Title
Control of Continuum Swarm Systems via Optimal Control and Optimal Transport Theory

Author
Emerick, Max

Publication Date
2022

Peer reviewed|Thesis/dissertation

University of California
Santa Barbara

Control of Continuum Swarm Systems via Optimal Control and Optimal Transport Theory

A thesis submitted in partial satisfaction
of the requirements for the degree

Master of Science
in
Mechanical Engineering

by

Max Joseph Emerick

Committee in charge:

Professor Bassam Bamieh, Chair
Professor Jeff Moehlis
Professor Francesco Bullo

September 2022

The thesis of Max Joseph Emerick is approved.

Professor Jeff Moehlis

Professor Francesco Bullo

Professor Bassam Bamieh, Committee Chair

September 2022

Control of Continuum Swarm Systems via Optimal Control and Optimal Transport Theory

by

Max Joseph Emerick

Acknowledgements

Thank you first and foremost to my advisor and mentor, Bassam Bamieh. Thank you to my committee members, Bassam Bamieh, Francesco Bullo, and Jeff Moehlis, and to my collaborators, Bassam Bamieh and Stacy Patterson. Thank you to my family, friends, and coworkers for the constant support, and particularly to those who provided feedback on my work: Connor Hughes, Daniella Walter, Alexander Davydov, Shannon Emerick, Craig Emerick, Karthik Chikmagalur, Poorva Shukla, and Pascal Ackermann-Karam. Thank you to UC Santa Barbara and the Department of Mechanical Engineering for the financial support and for the opportunity to work on this research. Last but not least, thank you to all of my teachers and mentors, both at UC Santa Barbara and at prior institutions, and in particular Eric Brussel and Russell Westphal. I would not be here if it were not for all of you.

Abstract

Control of Continuum Swarm Systems via Optimal Control and Optimal Transport Theory

by

Max Joseph Emerick

We consider problems of optimal motion control for multi-agent systems where assignments as well as motions have costs. In particular, we consider a demand distribution and a distribution of resource agents, which require and provide support respectively. We formulate a time-varying assignment problem which trades off two typically competing costs, namely an assignment cost which depends on distances between resource and demand, and a motion cost associated with moving resource agents to locations with lower assignment costs. We use the formalism of optimal transport theory, and the Wasserstein distance in particular, to capture assignment cost, while motion cost is captured by vehicular velocities over time. Both particle and continuum models for large-scale systems are considered, leading to infinite-dimensional nonlinear optimal control problems in general. We show how in the special case of one spatial dimension, the optimal control problem can be converted into an infinite dimensional Linear-Quadratic tracking problem by reparameterizing in terms of quantile functions, which can then be converted into a family of decoupled scalar Linear-Quadratic tracking problems. An analytic solution is provided in the general case. We investigate further two special cases where the demand distribution is static and where it is periodic in time. Explicit results and simulations are provided in each of these cases as well.

Table of Contents

Introduction

Low-cost sensing, processing, and communication hardware is driving the use of autonomous swarms of robotic agents in diverse settings, including emergency response, transportation, logistics, data collection, and defense [1, 2]. Large swarms can have significant advantages in efficiency and robustness, but as they scale in size, it becomes increasingly difficult to plan and coordinate motion. Thus, the development of effective motion planning and control strategies remains a problem of central importance.

A common problem setting in swarm control is that of *coverage* or *deployment*, where a swarm is deployed within some region and must move around in order to approximate some target distribution. These problems are well-studied in the literature [3] but have traditionally focused on swarms of *discrete* agents. However, as the number of agents grows large, it can become intractable to even store the positions of all the agents. One solution to this issue is to model the swarm as a *distribution* over the domain, which emphasizes the state of the overall swarm over the states of individual agents, providing a significant model reduction.

Optimal transport theory has provided many useful tools towards this end, that is, for working with distributions in a physical space. In recent years, several approaches have been taken to swarm control using tools from optimal transport theory, including works by Bandyopadhyay, Krishnan, Inoue, and their respective collaborators [4, 5, 6, 7, 8, 9]. However, these approaches have been limited in their treatment of the objective and constraints: they seek to minimize the transport distance while assuming convergence to the target distribution. In a real-world setting, it is desirable to have a model that can accommodate more general objectives and constraints, especially since convergence to a target distribution is not always possible.

The necessary generalization is provided by *optimal control theory*. Optimal control theory has emerged as the predominant framework in which to treat control problems with general objectives and constraints, and has been applied to continuum swarm models in numerous approaches, including those by Fornasier, Ferrari, Elamvazhuthi, Bonnet, Jimenez, Burger, and their respective collaborators [10, 11, 12, 13, 14, 15]. However, most of these approaches focus on the analytic aspects of the problem (e.g., well-posedness, regularity), and while providing useful theoretical tools, do not suggest particular models or algorithms for simulating real swarms.

This work aims to bridge that gap. That is, to propose a specific model that can accommodate general objectives yet is tractable enough for a practical implementation. In this work, I propose such a model, investigate the behavior and consequences of this model in depth, and provide simulations to demonstrate some of these theoretical results. The main contributions of this work are thus

- The formulation of a novel model for swarm tracking control which can accommodate more general objectives than existing continuum models
- Demonstration of a transformation which allows for an analytic solution in the one-dimensional case and provides insight into the structure of the problem
- Explicit solutions, further analysis, and simulation results for the special cases where the demand distribution is static or periodic in time

In chapter 1, I present the background necessary to support the main results. I begin with a discussion of notation, then give an introduction to continuum swarm models, before developing several of some of the key results from optimal transport theory that will be used throughout this work.

In chapter 2, I motivate and develop the proposed mathematical model. Two models are actually presented: a *general model* and a *specific model*. The general model is abstract and has the flexibility to accommodate different objectives, while the specific model makes particular choices for these objectives so that the model can be solved and analyzed.

In chapter 3, I focus on solving the specific model in the one-dimensional case. The model is solved in two main steps. The first step is a reparameterization in terms of quantile functions, which transforms the problem to an infinite-dimensional Linear-Quadratic (LQ) tracking problem. The second step is a "decoupling" of this problem into a family of scalar LQ tracking problems. The scalar LQ tracking problem has been studied in the literature before [16] and admits an analytic solution which is reproduced in this work. The cases where the demand distribution is static or time-periodic are then analyzed further, where explicit solutions and simulation results are provided.

Permissions and Attributions

Much of the content of this work (and in particular of Chapters 2 and 3) is the result of a collaboration with Bassam Bamieh and Stacy Patterson which has previously appeared at The 9th IFAC Conference on Networked Systems (NecSys22) and in IFAC-PapersOnLine [17]. The material is reproduced here under the IFAC Article Sharing Policy (`https://www.ifac-control.org/publications/copyright-conditions`), and the original article can be accessed via its DOI (`https://doi.org/10.1016/j.ifacol.2022.07.246`). To the material in that original article, I have added additional background, proofs, figures, and simulations, so that this work is now more readable and mostly self-contained. I have also extended the work slightly with additional results, interpretations of key findings, and the inclusion of the periodic case.

Chapter 1

Background

This chapter develops the background necessary to support the main results of this work, which are the focus of Chapters 2 and 3. We begin with a brief explanation of notation, discuss continuum swarm models, and then present several of the key results from optimal transport theory.

1.1 Notation

The material in this document relies heavily on measure-theoretic concepts. However, as the intended purpose of this work is to propose a model for engineering analysis and the main audience is largely composed of engineers, I will avoid talking about measures as much as possible, electing to talk about *distributions* instead. To many in this discipline, this is a much more intuitive concept, but the reader should be assured that this clarity does not come at the expense of rigor. It is a misconception that distributions are mathematically imprecise objects. For a formal development of distribution theory, see Gelfand and Shilov [18].

While reading this material, it is helpful to be able to transition smoothly between thinking of functions as maps and thinking of functions as points in function space. The notation used intends to emphasize these different perspectives in different contexts. For example, the notations R or $R(\cdot,\cdot)$ are used interchangeably, and refer to a whole function as an object, while notation like $R(x,t)$ refers to the evaluation of that function at the point (x,t). Also, notations like $R_t(\cdot)$ or $R(\cdot,t)$ are used interchangeably to suggest R either as a parameterized curve in function space or as a spatio-temporal field.

It is also important to point out that maps can act on different objects in different ways. For example, the flow map (1.1) can act on points by evaluation, maps by composition, or distributions by pushforward. These will all be defined later, but the important point is that these different actions are each given their own notation as well.

One last point is that there are a few places where standard notations conflict. For example, the star (*) is standard both for the adjoint and for the minimizing argument of a function. Here, star will be reserved to mean the adjoint and so an overbar ($\bar{\ }$) will be used for a minimizer instead. The following table describes all of the notation that I have ultimately settled upon, including some special variables that have been reserved for a particular purpose.

∂_x	–	the partial derivative with respect to x
D	–	the demand distribution, considered an external input
$\mathcal{D}$	–	the derivative of a mapping (e.g. Jacobian, Fréchet derivative)
$\mathfrak{D}(\Omega)$	–	the space of normalized distributions over Ω
f, g, h	–	scalar-valued functions of a scalar variable
F_μ	–	the cumulative distribution function of a distribution μ
$\mathcal{F}$	–	the Fourier transform of a function
$\mathcal{I}$	–	the identity operator
$\mathcal{J}$	–	an objective function

$\mathcal{K}$	–	a transport plan
$L^p_n(\Omega)$	–	L^p space of n-vector-valued functions over Ω
$\mathcal{M}$	–	a transport map
$\mathcal{N}$	–	the normal distribution
$\mathcal{P}$	–	a partition
Q_μ	–	the quantile function of a distribution μ
R	–	the resource distribution, considered the state
T	–	the final time in a time interval
u, U, v, V	–	velocity, considered a control input
$\mathcal{U}$	–	the uniform distribution over unit volume $[0,1]^n \subset \mathbb{R}^n$
$\mathcal{W}_p(\mu,\nu)$	–	p-Wasserstein distance between distributions μ and ν
$\mathbb{W}_p(\Omega)$	–	the p-Wasserstein space of distributions over Ω
x, y, z	–	spatial variables
α	–	a (constant) trade-off parameter
$\Gamma(\mu,\nu)$	–	the Wasserstein geodesic between distributions μ and ν
γ	–	a distribution in a curve or geodesic
δ	–	the Dirac distribution
ζ, η	–	position variables
μ, ν	–	arbitrary distributions
ξ	–	a dummy variable
Π_x	–	projection operator onto x
φ	–	a Kantorovich potential (Kantorovich dual variable)
ϕ, ψ	–	the flow of a vector field
Φ	–	a state-transition function
Ω	–	the domain, a compact convex subset of $\mathbb{R}^n$

∇	–	the gradient operator
$\lvert \cdot ' \rvert$	–	the metric derivative
$\circ$	–	the composition operator
$\#$	–	the pushforward operator
$\lvert \cdot \rvert$	–	the absolute value of a scalar
$\lvert\lvert \cdot \rvert\rvert$	–	the norm of a vector
$\times$	–	the cartesian product
$*$	–	the adjoint of an operator
$\langle \cdot, \cdot \rangle$	–	the inner product
$\bar{x}$	–	the minimizing argument x of an objective function

1.2 Continuum Swarm Models

Suppose we had a small swarm of k (discrete) robotic agents, labeled 1 to k, that roamed around on some subset of n-dimensional space. For simplicity, let's say that the state of each agent could be fully described by its position x_i, and so the state of the whole swarm could be described by a vector $X = (x_1, \ldots, x_k)$, where each $x_i \in \mathbb{R}^n$. As an input, we could specify the velocity of each agent as v_i, with the dynamics being $\dot{x}_i = v_i$, so that the input to the whole swarm was $V = (v_1, \ldots, v_k)$, with dynamics given by $\dot{X} = V$. This all works well for small swarms, but as the number of agents k increases to be become very large, it starts to become intractable to even store the positions of all of the agents. In this case, we gain an enormous model simplification by representing the state of the swarm by a distribution over the domain which describes the *density* of agents in each neighborhood. This effectively discards the *microscopic* information about each agent such as its state and identity, while retaining the *macroscopic* information describing the state of the swarm as a whole.

In the small-swarm case, we were able to describe the input to the swarm as the union of the inputs to each individual agent, but we are no longer able to do that for the large swarm, as we discarded the identities of each individual agent when we described the swarm as a distribution. The remedy is to instead specify the velocity input as a function of position rather than as a function of identity, so that instead of specifying $\dot{x}_i = v_i$ for agent i, we specify $\dot{x}(t) = v(x(t))$ for the agent at location x. However, this is confusing notation, as x is being used both as a coordinate and as a position variable. To avoid this sort of confusing notation, we define *the flow*.

Definition 1.1 (Flow of a Vector Field)**.** Let v be a vector field defined on $\mathbb{R}^n \times \mathbb{R}$. The *flow* of the vector field v is a map $\phi : \mathbb{R}^n \times \mathbb{R} \to \mathbb{R}^n$ satisfying

$$\begin{aligned} \partial_t \phi(x,t) &= v(\phi(x,t),t) \\ \phi(x_0, 0) &= x_0 \end{aligned} \tag{1.1}$$

By defining the flow, we reserve x as a coordinate by using ϕ for the position variable. The symbol $\phi(x_0, t)$ represents the position at time t of the agent that started from x_0 at time 0. Note that although we no longer keep track of the *identities* of individual agents, we can still *follow* individual agents over time via the flow. The flow gives us the swarm dynamics in the so-called *Lagrangian* perspective of fluid mechanics (that is, following individual agents as their position changes). There is also an expression for the swarm dynamics in the *Eulerian* viewpoint, describing the time-evolution of the entire density field. The "Eulerian" dynamics are given by the so-called *transport equation* (sometimes also referred to as the *advection equation* or *continuity equation*).

Definition 1.2 (Transport Equation)**.** Let μ_0 be an initial distribution defined on $\mathbb{R}^n$ and v be a velocity field defined on $\mathbb{R}^n \times \mathbb{R}$. The *transport equation* is the following partial differential equation

$$\partial_t \mu_t = -\nabla \cdot (v_t \mu_t). \tag{1.2}$$

Lemma 1.3. *Let μ_0 be an initial distribution defined on $\mathbb{R}^n$ and v be a velocity field defined on $\mathbb{R}^n \times \mathbb{R}$ that acts on particles in the distribution according to the dynamics* (1.1)*. The dynamics governing the resulting evolution of the distribution μ are given by the transport equation* (1.2)*.*

Proof. This is a foundational result in continuum mechanics – see e.g. Section 5.3 in Mase [19]. □

Equations (1.1) and (1.2) are the foundation of the model that we will use throughout this work.

1.3 Optimal Transport Theory

Optimal transport theory has provided many useful tools for modeling and analyzing the motions of distributions in a physical space. This section reviews many of the key results from this field that we will use throughout this work. We proceed rather informally and without proof, attempting to build a little intuition rather than give a rigorous development. Indeed, optimal transport theory can be a technical discipline and a rigorous development is well-outside the scope of this work. For references on the material in this section, one can consult any of the classic texts on the subject by Santambrogio [20], Ambrosio [21], or Villani [22, 23].

In short, the optimal transport problem seeks to find the most efficient way to transform one distribution of mass into another. The situation is formalized as follows. Suppose that we have an initial distribution of mass μ, where $\mu(x)$ specifies the density of mass at location x, and we wish to transform it into a target distribution ν, with a density $\nu(y)$ (we assume that the two distributions have the same total mass). One way to do this is by specifying a map $\mathcal{M} : x \mapsto y$ which takes the distribution μ and creates a new distribution by taking the mass from location x and moving it to location y. This new distribution is called the *pushforward distribution* and is defined as follows.

Definition 1.4 (Pushforward Distribution)**.** Let μ be a distribution defined on some subset of $\mathbb{R}^n$ and let $\mathcal{M} : \mathbb{R}^n \to \mathbb{R}^m$ be a map. Then the *pushforward* of μ by $\mathcal{M}$, denoted $\mathcal{M}_{\#}\mu$, is a distribution on a subset of $\mathbb{R}^m$ such that

$$\int_A (\mathcal{M}_{\#}\mu)(y)\, dy = \int_{\mathcal{M}^{-1}(A)} \mu(x)\, dx \tag{1.3}$$

for any measurable set $A \subset \mathbb{R}^m$, where $\mathcal{M}^{-1}$ denotes the preimage.

Figure 1.1 shows a pictorial representation of a pushforward distribution.

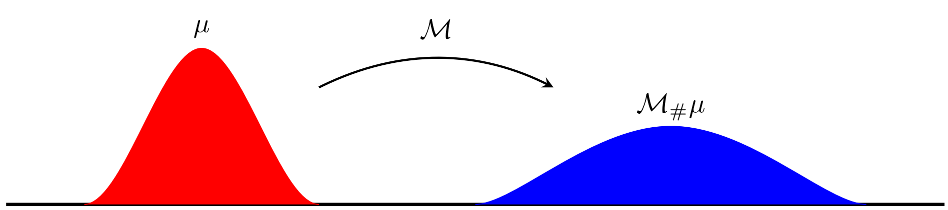

Figure 1.1: Pictorial representation of a pushforward distribution. The function $\mathcal{M}$ transforms μ into $\mathcal{M}_{\#}\mu$ by specifying where the mass at each point gets sent to.

Perhaps the most straightforward way to formalize the optimal transport problem is to define a function $c(x, y)$ representing the cost of moving mass from location x to location y, and then minimize the total cost over all maps $\mathcal{M}$ that transport μ to ν. This formulation is often referred to as the *Monge Problem*.

Definition 1.5 (Monge Problem)**.** Let μ and ν be two distributions defined on $\Omega \subset \mathbb{R}^n$ and $c(x, y)$ be a cost function. The *Monge Problem* is to find the map of minimum cost transporting μ to ν.

$$\begin{aligned} \bar{\mathcal{M}} \;&=\; \operatorname*{argmin}_{\mathcal{M}} \int_{\Omega} c\,(x, \mathcal{M}(x))\, \mu(x)\, dx \\ &\text{s.t.} \quad \mathcal{M}_{\#}\mu = \nu \end{aligned} \tag{1.4}$$

The minimizer $\bar{\mathcal{M}}$ is termed the *optimal transport map* and the value of the objective is termed the *transport cost* and is written $\mathcal{J}$.

While the Monge problem plays an important role in optimal transport theory, it is difficult to solve. It is nonconvex, and solutions are not even guaranteed to exist[1]. These issues are remedied by another formulation, the so-called *Kantorovich Problem*. Instead of seeking a map between the distributions, the Kantorovich problem seeks a more general coupling called a *plan*.

Definition 1.6 (Kantorovich Problem)**.** Let μ and ν be two distributions defined on $\Omega \subset \mathbb{R}^n$ and $c(x, y)$ be a cost function. The *Kantorovich Problem* is to find the plan of minimum cost transporting μ to ν.

$$\begin{aligned} \bar{\mathcal{K}} \;&=\; \operatorname*{argmin}_{\mathcal{K}} \int_{\Omega\times\Omega} c(x, y)\, \mathcal{K}(x, y)\, dx\, dy \\ &\text{s.t.} \quad \Pi_{x\#}\mathcal{K} = \mu \\ &\qquad\;\; \Pi_{y\#}\mathcal{K} = \nu \end{aligned} \tag{1.5}$$

where $\Pi_x : (x, y) \to x$ and $\Pi_y : (x, y) \to y$ are the projections onto x and y, respectively. The minimizer $\bar{\mathcal{K}}$ is termed the *optimal transport plan* and the value of the objective is again termed the *transport cost* and written $\mathcal{J}$.

[1]For example, consider μ to be a Dirac distribution. Since the mass in μ can only be sent to one location, it is obvious that there is no solution unless ν is also a Dirac distribution.

Whereas the Monge problem specifies *where* the mass at each point x is transported, the Kantorovich problem specifies *how much* mass is transported from each point x to each point y. In this way, the Kantorovich problem is more general, because one is allowed to “split mass” by sending it to multiple locations. A transport plan $\mathcal{K}$ can be interpreted as a joint distribution having the source and target distributions as marginals. A visual representation of a transport plan is shown in Figure 1.2.

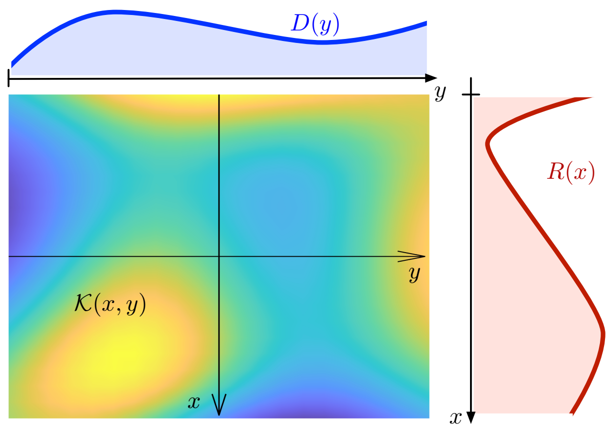

Figure 1.2: A Kantorovich transport plan between distributions R and D. The heat map shows the density of the transport plan, which determines how much mass at location x in R gets mapped to location y in D.

This formulation yields a number of advantages over the Monge problem. First, the Kantorovich problem consists of a linear objective function under linear constraints and is thus an infinite-dimensional linear program, which can be solved by established methods. Second, solutions are always guaranteed to exist. Third, if the Monge problem is solvable, its solutions can be recovered from the solution to the Kantorovich problem as follows.

Lemma 1.7. *When μ is absolutely continuous, the optimal transport plan $\bar{\mathcal{K}}$ solving* (1.5) *takes the form $(\mathcal{I}, \bar{\mathcal{M}})_{\#}\mu$, where $\mathcal{I}$ is the identity map on x and $\bar{\mathcal{M}}$ is the optimal transport map solving* (1.4).

Proof. See Theorem 1.17 in Santambrogio [20]. □

Fourth, when the cost function takes the form $c(x,y) = ||y-x||_p^p$, the resulting optimal transport cost forms a metric on the space of distributions, called the *p-Wasserstein distance.*

Definition 1.8 (p-Wasserstein Distance)**.** Let μ, ν be two distributions defined on $\Omega \subset \mathbb{R}^n$. The *p-Wasserstein distance* between μ and ν, written $\mathcal{W}_p(\mu,\nu)$, is given by

$$\begin{aligned} \mathcal{W}_p^p(\mu,\nu) \;:=\; & \min_{\mathcal{K}} \int_{\Omega\times\Omega} ||y-x||_p^p \, \mathcal{K}(x,y)\, dx\, dy \\ \text{s.t.} \quad & \Pi_{x\#}\mathcal{K} = \mu \\ & \Pi_{y\#}\mathcal{K} = \nu \end{aligned} \tag{1.6}$$

This quantity satisfies all the requirements of a metric, and turns $\mathfrak{D}(\Omega)$, the space of normalized distributions over Ω, into a complete metric space, which we call p-Wasserstein space[2] and denote $\mathbb{W}_p$. This proves to be extremely useful, as it endows the space of distributions with a rich geometry and provides many additional tools for solving problems. We will describe some aspects of this geometry shortly, but first, we introduce one more formulation of the optimal transport problem.

Since the Kantorovich problem is a linear program, it also has a dual formulation. After applying linear duality to (1.5) and eliminating a free variable using the convex conjugate transform, we obtain the following formulation.

[2] In this work, we will focus exclusively on the case $p = 2$.

Definition 1.9 (Dual Problem)**.** Let μ and ν be two distributions defined on $\Omega \subset \mathbb{R}^n$ and $c(x, y)$ be a cost function. The *Dual Problem* is to find the dual variable $\bar{\varphi}$ maximizing the objective

$$\begin{aligned} \bar{\varphi} \;&=\; \underset{\varphi}{\operatorname{argmax}} \int_{\Omega} \varphi(x)\mu(x)\, dx \;+\; \int_{\Omega} \varphi^c(y)\nu(y)\, dy \\ &\text{s.t.} \quad |x|^2/2 \;-\; \varphi(x) \text{ is convex} \end{aligned} \tag{1.7}$$

where $\varphi^c(y) := \inf_x c(x, y) - \varphi(x)$ is the convex conjugate of φ. The dual variables φ and φ^c are often referred to as *Kantorovich potentials.*

Solving the problems (1.5) and (1.7) are equivalent in the following sense.

Lemma 1.10. *The maximum value obtained in* (1.7) *is equal to the minimum value obtained in* (1.5)*. Furthermore, when μ is absolutely continuous and $c(x, y) = ||y - x||_2^2$, the solutions to the two problems are related by $\bar{\mathcal{K}} = (\mathcal{I}, \bar{\mathcal{M}})_{\#}\mu$ where*

$$\bar{\mathcal{M}}(x) \;=\; \nabla \left(|x|^2/2 \;-\; \bar{\varphi}(x) \right) \;=\; x \;-\; \nabla\bar{\varphi}(x) \tag{1.8}$$

where $\bar{\varphi}$ is dual variable solving (1.7)*.*

Proof. See Section 1.3.1 and Theorem 1.40 in Santambrogio [20]. □

In particular, notice that the optimal transport map is equal to the gradient of a convex function. Indeed, we have the following characterization of optimal transport maps.

Lemma 1.11. *Under the assumptions in Lemma 1.10, a transport map $\mathcal{M}$ is optimal if and only if $\mathcal{M} = \nabla\Psi$ for some convex function Ψ.*

Proof. See Section 1.3.1 and Theorem 1.48 in Santambrogio [20]. The key connection is the concept of a *cyclically monotone set.* □

The dual formulation also allows us to obtain a formula for the derivative of the squared 2-Wasserstein distance with respect to one of its arguments.

Lemma 1.12. *Let μ and ν be two distributions defined on $\Omega \subset \mathbb{R}^n$ with μ absolutely continuous. Then the derivative of the squared 2-Wasserstein distance at μ with respect to fixed ν is given by*

$$\mathcal{D}_\mu \mathcal{W}_2^2(\mu, \nu)(\cdot) \;=\; \langle \bar{\varphi}, \cdot \rangle \tag{1.9}$$

where $\bar{\varphi}$ is the maximizing variable in the dual problem (1.7)*.*

Proof. See Proposition 7.17 in Santambrogio [20]. □

Having given a brief review of the main results on the optimization side of the theory, we now build up the geometric picture. Here, we shift our thinking from maps and plans relating distributions to continuous curves valued in $\mathbb{W}_2$. Recall that the velocity of a parameterized curve is formalized as the *tangent vector* to the curve.

Definition 1.13 (Tangent Vector)**.** Let μ_t be a continuous curve valued in $\mathbb{W}_2$ and parameterized by t. The *tangent vector* of μ_t at $t = \tau$ is defined to be

$$v_\tau \;:=\; \lim_{s \to 0} \frac{\mu_{\tau+s} - \mu_\tau}{s} \tag{1.10}$$

Since such a tangent vector represents a velocity field on $\mathbb{R}^n$, we will also refer to this as a tangent velocity field. We have the following result.

Lemma 1.14. *Let* μ_t *be a continuous curve valued in* $\mathbb{W}_2$ *and parameterized by* t*. The tangent vector (or tangent velocity field) of* μ_t *at* $t = \tau$ *is given by*

$$\begin{aligned} \bar{v}_\tau \;&=\; \operatorname*{argmin}_{v_\tau} \int_\Omega ||v_\tau(x)||_2^2 \, \mu(x) \, dx \\ & s.t. \quad \partial_t \mu_t \Big|_{t=\tau} = -\nabla \cdot (v_\tau \mu_\tau) \end{aligned} \tag{1.11}$$

Proof. See Theorem 5.14 in Santambrogio [20]. □

That is, the tangent velocity field is the velocity field of minimum norm that generates the curve via the transport equation. The *speed* of the curve is formalized by the *metric derivative.*

Definition 1.15 (Metric Derivative)**.** Let μ_t be a continuous curve valued in $\mathbb{W}_2$ and parameterized by t. Then the *metric derivative* of the curve at $t = \tau$ is defined to be

$$|\mu'_\tau| \;:=\; \lim_{s \to 0} \frac{\mathcal{W}_2(\mu_\tau, \mu_{\tau+s})}{|s|} \tag{1.12}$$

As in Euclidean space, the speed of a curve is equal to the magnitude of its velocity.

Lemma 1.16. *Let* μ_t *be a continuous curve valued in* $\mathbb{W}_2$ *and parameterized by* t*. Then*

$$|\mu'_\tau|^2 \;=\; \int_\Omega ||v_\tau(x)||_2^2 \, \mu_\tau(x) \, dx \tag{1.13}$$

where v_τ *is the tangent vector at* $t = \tau$*.*

Proof. See Theorem 5.14 in Santambrogio [20]. □

Also as in Euclidean space, the length of a curve is equal to the integral of its speed.

Definition 1.17 (Length of a Curve)**.** Let μ_t be a continuous curve valued in $\mathbb{W}_2$ and parameterized by t. The *length* of the curve on the interval $[0, T]$ is defined to be

$$\text{length}(\mu_{[0,T]}) \ := \ \int_0^T |\mu_t'| \, dt \tag{1.14}$$

A critical result is that the space $\mathbb{W}_2$ is a *geodesic space*. This means that there exists a curve achieving a minimum length between any two distributions, and that this minimum length is equal to the distance between the two distributions. This curve is called a *geodesic*.

Definition 1.18 (Wasserstein Geodesic)**.** The *2-Wasserstein geodesic* between distributions μ and ν is defined to be the curve

$$\Gamma(\mu, \nu) \ := \ \{\gamma \in \mathbb{W}_2 : \mathcal{W}_2(\mu, \gamma) + \mathcal{W}_2(\gamma, \nu) = \mathcal{W}_2(\mu, \nu)\}. \tag{1.15}$$

It is known that given distributions μ and ν, the geodesic between them exists and is unique. Also notice that since $\mathcal{W}$ is a metric, we know that $\text{length}(\Gamma) < \text{length}(\Theta)$ for any other continuous curve Θ with endpoints μ and ν.

It is common to parameterize the Wasserstein geodesic as $\Gamma = \{\gamma_t\}$ where t spans some interval in $\mathbb{R}$, most often $[0, 1]$. By abuse of terminology, γ_t will sometimes be referred to as the geodesic. However, note that while the geodesic itself (that is, the set of points) is unique, the parameterization certainly is not.

Given an optimal transport map, we know how to compute the geodesic.

Lemma 1.19. *Let μ and ν be two distributions defined on $\mathbb{R}^n$ with μ absolutely continuous, and suppose $\bar{\mathcal{M}}$ is an optimal transport map transporting μ to ν. Then a constant-speed parameterization of the Wasserstein geodesic between μ and ν is given by*

$$\begin{aligned} \gamma_t &= \mathcal{M}_{t\#}\mu \\ \mathcal{M}_t &:= (1-t)\mathcal{I} + t\,\bar{\mathcal{M}} \end{aligned} \tag{1.16}$$

for $t \in [0,1]$, where $\mathcal{I}$ is the identity map.

Proof. See Theorem 5.27 in Santambrogio [20]. □

This concludes our review of optimal transport theory.

Chapter 2

The Proposed Model

In this chapter, we motivate and develop our proposed model, which is the first main contribution of this work. There are actually two models we develop – a "general" model and a "specific" model. The general model (2.11) is an abstract model which has the flexibility to accommodate different objectives. The specific model (2.13) makes particular choices for these objectives so that the model is fully defined and can be solved and analyzed. We start by introducing the features which compose our model.

2.1 Problem Formulation

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 either locations, facilities, or independently-controlled mobile robots that require supplies or services (e.g. data collection, communication). Resource agents are mobile and are able to service the needs of multiple demand components. Physical space is modeled as a compact convex subset $\Omega \subseteq \mathbb{R}^n$ where n is a positive integer (typically 1, 2, or 3). Both the demand and resource distributions are modeled as density functions over the domain, where discrete agents are modeled by Dirac distributions and continuum swarms are modeled by continuous distributions. In the most general setting, our model can accommodate both. The following are the five ingredients in our model.

(1) A *demand distribution*, $D_t(\cdot)$, defined over $\Omega \subset \mathbb{R}^n$, describing the state of the demand at time t which is of the form

$$D_t(x) \;=\; d(x,t) \;+\; \sum_{k=1}^{N_d} d_k(t)\;\delta\big(x-\zeta_k(t)\big). \tag{2.1}$$

In the purely discrete setting, the first term is zero, and the distribution represents N_d discrete demand components, each with demand $d_k(t)$ located at position $\zeta_k(t) \in \Omega$. In the purely continuum case, the second term is zero, and $d(\cdot,t) : \Omega \to \mathbb{R}$ is a non-negative function describing the density of demand. The expression above allows for a mixture of continuum and discrete components. The continuum model may be appropriate to describe a very large number of components, where only the spatial distribution, rather than the identity of each component, matters.

In classic coverage control problems, D is usually assumed to be constant in time, but we could choose to model D as time-varying, stochastic, or even having its own state-dependent dynamics, as in the case where the resource "satisfies" the demand and subsequently reduces it[3]. Here, we focus on the static and time-varying cases, and the stochastic and dynamically-coupled extensions will be presented elsewhere.

[3] It should be noted that this latter case is considerably more complex. In particular, either or both of the distributions may have mass that is changing in time, in which case a notion of "nonuniform optimal transport" can be applied (see e.g. [24]) as opposed to the setup which proceeds in this work.

(2) A *resource distribution*, $R_t(\cdot)$, also defined over $\Omega \subset \mathbb{R}^n$, describing the state of the resource agents at time t

$$R_t(x) \;=\; r(x,t) \;+\; \sum_{k=1}^{N_r} r_k(t)\; \delta\big(x - \eta_k(t)\big), \tag{2.2}$$

with similar interpretations of each term as those given to D_t in (2.1). In this paper, we assume that the total mass of each distribution is constant in time, and so without loss of generality, we assume that both distributions are normalized to integrate to 1.

$$\int_\Omega R_t(x)\,dx \;=\; \int_\Omega D_t(x)\,dx \;=\; 1 \tag{2.3}$$

A visual representation of the two distributions is shown in Figure 2.1.

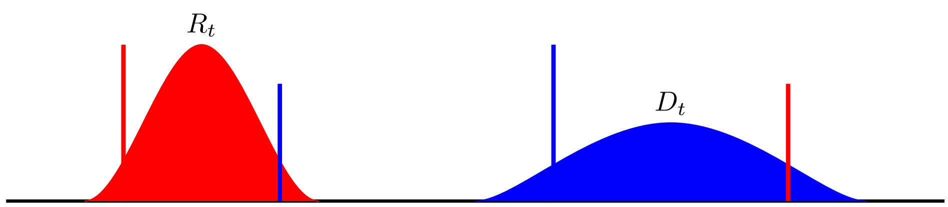

Figure 2.1: Resource distribution (red) and demand distribution (blue) over physical space.

(3) An *assignment kernel*, $\mathcal{K}_t(\cdot,\cdot)$. This kernel is a non-negative, scalar-valued (or possibly generalized) function where $\mathcal{K}_t(x,y)$ specifies the amount of service that resource agent at location x provides to demand component at location y at time t. Thus in the most general case, each resource agent can service multiple demand components, and similarly each demand component can be serviced by multiple resource agents.

With the distributions D_t and R_t normalized, the interpretation of the two-variable function $\mathcal{K}_t(\cdot,\cdot)$ as assignment of demands to resources gives the following "marginalization property"

$$R_t = \Pi_{x\#}\mathcal{K}_t, \qquad D_t = \Pi_{y\#}\mathcal{K}_t. \tag{2.4}$$

where Π_x and Π_y are the projection operators onto x and y respectively. When either R_t or D_t contain generalized functions, then the kernel $\mathcal{K}_t$ will include generalized functions as well. These two cases are shown in Figures 2.2 and 2.3.

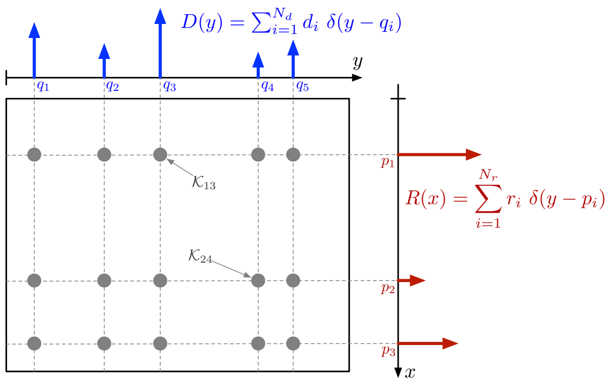

Figure 2.2: Example of a discrete assignment kernel.

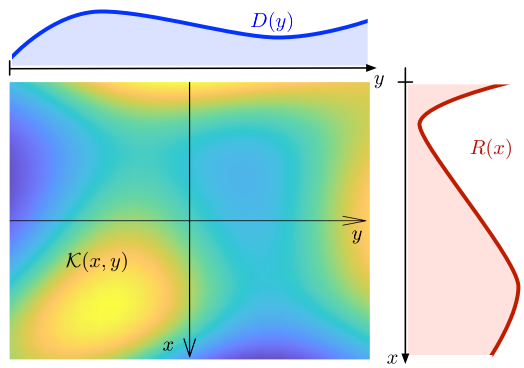

Figure 2.3: Example of a continuous assignment kernel.

The assignment kernel is one of the decision variables in our optimal control problem. Each assignment kernel incurs an assignment cost which will be defined shortly. The marginalization requirements (2.4) can be thought of as *constraints* that any valid assignment kernel must satisfy.

(4) A *dynamic model*, describing the equations of motion of the system and the role of the input function. Motivated by the need to consider large-scale continuum models of agents, we use the following *transport equation* to describe how a density moves in time

$$\partial_t R(x,t) = -\sum_{i=1}^{n} \partial_{x_i}\Big(v_i(x,t)\ R(x,t)\Big) = -\nabla \cdot \Big(V(x,t)\ R(x,t)\Big), \tag{2.5}$$

where $V_t(x) = V(x,t) := [v_1(x,t) \cdots v_n(x,t)]^*$ is a time-varying velocity field. The velocity field is one of the decision variables in our optimal control problem. The optimal velocity field "steers" the resource agents to move over optimal trajectories. We assume that $V = 0$ on the boundary of Ω and that V is continuous except on sets where R_t has vanishing mass to ensure existence and uniqueness of solutions on the support of R_t. Notice also there there is an *implicit constraint* in the fact that any two particles at the same location must move with the same velocity.

The transport equation (2.5) is used for both continuum and discrete models. Note that this velocity field is defined as a function of space, thus the velocity of each agent is determined by its location in space rather than its identity. Figure 2.4 shows a visual representation of how a velocity field acts to transport a distribution.

Figure 2.4: A velocity field V (green) acting to transport a resource distribution R (red).

(5) An *objective function*, $\mathcal{J}$, which trades off two competing cost components. The first is an *assignment cost* $\mathcal{C}_a$, which is a function of the assignment kernel and reflects that assignments between agents that are further apart are more costly. An example of such an assignment cost takes the form

$$\mathcal{C}_a(\mathcal{K}_t) \;=\; \langle C_a, \mathcal{K}_t \rangle \;:=\; \int_{\Omega\times\Omega} C_a(x,y)\, \mathcal{K}_t(x,y)\, dx\, dy, \tag{2.6}$$

where C_a is a "distance-like" function, that is, a monotonically increasing function of the "distance" between locations x and y. A mathematically convenient choice uses the squared Euclidean distance

$$C_a(x,y) \;:=\; ||y-x||_2^2. \tag{2.7}$$

The definition of the assignment cost (2.6) is "location aware" in the sense that assigning resources to demands that are far away incurs a high cost and vice versa. One interpretation of this is as the cost of "communication" between demand and assigned resource agents. In a simplistic communication model with no interference, the minimum transmission power required for error-free communication scales with the square of the distance, and therefore (2.6) with the choice (2.7) could be interpreted as the total communication power required for a given assignment $\mathcal{K}_t$.

In the static case where resource and demand distributions R and D are fixed in time, the optimal (static) assignment kernel $\bar{\mathcal{K}}$ is the one that minimizes (2.6) subject to the marginalization constraints (2.4), i.e.

$$\bar{\mathcal{K}} \;:=\; \underset{\substack{\Pi_{x\#}\mathcal{K}=R,\\ \Pi_{y\#}\mathcal{K}=D}}{\operatorname{argmin}} \int_{\Omega\times\Omega} \|y-x\|_2^2 \; \mathcal{K}(x,y)\; dx\, dy \tag{2.8}$$

Notice that the inner product structure of this cost function with the linear constraints makes this problem an infinite-dimensional linear program. When R and D are normalized, then this problem is precisely the *Kantorovich Problem* of optimal transport theory, where $\bar{\mathcal{K}}$ is the *optimal transport plan.* Furthermore, the minimum of the expression (2.8) is the squared *2-Wasserstein distance* $\mathcal{W}_2^2(R,D)$ between the distributions R and D.

We are interested however in a dynamic situation where both resource and demand distributions can be time varying. Clearly, if the resource distribution can freely move, then the optimal assignment is for the resource distribution to perfectly match the demand, resulting in a zero distance $\mathcal{W}_2^2(R_t,D_t)=0$ between them. To capture the idea that motion has its cost, we propose a cost for the motion of the resource field as

$$\mathcal{C}_m\big(R_t, V_t\big) \;=\; \int_\Omega \|V_t(x)\|_2^2 \, R_t(x) \; dx. \tag{2.9}$$

If we consider the resource agents to be flying drones, this quantity would be interpreted as the total aerodynamic drag on the swarm. Over a finite-time maneuver, the total cost of motion would be the combination of hover cost and the energetic cost of overcoming drag. However, since the hover cost is independent of maneuvers, it is the cost of overcoming drag that has the dominant effect on solutions.

We therefore propose the following combined cost function for maneuvers over a time horizon $[0, T]$

$$\mathcal{J}(\mathcal{K}, V) \;:=\; \int_0^T \Big(\mathcal{C}_a\big(\mathcal{K}_t\big) + \alpha\, \mathcal{C}_m\big(R_t, V_t\big)\Big)\, dt, \tag{2.10}$$

where $\alpha > 0$ is a "trade off" parameter. The two objectives $\mathcal{C}_a$ and $\mathcal{C}_m$ are clearly competing. If motion cost is negligible, then the optimal solution would be to move R_t quickly so that it matches D_t, and then $\mathcal{C}_a$ becomes small. However, if motion cost is expensive (i.e. large α), then the optimal solution would tolerate a high assignment cost while keeping motion cost small. The length of time T in which to carry out this maneuver will also be a factor in the tradeoff between the two costs.

2.2 Problem Statement

We finally present our model formally in its full generality.

Definition 2.1 (The General Model)**.** The *General Model* is defined to be the following problem. Given an initial resource distribution R_0, demand distribution D_t, weighting parameter α, time horizon T, and cost functions $\mathcal{C}_a$ and $\mathcal{C}_m$, solve

$$\begin{aligned} \{\bar{\mathcal{K}}, \bar{V}\} \;&=\; \operatorname*{argmin}_{\mathcal{K},V} \int_0^T \Big(\mathcal{C}_a\big(\mathcal{K}_t\big) + \alpha\; \mathcal{C}_m\big(R_t, V_t\big)\Big)\, dt \\ \text{s.t.}\quad & \Pi_{x\#}\mathcal{K}_t = R_t \\ & \Pi_{x\#}\mathcal{K}_t = D_t \\ & \partial_t R_t(x) = -\nabla \cdot \big(V_t(x)\; R_t(x)\big)\,. \end{aligned} \tag{2.11}$$

A *solution* of the problem is defined to be an optimal pair $\{\bar{\mathcal{K}}, \bar{V}\}$, and the *cost* is defined to be the value of the objective and is written $\mathcal{J}$.

Remark. We have not yet shown under what conditions the above model is well-posed, that is, that solutions exist, are unique, and depend continuously on initial conditions. Indeed, these problems can sometimes be subtle, and we have elected to leave these issues to later work.

Figures 2.5 and 2.6 show two different representations of the general model.

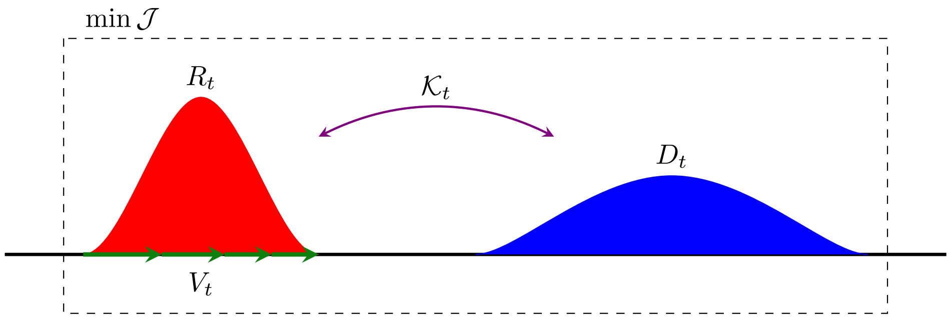

Figure 2.5: Pictorial representation of the general model. The resource distribution R (red) is paired to the demand distribution D (blue) by the assignment kernel $\mathcal{K}$ (purple) and is transported by the velocity field V (green). The objective $\mathcal{J}$ is minimized over the whole maneuver.

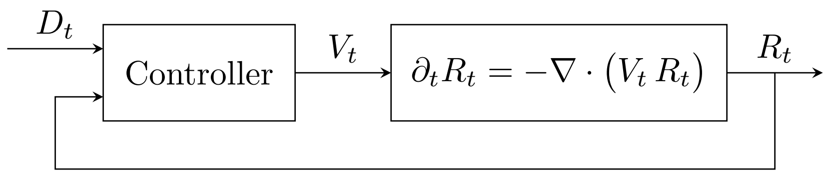

Figure 2.6: Block diagram representation of the general model. The controller sees the resource and demand distributions and specifies the optimal velocity field V to transport the resource distribution via the plant dynamics (i.e. the transport equation).

In this setting, the demand distribution D_t is considered as an external signal (which is known apriori), the resource distribution R_t is considered as the state, and the velocity field V_t as the control input. The assignment kernel $\mathcal{K}_t$ is also a decision variable, but in the cases which we examine in this paper, it plays a different role than V since it is an instantaneous (i.e. non-dynamic) variable[4]. This fact has an implication for the minimization over $\mathcal{K}$ that we state next.

[4]We remark that certain extensions may require that $\mathcal{K}_t$ be coupled to the dynamics, and so (2.11) describes the model in full generality, but we leave these extensions to later work and proceed to adapt this model to the cases presented.

The objective in problem (2.11) has two components, and since the motion and motion cost are independent of the assignment $\mathcal{K}$, we can split the optimization as follows

$$\inf_{V} \inf_{\substack{\Pi_{x\#}\mathcal{K}=R,\\ \Pi_{y\#}\mathcal{K}=D}} \int_0^T \Big(\mathcal{C}_a\big(\mathcal{K}_t\big) + \alpha\, \mathcal{C}_m\big(R_t, V_t\big)\Big)\, dt$$

$$= \inf_{V} \int_0^T \left(\left(\inf_{\substack{\Pi_{x\#}\mathcal{K}=R,\\ \Pi_{y\#}\mathcal{K}=D}} \mathcal{C}_a\big(\mathcal{K}_t\big)\right) + \alpha\, \mathcal{C}_m\big(R_t, V_t\big)\right) dt \quad (2.12)$$

Application of (2.6) with (2.7) makes the assignment cost equal to the 2-Wasserstein distance between R_t and D_t, and together with the choice of (2.9) for the motion cost we obtain the following "specific" model.

Definition 2.2 (The Specific Model)**.** The *Specific Model* is defined to be the following problem. Given an initial resource distribution R_0, demand distribution D_t, weighting parameter α, and time horizon T, solve

$$\begin{aligned} \bar{V} \;=\; & \operatorname*{argmin}_{V} \int_0^T \left(\mathcal{W}_2^2\big(R_t, D_t\big) + \alpha \int_\Omega \|V_t(x)\|_2^2\, R_t(x)\ dx\right) dt \\ & \text{s.t.} \quad \partial_t R_t(x) = -\nabla \cdot \big(V_t(x)\ R_t(x)\big) \end{aligned} \quad (2.13)$$

A *solution* of the problem is defined to be an optimal velocity field $\bar{V}$, and the *cost* is defined to be the value of the objective and is again written $\mathcal{J}$.

Since the quantity $\mathcal{W}_2^2(R_t, D_t)$ is a measure of distance between the state R_t and the value of the signal D_t, the first term in the objective can be thought of as a "tracking error". The second term can be interpreted as a "control cost", which in our context is actually the cost of motion of the resource agents.

With the interpretations given above, the optimal control problem (2.13) is analogous to an infinite-dimensional, generally nonlinear, "servo-mechanism" problem (see [16]). The necessary conditions for optimality for such problems lead to nonlinear two-point boundary value differential equations, which typically need to be solved numerically. However, at least when the spatial domain is one-dimensional, there is additional structure which allows us to simplify the problem considerably. We focus exclusively on this special case for the remainder of this work, while the higher-dimensional case will be presented elsewhere.

Chapter 3

The One-Dimensional Problem

In this chapter, we solve the specific model in the case where the domain is a subset of $\mathbb{R}$. We first provide some additional background for this case, then demonstrate two main results which simplify the problem considerably, before solving this simplified problem. In addition, two special cases are investigated: that where the demand is static, and that where the demand is periodic in time. Simulation results are included to support the analysis and to help provide visual intuition.

3.1 Transformation and Decoupling

The assignment cost component in (2.13) is expressed in terms of the 2-Wasserstein distance between two distributions. In general, there is not an explicit expression for this distance in terms of the distributions themselves. Rather, it must be found as the minimum of an optimization problem like (2.8). The problem (2.8) is an infinite-dimensional linear program, and some properties of the 2-Wasserstein distance can be obtained from this linear program and its dual. However, in the one-dimensional case (i.e. when $\Omega \subset \mathbb{R}$), there is an explicit expression for $\mathcal{W}_2^2(R_t, D_t)$ in terms of the *quantile* *functions* associated with these distributions. In this section, we show how to reformulate the problem (2.13) in the one-dimensional case using quantile functions as the state and external signals. It is interesting to note that with this reformulation, the infinite dimensional nonlinear optimal control problem (2.13) becomes a decoupled set of scalar Linear-Quadratic servomechanism problems, for which analytic solutions can be derived.

We begin with a well-known result about the 2-Wasserstein distance for distributions over one-dimensional space. First, recall the *cumulative distribution function* and the *quantile function* of a distribution.

Definition 3.1. Let μ be a distribution defined over $\Omega \subset \mathbb{R}$ with mass 1. The *cumulative distribution function (CDF)* $F_\mu : \Omega \to [0,1]$ and *quantile function* $Q_\mu : [0,1] \to \Omega$ of μ are defined by

$$F_\mu(x) \;:=\; \int_{\inf(\Omega)}^{x} \mu(\xi)\, d\xi, \tag{3.1}$$

$$Q_\mu(z) \;:=\; \inf\{x : F_\mu(x) \geq z\}. \tag{3.2}$$

Recall that F_μ and Q_μ are equivalent representations for the distribution μ, in that the associations are all 1-1. Recall also that in the case where F_μ is strictly increasing, the quantile function Q_μ is the *inverse function* of F_μ. We also have the additional relation

$$F_\mu(x,t) \;=\; \sup\{z : Q_\mu(z,t) \leq x\}. \tag{3.3}$$

Figures 3.1 and 3.2 show graphically the relationship between the three objects.

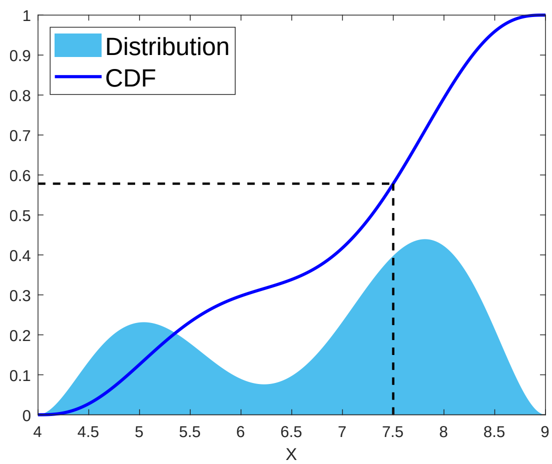

Figure 3.1: The CDF is the integral of the distribution.

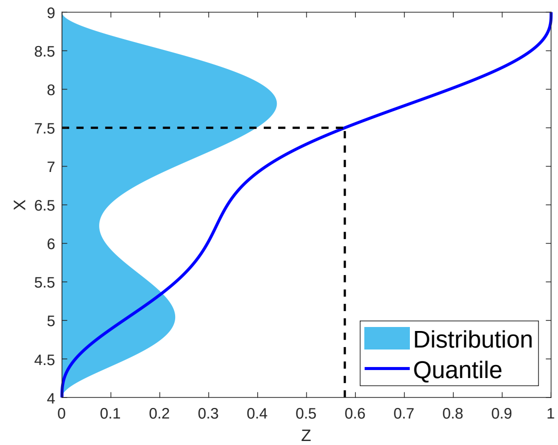

Figure 3.2: The quantile function is the inverse function of the CDF.

There is an explicit form for the Wasserstein distance in one dimension in terms of the quantile functions of the two distributions.

Lemma 3.2. *Let μ and ν be any two normalized distributions over $\Omega \subset \mathbb{R}$. The 2-Wasserstein distance between them is given by*

$$\mathcal{W}_2^2(\mu, \nu) \;=\; \int_0^1 (Q_\nu(z) - Q_\mu(z))^2 \, dz \tag{3.4}$$

where Q_μ, $Q_\nu : [0,1] \to \mathbb{R}$ are the quantile functions of the distributions μ and ν respectively.

Proof. The formula is well-known in the literature – see e.g. Proposition 2.17 in Santambrogio [20]. The critical fact is that the optimal assignment in 1D is monotone. □

Thus, in the one-dimensional case, the 2-Wasserstein distance between the distributions is exactly the L^2 distance between the respective quantile functions. Note that this implies that the space of normalized distributions on $\mathbb{R}$ equipped with $\mathcal{W}_2$ is isometric to the set of monotone nondecreasing functions on $[0, 1]$ equipped with the L^2 distance[5].

[5] However, this has strong implications (such as an inherited notion of projection) which are central

This gives a rather simple and familiar expression for the assignment cost in (2.13) in terms of quantile functions. The next natural question is whether the dynamics (i.e. the transport equation) also have a simple expression in terms of quantile functions. The answer is affirmative as we show next. First, we establish the evolution of the CDF.

Lemma 3.3. *Consider a one-dimensional distribution μ advected by (2.5). Its cumulative distribution F_μ evolves according to the non-uniform transport equation*

$$\partial_t F_\mu(x,t) \;=\; V(x,t)\;\partial_x F_\mu(x,t), \quad F_\mu(\inf(\Omega),t) = 0. \tag{3.5}$$

Proof. First, we have the property that $\mu(x,t) = \partial_x F_\mu(x,t)$. Substituting this into (2.5) yields

$$\partial_t \partial_x F_\mu(x,t) = \partial_x\big(V(x,t)\,\partial_x F_\mu(x,t)\big).$$

Integrating over x provides

$$\partial_t \int_{\inf(\Omega)}^{x} \partial_\xi F_\mu(\xi,t)\,d\xi = \int_{\inf(\Omega)}^{x} \partial_\xi\big(V(\xi,t)\,\partial_\xi F_\mu(\xi,t)\big)\,d\xi$$

and the fundamental theorem of calculus together with $F_\mu(\inf(\Omega),t) = 0$ gives

$$\partial_t F_\mu(x,t) \;=\; V(x,t)\;\partial_x F_\mu(x,t).$$

Note that the derivatives are all taken in the distributional sense. □

It turns out that the "bilinear dynamics" of R (2.5) and its CDF F_R (3.5) transform into *linear additive* dynamics of the quantile function Q_R. We state this precisely with the next definition and theorem.

to the special structure of the 1D case. This makes the geometry in 1D very different from the geometry in higher dimensions.

Definition 3.4 (Transformed Model)**.** The *Transformed Model* is defined to be the following problem. Given an initial state $Q_{R,0}$, signal $Q_{D,t}$, weighting parameter α, and time horizon T, solve

$$\begin{aligned} \bar{U} &= \min_{U} \int_0^T \int_0^1 \Big(\big(Q_D(z,t) - Q_R(z,t)\big)^2 + \alpha U^2(z,t) \Big)\, dz\, dt \\ \text{s.t.} \quad \partial_t Q_R(z,t) &= U(z,t), \\ \partial_z Q_R(z,t) &= 0 \quad \Rightarrow \quad \partial_z U(z,t) = 0, \end{aligned} \tag{3.6}$$

A *solution* of the problem is defined to be an optimal input $\bar{U}$, and the *cost* is defined to be the value of the objective and is again written $\mathcal{J}$.

The name "transformed model" as well as the suggestive choice of notation are justified by the following theorem.

Theorem 3.5. *The specific model (2.13) over one spatial dimension and the transformed model (3.6) are equivalent in that the solutions are 1-1 and attain the same cost. Specifically, the solutions are related by the equations*

$$U(z,t) = V\big(Q_R(z,t), t\big) \tag{3.7}$$

$$V(x,t) = U\big(F_R(x,t), t\big). \tag{3.8}$$

where F_R, F_D, and Q_R, Q_D are the CDFs and quantile functions of R, D respectively.

Proof. First, we wish to establish that $z = F_R(\phi(x,t), t)$ is constant, where ϕ is the flow of the vector field V. The defining property of the transport equation (2.5) is that it preserves the total mass of any transported volume element. More precisely, the total mass in any set A is preserved as A is transported by ϕ

$$\int_{\phi(A,t)} R(x,t)\ dx \ = \ \int_A R(x,0)\ dx.$$

In one dimension, the CDF value $F_R(x,t)$ is the total mass within the set $[\inf(\Omega), x]$. Such sets are transported as $[\inf(\Omega), \phi(x,t)]$ by the flow, and we can therefore conclude that $z = F_R(\phi(x,t),t) = \text{constant}$.

Now suppose that V satisfies the constraint in (2.13). Then we know

$$\partial_t \phi(x,t) \ = \ V(\phi(x,t),t).$$

Since $z = F_R(\phi(x,t),t)$ is constant, and Q_R and F_R are inverse functions, we have

$$Q_R(z,t) \ = \ Q_R(F_R(\phi(x,t),t),t) \ = \ \phi(x,t)$$

and therefore

$$\partial_t Q_R(z,t) \ = \ V(Q_R(z,t),t) \ =: \ U(z,t).$$

For the converse, the reasoning is almost the same, except that we require $U(z_1,t) = U(z_2,t)$ if $Q_R(z_1,t) = Q(z_2,t)$ so that $V(Q_R(z,t),t)$ is single-valued. Then V satisfies the constraint in (2.13) if and only if U satisfies the constraints in (3.6), and the solutions are 1-1.

To see that the costs are the same, first observe that (3.4) provides

$$\mathcal{W}_2^2(R_t, D_t) \ = \ \int_0^1 \left(Q_D(z,t) - Q_R(z,t)\right)^2 dz$$

and so it only remains to show that

$$\int_\Omega V^2(x,t)R(x,t)\,dx \;=\; \int_0^1 U^2(z,t)\,dz.$$

Applying the change of variables $z = F_R(x,t)$, $x = Q_R(z,t)$ to the instantaneous motion cost gives us

$$\int_\Omega V^2(x,t)R(x,t)\,dx \;=\; \int_0^1 V^2\Big(Q_R(z,t),t\Big)R\Big(Q_R(z,t),t\Big)\frac{\partial x}{\partial z}\,dz$$

and using $U(z,t) := V(Q_R(z,t),t)$ together with

$$\frac{\partial x}{\partial z} \;=\; \frac{1}{\frac{\partial z}{\partial x}} \;=\; \frac{1}{\partial_x F_R(x,t)} \;=\; \frac{1}{R(x,t)} \;=\; \frac{1}{R\big(Q_R(z,t),t\big)}$$

we have

$$\int_\Omega V^2(x,t)R(x,t)\,dx \;=\; \int_0^1 U^2(z,t)\,dz.$$

Thus the costs are identical, completing the proof. Again, note that the derivatives are taken in the distributional sense. □

The problem above is a distributed control problem with $Q_{R,t} : [0,1] \to \mathbb{R}$ as the state, and the signal $U_t : [0,1] \to \mathbb{R}$ as the new control input. The first constraint expresses the equivalent dynamics of the transformed model. Recall that in the original dynamics (2.5) there was an implicit constraint that particles at the same location had to move with the same velocity. This implicit constraint has now become explicit in the transformed model as the second constraint. A graphical representation of this transformation is shown in Figure 3.3.

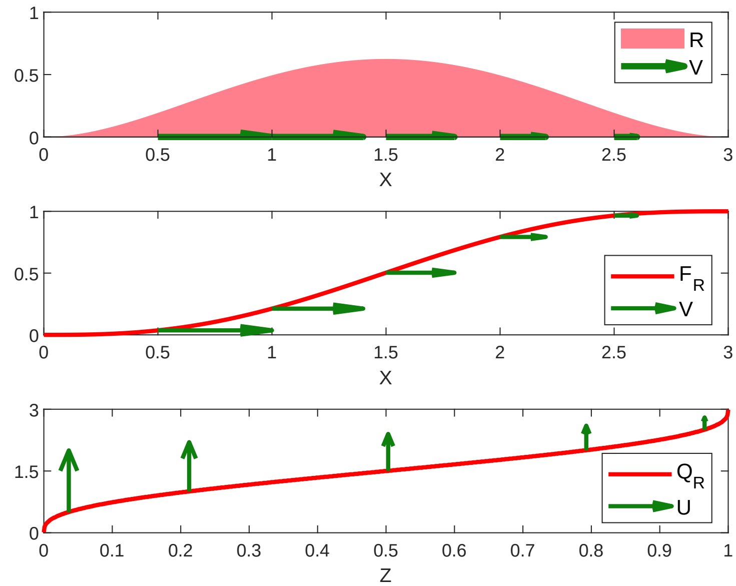

Figure 3.3: Equivalent dynamics for a distribution, the CDF, and quantile function. Whereas the state and dynamics of the specific model are akin to the top figure, those of the transformed model are akin to the bottom.

The distributed optimal control problem in Definition 3.4 has the following special features that greatly simplify its solution.

- The dynamics in (3.6) act pointwise in z, and are thus "decoupled" in z.
- The second constraint in (3.6) implies that U should be constant over regions in z where $Q_{R,t}$ is constant. This, together with the dynamics, implies that if the initial condition $Q_{R,0}$ is constant over some set $S \subseteq [0,1]$, then it remains constant over that set for all time. The control input U over S is also constant, and thus the solution can be parameterized by a single scalar for each set S.

The above observations together with the fact that the objective is an integral in z imply that the distributed optimal control problem decouples into a family of *scalar* optimal control problems. There is a distinction between the regions where $Q_{R,0}$ is constant versus the regions where it is strictly increasing. It is obvious for example, that when $Q_{R,0}$ is strictly increasing, the problem (3.6) decouples into an infinite number of scalar Linear-Quadratic (LQ) tracking problems, one for each value of z. The case where $Q_{R,0}$ is piecewise constant occurs when the resource distribution is composed of discrete agents. We present a unified solution to these problems below. We begin first with a definition.

Definition 3.6 (Partition According to Level Sets)**.** Given a function $Q : A \to \mathbb{R}$, a *partition of A according to level sets of Q* is a set $\mathcal{P}$ of sets $\mathcal{P}_i$ such that

$$A = \bigcup_i \mathcal{P}_i \tag{3.9}$$

$$\mathcal{P}_i \cap \mathcal{P}_j = \emptyset \qquad \forall\, i \neq j \tag{3.10}$$

$$Q(\mathcal{P}_i) = \text{constant} \qquad \forall\, i \tag{3.11}$$

$$Q(\mathcal{P}_i) = Q(\mathcal{P}_j) \quad \Leftrightarrow \quad i = j \tag{3.12}$$

Note that $\{i\}$ may be uncountable, but since the codomain of Q is $\mathbb{R}$, $\{i\}$ can always be well-ordered.

Definition 3.7 (Decoupled Model)**.** The *Decoupled Model* is defined to be the following problem. Given an initial state $Q_{R,0}$, signal $Q_{D,t}$, weighting parameter α, and time horizon T, first partition $[0, 1]$ according to level sets of $Q_{R,0}$. Then, for each element $\mathcal{P}_i$ in the partition, solve

$$\begin{aligned} \bar{u}_i \;=\; & \operatorname*{argmin}_{u_i} \int_0^T \left(\left(\zeta_i(t) - \eta_i(t)\right)^2 + \alpha u_i^2(t) \right) dt \\ & \text{s.t.} \quad \dot{\eta}_i(t) = u_i(t) \end{aligned} \tag{3.13}$$

where $\eta_{i,0} := Q_{R,0}(\mathcal{P}_i)$ and $\zeta_i(t) := \operatorname{avg}_{\mathcal{P}_i}(Q_{D,t})$. The *solution* is the set of all optimal inputs $\{\bar{u}_i\}$ and the *total cost* is defined by the integral

$$\mathcal{J}' \;:=\; \int_0^1 \int_0^T \left(\big(\zeta(z,t) - \eta(z,t)\big)^2 + \alpha u^2(z,t) \right) dt\, dz \tag{3.14}$$

where $\zeta(z,t) = \zeta_i(t)$ for $z \in \mathcal{P}_i$ and similarly for η and u.

The decoupled model (3.13) is equivalent to the transformed model (3.6) (and by extension the specific model (2.13)) in the following sense.

Theorem 3.8. *The problem (3.13) is equivalent to problem (3.6) in that the solutions are 1-1 on some subset which includes the optima and have costs differing by a constant depending only on $Q_{R,0}$ and $Q_{D,t}$. Specifically, the two solutions are related by $U(z,t) = u_i(t)$ for $z \in \mathcal{P}_i$.*

Proof. To see that the solutions are 1-1 on some subset including the optimum, first observe that the dynamics are identical. Since the dynamics are the only constraint in (3.13), clearly any solution U to (3.6) forms a feasible solution $\{u_i\}$ to (3.13) by partitioning according to level sets of $Q_{R,0}$. Clearly any solution $\{u_i\}$ to (3.13) satisfies the dynamics to (3.6) as well, and so it only remains to show that at least the optimal $\{\bar{u}_i\}$ satisfies the second constraint. Clearly the partition $\mathcal{P}$ is defined to satisfy this constraint *within* each element of the partition, but it still remains to be shown that the constraint is satisfied *between* elements of the partition. However, we assert that this is true while delaying the completion of the proof until Section 3.2. (We make no assumptions at this point and proceed to find that the unconstrained optimum takes the form (3.17) - (3.22). With this form, observe that $Q_R(z_1,0) < Q_R(z_2,0) \;\Rightarrow\; z_1 < z_2 \;\Rightarrow\; \zeta(z_1,t) \leq \zeta(z_2,t) \;\Rightarrow\; g(z_1,t) \leq g(z_2,t)$. Define a new variable $\beta(t) := \eta(z_2,t) - \eta(z_1,t)$ and find its dynamics to be

$$\dot{\beta} \;=\; -\,f(t)\beta/\alpha \;-\; \big(g(z_2,t) - g(z_1,t)\big)/\alpha.$$

We apply $g(z_1, t) \leq g(z_2, t)$ and solve the remaining equation by separation of variables to find

$$\beta(t) \;\geq\; \beta(0)\, \exp\left(\frac{1}{\alpha}\int_0^t f(\tau)\, d\tau\right)$$

from which we conclude that $\beta(t) > 0$ or that $\eta(z_2, t) > \eta(z_1, t)$. This ensures that the second constraint of (3.6) is satisfied between elements of the partition $\mathcal{P}$, completing the first part of the proof.)

To show equivalence of the objectives, first observe that on each set $\mathcal{P}_i$ the demand signal Q_D can be written in terms of an average and remainder component as

$$Q_D(z,t) \;=\; \mathrm{avg}_{\mathcal{P}_i}(Q_{D,t}) + \text{remainder} \;=:\; \zeta_i(t) + \tilde{\zeta}(z,t)$$

Notice that the remainder is by definition zero-mean, that is,

$$\int_{\mathcal{P}_i} \tilde{\zeta}(z,t)\, dz \;=\; 0.$$

Then we can split the assignment cost component as follows

$$\begin{aligned}
\int_{\mathcal{P}_i} \big(Q_D(z,t) - Q_r(z,t)\big)^2\, dz \;&=\; \int_{\mathcal{P}_i} \big(\zeta_i(t) + \tilde{\zeta}(z,t) - \eta_i(t)\big)^2\, dz \\
&=\; \int_{\mathcal{P}_i} \big(\zeta_i(t) - \eta_i(t)\big)^2 \;+\; \tilde{\zeta}(z,t)\big(\zeta_i(t) - \eta_i(t)\big) \\
&\qquad\qquad\qquad\qquad +\; \tilde{\zeta}^2(z,t)\, dz \\
&=\; \int_{\mathcal{P}_i} \big(\zeta_i(t) - \eta_i(t)\big)^2\, dz \;+\; \int_{\mathcal{P}_i} \tilde{\zeta}(z,t)\big(\zeta_i(t) - \eta_i(t)\big)\, dz \\
&\qquad\qquad\qquad\qquad +\; \int_{\mathcal{P}_i} \tilde{\zeta}^2(z,t)\, dz \\
&=\; \int_{\mathcal{P}_i} \big(\zeta_i(t) - \eta_i(t)\big)^2\, dz \;+\; 0 \;+\; c_i(t)
\end{aligned}$$

where $c_i(t)$ is a constant depending only on Q_D (and the partition $\mathcal{P}$). Integrating over

t and z we obtain

$$\int_0^1 \int_0^T \big(Q_D(z,t) - Q_r(z,t)\big)^2 \, dt \, dz \;=\; \int_0^1 \int_0^T \big(\zeta(z,t) - \eta(z,t)\big)^2 \, dt \, dz \;+\; C$$

where C is a constant depending only on Q_D and the partition $\mathcal{P}$. Then we can write

$$\int_0^T \int_0^1 \Big(\big(Q_D(z,t) - Q_R(z,t)\big)^2 + \alpha U^2(z,t)\Big) \, dz \, dt$$
$$= \int_0^1 \int_0^T \Big(\big(\zeta(z,t) - \eta(z,t)\big)^2 + \alpha u^2(z,t)\Big) \, dt \, dz \;+\; C \quad (3.15)$$

or that $\mathcal{J} = \mathcal{J}' + C$, where C is a constant depending only on Q_D and the partition $\mathcal{P}$, and thus only on Q_D and $Q_{R,0}$, completing the proof. □

The problem (3.13) is made up of a family of decoupled scalar LQ-tracking problems with the signals ζ_i as the reference signals. These problems can be solved independently of each other, and the optimal solution reconstructed from the independent solutions. Figure 3.4 shows this decoupling graphically.

The next section details the solution of a single one of these scalar LQ-tracking problems.

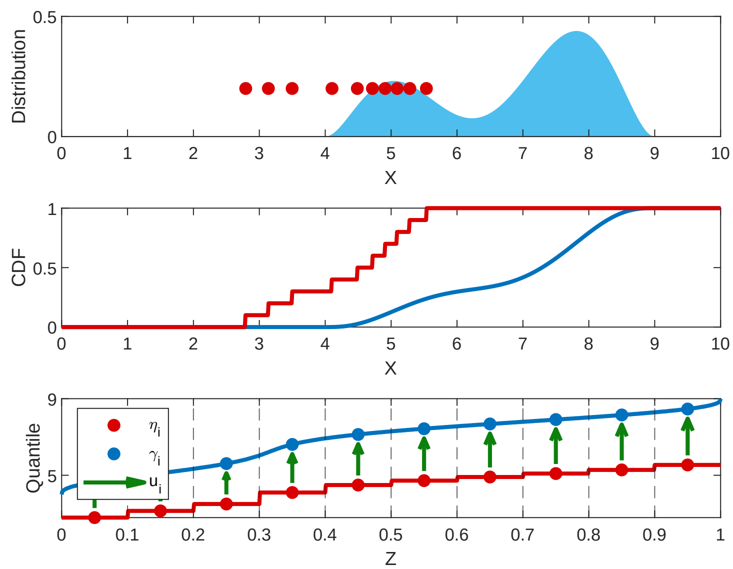

Figure 3.4: Decoupled model example. Each discrete agent corresponds to a constant region in the quantile function. The vertical dotted lines show the separation of these constant regions via partitioning. A single scalar LQ problem can then be written for each member of the partition.

3.2 The Scalar LQ Tracking Problem

We begin with a formal definition.

Definition 3.9 (Scalar LQ Tracking Problem)**.** The *scalar Linear-Quadratic (LQ) tracking problem* is defined to be the following problem. Given an initial state η_0, signal ζ_t, weighting parameter α, and time horizon T, solve

$$\begin{aligned} \bar{u} \;=\; & \min_{u} \int_0^T \big(\zeta(t) - \eta(t)\big)^2 + \alpha u^2(t)\, dt \\ & \text{s.t.} \quad \dot{\eta} = u \end{aligned} \tag{3.16}$$

The *solution* of the problem is defined to be the optimal input $\bar{u}$ and the *cost* is the value of the objective and is written $\mathcal{J}'$.

Figures 3.5 and 3.6 show two pictorial representations of the scalar LQ-tracking problem, as well as its relation to the overall problem we wish to solve.

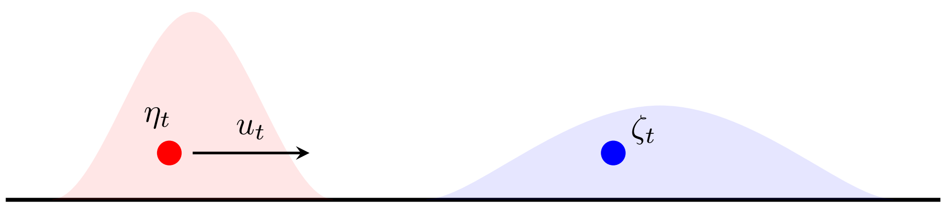

Figure 3.5: A single scalar LQ-tracking problem is written for two elements in the original distributions. The variables η_t and ζ_t represent the positions of the resource agent and demand component respectively and u_t represents the velocity of the resource agent.

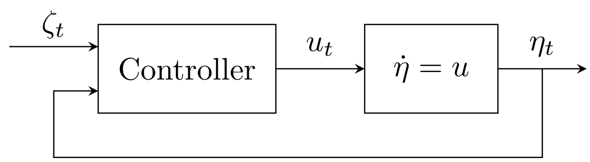

Figure 3.6: Block diagram representation of the scalar LQ-tracking problem. The variable η is considered as the state, ζ as the reference signal, and u as the control input.

This problem has been studied in the literature before (see e.g. [16]), but we reproduce its solution here. The optimal solution has two components, a feedback and feedforward term

$$u(t) \;=\; -\,f(t)\;\eta(t)/\alpha \;-\; g(t)/\alpha, \tag{3.17}$$

where f solves a differential Riccati equation and g is the output of a linear time-varying system driven by the "reference signal" ζ

$$\dot{f}(t) \;=\; f^2(t)/\alpha - 1, \qquad f(T) = 0, \tag{3.18}$$

$$\dot{g}(t) \;=\; f(t)\;g(t)/\alpha \;+\; \zeta(t), \qquad g(T) = 0. \tag{3.19}$$

As usual, the feedback gain arises out of a differential equation with a final condition, and can be computed ahead of time. The same is true for the feedforward term, which requires that in this problem setting, the demand signal must be known ahead of time. Figure 3.7 shows the structure of the optimal controller.

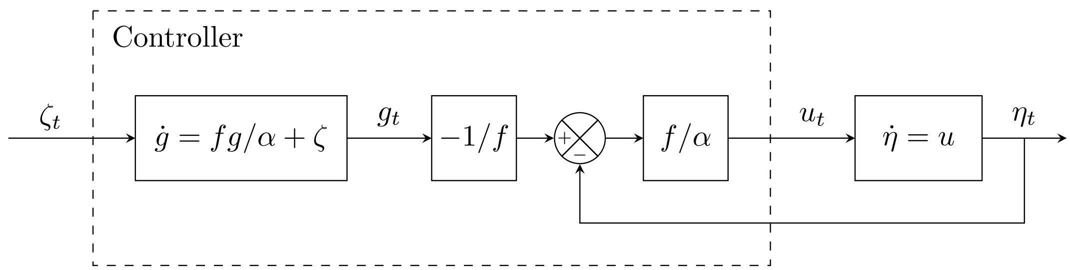

Figure 3.7: Structure of the optimal controller. The variable f is a feedback gain that arises out of a differential Riccati equation and g is an internal controller state. Note that the equation for g is solved backwards in time and so ζ must be known ahead of time. In other words, the optimal controller is noncausal.

The differential Riccati equation (3.18) can be solved analytically by separation of variables to give

$$f(t) = \sqrt{\alpha}\tanh\left((T-t)/\sqrt{\alpha}\right). \tag{3.20}$$

The differential equation (3.19) for g can in principle be solved backwards using its state transition function Φ_g (which is in turn determined by f/α)

$$g(t) = \int_T^t \Phi_g(t,\tau)\ \zeta(\tau)\,d\tau \tag{3.21}$$

where an explicit form for Φ_g can be derived from (3.20)

$$\Phi_g(t,\tau) = \cosh\left((T-\tau)/\sqrt{\alpha}\right)\cdot \operatorname{sech}\left((T-t)/\sqrt{\alpha}\right). \tag{3.22}$$

This is as much as we can say in the general case. We remark once again that the form of (3.21) requires the entire demand signal to be known ahead of time, and so this is a noncausal controller. However, there are at least two cases where this could reasonably be assumed. The first case is that where the demand is static, and the second is where the demand is periodic in time. In both these cases, we can use past measurements to infer the future state of the signal. The next two sections are dedicated to investigating each of these scenarios.

3.3 The Static Case

In the special case where $\zeta(t) = \zeta$ is constant in time, we can obtain more explicit expressions for the optimal control and trajectory. First, the integral (3.21) evaluates to

$$g(t) \;=\; -\,\zeta\sqrt{\alpha}\tanh\left((T-t)/\sqrt{\alpha}\right). \tag{3.23}$$

Combining this with (3.20) gives the optimal control (3.17) in the following "error feedback" form

$$u(t) \;=\; \left(\zeta-\eta(t)\right)\tanh\left((T-t)/\sqrt{\alpha}\right)/\sqrt{\alpha} \tag{3.24}$$

which provides a feedback control law. With this control law, we obtain the trajectory

$$\eta(t) \;=\; \Phi(t,0)\,\eta_0 \;+\; (1-\Phi(t,0))\,\zeta \tag{3.25}$$

and can find the open-loop velocity solving (3.16) as

$$\bar{u}(t) \;=\; -\,\dot{\Phi}(t,0)\,(\zeta-\eta_0) \tag{3.26}$$

where

$$\Phi(t,\tau) := \cosh\left((T-t)/\sqrt{\alpha}\right)\cdot \operatorname{sech}\left((T-\tau)/\sqrt{\alpha}\right) \tag{3.27}$$

$$\dot{\Phi}(t,\tau) = -\sinh\left((T-t)/\sqrt{\alpha}\right)\cdot \operatorname{sech}\left((T-\tau)/\sqrt{\alpha}\right)/\sqrt{\alpha}. \tag{3.28}$$

Furthermore, the optimal solution obtains a cost of

$$\mathcal{J}'(\eta_0,\zeta;\alpha;T) = (\zeta-\eta_0)^2\sqrt{\alpha}\tanh\left(T/\sqrt{\alpha}\right)/2. \tag{3.29}$$

Notice that $\Phi(t,0)$ decreases monotonically from 1 towards 0 in the limiting case where $T \to \infty$, so the trajectory (3.25) is an interpolation between η_0 and ζ. This implies a rather nice form for the overall solution to (2.13): the solution moves along the Wasserstein geodesic between R_0 and some distribution R^*, which turns out to be the nearest reachable distribution to D in the $\mathcal{W}_2$ sense. We formalize this in the following theorem.

Theorem 3.10. *Let R_0 be an initial resource distribution, D a static demand distribution, α a weighting parameter, and T a time horizon. Let $F_{R,t}$, F_D and $Q_{R,t}$, Q_D be the CDFs and quantile functions of R_t, D respectively. Now let $\Pi_{\mathcal{P}}$ be the L^2 projection operator onto the subspace of functions which are constant on sets $\mathcal{P}_i$ where $Q_{R,0}$ is constant. Define R^* to be the distribution and F_{R^*} the CDF of the quantile function $Q_{R^*} := \Pi_{\mathcal{P}} Q_D$. Then the optimal velocity is given in closed-loop form by*

$$\bar{V}_t \;=\; \Big(Q_{R^*} \circ F_{R,t} \;-\; \mathcal{I}\Big) \tanh\left((T-t)/\sqrt{\alpha}\right)/\sqrt{\alpha} \tag{3.30}$$

or in open-loop form by

$$\bar{V}_t \;=\; -\,\dot{\Phi}_t \Big[(Q_{R^*} - Q_{R,0}) \circ F_{R,t}\Big] \tag{3.31}$$

solving (2.13) and resulting in the trajectory

$$R_t \;=\; \Big[\Phi_t \mathcal{I} \;+\; (1-\Phi_t)\left(Q_{R^*} \circ F_{R,0}\right)\Big]_{\#} R_0 \tag{3.32}$$

where

$$\Phi_t \;:=\; \cosh\left((T-t)/\sqrt{\alpha}\right) \cdot \operatorname{sech}\left(T/\sqrt{\alpha}\right) \tag{3.33}$$

$$\dot{\Phi}_t \;=\; -\sinh\left((T-t)/\sqrt{\alpha}\right) \cdot \operatorname{sech}\left(T/\sqrt{\alpha}\right)/\sqrt{\alpha}. \tag{3.34}$$

Furthermore, the solution obtains the cost

$$\mathcal{J}(R_0, D; \alpha; T) \;=\; \mathcal{W}_2^2(R_0, D)\,\sqrt{\alpha}\,\tanh(T/\sqrt{\alpha})\,/\,2. \tag{3.35}$$

Proof. First, we have the velocity solving (3.17) given in closed-loop form by (3.24). By applying Theorem 3.8, we know the solution to (3.6) is given by

$$\bar{U}_t \;=\; \big(Q_{R^*} - Q_{R,t}\big) \tanh\big((T-t)/\sqrt{\alpha}\big)\,/\sqrt{\alpha}.$$

Then by applying Theorem 3.5, we have

$$\begin{aligned}
\bar{V}_t \;&=\; U_t \circ F_{R,t} \\
&=\; \big(Q_{R^*} \circ F_{R,t} - Q_{R,t} \circ F_{R,t}\big) \tanh\big((T-t)/\sqrt{\alpha}\big)\,/\sqrt{\alpha} \\
&=\; \big(Q_{R^*} \circ F_{R,t} - \mathcal{I}\big) \tanh\big((T-t)/\sqrt{\alpha}\big)\,/\sqrt{\alpha}.
\end{aligned}$$

Similarly, for the open-loop form, by applying Theorems 3.8 and 3.5 to (3.26), we obtain

$$\begin{aligned}
\bar{U}_t \;&=\; -\,\dot{\Phi}_t\,\big(Q_{R^*} - Q_{R,0}\big), \\
\bar{V}_t \;&=\; -\,\dot{\Phi}_t\,\Big[(Q_{R^*} - Q_{R,0}) \circ F_{R,t}\Big].
\end{aligned}$$

To find the trajectory R_t, recall that $R_t = \phi_{t\#} R_0$ where ϕ is the flow generated by $\bar{V}$. Using the identity

$$\phi(x,t) \;=\; Q_R(F_R(x,0),t)$$

from the proof of 3.5 together with the solution

$$Q_{R,t} \;=\; \Phi_t\, Q_{R,0} \;+\; (1-\Phi_t)\, Q_{R^*}$$

yields

$$R_t \;=\; \Big[\Phi_t \mathcal{I} \;+\; (1-\Phi_t)\,\big(Q_{R^*} \circ F_{R,0}\big)\Big]_{\#} R_0.$$

To find the cost, we apply Theorem 3.8 and integrate (3.29) over z to find

$$\mathcal{J}' \;=\; \left\|\zeta - Q_{R,0}\right\|_2^2 \sqrt{\alpha}\,\tanh(T/\sqrt{\alpha})\,/\,2$$

so that

$$\mathcal{J} \;=\; \left\|Q_D - Q_{R,0}\right\|_2^2 \sqrt{\alpha}\,\tanh(T/\sqrt{\alpha})\,/\,2$$

and by applying Theorem 3.5 we have

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

completing the proof. □

This surprisingly straightforward solution can be attributed to the consequences of the isometry between the set of quantile functions and the space of distributions described in Footnote 5. First, the isometry with L^2 provides an inherited notion of projection for distributions over $\mathbb{R}$. Second, the isometry ensures that geodesics in L^2 (i.e. straight lines) map to geodesics in $\mathbb{W}_2$ and vice versa. Figure 3.8 shows the situation pictorially.

Lastly, we remark also that the optimal cost function (3.35) solves the Hamilton-Jacobi-Bellman equation for this system. This concludes our analysis of the static case.

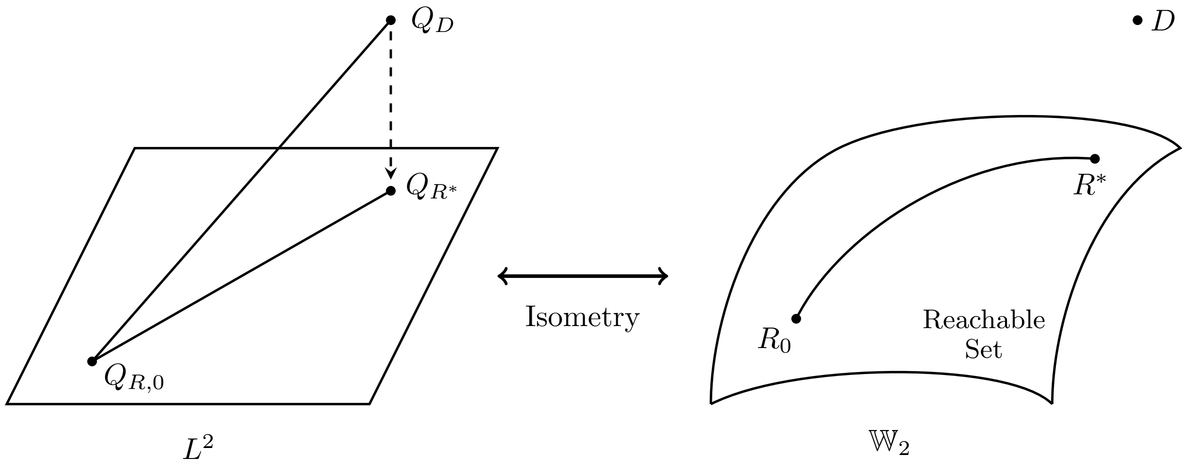

Figure 3.8: Consequences of isometry in the 1D case. We can obtain the solution of the transformed problem via projection in L^2 and then pull the solution back into Wasserstein space. Since the solution in L^2 is a straight line, the resulting solution in Wasserstein space is a Wasserstein geodesic.

3.3.1 Simulations

We simulated a simple example with a discrete resource distribution and continuous demand distribution. The resource distribution consists of ten agents with equal weight, initialized evenly between 0 and 0.9. The demand distribution is given by a relatively arbitrary static bimodal distribution. We used the time horizon $T = 10$ and the parameter $\alpha = 2$. Figures 3.9 - 3.11 show the results.

Theorem 3.10 tells us that changing the parameters α and T does not change the path that the resource distribution takes, only the rate at which it traverses that path. We know that this simulation used a relatively large value for $T/\sqrt{\alpha}$ since the final distribution appears to converge to R^*. However, if we had used a much smaller value for $T/\sqrt{\alpha}$ then the resource distribution would have moved much more slowly, and the final condition at $t = T$ may have looked more similar to Figure 3.10.

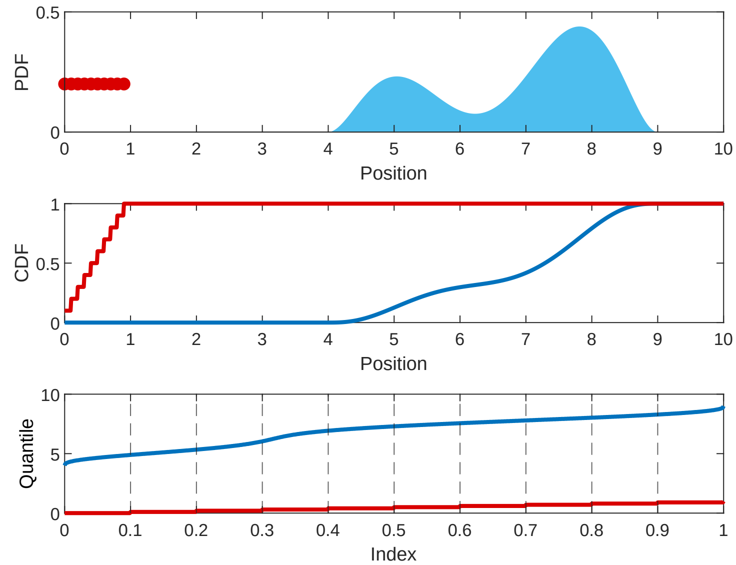

Figure 3.9: Initial conditions of the distributions R_0 (red), D (blue), and the corresponding CDFs and quantile functions. The vertical dotted lines on the quantile function plot separate members of the partition $\mathcal{P}$ of Definition 3.4.

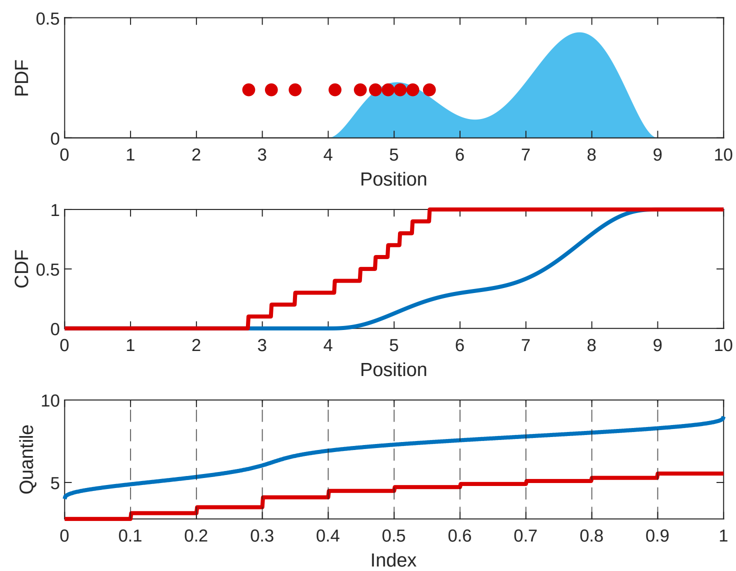

Figure 3.10: Intermediate conditions of the distributions R_0, D, their CDFs, and quantile functions at $t \approx 1.2$.

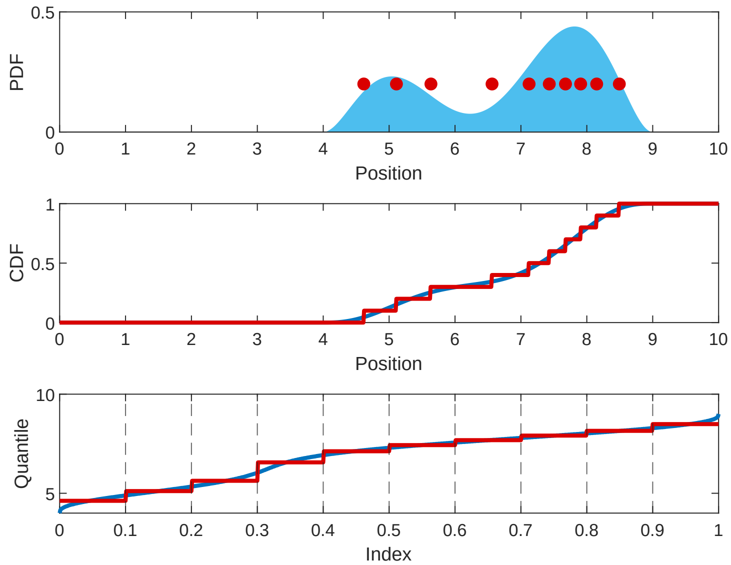

Figure 3.11: Final conditions of the distributions R_0, D, their CDFs, and quantile functions. Notice how the distribution R appears to converge to the distribution R^*, the closest reachable distribution to D.

3.4 The Periodic Case

Here, we are interested in the steady-state response of the system in the infinite-horizon case where the demand is periodic in time. In this case, we find that

$$f(t) \;\to\; \sqrt{\alpha}, \tag{3.36}$$

$$\dot{g}(t) \;\to\; g(t)/\sqrt{\alpha} \;+\; \zeta(t), \tag{3.37}$$

$$\dot{\eta}(t) \;\to\; -\,\eta(t)/\sqrt{\alpha} \;-\; g(t)/\alpha. \tag{3.38}$$

This is a forced linear system and so we can solve for the overall transfer function. We find the frequency response to be

$$\frac{\mathcal{F}(\eta)}{\mathcal{F}(\zeta)} \;=\; \frac{-1/\alpha}{(j\omega)^2 - 1/\alpha} \;=\; \frac{1/\alpha}{\omega^2 + 1/\alpha} \tag{3.39}$$

where $\mathcal{F}$ denotes the Fourier transform. Note that this looks like a low-pass filter with cutoff frequency $1/\sqrt{\alpha}$. Also note that the phase response is identically zero, so that the state is perfectly in-phase with the reference signal. While this is expected for the optimal behavior, it also reaffirms that the optimal controller is noncausal, as we would expect any causal low-pass filter to have some amount of delay. We also remark that there are two roots of the transfer function: one in the right half plane (which stable solving backward for g) and one in the left half plane (which is stable solving forwards for η).

3.4.1 Simulations

We simulated an example with a discrete resource distribution and continuous demand distribution. The resource distribution again consists of ten discrete agents with equal weight, and the demand distribution is given by two periodically alternating Gaussian distributions. The equation for the demand distribution is given by

$$d(x,t) \;=\; (1+\sin(2\pi t))\,\mathcal{N}(2.5,1) \;+\; (1-\sin(2\pi t))\,\mathcal{N}(7.5,1) \tag{3.40}$$

where $\mathcal{N}(x,\sigma^2)$ denotes the normal distribution with mean x and variance σ^2. A time-series of the demand distribution is shown in Figure 3.12, and the resulting tracking signals ζ_i are shown in Figure 3.13. Notice that even though the distribution is sinusoidally time-varying, the quantile function (and thus the tracking signals) have higher harmonics.

We ran simulations with three different values of α to demonstrate behavior in three different regimes. Figures 3.14 - 3.19 show the results.

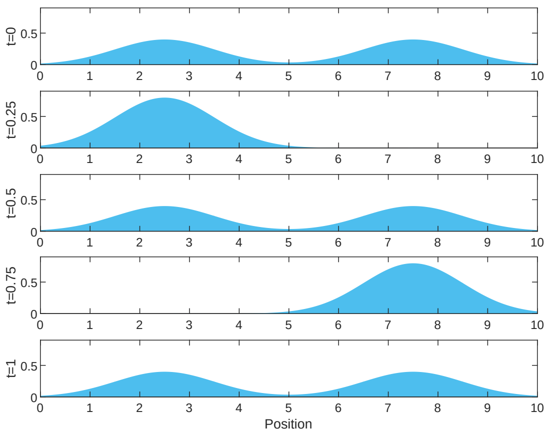

Figure 3.12: Demand signal D consisting of two periodically alternating Gaussian distributions.

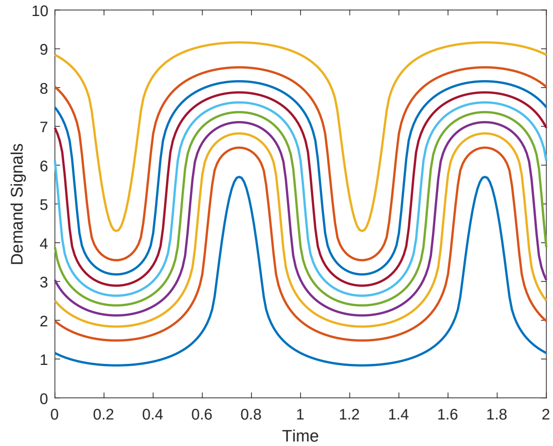

Figure 3.13: Resulting reference signals ζ_i for periodic demand field.

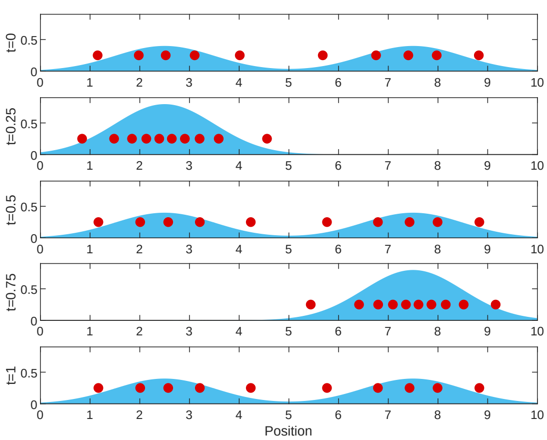

Figure 3.14: Trajectory of resource distribution in small-α case. Notice that since α is small, motion is penalized very little, and the optimal trajectory follows the demand field closely.

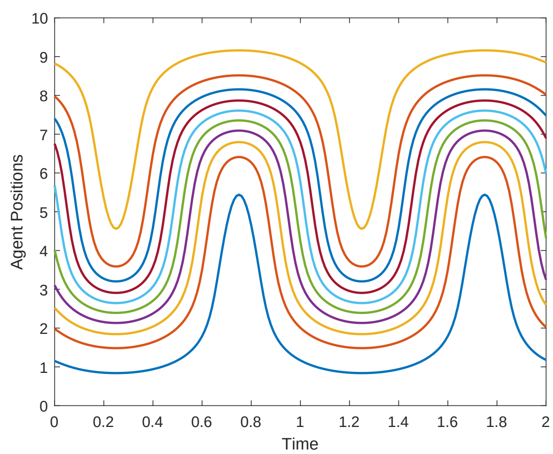

Figure 3.15: Positions η_i for ten resource agents in small-α case. Notice that since α is small, almost all the frequencies in the reference signals are passed by the filter.

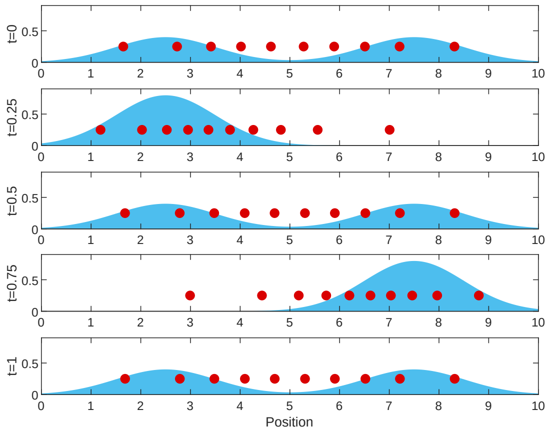

Figure 3.16: Trajectory of resource distribution in moderate-α case. Notice that since α is moderate, motion is penalized somewhat, and the optimal trajectory follows the demand field, but less closely.

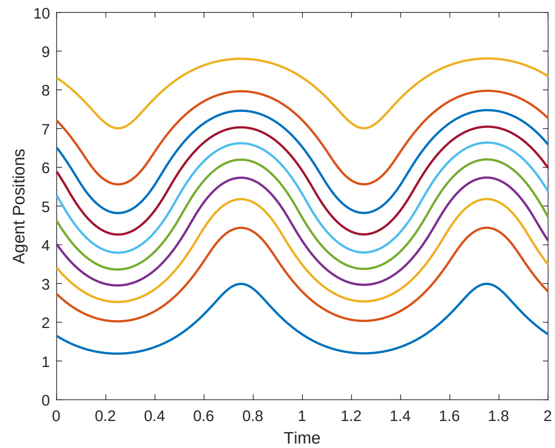

Figure 3.17: Positions η_i for ten resource agents in moderate-α case. Since α was chosen to be about twice the base frequency of the demand field, some of the frequencies in the reference signals are passed by the filter while some are attenuated.

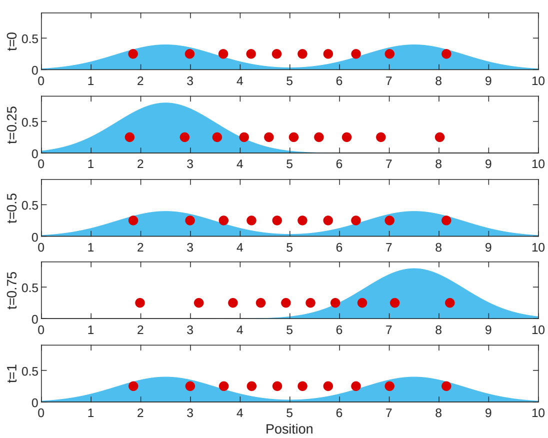

Figure 3.18: Trajectory of resource distribution in large-α case. Notice that since α is large, motion is penalized heavily, and the optimal trajectory is to move very little.

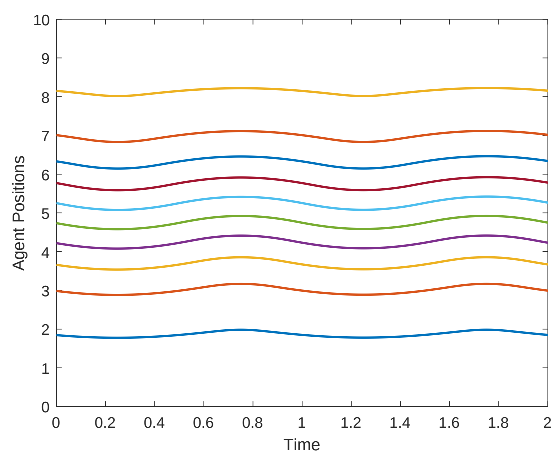

Figure 3.19: Positions η_i for ten resource agents in large-α case. Notice that since α is large, almost all the frequencies in the reference signals are attenuated by the filter.

Conclusion

In this work, I formulated and analyzed a novel model for control of large swarms of autonomous agents. I first provided some background, then motivated and developed the proposed model. Two models were introduced: a *general model* which was abstract and able to accommodate different objectives, and a *specific model* which made particular choices for the objectives so that the model was fully defined. I then solved the specific model in the case of one spatial dimension by means of a *transformation* and then a *decoupling*, which reduced the problem to a well-understood form. An analytic solution was reproduced in the general case. Then, the cases where the demand distribution was static and time-periodic were analyzer further, where explicit solutions and simulation results were provided.

Main Contributions

The main contributions of this work are several definitions and theorems which describe the proposed model, its structure, solution, and provide new perspectives and techniques for solving these types of problems. Specifically:

- Definitions 2.1 and 2.2 propose two new models for large-scale swarm control

- Definitions 3.4 and 3.7 and Theorems 3.5 and 3.8 provide new tools for solving these types of problems and lend insight into the structure of the proposed model in one spatial dimension
- Analytic solutions to the general case, static case, and time-periodic case in 1D are presented in Sections 3.2 - 3.4, characterizing optimal swarm behavior and demonstrating that the proposed model is analytically tractable and that the Wasserstein distance cooperates well with the framework of optimal control
- Simulation results are provided for the static case and the time-periodic case in Subsections 3.3.1 and 3.4.1, helping support the theory in addition to providing visual intuition

There are many practical merits to this approach to swarm control as well. The framework is relatively abstract and could be applied to a wide range of swarm control scenarios, including, for example, drone delivery, autonomous taxi services, data collection, attacker/defender scenarios, and emergency response. The ability to treat general objectives and constraints is also desirable for many of these real-world applications. Optimal control also allows swarm behavior to be tuned by an operator through high-level parameters without knowledge of the underlying mathematics. For these reasons, this seems to be a compelling model for many types of real-world swarm control problems.

Limitations and Future Work

The foremost limitation of this work is that the control strategy developed is *centralized* as opposed to *distributed*. For large-scale swarms in the real world, a centralized controller will suffer from the same drawbacks that motivated the continuum model in the first place. Namely, it is infeasible for a single controller to communicate with and plan motion for every agent. Thus, the control strategy employed on any large-scale real-world swarm will need to be distributed. It is not immediately clear if this work could be extended to distributed controllers, and exploring this is the first main thrust for future work. However, even as is, these results may still be used in a motion planning phase to find a trajectory which could then be followed by a distributed controller. Additionally, this work characterizes best-possible performance both qualitatively and quantitatively by providing behavior objectives for distributed controllers and a metric by which to compare them. Thus, the work does ultimately provide useful tools for developing better distributed controllers regardless of whether it can produce them directly.

The second limitation of this work is that the control strategy developed is noncausal, that is, the entire demand signal needs to be known ahead of time in order to use this framework. The two cases where this is a reasonable assumption are where the demand is static and where it is periodic in time, since in both of these cases, information about the future signal can be inferred from past measurements. However, this limitation can be overcome by treating the problem within the framework of robust control. Developing a version of the controller which can handle unknown time-varying signals is the second main thrust for future work.

The third main thrust is the development of a model that can treat the case where the demand is *dynamically coupled* to the resource, as in the case where the resource "satisfies" the demand and subsequently reduces it. As mentioned earlier in Footnote 3, this is a more interesting (albeit much more complex) problem, and can be treated using the notion of "nonuniform optimal transport" [24].

The last main thrust we mention is of course exploration of the higher-dimensional case. The degree to which these results may generalize to higher dimensions is still unclear. In particular, it is known that the Wasserstein distance does not admit a general closed-form solution in higher dimensions, and it is not immediately obvious what analogue might be used in place of the quantile function in order to transform the problem (nor whether there exists a similar isometry). Nevertheless, even if these properties do not generalize, certain analysis and optimization techniques may still make optimal control based on the Wasserstein distance compelling. Further investigation must be done.

References

[1] A. Gautam and S. Mohan, *A review of research in multi-robot systems*, in *2012 IEEE 7th international conference on industrial and information systems (ICIIS)*, pp. 1–5, IEEE, 2012.

[2] J. Cortés and M. Egerstedt, *Coordinated control of multi-robot systems: A survey, SICE Journal of Control, Measurement, and System Integration* **10** (2017), no. 6 495–503.

[3] S. Huang, R. S. H. Teo, and W. L. Leong, *Review of coverage control of multi unmanned aerial vehicles*, in *2017 11th Asian Control Conference (ASCC)*, pp. 228–232, IEEE, 2017.

[4] 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.

[5] S. Bandyopadhyay, S.-J. Chung, and F. Y. Hadaegh, *Probabilistic and distributed control of a large-scale swarm of autonomous agents*, *IEEE Transactions on Robotics* **33** (2017), no. 5 1103–1123.

[6] V. Krishnan and S. Martínez, *Distributed optimal transport for the deployment of swarms*, in *2018 IEEE Conference on Decision and Control (CDC)*, pp. 4583–4588, IEEE, 2018.

[7] V. Krishnan and S. Martínez, *Distributed online optimization for multi-agent optimal transport*, *arXiv preprint arXiv:1804.01572* (2018).

[8] V. Krishnan and S. Martinez, *Distributed control for spatial self-organization of multi-agent swarms*, *SIAM Journal on Control and Optimization* **56** (2018), no. 5 3642–3667.

[9] 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* **5** (2021), no. 4 1483–1488.

[10] M. Fornasier and F. Solombrino, *Mean-field optimal control*, *ESAIM: Control, Optimisation and Calculus of Variations* **20** (2014), no. 4 1123–1152.

[11] S. Ferrari, G. Foderaro, P. Zhu, and T. A. Wettergren, *Distributed optimal control of multiscale dynamical systems: a tutorial*, *IEEE Control Systems Magazine* **36** (2016), no. 2 102–116.

[12] K. Elamvazhuthi, H. Kuiper, and S. Berman, *Pde-based optimization for stochastic mapping and coverage strategies using robotic ensembles*, *Automatica* **95** (2018) 356–367.

[13] B. Bonnet and F. Rossi, *The pontryagin maximum principle in the wasserstein space*, *Calculus of Variations and Partial Differential Equations* **58** (2019), no. 1 1–36.

[14] C. Jimenez, A. Marigonda, and M. Quincampoix, *Optimal control of multiagent systems in the wasserstein space*, *Calculus of Variations and Partial Differential Equations* **59** (2020), no. 2 1–45.

[15] 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* **59** (2021), no. 2 977–1006.

[16] A. P. Sage and C. C. White, *Optimum systems control*. Prentice Hall, 1977.

[17] M. Emerick, S. Patterson, and B. Bamieh, *Optimal combined motion and assignments with continuum models*, *IFAC-PapersOnLine* **55** (2022), no. 13 121–126.

[18] I. M. Gelfand and G. E. Shilov, *Generalized Functions, Volume 1: Properties and Operations*. Academic Press, 1964.

[19] G. T. Mase, R. E. Smelser, and G. E. Mase, *Continuum mechanics for engineers*. CRC press, 2009.

[20] F. Santambrogio, *Optimal transport for applied mathematicians*, *Birkäuser, NY* **55** (2015), no. 58-63 94.

[21] L. Ambrosio, N. Gigli, and G. Savaré, *Gradient flows: in metric spaces and in the space of probability measures*. Springer Science & Business Media, 2005.

[22] C. Villani, *Topics in optimal transportation*, vol. 58. American Mathematical Soc., 2021.

[23] C. Villani, *Optimal transport: old and new*, vol. 338. Springer, 2009.

[24] W. Gangbo, W. Li, S. Osher, and M. Puthawala, *Unnormalized optimal transport*, *Journal of Computational Physics* **399** (2019) 108940.