – we still need to recover the feedback form of the optimal controller. To do this, we need to find the velocity field \bar{V} in terms of \bar{R}_t and D . Observe that we already have the optimal trajectory

$$\bar{R}_t = [(1 - \sigma(t))\mathcal{I} + \sigma(t)\bar{M}]_{\#} R_0 \quad (48)$$

in \mathbb{W}_2 , which corresponds to the trajectory

$$\bar{\phi}_t = (1 - \sigma(t))\mathcal{I} + \sigma(t)\bar{M} \quad (49)$$

in \mathbb{O} . We find the tangent of this trajectory to be

$$u_t := \phi'_t = \sigma'(t)(\bar{M} - \mathcal{I}). \quad (50)$$

Now, we want to write this tangent in terms of \bar{M} and $\bar{\phi}_t$ to get it in feedback form. Observe that from (49) we have

$$\mathcal{I} = \frac{-\sigma(t)}{1 - \sigma(t)}\bar{M} + \frac{1}{1 - \sigma(t)}\bar{\phi}_t, \quad (51)$$

which we substitute into (50) to obtain

$$u_t = \sigma'(t) \left(\frac{1}{1 - \sigma(t)} \right) (\bar{M} - \bar{\phi}_t). \quad (52)$$

Using our coordinate transformation $u_t = v_t \circ \bar{\phi}_t$, we have

$$v_t = \sigma'(t) \left(\frac{1}{1 - \sigma(t)} \right) (\bar{M} \circ \bar{\phi}_t^{-1} - \mathcal{I}). \quad (53)$$

Since $\bar{\phi}$ is a geodesic in \mathbb{O} , we know that $\bar{M}_t = \bar{M} \circ \bar{\phi}_t^{-1}$ is the optimal transport map taking \bar{R}_t to D . Using this fact, we get the optimal controller in error-feedback form:

$$v_t = -f(t)(\mathcal{I} - \bar{M}_t)/\alpha, \quad (54)$$

where

$$f(t)/\alpha = \sigma'(t) \left(\frac{1}{1 - \sigma(t)} \right). \quad (55)$$

This completes the development of the results presented in Theorem 5.

VI. DISCUSSION/CONCLUSION

In this paper, we presented a geometric perspective on continuum swarm tracking control and used this perspective to characterize solutions to our model in the n -dimensional case where the resource distribution is absolutely continuous and the demand distribution is static.

Where exactly do these results fail under more general assumptions? When the resource distribution is not absolutely continuous (i.e. if it has discrete components), then we can still formulate both the original problem (6) and the geometric problem (35), and we can still solve the geometric problem in the same manner as was done here. However, the failure is in the equivalence of these models: if the resource distribution has discrete components, then the constraint that the trajectory be absolutely continuous is strictly weaker than the dynamic constraint in the original problem. In particular, absolute continuity of the curve allows discrete masses to be broken up. When the demand is time-varying, we can also still formulate the original problem (6) and the geometric problem (35), and here they are still equivalent. However, the failure in this case is in our solution technique: the resource distribution no longer traverses a single geodesic, having different and more complicated necessary conditions for optimality instead. The exploration of both of these cases will be the subject of future work.

We also briefly mentioned how our solution based on a static demand distribution can be used as the basis for a

model-predictive control scheme for tracking time-varying, stochastic, and apriori unknown demand trajectories. Developing a numerical control scheme along these lines will also be treated in future work.

Lastly, the robustness of these solutions with respect to noise and approximation error (in particular, that which results from using a continuum model) will be investigated in future work as well.

APPENDIX I

THE SCALAR LQ TRACKING PROBLEM

In our previous work [13], we showed that in the special case of one spatial dimension (i.e. distributions on \mathbb{R}), Problem 1 could be transformed into an infinite-dimensional linear-quadratic (LQ) tracking problem which could then be decoupled into a superposition of scalar LQ tracking problems. The scalar LQ tracking problem plays an important role in this paper, and so we review its solution here. The problem takes the following form.

Problem 15 (Scalar LQ Tracking Problem). Given an initial state $\eta_0 \in \mathbb{R}$ and tracking signal $\zeta : [0, T] \rightarrow \mathbb{R}$, solve

$$\begin{aligned} \inf_u \int_0^T (\zeta(t) - \eta(t))^2 + \alpha u^2(t) dt \\ \text{s.t. } \dot{\eta}(t) = u(t), \end{aligned} \quad (56)$$

where $\alpha > 0$ is the trade-off parameter and T is the time horizon.

This problem has been addressed in [13] with the following solution.

Proposition 16 ([13]). *The solution to Problem 15 is given implicitly by the non-causal feedback controller*

$$\bar{u}(t) = -f(t)\eta(t)/\alpha - g(t)/\alpha \quad (57)$$

where

$$f(t) = \sqrt{\alpha} \tanh((T-t)/\sqrt{\alpha}) \quad (58)$$

$$g(t) = \int_T^t \Phi^{-1}(t, \tau) \zeta(\tau) d\tau \quad (59)$$

$$\Phi(t, 0) = \cosh((T-t)/\sqrt{\alpha}). \quad (60)$$

In the special case where the demand distribution is static, we can sharpen the result as follows.

Corollary 17 ([13]). *When $\zeta(t) = \zeta$ is constant in time, the solution to Problem 15 takes the form*

$$u(t) = -f(t)(\eta(t) - \zeta)/\alpha \quad (61)$$

where

$$f(t) = \sqrt{\alpha} \tanh((T-t)/\sqrt{\alpha}). \quad (62)$$

The trajectory generated by this control law is

$$\eta(t) = \Phi(t, 0)\eta_0 + (1 - \Phi(t, 0))\zeta \quad (63)$$

where

$$\Phi(t, 0) = \cosh((T-t)/\sqrt{\alpha}), \quad (64)$$

and the cost attained with this control law is

$$\mathcal{J}^1(\eta_0, \zeta; \alpha; T) = (\zeta - \eta_0)^2 \sqrt{\alpha} \tanh(T/\sqrt{\alpha}) / 2. \quad (65)$$

REFERENCES

- [1] S. Bandyopadhyay, S.-J. Chung, and F. Y. Hadaegh, "Probabilistic swarm guidance using optimal transport," in *2014 IEEE Conference on Control Applications (CCA)*, pp. 498–505, IEEE, 2014.
- [2] V. Krishnan and S. Martínez, "Distributed online optimization for multi-agent optimal transport," *arXiv preprint arXiv:1804.01572*, 2018.
- [3] D. Inoue, Y. Ito, and H. Yoshida, "Optimal transport-based coverage control for swarm robot systems: Generalization of the voronoi tessellation-based method," *IEEE Control Systems Letters*, vol. 5, no. 4, pp. 1483–1488, 2021.
- [4] M. Fornasier and F. Solombrino, "Mean-field optimal control," *ESAIM: Control, Optimisation and Calculus of Variations*, vol. 20, no. 4, pp. 1123–1152, 2014.
- [5] J. A. Carrillo, Y.-P. Choi, and M. Hauray, "The derivation of swarming models: mean-field limit and wasserstein distances," in *Collective dynamics from bacteria to crowds*, pp. 1–46, Springer, 2014.
- [6] B. Bonnet and F. Rossi, "The pontryagin maximum principle in the wasserstein space," *Calculus of Variations and Partial Differential Equations*, vol. 58, no. 1, pp. 1–36, 2019.
- [7] B. Bonnet, "A pontryagin maximum principle in wasserstein spaces for constrained optimal control problems," *ESAIM: Control, Optimisation and Calculus of Variations*, vol. 25, no. 1, p. 52, 2019.
- [8] B. Bonnet and H. Frankowska, "Necessary optimality conditions for optimal control problems in wasserstein spaces," *Applied Mathematics & Optimization*, vol. 84, no. 2, pp. 1281–1330, 2021.
- [9] M. Bongini, M. Fornasier, F. Rossi, and F. Solombrino, "Mean-field pontryagin maximum principle," *Journal of Optimization Theory and Applications*, vol. 175, no. 1, pp. 1–38, 2017.
- [10] W. Gangbo, T. Nguyen, and A. Tudorascu, "Hamilton-jacobi equations in the wasserstein space," *Methods and Applications of Analysis*, vol. 15, no. 2, pp. 155–184, 2008.
- [11] C. Jimenez, A. Marigonda, and M. Quincampoix, "Optimal control of multiagent systems in the wasserstein space," *Calculus of Variations and Partial Differential Equations*, vol. 59, no. 2, pp. 1–45, 2020.
- [12] M. Burger, R. Pinnau, C. Totzeck, and O. Tse, "Mean-field optimal control and optimality conditions in the space of probability measures," *SIAM Journal on Control and Optimization*, vol. 59, no. 2, pp. 977–1006, 2021.
- [13] M. Emerick, S. Patterson, and B. Bamieh, "Optimal combined motion and assignments with continuum models," *IFAC-PapersOnLine*, vol. 55, no. 13, pp. 121–126, 2022.
- [14] M. Emerick, "Control of continuum swarm systems via optimal control and optimal transport theory," Master's thesis, UC Santa Barbara, 2022.
- [15] F. Santambrogio, "Optimal transport for applied mathematicians," *Birkhäuser, NY*, vol. 55, no. 58-63, p. 94, 2015.
- [16] D. Limón, I. Alvarado, T. Alamo, and E. F. Camacho, "Mpc for tracking piecewise constant references for constrained linear systems," *Automatica*, vol. 44, no. 9, pp. 2382–2387, 2008.
- [17] F. Otto, "The geometry of dissipative evolution equations: the porous medium equation," 2001.
- [18] L. Ambrosio, N. Gigli, and G. Savaré, *Gradient flows: in metric spaces and in the space of probability measures*. Springer Science & Business Media, 2005.