UNIVERSITY OF CALIFORNIA SAN DIEGO

Large-Scale Multi-Agent Transport: Theory, Algorithms and Analysis

A dissertation submitted in partial satisfaction of the
requirements for the degree of Doctor of Philosophy

in

Engineering Sciences (Mechanical Engineering)

by

Vishaal Krishnan

Committee in charge:

Professor Sonia Martínez, Chair
Professor Jorge Cortés
Professor Miroslav Krstić
Professor Lei Ni
Professor Andrej Zlatoš

2019

The Dissertation of Vishaal Krishnan is approved, and it is acceptable in quality and form for publication on microfilm and electronically:

Chair

University of California San Diego

2019

DEDICATION

To the memory of my late grandfather, S. Balasubramanian

TABLE OF CONTENTS

LIST OF FIGURES

ACKNOWLEDGEMENTS

I would first like to express my deepest gratitude to my advisor, Prof. Sonia Martínez, whose mentorship has enabled me to broadly develop my skills as a researcher. My work has benefited immensely from her expertise and rigorous approach. I am forever thankful to her for providing me the space to explore new ideas and for her patience when results were hard to come by, while her encouragement has allowed me to persist with my ideas in the face of mounting challenges. I also take great inspiration from her role as teacher and mentor, and although I have yet to fully realize the true extent of her contribution to my growth as a researcher, I know that I will continue to reap the fruits of these formative interactions for as long as I am able.

I am also very thankful to the rest of my doctoral committee, Prof. Jorge Cortés, Prof. Miroslav Krstić, Prof. Andrej Zlatoš and Prof. Lei Ni, whose valuable inputs have guided me through the course of my research. I am indebted to Prof. Cortés and Prof. Krstić, who have helped me along in numerous ways ever since I began my work, from helping me understand the fundamentals of control theory to their support through the course of my studies. My research can be viewed as an attempt to build on some of their past work, and their contributions have provided clarifying insight into the problems I have worked on. I am particularly thankful to Prof. Cortés for sharing his knowledge in distributed optimization, control and a variety of other topics during the weekly lab meetings over the years that have vastly improved my understanding of the subjects. I am immensely grateful to Prof. Ni for sharing his expertise in Differential Geometry through a year-long graduate course, which was an intellectually rewarding experience. My consultations with him on aspects of my research involving the application of differential geometry and spectral theory have greatly helped clarify my thinking and exposed me to interesting ideas. I am very thankful to Prof. Zlatoš for his insightful and probing comments, and for helping me understand the mathematics of fluid dynamics, which have contributed greatly to improving this work.

I would like to thank my lab mates, graduate students and postdocs in Prof. Martínez's and Prof. Cortés' groups over the years, for all the interesting discussions, academic and otherwise. The shared sense of community has been a constant source of support over the years.

I am, most of all, thankful to my family for their support over the years. I have learnt over the course of my doctoral work that the challenges in research are often not merely of a technical nature, and I have my family to thank for providing the necessary wherewithal to undertake my work. To Sahana, your love and unflinching support have been the wellspring of my courage and well-being. On whether navigating through the doctoral degree would be possible without them, I do not know another way. To my parents, your numerous sacrifices and the self-belief you have instilled in me early in life continue to carry me through. I take great inspiration from your lives and the challenges you have had to face. I would be remiss to not mention Prakrit, Mayank, Samuel, Aamodh and others, friends whose contributions have been immense and inarticulable.

The material in Chapter 2 is currently being prepared for submission as *Multiscale Analysis of Multi-Agent Transport by Gradient Descent*, V. Krishnan and S. Martínez. The dissertation author was the primary investigator and author of this paper.

The material in Chapter 3, in full, has been submitted for publication to the SIAM Journal on Control and Optimization and is under review. It may appear as *Distributed Online Optimization for Multi-Agent Optimal Transport*, V. Krishnan and S. Martínez. A preliminary version of the work appeared in the proceedings of the IEEE Conference on Decision and Control, Miami Beach, USA, pp. 4583–4588, December 2018, as *Distributed optimal transport for the deployment of swarms*, V. Krishnan and S. Martínez. The dissertation author was the primary investigator and author of these papers.

Chapter 4, in full, is a reprint of the material as it appears in the publication *Distributed Control for Spatial Self-Organization of Multi-Agent Swarms*, V. Krishnan and S. Martínez, SIAM Journal on Control and Optimization, 56(5), pp. 3642–3667, 2018. A preliminary version of the work appeared in the proceedings of the International Symposium on Mathematical Theory of Networks and Systems, Minneapolis, USA, pp. 706–713, July 2016, as *Self-Organization in Multi-Agent Swarms via Distributed Computation of Diffeomorphisms*, V. Krishnan and S. Martínez. The dissertation author was the primary investigator and author of these papers.

The material in Chapter 5, in full, has been submitted for publication to the IEEE Transactions on Automatic Control and is under review. It may appear as *A Distributional Framework for Moving Horizon Estimation: Stability and Privacy Considerations*, V. Krishnan and S. Martínez. A preliminary version of the work appeared in the proceedings of the American Control Conference, Philadelphia, USA, pp. 459–464, July 2019, as *On Observability and Stability of Moving-Horizon Estimation in a Distributional Framework*, V. Krishnan and S. Martínez. The dissertation author was the primary investigator and author of these papers.

Chapter 6, in full, is a reprint of the material in *Identification of Critical Nodes in Large-Scale Spatial Networks*, V. Krishnan and S. Martínez, IEEE Transactions on Control of Network Systems, 6(2), pp. 842–851, 2019. A preliminary version of the work appeared in the proceedings of the IFAC World Congress, Toulouse, France, pp. 14721–14726, July 2017, as *Identification of Critical Node Clusters for Consensus in Large-Scale Spatial Networks*, V. Krishnan and S. Martínez. The dissertation author was the primary investigator and author of these papers.

VITA

2019	Ph.D, University of California, San Diego
2014	M. Tech, Indian Institute of Technology, Madras
2014	B. Tech, Indian Institute of Technology, Madras

PUBLICATIONS

1. *Multiscale Analysis of Multi-Agent Transport by Gradient Descent*, V. Krishnan and S. Martínez, In preparation.

2. *Distributed Online Optimization for Multi-Agent Optimal Transport*, V. Krishnan and S. Martínez, SIAM Journal on Control and Optimization, Submitted.

3. *A Distributional Framework for Moving Horizon Estimation: Stability and Privacy Considerations*, V. Krishnan and S. Martínez, IEEE Transactions on Automatic Control, Submitted.

4. *On Observability and Stability of Moving-Horizon Estimation in a Distributional Framework*, V. Krishnan and S. Martínez, American Control Conference, Philadelphia, USA, pp. 459–464, July 2019.

5. *Identification of Critical Nodes in Large-Scale Spatial Networks*, V. Krishnan and S. Martínez, IEEE Transactions on Control of Network Systems, 6(2), pp. 842–851, 2019.

6. *Distributed optimal transport for the deployment of swarms*, V. Krishnan and S. Martínez, IEEE Conference on Decision and Control, Miami Beach, USA, pp. 4583–4588, December 2018.

7. *Distributed Control for Spatial Self-Organization of Multi-Agent Swarms*, V. Krishnan and S. Martínez, SIAM Journal on Control and Optimization, 56(5), pp. 3642–3667, 2018.

8. *Identification of Critical Node Clusters for Consensus in Large-Scale Spatial Networks*, V. Krishnan and S. Martínez, IFAC World Congress, Toulouse, France, pp. 14721–14726, July 2017.

9. *Self-Organization in Multi-Agent Swarms via Distributed Computation of Diffeomorphisms*, V. Krishnan and S. Martínez, Int. Symposium on Mathematical Theory of Networks and Systems, Minneapolis, USA, pp. 706–713, July 2016.

ABSTRACT OF THE DISSERTATION

Large-Scale Multi-Agent Transport: Theory, Algorithms and Analysis

by

Vishaal Krishnan

Doctor of Philosophy in Engineering Sciences (Mechanical Engineering)

University of California San Diego, 2019

Professor Sonia Martínez, Chair

The problem of transport of multi-agent systems has received much attention in a wide range of engineering and biological contexts, such as spatial coverage optimization, collective migration, estimation and mapping of unknown environments. In particular, the emphasis has been on the search for scalable decentralized algorithms that are applicable to large-scale multi-agent systems.

For large multi-agent collectives, it is appropriate to describe the configuration of the collective and its evolution using macroscopic quantities, while actuation rests at the microscopic scale at the level of individual agents. Moreover, the control problem faces a multitude of information constraints imposed by the multi-agent setting, such as limitations in sensing,

communication and localization. Viewed in this way, the problem naturally extends across scales and this motivates a search for algorithms that respect information constraints at the microscopic level while guaranteeing performance at the macroscopic level.

We address the above concerns in this dissertation on three fronts: theory, algorithms and analysis. We begin with the development of a multiscale theory of gradient descent-based multi-agent transport that bridges the microscopic and macroscopic perspectives and sets out a general framework for the design and analysis of decentralized algorithms for transport. We then consider the problem of optimal transport of multi-agent systems, wherein the objective is the minimization of the net cost of transport under constraints of distributed computation. This is followed by a treatment of multi-agent transport under constraints on sensing and communication, in the absence of location information, where we study the problem of self-organization in swarms of agents. Motivated by the problem of multi-agent navigation and tracking of moving targets, we then present a study of moving-horizon estimation of nonlinear systems viewed as a transport of probability measures. Finally, we investigate the robustness of multi-agent networks to agent failure, via the problem of identifying critical nodes in large-scale networks.

Introduction

Multi-agent systems are, broadly, collections of autonomous agents with sensing, communication and computational capabilities. Often to ensure scalability, these systems are characterized by the absence of centralized decision-making, where the agents use information locally available to them to make decisions. An important class of problems in the context of multi-agent systems is that of transport, which arises in a wide range of scenarios in engineering and biology. This includes (i) seeking the extrema of a scalar field, for instance, in scenarios such as collective chemotaxis in biology and estimating and controlling wildfires; (ii) achieving a desired coverage of a spatial region, as in the case of mobile sensing networks [27, 35, 88] and emerging applications such as autonomous mobility-on-demand; (iii) mapping, navigation and tracking of moving targets; (iv) collective manipulation.

In problems of multi-agent transport such as the above, particularly when they involve large collectives of agents, it is more appropriate to track the configuration of the collective and its evolution by macroscopic quantities (such as the distribution of agents over a spatial region in coverage problems). This involves a scale transformation, wherein a description of the system at the microscopic scale is mapped onto a macroscopic description, and the objectives of transport are specified at the macroscopic scale. However, the actuation still rests at the microscopic scale at the level of individual agents, and the control problem faces a multitude of constraints imposed by the multi-agent setting, broadly catergorized into information constraints (such as the need for online, decentralized algorithms for scalability; limitations in sensing, communication and localization) and physical constraints (collision and obstacle avoidance, to name a few). In the macroscopic scale, theoretical tools from infinite-dimensional analysis are often more appropriate, while at the microscopic scale, multi-agent control problems with the above constraints have been more appropriately dealt with using tools of finite-dimensional analysis. However, a formal theory bridging the two scales in the context of multi-agent transport has been elusive, and this poses a challenge to control and algorithm design and analysis for large-scale multi-agent systems. There is a need for such a bridge theory because it is important to understand how macroscopic transport objectives translate into the microscopic scale and conversely, how the microscopic control laws and algorithms affect macroscopic behavior and scale as the number of agents $N \to \infty$.

Models of multi-agent transport

We begin the study of multi-agent transport in Chapter 2 with the formulation of an iterative optimization-based transport scheme. We model the transport of agents as an iterative proximal descent scheme in a compact Euclidean domain $\Omega \subset \mathbb{R}^d$, of the form:

$$x^{+} = \arg\min_{z \in \Omega} \frac{1}{2\tau}|x - z|^2 + \varphi(z). \tag{1}$$

Macroscopically, the objective of multi-agent transport is formulated as the minimization of a (strictly) convex functional $F : \mathscr{P}(\Omega) \to \mathbb{R}$ on the space of probability measures $\mathscr{P}(\Omega)$ over Ω. We then address the question of how this macroscopic objective can be achieved by the scheme (1). With an appropriate choice of φ, we show that the "lift" of the scheme (1) to the space of probability measures $\mathscr{P}(\Omega)$ corresponds to a transport scheme that minimizes F. This establishes the connection between the microscopic and macroscopic perspectives, i.e., between an iterative proximal descent scheme for the agents in the Euclidean space and the minimization of F in the space of probability measures that describes the transport of the collective.

We then propose an implementable multi-agent transport scheme for a finite N number of agents as a proximal descent w.r.t. a discretization of the functional F, and show convergence to critical points, and in some cases the local minimizers, of such a scheme. In the limit $N \to \infty$, we show that we recover the scheme (1) and convergence to the global minimizer of F. We also investigate the asymptotic stability of the continuous-time gradient flow in the space of absolutely continuous probability measures, obtained from the lift of (1) in the limit $\tau \to 0$, and this serves as the candidate continuous-time model for multi-agent transport. We then use these results to shed light on some multi-agent coverage control algorithms from the literature.

In summary, Chapter 2 sets up the theory and framework for the design of algorithms for multi-agent transport in the rest of the thesis.

Multi-agent optimal transport

In multi-agent transport scenarios there is, typically, an associated cost of transport owing to energy considerations. Optimal transport theory [117], which deals with the problem of rearranging probability measures while minimizing the cost of transport, presents the appropriate theoretical tools. Continuing in the spirit of Chapter 2, the objective of Chapter 3 is to set up a scheme similar to (1) for the problem of multi-agent optimal transport with the goal of minimizing the net cost of transport of the collective. We consider costs $c : \Omega \times \Omega \to \mathbb{R}$ that are continuous and satisfy the properties of a metric. Lifting the cost c using the optimal transport formulation, we obtain a metric $C : \mathscr{P}(\Omega) \times \mathscr{P}(\Omega) \to \mathbb{R}$ in the space of probability measures. Fixing a target probability measure $\mu^* \in \mathscr{P}(\Omega)$, we define an objective functional $F(\mu) = C(\mu, \mu^*)$, as the cost of optimal transport from μ^*. We then set up an iterative proximal descent-based transport scheme to minimize F. The particular challenges to multi-agent optimal transport come from the contraints of online and distributed computation. In other words, the estimates of the local objective functions for the agents, to be used in the proximal descent scheme, are to be computed online by a distributed algorithm. This results in a coupling between the transport and the distributed online computation of the local objective functions, and we investigate in Chapter 3 the asymptotic stability of the transport under such a coupling.

Self-organizing multi-agent transport

Self-organization in swarms refers broadly to the emergence of patterns of long-range order in large groups of dynamic agents which interact locally with each other. It is a pervasive phenomenon in nature, observed in biological [28] and other natural systems [120]. These instances are characterized by primitive agents functioning under severe constraints on information, in the form of limitations on sensing and communication. Moreover, the agents are constrained to operate in the absence of location information. Formulating self-organization as a problem of multi-agent transport, we explore in Chapter 4 mechanisms that enable large-scale multi-agent transport towards a target measure in the absence of location information. In examples of biological transport and pattern formation, the absence of location information is typically mitigated, if only partially, by a process of (cellular) differentiation that assigns to every agent an identifier which modulates the behavior of the agent. Adopting this perspective in our context, the transport scheme is to be accompanied by an identifier-assignment algorithm that serves to modulate the local objective function that the agent seeks to minimize via a descent scheme. We call this identifier-assignment algorithm pseudo-localization, as it serves to partially mitigate the absence of location information.

To further illustrate the problem, we refer again to the proximal descent scheme (1). In order to implement such a scheme, every agent must be able to evaluate its local objective function φ at a point $z \in \Omega$, typically in its vicinity (from its location $x \in \Omega$). In the absence of access to its location $x \in \Omega$, such a scheme is not implementable. We therefore reformulate the problem as one of composite optimization, where we replace the local objective function φ with $\tilde{\varphi} \circ X$, where X is the identifier-assignment map resulting from the pseudo-localization algorithm. In other words, the agents do not have access to their locations in Ω but to the image of their locations in $X(\Omega)$, i.e., their identifiers.

We note that in Chapter 4, we work entirely with a macroscopic model of transport in continuous-time, described by a coupled system of PDEs, with the continuity equation describing the motion of the collective, and a PDE describing the process of pseudo-localization (differentiation). We target setting up the transport to minimize the squared L^2-distance $\int_\Omega |\rho - \rho^*|^2 \,\mathrm{dvol}$ from the (absolutely continuous) target measure. Using Lyapunov-based methods, we derive control laws for the coupled system of PDEs, in order that the transport converges asymptotically to the target measure.

State estimation for tracking and navigation

Chapters 2, 3 and 4 consider the problem of multi-agent transport where the target probability measure, or the objective functional $F : \mathscr{P}(\Omega) \to \mathbb{R}$, is fixed. We present a version of the time-varying case in Chapter 5, motivated by applications of navigation and tracking of moving targets. We let a discrete-time nonlinear system of the form $x_{k+1} = f(x_k, w_k)$ (with process noise w_k) describe a moving point target, with measurements of the underlying state given by $y_k = h(x_k) + v_k$ (with measurement noise v_k), where the objective of the multi-agent transport is to track the true underlying state x_k. We consider the case where all the agents have access to the measurements y_k, and formulate the multi-agent tracking problem as one of transport by defining an appropriate objective functional using a finite moving window of measurements. Owing to the underlying discrete-time nonlinear system, this objective functional is time-varying.

Viewed another way, this problem is essentially one of optimization-based state estimation, in particular, moving-horizon estimation formulated as a transport of probability measures, and the material in Chapter 5 is presented entirely from this point of view. We also investigate the allied concern of guaranteeing privacy in state estimation, and design differentially private moving-horizon estimation schemes via an entropy regularization of the objective functional.

Robustness of multi-agent networks

Chapter 6 contains an investigation of robustness of multi-agent networks to failure of agents, posed as a problem of identifying the critical nodes in a large-scale spatial network. The identification of critical nodes in a network, motivated by the question of network robustness, is crucial to improving its resilience to attacks and failures. The notion of critical nodes refers to the subset of nodes in the network whose removal results in the maximum deterioration of a given performance metric. In the context of robustness of networks/graphs, a widely studied metric [58, 75] is the second smallest eigenvalue of the graph Laplacian matrix (also called the algebraic connectivity of the graph). In addition to being an indicator of how well connected the graph is, it is typically of significance in the context of agreement dynamics on networks (such as consensus and synchronization), as it governs the convergence rate of the dynamics. The problem of identifying critical nodes in a network graph leads to combinatorial optimization problems. Thus, for large-scale networks any algorithm that solves the problem exactly is of high complexity. Motivated by this, we study a relaxation of the problem through a continuum approximation of the network to the spatial domain where the nodes are distributed.

Chapter 1

Notation and Preliminaries

Let $\|\cdot\| : \mathbb{R}^d \to \mathbb{R}_{\geq 0}$ denote the Euclidean norm on $\mathbb{R}^d$ and $|\cdot| : \mathbb{R} \to \mathbb{R}_{\geq 0}$ the absolute value function. We denote by $\nabla = \left(\frac{\partial}{\partial x_1}, \dots \frac{\partial}{\partial x_n}\right)$ the gradient operator in $\mathbb{R}^d$. As a shorthand, we let $\frac{\partial}{\partial z}(\cdot) = \partial_z(\cdot)$ for a variable z.

Let $\partial\Omega$ denote the boundary of Ω, $\bar{\Omega} = \Omega \cup \partial\Omega$ its closure and $\mathring{\Omega} = \Omega \setminus \partial\Omega$ its interior with respect to the standard Euclidean topology. For $M \subseteq \Omega$, let the distance $d(x, M)$ of a point $x \in \Omega$ to the set M be given by $d(x, M) = \inf_{y \in M} \|x - y\|$. For any $x \in \Omega \subset \mathbb{R}^d$, we denote by $B_r^m(x)$ the closed d-ball of radius $r > 0$, with respect to a metric m, centered at x. Let $1_M : \Omega \to \{0, 1\}$ be the indicator function on Ω for the subset M. We denote by $\langle f, g \rangle$ the inner product of functions $f, g : \Omega \to \mathbb{R}$ with respect to the Lebesgue measure, given by $\langle f, g \rangle = \int_\Omega f g \,\mathrm{dvol}$.

Let $\mu \in \mathscr{P}(\Omega)$ be an absolutely continuous probability measure on $\Omega \subset \mathbb{R}^d$, with ρ the corresponding density function (where $\mathrm{d}\mu = \rho \,\mathrm{dvol}$), with vol being the Lebesgue measure. We denote by $\mathbb{E}_\mu$ the expectation w.r.t. the measure μ. Given a map $\mathscr{T} : \Omega \to \Gamma$ and a measure $\mu \in \mathscr{P}(\Omega)$, we let $\nu = \mathscr{T}_\# \mu$ denote the pushforward measure of μ by $\mathscr{T}$, where for a measurable set $\mathscr{B} \subset \mathscr{T}(\Omega)$, we have $\nu(\mathscr{B}) = \mathscr{T}_\# \mu(\mathscr{B}) = \mu(\mathscr{T}^{-1}(\mathscr{B}))$. Let $F : \mathscr{P}(\Omega) \to \mathbb{R}$ be a smooth real-valued function on the space of probability measures on $\Omega \subset \mathbb{R}^d$. We denote by $\frac{\delta F}{\delta \mu}(x)$ the first variation of F with respect to the measure μ, such that a perturbation $\delta\mu$ of the measure results in a perturbation $\delta F = \int_{\mathscr{X}} \frac{\delta F}{\delta \mu} d(\delta\mu)$.

We denote by $C^k(\Omega)$ the space of k-times continuously differentiable functions on Ω, and by $Lip(\Omega)$ the space of Lipschitz continuous functions on Ω. The L^p space of functions on a measurable space U is given by $L^p(U) = \{f : U \to \mathbb{R} \,|\, \|f\|_{L^p(U)} = (\int_U |f|^p \,\mathrm{dvol})^{1/p} < \infty\}$, where $\|\cdot\|_{L^p(U)}$ is the L^p norm. Of particular interest is the L^2 space, or the space of square-integrable functions. In this paper, we denote by $\|f\|_{L^2(\Omega)}$ the L^2 norm of f with respect to the Lebesgue measure, and by $\|f\|_{L^2(\Omega,\mu)} = \left(\int_\Omega |f|^2 \,\mathrm{d}\mu\right)^{1/2}$ the weighted L^2 norm. The Sobolev space $W^{1,p}(\Omega)$ is defined as $W^{1,p}(\Omega) = \{f : \Omega \to \mathbb{R} \,|\, \|f\|_{W^{1,p}} = (\int_\Omega |f|^p + \int_\Omega |\nabla f|^p)^{1/p} < \infty\}$. For two functions $f : \mathbb{R} \times \Omega \to \mathbb{R}$ and $g : \Omega \to \mathbb{R}$, denote $f(t,\cdot) \equiv f_t$ and further denote $f \to_{L^2} g$ the convergence in L^2 norm of f_t to g as $t \to \infty$, that is, $\lim_{t\to\infty} \|f_t - g\|_{L^2} = 0$. Convergence in H^1 norm is denoted similarly by $f \to_{H^1} g$.

We now state some well-known results that will be used in the subsequent chapters in this thesis.

Lemma 1 (Divergence Theorem [32])**.** *For a smooth vector field* $\mathbf{F}$ *over a bounded open set* $\Omega \subseteq \mathbb{R}^d$ *with boundary* $\partial\Omega$*, the volume integral of the divergence* $\nabla \cdot \mathbf{F}$ *of* $\mathbf{F}$ *over* Ω *is equal to the surface integral of* $\mathbf{F}$ *over* $\partial\Omega$*:*

$$\int_\Omega (\nabla \cdot \mathbf{F}) \;\mathrm{dvol} = \int_{\partial\Omega} \mathbf{F} \cdot \mathbf{n} \, dS, \tag{1.1}$$

where $\mathbf{n}$ *is the outward normal to the boundary and* dS *the surface measure on the boundary. For a scalar field* ψ *and a vector field* $\mathbf{F}$ *defined over* $\Omega \subseteq \mathbb{R}^d$*:*

$$\int_\Omega (\mathbf{F} \cdot \nabla\psi) \, d\,\mathrm{dvol} = \int_{\partial\Omega} \psi(\mathbf{F} \cdot \mathbf{n}) \, dS - \int_\Omega \psi(\nabla \cdot \mathbf{F}) \;\mathrm{dvol}.$$

Lemma 2. ***(Leibniz Integral Rule [32]).*** *Let* $f \in \mathscr{C}^\infty(\mathbb{R} \times \mathbb{R}^n)$ *and* $\Omega : \mathbb{R} \rightrightarrows \mathbb{R}^n$ *be a smooth one-parameter family of bounded open sets in* $\mathbb{R}^n$ *generated by the flow corresponding to the smooth vector field* $\mathbf{v}$ *on* $\mathbb{R}^n$*. Then:*

$$\frac{d}{dt}\left(\int_{\Omega(t)} f(t,\mathbf{r}) \, d\mu\right) = \int_{\Omega(t)} \partial_t(f(t,\mathbf{r})) \, d\mu + \int_{\partial\Omega(t)} f(t,\mathbf{r})\mathbf{v} \cdot \mathbf{n} \, dS.$$

Corollary 1. ***(Derivative of Energy Functional).*** *Let U be an energy functional defined as follows:*

$$U = \frac{1}{2}\int_{\Omega} |f|^2 \, d\mu,$$

for some function $f : \Omega \to \mathbb{R}$. Then,

$$\dot{U} = \int_{\Omega} f \left(\frac{df}{dt}\right) \, d\mu + \frac{1}{2}\int_{\Omega} |f|^2 \nabla \cdot \mathbf{v} \, d\mu.$$

where $\frac{d}{dt} = \partial_t + \mathbf{v} \cdot \nabla$ is the total derivative.

Proof. We have included the proof for this corollary for the sake of completeness. Using the Leibniz integral rule and the Divergence theorem, we have (it is understood that the integrations are with respect to the measure μ):

$$\begin{aligned}
\frac{\partial U}{\partial t} &= \int_{\Omega} f \partial_t f + \frac{1}{2}\int_{\partial\Omega} |f|^2 \mathbf{v} \cdot \mathbf{n} \\
&= \int_{\Omega} f \partial_t f + \frac{1}{2}\int_{\Omega} \nabla \cdot (|f|^2 \mathbf{v}) \\
&= \int_{\Omega} f \partial_t f + \int_{\Omega} f \cdot (\mathbf{v} \cdot \nabla) f + \frac{1}{2}\int_{\Omega} |f|^2 \nabla \cdot \mathbf{v} \\
&= \int_{\Omega} f (\partial_t f + (\mathbf{v} \cdot \nabla) f) + \frac{1}{2}\int_{\Omega} |f|^2 \nabla \cdot \mathbf{v} \\
&= \int_{\Omega} f \left(\frac{df}{dt}\right) + \frac{1}{2}\int_{\Omega} |f|^2 \nabla \cdot \mathbf{v}.
\end{aligned}$$

□

Lemma 3 (Rademacher's Theorem [84])**.** *Let $\Omega \subset \mathbb{R}^d$ be open and $f : \Omega \to \mathbb{R}^m$ be Lipschitz continuous. Then f is differentiable at almost every $x \in \Omega$.*

Lemma 4 (Poincaré-Wirtinger Inequality [84])**.** *For $p \in [1, \infty]$ and Ω, a bounded connected open subset of $\mathbb{R}^d$ with a Lipschitz boundary, there exists a constant C depending only on Ω and* *p such that for every function u in the Sobolev space* $W^{1,p}(\Omega)$*:*

$$\|u - u_\Omega\|_{L^p(\Omega)} \leq C\|\nabla u\|_{L^p(\Omega)},$$

where $u_\Omega = \frac{1}{|\Omega|}\int_\Omega u d\mu$*, and* $|\Omega|$ *is the Lebesgue measure of* Ω*.*

Lemma 5 (Rellich-Kondrachov Compactness Theorem [55])**.** *Let* $\Omega \subset \mathbb{R}^d$ *be open, bounded and such that* $\partial\Omega$ *is* C^1*. Suppose* $1 \leq p < n$*, then* $W^{1,p}(\Omega)$ *is compactly embedded in* $L^q(\Omega)$ *for each* $1 \leq q < \frac{pn}{n-p}$*. In particular, we have* $W^{1,p}(\Omega)$ *is compactly contained in* $L^p(\Omega)$*.*

We now state the following version of the LaSalle invariance principle for Banach spaces, which will be used later:

Lemma 6 (LaSalle Invariance Principle [70, 118, 119])**.** *Let* $\{\mathscr{C}(t)\,|\,t \in \mathbb{R}_{\geq 0}\}$ *be a continuous semigroup of operators on a Banach space U (closed subset of a Banach space with norm* $\|\cdot\|_U$*), and for any* $u \in U$*, define the positive orbit starting from u at* $t = 0$ *as* $\Gamma_+(u) = \{\mathscr{C}(t)u\,|\,t \in \mathbb{R}_{\geq 0}\} \subseteq U$*. Let* $V : U \to \mathbb{R}$ *be a continuous Lyapunov functional on* $\mathscr{G} \subset U$ *for* $\mathscr{C}$ *(such that* $\dot{V}(u) = \frac{d}{dt}V(\mathscr{C}(t)u) \leq 0$ *in* $\mathscr{G}$*). Define* $E = \{u \in \bar{\mathscr{G}}\,|\,\dot{V}(u) = 0\}$*, and let* $\tilde{E}$ *be the largest invariant subset of E. If for* $u_0 \in \mathscr{G}$*, the orbit* $\Gamma_+(u_0)$ *is pre-compact (lies in a compact subset of U), then* $\lim_{t\to+\infty} d_U(\mathscr{C}(t)u_0, \tilde{E}) = 0$*, where* $d_U(y, \tilde{E}) = \inf_{x\in\tilde{E}} \|y - x\|_U$ *(where* d_U *is the distance in U).*

Harmonic diffeomorphisms

Let (M, g) and (N, h) be two Riemannian manifolds of dimensions m and n, and Riemannian metrics g and h, respectively. A map $\phi : M \to N$ is called harmonic if it minimizes the functional:

$$E(\phi) = \int_M |\nabla\phi|^2 dv_g, \tag{1.2}$$

where dv_g is the Riemannian volume form on M. The Euler-Lagrange equation for the functional E, which also yields the minimum energy, is given by $\Delta\phi = 0$, the Laplace equation [74]. It is useful to note that the solutions to the heat equation, in the limit $t \to \infty$, approach the harmonic map. We now state a lemma on harmonic diffeomorphisms of Riemann surfaces (i.e., $m = n = 2$ above).

Lemma 7. ***(Harmonic diffeomorphism [48]).*** *Let (M, g) be a compact surface with boundary and (N, h) a compact surface with non-positive curvature. Suppose that $\psi : M \to N$ is a diffeomorphism onto $\psi(M)$. Assume that $\psi(M)$ is convex. Then there is a unique harmonic map $\phi : M \to N$ with $\phi = \psi$ on ∂M, such that $\phi : M \to \phi(M)$ is a diffeomorphism.*

We note that the non-positive curvature constraint in the lemma is essentially a constraint on the metric h on N, and the curvature is zero for the Euclidean metric.

The space of probability measures and its weak topology

Let $\Omega = \bar{D}$, with $D \subset \mathbb{R}^d$ an open, bounded set in the d-dimensional Euclidean space $\mathbb{R}^d$. Let $\mathscr{B}(\Omega)$ be the collection of Borel sets (the Borel σ-algebra) in Ω, which we take in this paper to be the collection of measurable sets. The space of probability measures, $\mathscr{P}(\Omega)$, is the collection of functions $\mu \in \mathscr{P}(\Omega)$ satisfying:

1. Values in the unit interval: $\mu : \mathscr{B}(\Omega) \to [0,1]$, with $\mu(\emptyset) = 0$ and $\mu(\Omega) = 1$.

2. Countable additivity: For a pairwise disjoint sequence $\{A_i\}_{i\in\mathbb{N}}$ of measureable sets, $\mu(\cup_{i\in\mathbb{N}} A_i) = \sum_{i\in\mathbb{N}} \mu(A_i)$.

We denote by $\mathscr{P}^r(\Omega) \subset \mathscr{P}(\Omega)$ the space of atomless probability measures, where a measure $\mu \in \mathscr{P}(\Omega)$ is said to be atomless if for any $A \in \mathscr{B}(\Omega)$ with $\mu(A) > 0$, there exists $B \in \mathscr{B}(\Omega)$, $B \subset A$, such that $\mu(A) > \mu(B) > 0$. It follows that for an atomless measure μ, we will have $\mu(\{x\}) = 0$ for all $x \in \Omega$. We consider this a notion of regularity of probability measures in this paper, in that atomless measures are regular, and hence the superscript r in $\mathscr{P}^r(\Omega)$.

We now define absolutely continuous probability measures over Ω.

Definition 1 (Absolutely continuous probability measures)**.** *A probability measure* $\mu \in \mathscr{P}(\Omega)$ *is said to be absolutely continuous if for any* $A \in \mathscr{B}(\Omega)$*, we have* $\mu(A) = 0$ *if* $\mathrm{vol}(A) = 0$ *(where* vol *is the Lebesgue measure).*

This allows us to define a density function ρ corresponding to μ (where $\mathrm{d}\mu = \rho\,\mathrm{dvol}$).

We now introduce the notion of pushforward of a measure under a mapping $T : \Omega \to \Omega$.

Definition 2 (Pushforward measures)**.** *Given a map* $\mathscr{T} : \Omega \to \Omega$ *and a measure* $\mu \in \mathscr{P}(\Omega)$*, we let* $\nu = \mathscr{T}_{\#}\mu$ *denote the* pushforward measure *of* μ *by* $\mathscr{T}$*, where for a measurable set* $\mathscr{B} \subset \mathscr{T}(\Omega)$*, we have* $\nu(\mathscr{B}) = \mathscr{T}_{\#}\mu(\mathscr{B}) = \mu(\mathscr{T}^{-1}(\mathscr{B}))$.

We now introduce the notions of weak convergence in $\mathscr{P}(\Omega)$, the topology of weak convergence, the metrizability of this topology, the compactness of collections of probability measures and the underlying connections between them. These will be crucial for the development of the theoretical results contained in this paper, and a detailed account can be found in [24].

Definition 3 (Weak convergence)**.** *A sequence of measures* $\{\mu_k\}_{k\in\mathbb{N}}$ *in* $\mathscr{P}(\Omega)$ *is said to* converge weakly *to* $\mu \in \mathscr{P}(\Omega)$ *if for any bounded and continuous function* f *on* Ω*, it holds that* $\lim_{k\to\infty} \int_\Omega f d\mu_k = \int_\Omega f d\mu$.

Equivalently, in the definition above, the sequence $\{\mu_k\}_{k\in\mathbb{N}}$ in $\mathscr{P}(\Omega)$ is said to *converge to* μ *in* $\mathscr{P}(\Omega)$ *equipped with the weak topology*. The space of probability measures $\mathscr{P}(\Omega)$ equipped with the weak topology is *metrizable* [24]. In other words, there exists a metric on $\mathscr{P}(\Omega)$ such that the weak topology is obtained as the topology induced by the metric. One such metric is the Wasserstein distance, which will be defined below. We now state Prokhorov's theorem [24] on the equivalence between tightness and precompactness of a collection of probability measures over a separable and complete metric (Polish) space.

Lemma 8 (Prokhorov's theorem)**.** *A set $\mathscr{S} \subseteq \mathscr{P}(\Omega)$ (with Ω a separable, complete metric space) is* precompact *w.r.t. the topology of weak convergence if and only if it is* tight*; i.e., for any $\varepsilon > 0$ there exists a compact $K_\varepsilon \subseteq \Omega$ such that $\mu(K_\varepsilon) > 1 - \varepsilon$ for all $\mu \in \mathscr{S}$.*

Corollary 2 (Compactness of $\mathscr{P}(\Omega)$)**.** *Prokhorov's theorem in Lemma 8 implies that, for compact set Ω, $\mathscr{P}(\Omega)$ is precompact since it is tight (where for any $\varepsilon > 0$, we choose Ω itself as the compact set and have $\mu(\Omega) = 1 > 1 - \varepsilon$ for any $\mu \in \mathscr{P}(\Omega)$). Moreover, since $\mathscr{P}(\Omega)$ is also closed, it is therefore compact.*

The L^2-Wasserstein distance

The L^2-Wasserstein distance between two probability measures $\mu, \nu \in \mathscr{P}(\Omega)$ is given by:

$$W_2^2(\mu,\nu) = \min_{\pi \in \Pi(\mu,\nu)} \int_{\Omega\times\Omega} |x-y|^2 \, d\pi(x,y), \tag{1.3}$$

where $\Pi(\mu,\nu)$ is the space of joint probability measures over $\Omega \times \Omega$ with marginals μ and ν. The definition of L^2-Wasserstein distance in (1.3) follows from the so-called Kantorovich formulation of optimal transport. An alternative formulation, called the Monge formulation of optimal transport, is given below:

$$W_2^2(\mu,\nu) = \min_{\substack{T:\Omega\to\Omega \\ T_\#\mu=\nu}} \int_{\Omega} |x-T(x)|^2 \, d\mu(x). \tag{1.4}$$

In the Monge formulation (1.4), the minimization is carried out over the space of maps $T : \Omega \to \Omega$ for which the probability measure ν is obtained as the pushforward of μ. This can be viewed as a deterministic formulation of optimal transport, where the transport is carried out by a map, whereas the Kantorovich formulation (1.3) can be seen as a relaxation where the transport plan is described by a joint probability measure π over $\Omega \times \Omega$, with μ and ν as its marginals. It is to be noted that the Monge formulation does not admit a solution for all probability measures μ

and ν, while the Kantorovich formulation does. However, the two formulations (1.3) and (1.4) are equivalent under certain conditions and in the sense laid out in the ensuing lemma. We refer the reader to [106] for detailed proofs.

Lemma 9 (Existence and Uniqueness)**.** *There exists a unique minimizer π^* to the Kantorovich formulation* (1.3) *of the L^2-Wasserstein distance. Moreover, if the measure μ is atomless, the Monge formulation* (1.4) *has a unique minimizer T^* and it holds that $\pi^* = (id, T^*)_{\#}\mu$.*

The Kantorovich formulation (1.3) admits a dual formulation for which strong duality holds, so that the L^2-Wasserstein distance is also be given by:

$$\begin{aligned} W_2^2(\mu,\nu) = \sup_{\phi \in L^1(\Omega); \psi \in L^1(\Omega)} & \int_\Omega \phi \, d\mu + \int_\Omega \psi \, d\nu \\ & \phi(x) + \psi(y) \leq |x-y|^2. \end{aligned} \tag{1.5}$$

Equivalently, the above can be formulated as:

$$W_2^2(\mu,\nu) = \sup_{\substack{\phi \in L^1(\Omega) \\ s.t. \ \phi^c \in L^1(\Omega)}} \int_\Omega \phi \, d\mu + \int_\Omega \phi^c \, d\nu, \tag{1.6}$$

where $\phi^c(y) = \inf_{x\in\Omega}\{|x-y|^2 - \phi(x)\}$. The space of probability measures $\mathscr{P}(\Omega)$ endowed with the L^2-Wasserstein distance W_2 will equivalently be referred to as the L^2-Wasserstein space $(\mathscr{P}(\Omega), W_2)$ over Ω.

We now present the following lemma, which follows from Theorem 6.9 in [117], on the equivalence between convergence in the weak topology sense and convergence in the L^2-Wasserstein metric sense.

Lemma 10 (Convergence in $(\mathscr{P}(\Omega), W_2)$)**.** *For bounded $\Omega \in \mathbb{R}^d$, the L^2-Wasserstein distance W_2 metrizes the weak convergence in $\mathscr{P}(\Omega)$, i.e., a sequence of measures $\{\mu_k\}_{k\in\mathbb{N}}$ in $\mathscr{P}(\Omega)$ converges weakly to $\mu \in \mathscr{P}(\Omega)$ if and only if $\lim_{k\to\infty} W_2(\mu_k, \mu) = 0$.*

Convexity of functionals on the Wasserstein space

Before we can define any notion of convexity, we need to introduce an appropriate notion of interpolation. We first recall that a function $f : \mathbb{R}^d \to \mathbb{R}$ is convex if $f((1-t)x+ty) \leq (1-t)f(x)+tf(y)$, for any $x,y \in \mathbb{R}^d$ and $t \in [0,1]$. In this way, the definition of convexity involves interpolation along the straight line segment γ between x and y, or the geodesic $\gamma(t) = (1-t)x+ty$, for $t \in [0,1]$, connecting x and y in $\mathbb{R}^d$. The generalization of convexity to the space of probability measures also requires a notion of interpolation between probability measures, such as the following. Given $x,y \in \Omega$ and Dirac measures $\delta_x, \delta_y \in \mathscr{P}(\Omega)$, we define its *linear interpolation* as $(1-t)\delta_x + t\delta_y$, for $t \in [0,1]$. Similarly, the *displacement interpolation* of δ_x and δ_y is given by $\delta_{(1-t)x+ty}$, for $t \in [0,1]$. Notice that the support of the linear interpolation of δ_x, δ_y, is given by $\{x,y\}$, for each $t \in [0,1]$. However, support of the displacement interpolation is located along the geodesic segment connecting x to y, for $t \in [0,1]$. The displacement interpolation is more appropriate for defining convexity of functionals over the space of probability measures and the transport schemes we construct in our work. More generally, we provide the following definition:

Definition 4 (Displacement interpolation)**.** *For $\mu, \nu \in \mathscr{P}(\Omega)$ such that there exists an optimal transport map $T : \Omega \to \Omega$ from μ to ν in the L^2-Wasserstein space over Ω, the* displacement interpolant *of μ and ν is given by $\gamma_t = ((1-t)id + tT)_\# \mu$, for $t \in [0,1]$.*

Observe that the displacement interpolant of μ and $\nu \in \mathscr{P}(\Omega)$ corresponds to the notion of *geodesic* from μ to ν in the L^2-Wasserstein space. Moreover, in the definition of the displacement interpolant above, we have assumed the existence of an optimal transport map from μ to ν, which restricts the class of probability measures considered. The following lemma, which is a consequence of Proposition 9.1.11 in [25], states that the existence of a transport map is guaranteed when the source measure is atomless.

Lemma 11 (Existence of pushforward map)**.** *For two probability measures $\mu, \nu \in \mathscr{P}(\Omega)$, there* *exists a map $T : \Omega \to \Omega$ such that $\nu = T_{\#}\mu$ if μ is atomless.*

In order for a definition of interpolation to apply to the entire space $\mathscr{P}(\Omega)$, the interpolants must be constructed using an optimal transport plan, which is a joint probability measure with marginals μ and ν. We refer to [8], where this notion is explored in greater detail. However, for the purposes of this work, we do not need such a general definition, because splitting of masses by transport plans is not relevant to the setting of multi-agent transport. Alternatively, we can work with the notion of generalized geodesics introduced below:

Definition 5 (Generalized displacement interpolation)**.** *Let $\mu, \nu \in \mathscr{P}(\Omega)$, and $\theta \in \mathscr{P}^r(\Omega)$ be an atomless probability measure, such that $T_{\theta\to\mu} : \Omega \to \Omega$ and $T_{\theta\to\nu} : \Omega \to \Omega$ are optimal transport maps from θ to μ, and θ to ν resp. in the L^2-Wasserstein space over Ω. A* (generalized) displacement interpolant *of μ and ν* w.r.t. *θ is given by $\gamma_t = \big((1-t)T_{\theta\to\mu} + tT_{\theta\to\nu}\big)_{\#}\theta$, for $t \in [0,1]$.*

We first state the following lemma before introducing the notion of geodesic convexity.

Lemma 12 (Geodesic convexity of $\mathscr{P}(\Omega)$)**.** *The L^2-Wasserstein space $(\mathscr{P}(\Omega), W_2)$ is geodesically convex (w.r.t. generalized displacement interpolations or geodesics) if Ω is convex.*

Proof. Let $\mu, \nu \in \mathscr{P}(\Omega)$ be two probability measures over Ω and let $\theta \in \mathscr{P}(\Omega)$ be an atomless probability measure. Moreover, let $T_{\theta\to\mu} : \Omega \to \Omega$ and $T_{\theta\to\nu} : \Omega \to \Omega$ be the optimal transport maps from θ to μ and θ to ν respectively, such that ${T_{\theta\to\mu}}_{\#}\theta = \mu$ and $T_{\theta\to\nu\#}\theta = \nu$. The L^2-Wasserstein generalized geodesic from μ to ν is generated as the pushforward by the one-parameter family of maps $T_t = (1-t)T_{\theta\to\mu} + tT_{\theta\to\nu}$ with $t \in [0,1]$, given by $\mu_t = T_{t\#}\theta$. From the convexity of Ω, for any $x \in \Omega$, we have that $T_t(x) \in \Omega$ since $T_t(x)$ is a convex combination of $T_{\theta\to\mu}(x) \in \Omega$ and $T_{\theta\to\nu}(x) \in \Omega$, and therefore lies on the straight line segment between them. This implies that $T_t : \Omega \to \Omega$. Moreover, since $\mu_t = T_{t\#}\theta$, the mass of μ_t is concentrated on

$T_t(\Omega)$, and we have:

$$\mu_t(\Omega) = \mu_t(T_t(\Omega)) = \int_{T_t(\Omega)} d\mu_t = \int_{T_t(\Omega)} d\,(T_{t\#}\theta) = \int_{\Omega} d\theta = \theta(\Omega) = 1.$$

Therefore, we have $\mu_t \in \mathscr{P}(\Omega)$ and that $(\mathscr{P}(\Omega), W_2)$ is convex. □

Now, under the assumption of convexity of Ω, we have:

Definition 6 (Geodesic convexity). *A functional $F : \mathscr{P}(\Omega) \to \mathbb{R}$ is* geodesically convex *if for $\mu, \nu \in \mathscr{P}(\Omega)$ such that there exists an optimal transport map $T : \Omega \to \Omega$ from μ to ν ($T_{\#}\mu = \nu$) in the L^2-Wasserstein space over Ω, we have:*

$$F(((1-t)id + tT)_{\#}\mu) \leq (1-t)F(\mu) + tF(\nu), \qquad \forall t \in [0,1].$$

It is useful to generalize the notion of geodesic convexity to accommodate functionals that are not convex in the sense of Definition 6. An example of a functional that is not convex in the sense of Definition 6 is $F(\mu) = W_2^2(\mu_{\text{ref}}, \mu)$ [8], defined as the squared L^2-Wasserstein distance from a reference measure μ_{ref}, which is nevertheless an attractive candidate for gradient flow-based transport as will be seen later. This motivates a definition of convexity along generalized geodesics, as given below:

Definition 7 (Generalized geodesic convexity). *Let $\mu, \nu \in \mathscr{P}(\Omega)$ and $\theta \in \mathscr{P}(\Omega)$ be an atomless probability measure, such that $T_{\theta\to\mu} : \Omega \to \Omega$ and $T_{\theta\to\nu} : \Omega \to \Omega$ are the optimal transport maps from θ to μ and from θ to ν respectively, in the L^2-Wasserstein space over Ω. A functional $F : \mathscr{P}(\Omega) \to \mathbb{R}$ is* (generalized) geodesically convex *if:*

$$F(\big((1-t)T_{\theta\to\mu} + tT_{\theta\to\nu}\big)_{\#}\theta) \leq (1-t)F(\mu) + tF(\nu), \qquad \forall t \in [0,1].$$

We note that $F(\mu) = W_2^2(\mu_{ref}, \mu)$ is convex [8] in the sense of Definition 7.

Derivatives of functionals on the space of atomless measures

We first introduce the notion of first variation of a functional as follows:

Definition 8 (First variation of a functional on $\mathscr{P}(\Omega)$)**.** *Consider a functional $F : \mathscr{P}(\Omega) \to \mathbb{R}$, and a $\mu_0 \in \mathscr{P}(\Omega)$. Let $\{\mu_\varepsilon\}_{\varepsilon\in\mathbb{R}}$ be any a smooth one-parameter family of probability measures such that the limit $\partial_\varepsilon \mu|_{\varepsilon=0} = \lim_{\varepsilon\to 0} \frac{\mu_\varepsilon - \mu_0}{\varepsilon}$ exists. Suppose that there exists a unique $\frac{\delta F}{\delta \mu}(\mu_0)$ such that $\frac{d}{d\varepsilon} F(\mu_\varepsilon)\big|_{\varepsilon=0} = \int_\Omega \frac{\delta F}{\delta \mu}(\mu_0) d\left(\partial_\varepsilon \mu|_{\varepsilon=0}\right)$, for any $\{\mu_\varepsilon\}_{\varepsilon\in\mathbb{R}}$. Then, $\frac{\delta F}{\delta \mu}(\mu_0)$ is called the first variation of F evaluated at μ_0.*

With the above definition on first variation of functionals in place, we are ready to introduce the notion of Fréchet derivative of a functional on the L^2-Wasserstein space $(\mathscr{P}^r(\Omega), W_2)$:

Definition 9 (Derivative of a functional on $(\mathscr{P}^r(\Omega), W_2)$)**.** *The* Fréchet derivative ξ *of a differentiable (over the space of atomless probability measures) functional $F : \mathscr{P}(\Omega) \to \mathbb{R}$ at an atomless measure $\mu \in \mathscr{P}(\Omega)$, is given implicitly by:*

$$\lim_{\nu\to\mu} \frac{F(\nu) - F(\mu) - \int_\Omega \langle \xi, T_{\mu\to\nu} - id \rangle \, d\mu}{W_2(\mu,\nu)} = 0,$$

where $\xi = \nabla \varphi$ and $\varphi = \frac{\delta F}{\delta \mu}$.

We now introduce the notion of directional derivative of a functional over probability measures. For this, let $\mathbf{v} = \frac{T_{\mu\to\nu} - id}{t}$ which implies that $\nu = (id + t\mathbf{v})_\# \mu$. We also have:

$$W_2(\mu,\nu) = \sqrt{\int_\Omega \left|T_{\mu\to\nu} - id\right|^2 d\mu} = t\sqrt{\int_\Omega |\mathbf{v}|^2 d\mu},$$

and we get:

$$\lim_{t\to 0} \frac{F((id + t\mathbf{v})_\# \mu) - F(\mu) - t\int_\Omega \langle \xi, \mathbf{v} \rangle \, d\mu}{t\sqrt{\int_\Omega |\mathbf{v}|^2 d\mu}} = 0.$$

Therefore, the directional derivative of F along $\mathbf{v}$ is

$$\frac{d}{dt}\bigg|_{\mathbf{v}} F(\mu) = \lim_{t\to 0} \frac{F((id+t\mathbf{v})_{\#}\mu) - F(\mu)}{t} = \int_{\Omega} \langle \xi, \mathbf{v} \rangle \, d\mu,$$

where ξ is the Fréchet derivative of F.

Lipschitz-continuous derivatives

We now introduce the notion of l-smoothness that will be useful for the development of gradient descent-based transport schemes later in the paper. We begin with a definition of this notion in the Euclidean space and then generalize it in the Wasserstein sense.

Definition 10 (l-smoothness)**.** *A function $f : \Omega \to \mathbb{R}$ is called l-*smooth *(or* Lipschitz differentiable*) if for any $x, y \in \Omega$, we have $|\nabla f(y) - \nabla f(x)| \leq l\|y-x\|$.*

It is easy to prove the following lemma on l-smooth functions, which is used later to define the notion of l-smoothness of functionals over $(\mathscr{P}(\Omega), W_2)$:

Lemma 13 (l-smooth functions)**.** *For an l-smooth function $f : \Omega \to \mathbb{R}$ and any $x, y \in \Omega$, we have $|f(y) - f(x) - \langle \nabla f(x), y-x \rangle| \leq \frac{l}{2}\|y-x\|^2$.*

We now generalize the above notion to functionals over the space of probability measures using Lemma 13.

Definition 11 (l-smoothness of functionals on $(\mathscr{P}(\Omega), W_2)$)**.** *A functional $F : \mathscr{P}(\Omega) \to \mathbb{R}$ is called l-*smooth *(or* Lipschitz differentiable*) if for any $\mu, \nu \in \mathscr{P}(\Omega)$, we have:*

$$\sqrt{\int_{\Omega} \left|\xi_\mu - \xi_\nu\right|^2 d\nu} \leq l W_2(\mu, \nu),$$

where $\xi_\mu = \nabla\left(\frac{\delta F}{\delta \bar{\mu}}\Big|_{\mu}\right)$, $\xi_\nu = \nabla\left(\frac{\delta F}{\delta \bar{\mu}}\Big|_{\nu}\right)$ and $T_{\nu\to\mu}$ is the optimal transport map from ν to μ.

Lemma 14 (l-smooth functionals)**.** *A functional $F : \mathscr{P}(\Omega) \to \mathbb{R}$ that is l-smooth on $(\mathscr{P}(\Omega), W_2)$ satisfies:*

1. $\left|F(\nu)-F(\mu)-\int_\Omega \left\langle \xi_\mu, T_{\mu\to\nu}-id\right\rangle d\mu\right| \leq \frac{l}{2}W_2^2(\mu,\nu)$,

2. $\left|\int_\Omega \left\langle \xi_\mu-\xi_\nu, T_{\nu\to\mu}-id\right\rangle d\nu\right| \leq lW_2^2(\mu,\nu)$,

for atomless probability measures $\mu,\nu\in\mathscr{P}^r(\Omega)$.

We also define the proximal operator on Ω with respect to a function $f:\Omega\to\mathbb{R}$ as follows:

$$\operatorname{prox}_f(x) = \arg\min_{z\in\Omega} \frac{1}{2}\|z-x\|^2 + f(z).$$

Chapter 2

A multiscale theory of multi-agent transport by gradient descent

In this chapter, we set out to establish a multiscale theory of gradient descent-based transport of multi-agent systems, with three main goals:

1. To present a macroscopic description of the behavior of multi-agent gradient descent algorithms as transport in the space of probability measures.
2. To shed new light on the behavior of coverage optimization algorithms as the number of agents $N \to \infty$.
3. To provide a framework for the development of algorithms based on iterative, gradient-based transport schemes in the space of probability measures.

2.1 Bibliographical comments

In the context of robotic systems, problems of deployment and formation control of groups of robots have been extensively studied [27, 35, 72, 88, 108]. More recently, research efforts have been undertaken to massively increase the scale of these robotic systems [104]. In the context of robotic swarms, programmable self-assembly of two-dimensional shapes with a thousand-robot swarm is demonstrated in [105]. These robots are capable of measuring distances to nearby neighbors which they use to localize themselves relative to other localized robots. Each robot then uses its position to implement an edge-following algorithm.

From a theoretical perspective, as the number of agents increases, the design and analysis of efficient distributed transport laws poses new challenges, starting with the choice of appropriate mathematical abstractions. The need for parsimonious descriptions of the collectives, along with the fact that tasks for these systems are more likely to be specified at a high level, calls for the use of macroscopic models. Among the approaches to the coverage control and deployment problem for large-scale multi-agent systems are transport by synthesis of Markov transition matrices [14, 16, 41], the use of continuum models [53, 80] for transport, coverage control by parameter tuning and/or boundary control of the reaction-advection-diffusion PDE [52, 60, 123], and mean-field stabilization [42, 50, 51].

2.2 Models for large-scale multi-agent transport

We consider a collection of $N \in \mathbb{N}$ identical agents with indices $i \in \mathscr{I} = \{1, \ldots, N\}$, distributed across a spatial region, we recall, $\Omega = \bar{D}$, with $D \subset \mathbb{R}^d$ an open, bounded set in the d-dimensional Euclidean space $\mathbb{R}^d$. Let $x_i \in \Omega$ be the position of the i-th agent, and let $\mathbf{x} = (x_1, \ldots, x_N)$. Since the problem of transport additionally involves the dimension of time, we consider the positions of the agents (and of other variables) as parametrized by time steps $k \in \mathbb{N}$ in the discrete-time case, as $x_i(k)$, or time $t \in \mathbb{R}_{\geq 0}$ in the continuous-time case, as $x_i(t)$

We view the problem of multi-agent transport as one of updating the positions $x_i(k)$ (or $x_i(t)$ in continuous-time) according to a specified rule. Moreover, the specified rule is such that the motion of an agent is seen as the result of computations that generate the successive positions of the agent (or its instantaneous velocity). In other words, the physical motion is is merely the physical realization of the results of underlying computations. Once we take this computational perspective, that the problem of collective motion is essentially one of collective computation or information processing, characterizing the flow of information within the collective becomes crucial. This is essentially due to the fact that in the absence of a centralized decision maker, collective computation is enabled by information flow between the agents, and it is important to characterize the nature of this information flow. To this end, we define a graph $\mathscr{G} = (\mathscr{I}, \mathscr{E})$, where the edge-set $\mathscr{E}$ specifies the sensing/communication model for the collective and characterizes the information flow between the agents. In the ensuing applications, we often consider $\mathscr{G}$ to be a proximity graph over the set of agents (vertices) $\mathscr{I}$, given the corresponding set of positions of the agents $\{x_i\}_{i=1}^N$.

We recall that the (microscopic) configuration of the collective is specified as $\mathbf{x} = (x_1, \ldots, x_N) \in \Omega^N$, and for an $N \times N$ permutation matrix P, we take that the configuration $(P \otimes I_d)\mathbf{x}$ is equivalent to $\mathbf{x}$, since the agents are assumed to be identical. We thereby look for a permutation-invariant description of the collective, which leads us to specifying its configuration (macroscopically) as a probability distribution over Ω, as $\widehat{\mu}_{\mathbf{x}}^N = \frac{1}{N}\sum_{i=1}^N \delta_{x_i}$. If we are further given that the positions x_i of the agents are independently and identically distributed according to a (absolutely continuous) probability measure $\mu \in \mathscr{P}(\Omega)$ (where $\mathscr{P}(\Omega)$ is the space of probability measures over Ω), it follows from the Glivenko-Cantelli theorem [23] that in the limit $N \to \infty$ the discrete probability measure $\widehat{\mu}_{\mathbf{x}}^N$ converges uniformly, almost surely to the probability measure μ. In this way, $\widehat{\mu}_{\mathbf{x}}^N$ is seen as a discretization of the underlying measure μ. In particular, this constitutes the sampling perspective, in that $\widehat{\mu}_{\mathbf{x}}^N$ is generated by N i.i.d. samples of the probability measure μ. Alternatively, we can view $\widehat{\mu}_{\mathbf{x}}^N$ as being obtained by quantization of the measure μ, wherein μ is discretized over an equitable partition of Ω (i.e., the individual cells are of equal mass) to obtain the discrete measure $\widehat{\mu}_{\mathbf{x}}^N$. We explore these ideas in greater detail later in this chapter, but it suffices to say at present that for a large N number of agents, the probability measure μ approximates closely the macroscopic configuration $\widehat{\mu}_{\mathbf{x}}^N$ of the collective.

We begin by considering a (discrete-time) deterministic update rule for the transport, specified by a map $T_k : \Omega \to \Omega$ at time instant k, such that the position update for the i-th agent is given by $x_i(k+1) = T(x_i(k))$. While the microscopic configuration undergoes the update $\mathbf{x}(k) = (x_1(k), \ldots, x_N(k)) \mapsto (T_k(x_1(k)), \ldots, T_k(x_N(k))) = \mathbf{x}(k+1)$, the updated macroscopic configuration is obtained as the pushforward of the measure $\widehat{\mu}_{\mathbf{x}}^N(k)$ by the map T_k, given by $\widehat{\mu}_{\mathbf{x}}^N(k+1) = T_{k\#}\widehat{\mu}_{\mathbf{x}}^N(k)$.

2.3 Iterative proximal descent schemes in the space of probability measures

In this section, we set up iterative descent schemes in the space of probability measures $\mathscr{P}(\Omega)$ to converge to the minimizer of a convex functional. We then obtain descent schemes in Ω that transport probability measures in accordance with the descent schemes in the space of probability measures, and establish that they result in weak convergence to the minimizer.

Assumption 1. *$\Omega \subset \mathbb{R}^d$ is the closure of an open, bounded subset of $\mathbb{R}^d$ and is convex.*

Under Assumption 1, we have from Lemma 12 that the space of probability measures $\mathscr{P}(\Omega)$ equipped with the L^2-Wasserstein metric is geodesically convex (w.r.t. generalized geodesics). We now construct an l-smooth and strictly geodesically convex functional $F : \mathscr{P}(\Omega) \to \mathbb{R}$ with $\nabla\left(\frac{\delta F}{\delta \nu}\right) \cdot \mathbf{n} = 0$ on $\partial\Omega$ (where $\mathbf{n}$ is the outward normal to $\partial\Omega$) for all $\nu \in \mathscr{P}(\Omega)$, such that $\mu^* = \arg\min_{\nu \in \mathscr{P}(\Omega)} F(\nu)$ is absolutely continuous, and set up the following proximal recursion in $\mathscr{P}(\Omega)$ to converge to μ^* from any absolutely continuous $\mu_0 \in \mathscr{P}(\Omega)$:

$$\mu_{k+1} \in \arg\min_{\nu \in \mathscr{P}(\Omega)} \frac{1}{2\tau} W_2^2(\mu_k, \nu) + F(\nu). \tag{2.1}$$

Remark 1 (Neumann boundary condition)**.** *The Neumann boundary condition on the derivative of the functional F, $\nabla\left(\frac{\delta F}{\delta \nu}\right) \cdot \mathbf{n} = 0$ on $\partial\Omega$ (where $\mathbf{n}$ is the outward normal to $\partial\Omega$), yields a zero-flux boundary condition in the context of gradient descent w.r.t. F. This ensures conservation of mass and that the solutions of gradient descent w.r.t. F, which are sequences of measures, are contained in $\mathscr{P}(\Omega)$ as probability measures over Ω.*

Lemma 15 (Compactness and convexity of sublevel sets)**.** *The sublevel sets of the functional F are compact and convex.*

Proof. For any $\mu \in \mathscr{P}(\Omega)$ (with the Wasserstein metric W_2 on $\mathscr{P}(\Omega)$), we have that the sublevel set $\mathscr{S}(\mu) = \{\nu \in \mathscr{P}(\Omega) | F(\nu) \leq F(\mu)\}$ is closed, since F is continuous and $\mathscr{P}(\Omega)$ is closed and compact (from Prokhorov's theorem in Lemma 8 and Corollary 2 on the compactness of $\mathscr{P}(\Omega)$), which implies that $\mathscr{S}(\mu)$ is also compact since it is a closed subset of a compact set.

Moreover, for any $\nu_0, \nu_1 \in \mathscr{S}(\mu)$, and $\nu_t \in \mathscr{P}(\Omega)$ on the generalized geodesic between ν_0 to ν_1 with $t \in [0,1]$ (which follows from Lemma 12), we have from the (generalized) geodesic convexity of F that $F(\nu_t) \leq (1-t)F(\nu_0) + tF(\nu_1) \leq F(\mu)$ (since $F(\nu_0) \leq F(\mu)$ and $F(\nu_1) \leq F(\mu)$ by definition of $\mathscr{S}(\mu)$). This implies that $\nu_t \in \mathscr{S}(\mu)$ for any $t \in [0,1]$, from which we infer convexity of $\mathscr{S}(\mu)$. □

Corollary 3 (Completeness of sublevel sets). *It follows from Lemma 15 that the sublevel sets of F in the L^2-Wasserstein space are complete, in that every Cauchy sequence in a sublevel set of F is convergent.*

Lemma 16 (Strong convexity of objective functional). *The objective functional in* (2.1) *is $\left(\frac{1}{\tau} - l\right)$-strongly geodesically convex in $\mathscr{P}^r(\Omega)$ for $\tau < 1/l$.*

Proof. Since F is l-smooth, by applying Lemma 14 we get:

$$\left| \int_\Omega \langle \xi_2 - \xi_1, T_{\nu_1 \to \nu_2} - id \rangle \, d\nu_1 \right| \leq l W_2^2(\nu_1, \nu_2). \tag{2.2}$$

Let $G(\nu) = \frac{1}{2\tau} W_2^2(\mu, \nu) + F(\nu)$ and $\eta = \nabla \left(\frac{\delta G}{\delta \nu} \right) \Big|_\nu$. Also, let $\phi = \frac{1}{2} \left. \frac{\delta W_2^2(\mu,\nu)}{\delta \nu} \right|_\nu$ be the Kantorovich potential for the transport from ν to μ. We now have:

$$\begin{aligned}
\int_\Omega \langle \eta_2 - \eta_1, T_{\nu_1 \to \nu_2} - id \rangle \, d\nu_1 &= \int_\Omega \left\langle \frac{1}{\tau} \nabla \phi_2 - \frac{1}{\tau} \nabla \phi_1 - \xi_1 + \xi_2, T_{\nu_1 \to \nu_2} - id \right\rangle d\nu_1 \\
&= \frac{1}{\tau} \int_\Omega \langle \nabla \phi_2 - \nabla \phi_1, T_{\nu_1 \to \nu_2} - id \rangle \, d\nu_1 + \int_\Omega \langle \xi_2 - \xi_1, T_{\nu_1 \to \nu_2} - id \rangle \, d\nu_1 \\
&\geq \frac{1}{\tau} \int_\Omega \langle \nabla \phi_2 - \nabla \phi_1, T_{\nu_1 \to \nu_2} - id \rangle \, d\nu_1 - l W_2^2(\nu_1, \nu_2) \\
&= \left(\frac{1}{\tau} - l \right) W_2^2(\nu_1, \nu_2),
\end{aligned}$$

where the inequality above follows from (2.2). Moreover, we have used the fact that $\nabla\phi_2 - \nabla\phi_1 = T_{\nu_1\to\nu_2} - id$, which implies that $\int_\Omega \langle \nabla\phi_2 - \nabla\phi_1, T_{\nu_1\to\nu_2} - id\rangle\, d\nu_1 = W_2^2(\nu_1,\nu_2)$. Since $\tau < \frac{1}{l}$, we get that the functional G, which is the objective functional in (2.1), is strongly-convex with parameter $\frac{1}{\tau} - l$. □

Assumption 2 (Atomless sequence)**.** *We assume that for any $\tau < \frac{1}{l}$, the sequence $\{\mu_k\}_{k\in\mathbb{N}}$ generated by* (2.1) *is such that $\mu_k \in \mathscr{P}^r(\Omega)$ for all $k \in \mathbb{N}$.*

We remark here that sufficient regularity of the functional F would guarantee validity of Assumption 2, and we conjecture that this is indeed the case if F is twice continuously differentiable. Since we do not offer a proof for this claim, we retain Assumption 2 in establishing the following theorem:

Theorem 1 (Convergence of proximal recursion (2.1))**.** *Under Assumption 2, the proximal recursion* (2.1) *converges weakly to μ^* as $k \to \infty$.*

Proof. It follows from (2.1) that:

$$\frac{1}{2\tau}W_2^2(\mu_k,\mu_{k+1}) + F(\mu_{k+1}) \le F(\mu_k)$$
$$\Rightarrow F(\mu_{k+1}) \le F(\mu_k) - \frac{1}{2\tau}W_2^2(\mu_k,\mu_{k+1})$$

This implies that for $\mu_k \neq \mu_{k+1}$, we have $F(\mu_{k+1}) < F(\mu_k)$. Therefore, the sequence $\{F(\mu_k)\}$ is decreasing, and given an initial μ_0, the sequence $\{\mu_k\}_{k\in\mathbb{N}}$ is contained in the sublevel set $\mathscr{S}(\mu_0)$ of $F(\mu_0)$. From Lemma 15, we have that $\mathscr{S}(\mu_0)$ is convex and compact in the L^2-Wasserstein space $(\mathscr{P}(\Omega), W_2)$, and by Corollary 3, complete. Moreover, we have:

$$\frac{1}{2\tau}W_2^2(\mu_k,\mu_{k+1}) \le F(\mu_k) - F(\mu_{k+1}),$$

and by summing over $k \in \{0, \ldots, K\}$, we get:

$$\frac{1}{2\tau} \sum_{k=0}^{K} W_2^2(\mu_k, \mu_{k+1}) \leq F(\mu_0) - F(\mu_{K+1}).$$

Since $F(\mu_{K+1}) \leq F(\mu_0)$ and is bounded (as the sequence $\{\mu_k\}_{k\in\mathbb{N}}$ is contained in the sublevel set $\mathscr{S}(\mu_0)$ and F is bounded below by $F(\mu^*)$), we get that $\lim_{K\to\infty} \sum_{k=0}^{K} W_2^2(\mu_k, \mu_{k+1})$ is bounded, which implies that $\lim_{K\to\infty} W_2^2(\mu_K, \mu_{K+1}) = 0$. Therefore, the sequence $\{\mu_k\}_{k\in\mathbb{N}}$ is Cauchy, and since $\mathscr{S}(\mu_0)$ is complete, it is also convergent. Thus we have $\lim_{K\to\infty} W_2^2(\mu_K, \bar{\mu}) = 0$ for some $\bar{\mu} \in \mathscr{S}(\mu_0)$.

Now, since the sequence $\{\mu_k\}_{k\in\mathbb{N}}$ is generated by the iterative proximal descent scheme (2.1), we must have that $\bar{\mu}$ is a fixed point of (2.1). By the strong convexity of the objective functional in (2.1) and since the functional F is strictly convex in $(\mathscr{P}(\Omega), W_2)$, with μ^* as the only minimizer, we infer that μ^* is the only fixed point of (2.1) and $\bar{\mu} = \mu^*$. □

We now consider the following proximal recursion in Ω from an initial condition $x_0 \in \Omega$:

$$x_{k+1} = \arg\min_{z\in\Omega} \frac{1}{2\tau} |x_k - z|^2 + g_k(z), \tag{2.3}$$

where $\{g_k\}_{k\in\mathbb{N}}$ is a sequence of functions on Ω. We now let μ_0 be the probability distribution of the initial condition x_0 (denoted $x_0 \sim \mu_0$) and obtain an assignment for the sequence $\{g_k\}_{k\in\mathbb{N}}$ to target the recursion (2.1). In other words, we are interested in defining the dynamics in Ω that would result in the transport of the initial measure μ_0 according to the recursion (2.1).

Theorem 2 (Target dynamics in Ω)**.** *The proximal recursion* (2.1) *from* $\mu_0 \in \mathscr{P}^r(\Omega)$ *is obtained as the transport of* μ_0 *by* (2.3) *with* $x_0 \sim \mu_0$ *and the choice* $g_k = \frac{\delta F}{\delta \nu}\Big|_{\mu_{k+1}}$.

Proof. We rewrite the single-step update in (2.1) from a probability measure $\mu \in \mathscr{P}(\Omega)$ as follows, for the purposes of this proof:

$$\mu^+ = \arg \min_{\nu \in \mathscr{P}(\Omega)} \frac{1}{2\tau} W_2^2(\mu, \nu) + F(\nu). \tag{2.4}$$

It follows from Lemma 16 that the minimizer μ^+ in (2.4) is unique. Thus, for a one-parameter family of absolutely continuous probability measures $\{\nu_\varepsilon\}_{\varepsilon \in \mathbb{R}}$ generated by a transport vector field $\mathbf{v}_\varepsilon$ (according to $\partial_\varepsilon \nu_\varepsilon + \nabla \cdot (\nu_\varepsilon \mathbf{v}_\varepsilon) = 0$), with $\nu_0 = \mu^+$, we have:

$$\begin{aligned}
0 &= \frac{d}{d\varepsilon} \left(\frac{1}{2\tau} W_2^2(\mu, \nu_\varepsilon) + F(\nu_\varepsilon) \right) \Big|_{\varepsilon=0} \\
&= \frac{1}{\tau} \int_\Omega \langle \nabla \phi_{\mu^+ \to \mu}, \mathbf{v} \rangle \, d\mu^+ + \int_\Omega \langle \xi, \mathbf{v} \rangle \, d\mu^+ \\
&= \int_\Omega \left\langle \frac{1}{\tau} \nabla \phi_{\mu^+ \to \mu} + \xi, \mathbf{v} \right\rangle d\mu^+.
\end{aligned}$$

where $\xi = \nabla \left(\frac{\delta F}{\delta \nu} \right) \Big|_{\nu = \mu^+}$ and $\nabla \phi_{\mu^+ \to \mu} = id - T^{\mu^+ \to \mu}$, with $T^{\mu^+ \to \mu} : \Omega \to \Omega$ being the optimal transport map from μ^+ to μ. Since $\int_\Omega \langle \frac{1}{\tau} \nabla \phi_{\mu^+ \to \mu} + \xi, \mathbf{v} \rangle \, d\mu^+ = 0$ for all $\mathbf{v}$, it implies that $\frac{1}{\tau} \nabla \phi_{\mu^+ \to \mu} + \xi = 0$ (μ^+ a.e. in Ω), and we get:

$$\frac{1}{\tau} \nabla \phi_{\mu^+ \to \mu} + \xi = \frac{1}{\tau} \left(id - T^{\mu^+ \to \mu} \right) + \xi = 0,$$

which implies that:

$$T^{\mu^+ \to \mu} = id + \tau \xi.$$

Let $\varphi = \left(\frac{\delta F}{\delta \nu} \right) \Big|_{\nu = \mu^+}$. For any $y \in \Omega$ and $\tau < 1/l$, we have a unique y^+ defined as follows (which corresponds to the single-step update in (2.3)):

$$y^+ = \arg \min_{z \in \Omega} \frac{1}{2\tau} |y - z|^2 + \varphi(z). \tag{2.5}$$

If $y^+ \in \mathring{\Omega}$, it is a critical point of (2.5) and satisfies $y^+ = y - \tau \nabla \varphi(y^+)$. Since $\xi = \nabla \varphi$, we therefore have $y^+ = (id + \tau \xi)^{-1}(y)$. We note here that when the image of a $y \in \Omega$ under the $\arg\min$ map in (2.5) is a critical point in the interior of Ω, it is also the inverse image of y under the optimal transport map $T^{\mu^+ \to \mu}$.

Now, for a $y \in \mathring{\Omega}$, consider the objective function in (2.5) $\beta(z) = \frac{1}{2\tau}|y - z|^2 + \varphi(z)$. The inner product of its gradient at any point $z \in \partial\Omega$ on the boundary of Ω with the outward normal $\mathbf{n}$ to $\partial\Omega$ at z is given by $\nabla\beta \cdot \mathbf{n} = \left(\frac{1}{\tau}(z - y) + \nabla\varphi(z)\right) \cdot \mathbf{n} = \frac{1}{\tau}(z - y) \cdot \mathbf{n} > 0$, since $\nabla\varphi \cdot \mathbf{n} = 0$ and $z - y$ points outward to Ω (as $z \in \partial\Omega$ and $y \in \mathring{\Omega}$ and Ω is convex). This implies that there exists a point $\tilde{z}$ in the interior of Ω in a neighborhood of z such that $\beta(\tilde{z}) < \beta(z)$, which implies that z cannot be the minimizer. Thus, for any $y \in \mathring{\Omega}$, the minimizer of $\beta(z) = \frac{1}{2\tau}|y - z|^2 + \varphi(z)$ cannot lie on the boundary $\partial\Omega$, and must therefore lie in the interior of Ω and be a critical point of the objective function β. Now, when $y \in \partial\Omega$, if $y^+ \notin \mathring{\Omega}$, it must be that $y^+ = y$ (otherwise we obtain a contradiction for the same reason as above, that the inner product of $\nabla\beta$ with the outward normal would be strictly positive) and the argmin map (and the optimal transport map) is an identity in this case.

It therefore follows that for any $y \in \Omega$, its image y^+ under the argmin map is also its inverse image under the optimal transport map $T^{\mu^+ \to \mu}$. Therefore, we get that the $\arg\min$ map in (2.5) is also the inverse of the optimal transport map $T^{\mu^+ \to \mu}$. Thus, we have that the map $T^{\mu^+ \to \mu} = id + \tau\xi$ is well-defined and so is its inverse, we have that $\left(T^{\mu^+ \to \mu}\right)^{-1}_{\#} \mu = (id + \tau\xi)^{-1}_{\#} \mu = \mu^+$, and (2.4) is the lift to the space of probability measures of (2.5).

We therefore conclude that the proximal recursion (2.1) starting from μ_0 is the transport of μ_0 by (2.3) with $x_0 \sim \mu_0$. □

From a computational perspective, we note from Theorem 2 that to implement the proximal recursion (2.1) by the dynamics (2.3), we need to evaluate at a given time instant k the first variation $\frac{\delta F}{\delta \nu}$ at μ_{k+1}, the transported measure at the time instant $k+1$. To circumvent the need to evaluate the first variation one time step ahead, we alternatively consider the dynamics (2.3) with the choice $g_k = \frac{\delta F}{\delta \nu}\Big|_{\mu_k}$, where for a given time instant k, the function g_k is obtained by evaluating the first variation of F at the transported measure μ_k.

Theorem 3. *The sequence $\{\mu_k\}_{k\in\mathbb{N}}$ obtained as the transport of measure $\mu_0 \in \mathscr{P}^r(\Omega)$ by (2.3) with $\tau < 1/l$, $x_0 \sim \mu_0$ and the choice $g_k = \frac{\delta F}{\delta \nu}\Big|_{\mu_k}$, converges weakly to μ^* as $k \to \infty$.*

Proof. By the l-smoothness of the functional F and Lemma 14, we have:

$$\int_\Omega \left\langle \nabla \left(\frac{\delta F}{\delta \nu}\bigg|_{\mu_k} - \frac{\delta F}{\delta \nu}\bigg|_{\mu_{k+1}} \right), T_{\mu_{k+1}\to\mu_k} - id \right\rangle d\mu_{k+1} \leq l W_2^2(\mu_k, \mu_{k+1}).$$

We have that $T_{\mu_{k+1}\to\mu_k} = id + \tau \nabla \left(\frac{\delta F}{\delta \nu}\Big|_{\mu_k} \right)$, which from the above implies:

$$\tau \int_\Omega \left\langle \nabla \left(\frac{\delta F}{\delta \nu}\bigg|_{\mu_k} - \frac{\delta F}{\delta \nu}\bigg|_{\mu_{k+1}} \right), \nabla \left(\frac{\delta F}{\delta \nu}\bigg|_{\mu_k} \right) \right\rangle d\mu_{k+1} \leq l W_2^2(\mu_k, \mu_{k+1}).$$

and we therefore have:

$$\tau \int_\Omega \left\langle \nabla \left(\frac{\delta F}{\delta \nu}\bigg|_{\mu_{k+1}} \right), \nabla \left(\frac{\delta F}{\delta \nu}\bigg|_{\mu_k} \right) \right\rangle d\mu_{k+1} \geq \left(\frac{1}{\tau} - l \right) W_2^2(\mu_k, \mu_{k+1}).$$

Moreover, by convexity of the functional F, we have:

$$F(\mu_k) \geq F(\mu_{k+1}) + \int_\Omega \left\langle \nabla \left(\frac{\delta F}{\delta \nu}\bigg|_{\mu_{k+1}} \right), T_{\mu_{k+1}\to\mu_k} - id \right\rangle d\mu_{k+1}.$$

Subsituting in the above inequality, we get:

$$F(\mu_k) \geq F(\mu_{k+1}) + \left(\frac{1}{\tau} - l \right) W_2^2(\mu_k, \mu_{k+1}).$$

From the above inequality we get that μ_{k+1} belongs to the F-sublevel set of μ_k, and consequently that the sequence $\{\mu_k\}_{k\in\mathbb{N}}$ is contained in $\mathscr{S}(\mu_0)$, the F-sublevel set of μ_0. Following the same arguments as in the proof of Theorem 1, we get that the sequence $\{\mu_k\}_{k\in\mathbb{N}}$ is convergent and

$\lim_{K\to\infty} W_2^2(\mu_K, \bar{\mu}) = 0$ for some $\bar{\mu} \in \mathscr{S}(\mu_0)$.

Now, since $\{\mu_k\}_{k\in\mathbb{N}}$ is a sequence of measures generated by (2.3) with the choice $g_k = \frac{\delta F}{\delta \nu}\Big|_{\mu_k}$ and initial condition $x_0 \sim \mu_0$, in order to characterize the limit $\bar{\mu}$, we first formulate the corresponding iterative descent scheme in $\mathscr{P}(\Omega)$. The descent in $\mathscr{P}(\Omega)$ corresponding to (2.3) with the choice $g_k = \frac{\delta F}{\delta \nu}\Big|_{\mu_k}$ is given by:

$$\mu_{k+1} = \arg\min_{\nu\in\mathscr{P}(\Omega)} \frac{1}{2\tau} W_2^2(\mu_k, \nu) + \mathbb{E}_\nu \left[\frac{\delta F}{\delta \mu}\Big|_{\mu_k} \right]. \tag{2.6}$$

As the sequence $\{\mu_k\}_{k\in\mathbb{N}}$ is generated by (2.6), the limit $\bar{\mu}$ must be one of its fixed points. It again simply follows from the properties of F that the objective functional in (2.6) is strongly convex with $\mathbb{E}_\nu \left[\frac{\delta F}{\delta \mu}\Big|_{\mu_k} \right]$ being convex, and that the only fixed point of (2.6) is μ^*. We therefore have that $\bar{\mu} = \mu^*$. □

Theorem 3 allows us to consider the transport in $\mathscr{P}(\Omega)$ by the following proximal scheme in Ω for minimizing F:

$$x^+ = \arg\min_{z\in\Omega} \frac{1}{2\tau} |x - z|^2 + g(z), \tag{2.7}$$

where $x \sim \mu$ and $g = \frac{\delta F}{\delta \nu}\Big|_{\mu}$.

2.4 Multi-agent transport

We recall that the configuration of the collective is given by $\mathbf{x} = (x_1, \ldots, x_N)$, with $x_i \in \Omega$ for $i \in \{1, \ldots, N\}$. Let $\widehat{\mu}_{\mathbf{x}}^N = \frac{1}{N}\sum_{i=1}^N \delta_{x_i}$, be the discrete measure in $\mathscr{P}(\Omega)$ corresponding to the configuration $\mathbf{x}$. For a macroscopic description of the transport, we first let the macroscopic configuration be specified by an absolutely continuous probability measure, and since $\widehat{\mu}_{\mathbf{x}}^N$ is is not absolutely continuous, we define an absolutely continuous probability measure $\widehat{\mu}_{\mathbf{x}}^{h,N}$ through its density function using a smooth kernel K, as follows:

$$\widehat{\mu}_{\mathbf{x}}^{h,N}(x) = \frac{1}{N}\sum_{i=1}^{N} K(x - x_i, h), \tag{2.8}$$

where $h > 0$ is the bandwidth of the kernel. We allow $\widehat{\mu}_{\mathbf{x}}^{h,N}$ to denote both the absolutely continuous measure and its corresponding density function. We also denote, for $x \in \Omega$, $\widehat{\mu}_{x}^{h,1}$ simply by $\widehat{\mu}_{x}^{h}$. Thus, for $\mathbf{x} \in \Omega^N$, we have $\widehat{\mu}_{\mathbf{x}}^{h,N} = \sum_{i=1}^{N} \widehat{\mu}_{x_i}^{h}$.

Assumption 3 (Properties of K)**.** *For $h > 0$ and $z \in \Omega$, the probability measures $\widehat{\mu}_{z}^{h}$ defined using the kernel K as in* (2.8)*, we have:*

1. *Smoothness: The kernel $K(\cdot, h) \in C^{\infty}(\Omega)$ for every $h > 0$.*
2. *Monotonicity of support: For any $z \in \Omega$ and $h_1 < h_2$, we let $supp\left(\widehat{\mu}_{z}^{h_1}\right) \subset supp\left(\widehat{\mu}_{z}^{h_2}\right)$.*
3. *Containment: For every $h > 0$, there exists a set $\tilde{\Omega}_h \subset \Omega$ such that for $z \in \tilde{\Omega}_h$, the support of the measure $\widehat{\mu}_{z}^{h}$ satisfies $supp(\widehat{\mu}_{z}^{h}) \subset \Omega$. Moreover, $\lim_{h\to 0} \tilde{\Omega}_h = \Omega$ in Hausdorff distance.*
4. *Total variation convergence: With $\mathscr{M}$ being the space of measureable functions over Ω, we have $\lim_{h\to 0} \sup_{f\in\mathscr{M}} \left\{\int_{\Omega} f(z)K(x - z, h)\,\mathrm{dvol}(z) - f(x)\right\} = 0$.*

2.4.1 Discretization of $F : \mathscr{P}(\Omega) \to \mathbb{R}$

We now define a function $F^{h,N}$ as the discretization of the functional F, for $h > 0$, as follows:

$$F^{h,N}(\mathbf{x}) = F(\widehat{\mu}_{\mathbf{x}}^{h,N}). \tag{2.9}$$

We note that, clearly, $F^{h,N}$ is invariant under permutation, in that, for $\mathbf{x} \in \tilde{\Omega}_h^N$ and $P \in \mathbb{R}^N$ a permutation, we have $F^{h,N}(\mathbf{x}) = F^{h,N}((P \otimes I_d)\,\mathbf{x})$.

Lemma 17 (Convergence as $h \to 0,\ N \to \infty$)**.** *Under Assumption 3, for $x_i \sim \mu$ independent and identically distributed, we have $\lim_{h\to 0} \lim_{N\to\infty} F^{h,N}(x_1, \ldots, x_N) = F(\mu)$, μ-almost surely.*

Proof. We have $F^{h,N}(\mathbf{x}) = F(\widehat{\mu}_{\mathbf{x}}^{h,N})$ and that $\lim_{h\to 0}\lim_{N\to\infty}\widehat{\mu}_{\mathbf{x}}^{h,N} = \mu$ uniformly, almost surely (u.a.s) by the Glivenko-Cantelli theorem and Assumption 3. Therefore, by continuity of F we have:

$$\begin{aligned}
\lim_{h\to 0}\lim_{N\to\infty} F^{h,N}(\mathbf{x}) &= \lim_{h\to 0}\lim_{N\to\infty} F(\widehat{\mu}_{\mathbf{x}}^{h,N}) \\
&= \lim_{\widehat{\mu}_{\mathbf{x}}^{h,N}\to_{u.a.s}\,\mu} F(\widehat{\mu}_{\mathbf{x}}^{h,N}) \\
&= F(\mu),\ \mu-\text{almost surely.}
\end{aligned}$$

□

On the derivative of $F^{h,N}$

We begin by relating, through the following lemma, the derivative $\partial_1 F^{h,N}$ of the function $F^{h,N}$ to the derivative of the functional F:

Lemma 18 (Derivative of $F^{h,N}$)**.** *The derivative of the function $F^{h,N}$ satisfies:*

$$\partial_1 F^{h,N}(z,\xi) = \int_{supp\left(\widehat{\mu}_z^h\right)} \nabla\varphi^{h,N}\, d\widehat{\mu}_z^h,$$

where $d\widehat{\mu}_z^h = \rho_z^h$ dvol with $\rho_z^h(x) = K(x-z,h)$, and $\varphi^{h,N} = \frac{\delta F}{\delta \nu}\mid_{\widehat{\mu}^{h,N}}$.

Proof. Let $\mathbf{x}(t) = (x_1(t),\dots,x_N(t))$ be a curve parametrized by $t\in\mathbb{R}$ and $\dot{\mathbf{x}}(0) = \mathbf{v} = (\mathbf{v}_1,\dots,\mathbf{v}_N)$.

$$\frac{d}{dt}F^{h,N}(\mathbf{x}(0)) = \sum_{i=1}^{N}\left\langle \partial_i F^{h,N}, \mathbf{v}_i\right\rangle,$$

and using the fact that $F^{h,N}(\mathbf{x}) = F(\widehat{\mu}_{\mathbf{x}}^{h,N})$, we also get:

$$\begin{aligned}
\frac{d}{dt}F^{h,N}(\mathbf{x}(0)) &= \sum_{i=1}^{N}\int_{\Omega}\left\langle \nabla\varphi^{h,N}, \mathbf{v}_i\right\rangle\, d\widehat{\mu}_{x_i(0)}^h \\
&= \sum_{i=1}^{N}\left\langle \int_{\Omega}\nabla\varphi^{h,N}\, d\widehat{\mu}_{x_i(0)}^h, \mathbf{v}_i\right\rangle.
\end{aligned}$$

Since the above applies for all $\mathbf{v} = (\mathbf{v}_1, \dots, \mathbf{v}_N)$, we get:

$$\partial_i F^{h,N}(\mathbf{x}) = \int_\Omega \nabla \varphi^{h,N} \, d\widehat{\mu}^h_{x_i}.$$

The above can be rewritten as:

$$\partial_1 F^{h,N}(z, \xi) = \int_\Omega \nabla \varphi^{h,N} \, d\widehat{\mu}^h_z = \int_{\operatorname{supp}(\widehat{\mu}^h_z)} \nabla \varphi^{h,N} \, d\widehat{\mu}^h_z,$$

where $d\widehat{\mu}^h_z = \rho^h_z$ dvol with $\rho^h_z(x) = K(x - z, h)$, and $\varphi^{h,N} = \frac{\delta F}{\delta \nu} \mid_{\widehat{\mu}^{h,N}}$, which proves the statement of the lemma. □

Lemma 19 (α-smoothness of $F^{h,N}$). *If the function $\varphi = \frac{\delta F}{\delta \tilde{\mu}}\Big|_\mu$ is continuously differentiable on Ω for all μ, then there exists an $\alpha > 0$ such that $F^{h,N}$ is α-smooth.*

Proof. We have:

$$\begin{aligned}
\|\nabla F(\mathbf{y}) - \nabla F(\mathbf{x})\| &= \sqrt{\sum_{i=1}^N |\partial_1 F(y_i, \mathbf{y}_{-i}) - \partial_1 F(x_i, \mathbf{x}_{-i})|^2} \\
&= \sqrt{\sum_{i=1}^N \left| \int_\Omega \nabla \varphi^{h,N}_{\mathbf{y}} d\widehat{\mu}^h_{y_i} - \int_\Omega \nabla \varphi^{h,N}_{\mathbf{x}} d\widehat{\mu}^h_{x_i} \right|^2} \\
&= \sqrt{\sum_{i=1}^N \left| \int_\Omega \left[\nabla \varphi^{h,N}_{\mathbf{y}}(z + (y_i - x_i)) - \nabla \varphi^{h,N}_{\mathbf{x}}(z) \right] d\widehat{\mu}^h_{x_i}(z) \right|^2} \\
&\leq \int_\Omega \left| \nabla \varphi^{h,N}_{\mathbf{y}}(z) - \nabla \varphi^{h,N}_{\mathbf{x}}(z) \right| d\widehat{\mu}^{h,N}_{\mathbf{x}}(z) \\
&\quad + \sum_{i=1}^N \int_\Omega \left| \nabla \varphi^{h,N}_{\mathbf{y}}(z + (y_i - x_i)) - \nabla \varphi^{h,N}_{\mathbf{y}}(z) \right| d\widehat{\mu}^h_{x_i}(z) \\
&\leq l W_2(\widehat{\mu}^{h,N}_{\mathbf{x}}, \widehat{\mu}^{h,N}_{\mathbf{y}}) + M\|\mathbf{y} - \mathbf{x}\| \\
&\leq \alpha \|\mathbf{y} - \mathbf{x}\|,
\end{aligned}$$

where the penultimate inequality results from the l-smoothness of F and the continuous differentiability of φ over compact Ω (which implies φ has a Lipschitz-continuous gradient). Moreover, the final inequality results from the fact that $W_2(\widehat{\mu}_{\mathbf{x}}^{h,N}, \widehat{\mu}_{\mathbf{y}}^{h,N}) \leq \|\mathbf{y}-\mathbf{x}\|$. □

We now characterize the behavior of the discretization $F^{h,N}$ along the boundary through the following assumption:

Assumption 4 (Boundary conditions)**.** *The function $F^{h,N}$ satisfies the boundary condition $\partial_1 F^{h,N}(z,\xi)\cdot\mathbf{n}(z) = 0$ for $z \in \partial\tilde{\Omega}_h$ and all $\xi \in \tilde{\Omega}_h^{N-1}$.*

On the (non)convexity of $F^{h,N}$

The function $F^{h,N}:\Omega^N \to \mathbb{R}$ is in general non-convex, although it is the discretization of a strictly geodesically convex functional $F:\mathscr{P}(\Omega)\to\mathbb{R}$. This is because the notion of convexity of functions over Ω^N, which is the domain of the function $F^{h,N}$, is not equivalent to the notion of geodesic convexity over the space of probability measures over Ω, in that for $\mathbf{x},\mathbf{y}\in\Omega^N$ with $\sum_{i=1}^N \frac{1}{N}\delta_{x_i}, \sum_{i=1}^N \frac{1}{N}\delta_{y_i} \in \mathscr{P}(\Omega)$ being the corresponding discrete measures, the supports of the geodesics (when they exist) between $\sum_{i=1}^N \frac{1}{N}\delta_{x_i}$ and $\sum_{i=1}^N \frac{1}{N}\delta_{y_i}$ in $\mathscr{P}(\Omega)$ do not correspond to the straight line segment between $\mathbf{x}$ and $\mathbf{y}$ in Ω^N. This is a source of non-convexity of F^N.

Definition 12 (Cyclical monotonicity)**.** *A set $\Gamma\subset\Omega\times\Omega$ is cyclically monotone if any sequence $\{(x_i,y_i)\}_{i=1}^N$, with $(x_i,y_i)\in\Gamma$, satisfies:*

$$\sum_{i=1}^N |x_i-y_i|^2 \leq \sum_{i=1}^N |x_i-y_{\sigma(i)}|^2,$$

where σ is any permutation.

We define a subset $\Delta\subset\Omega^N$ (with $\delta>0$) as follows:

$$\Delta = \left\{ \mathbf{z} = (z_1,\ldots,z_N)\in\mathring{\Omega}^N \,\middle|\, |z_i-z_j| > \delta,\ \forall\, i\neq j \right\}.$$

For every $\mathbf{x} \in \Delta$, we now define a set $\Gamma_{\mathbf{x}} \subset \Omega^N$ such that for all $\mathbf{y} \in \Gamma_{\mathbf{x}}$, we have:

$$\sum_{i=1}^{N} |x_i - y_i|^2 \leq \sum_{i=1}^{N} |x_i - y_{\sigma(i)}|^2,$$

for any permutation σ. In other words, $\Gamma_{\mathbf{x}}$ is the subset of Ω^N such that for any $\mathbf{y} \in \Gamma_{\mathbf{x}}$, we have $(x_i, y_i) \in \Gamma \subset \Omega \times \Omega$, i.e., $\{(x_i, y_i)\}_{i=1}^N$ is cyclically monotone. We now establish through the following lemma that the set $\Gamma_{\mathbf{x}}$ contains an open neighborhood of $\mathbf{x}$:

Lemma 20 ($\Gamma_{\mathbf{x}}$ contains an open neighborhood of $\mathbf{x}$)**.** *For any $\mathbf{x} \in \Delta$, there exists an open neighborhood $\mathscr{N}(\mathbf{x}) \subset \Omega^N$ of $\mathbf{x}$ such that $\mathscr{N}(\mathbf{x}) \subset \Gamma_{\mathbf{x}}$.*

Proof. For $\mathbf{x} \in \Delta \subset \Omega^N$, let $\mathbf{y} \in \mathring{\Omega}^N$ such that for all $i \in \{1, \dots, N\}$, we have $y_i \in B_{\delta/2}(x_i)$, where $B_{\delta/2}(x_i)$ is the open $\delta/2$-ball centered at $x_i \in \Omega$. Thus, there exists an open neighborhood $\mathscr{N}(\mathbf{x}) \subset \Omega^N$ of $\mathbf{x}$, such that $\mathbf{y} \in \mathscr{N}(\mathbf{x})$. Now for any $j \in \{1, \dots, N\}$ with $j \neq i$, we have $|y_i - x_j| = |y_i - x_i + x_i - x_j| \geq |x_i - x_j| - |y_i - x_i| > \delta - \delta/2 > \delta/2$, since $|x_i - x_j| > \delta$ as $\mathbf{x} \in \Delta$ and $|y_i - x_i| < \delta/2$. Thus, among all (non-identity) permutations σ, we have:

$$\frac{1}{N} \sum_{i=1}^{N} |x_i - y_{\sigma(i)}|^2 > \frac{\delta^2}{4} > \frac{1}{N} \sum_{i=1}^{N} |x_i - y_i|^2.$$

Thus, we infer that $\mathbf{y} \in \Gamma_{\mathbf{x}}$. Since the same holds for any $\tilde{\mathbf{y}} \in \mathscr{N}(\mathbf{x})$, we get that $\mathscr{N}(\mathbf{x}) \subset \Gamma_{\mathbf{x}}$. □

It follows from Lemma 20 that for an $\mathbf{x} \in \Delta$ with a given $\delta > 0$, under an appropriate choice of $h > 0$, the supports of the components $\widehat{\mu}_{x_i}^h$ of the measure $\widehat{\mu}_{\mathbf{x}}^{h,N}$ can be made disjoint.

Lemma 21 (Relaxation to atomless measures)**.** *For any $\mathbf{x} \in \Delta$ and $\mathbf{y} \in \Gamma_{\mathbf{x}}$, there exists $\bar{h} > 0$ such that for $0 \leq h \leq \bar{h}$ and the probability measures $\widehat{\mu}_{\mathbf{x}}^{h,N}, \widehat{\mu}_{\mathbf{y}}^{h,N}$ defined in* (2.8)*, the optimal transport map $T_{\widehat{\mu}_{\mathbf{x}}^{h,N} \to \widehat{\mu}_{\mathbf{y}}^{h,N}}$ from $\widehat{\mu}_{\mathbf{x}}^{h,N}$ to $\widehat{\mu}_{\mathbf{y}}^{h,N}$ satisfies:*

$$\left(T_{\widehat{\mu}_{\mathbf{x}}^{h,N} \to \widehat{\mu}_{\mathbf{y}}^{h,N}} - id\right)(z) = y_i - x_i, \quad \forall\, z \in supp\left(\widehat{\mu}_{x_i}^h\right).$$

Proof. The proof is based on a generalization of Brenier's Theorem [87]. We consider convex functions $\chi_i : \Omega \to \mathbb{R}$ for $i \in \{1, \ldots, N\}$ defined by:

$$\chi_i(z) = \frac{1}{2}|z + y_i - x_i|^2.$$

We note that the gradient of χ_i, $\nabla\chi_i(z) = z + y_i - x_i$ defines a map that transports the measure $\widehat{\mu}^h_{x_i}$ to $\widehat{\mu}^h_{y_i}$ simply by translation. Since by the generalized Brenier's Theorem [87] such a transport map (defined by the gradient of a convex function) is unique and is also the optimal transport map, the statement of the lemma follows. □

Lemma 21 essentially establishes that for $\mathbf{x} \in \Delta$ and any $\mathbf{y} \in \Gamma_{\mathbf{x}}$, the optimal transport from $\widehat{\mu}^{h,N}_{\mathbf{x}}$ to $\widehat{\mu}^{h,N}_{\mathbf{y}}$ is simply achieved by the translation of components $\widehat{\mu}^h_{x_i}$ along the rays $y_i - x_i$ to $\widehat{\mu}^h_{y_i}$ for each $i \in \{1, \ldots, N\}$.

Corollary 4 (L^2-Wasserstein distance)**.** *For any* $\mathbf{x} \in \Delta$ *and* $\mathbf{y} \in \Gamma_{\mathbf{x}}$*:*

$$W_2^2\left(\widehat{\mu}^{h,N}_{\mathbf{x}}, \widehat{\mu}^{h,N}_{\mathbf{y}}\right) = \frac{1}{N}\sum_{i=1}^{N}|x_i - y_i|^2,$$

for any $0 < h \leq \bar{h}$*.*

With the above results and the convexity of the functional F, we establish the following comparison lemma:

Lemma 22 (Comparison lemma for $F^{h,N}$ on cyclically monotone sets)**.** *For any* $\mathbf{x} \in \Delta$*,* $h \in (0, \bar{h}]$ *and* $\mathbf{y} \in \Gamma_{\mathbf{x}}$*, we have:*

$$F^{h,N}(\mathbf{y}) \geq F^{h,N}(\mathbf{x}) + \left\langle \nabla F^{h,N}(\mathbf{x}), \mathbf{y} - \mathbf{x} \right\rangle.$$

Proof. From convexity of the functional F and the fact that $\mathbf{x} \in \Delta$ and $\mathbf{y} \in \Gamma_{\mathbf{x}}$, it follows that:

$$\begin{aligned}
F^{h,N}(\mathbf{y}) = F(\widehat{\mu}_{\mathbf{y}}^{h,N}) &\geq F(\widehat{\mu}_{\mathbf{x}}^{h,N}) + \int_{\Omega} \left\langle \nabla \varphi_{\mathbf{x}}^{h,N}, T_{\widehat{\mu}_{\mathbf{x}}^{h,N} \to \widehat{\mu}_{\mathbf{y}}^{h,N}} - id \right\rangle d\mu_{\mathbf{x}}^{h,N} \\
&= F(\widehat{\mu}_{\mathbf{x}}^{h,N}) + \sum_{i=1}^{N} \int_{\Omega} \left\langle \nabla \varphi_{\mathbf{x}}^{h,N}, T_{\widehat{\mu}_{\mathbf{x}}^{h,N} \to \widehat{\mu}_{\mathbf{y}}^{h,N}} - id \right\rangle d\mu_{x_i}^{h} \\
&= F(\widehat{\mu}_{\mathbf{x}}^{h,N}) + \sum_{i=1}^{N} \int_{\mathrm{supp}(\widehat{\mu}_{x_i}^{h})} \left\langle \nabla \varphi_{\mathbf{x}}^{h,N}, T_{\widehat{\mu}_{\mathbf{x}}^{h,N} \to \widehat{\mu}_{\mathbf{y}}^{h,N}} - id \right\rangle d\mu_{x_i}^{h} \\
&= F(\widehat{\mu}_{\mathbf{x}}^{h,N}) + \sum_{i=1}^{N} \int_{\mathrm{supp}(\widehat{\mu}_{x_i}^{h})} \left\langle \nabla \varphi_{\mathbf{x}}^{h,N}, y_i - x_i \right\rangle d\mu_{x_i}^{h} \\
&= F(\widehat{\mu}_{\mathbf{x}}^{h,N}) + \sum_{i=1}^{N} \left\langle \int_{\mathrm{supp}(\widehat{\mu}_{x_i}^{h})} \nabla \varphi_{\mathbf{x}}^{h,N} d\mu_{x_i}^{h} \, , \, y_i - x_i \right\rangle \\
&= F^{h,N}(\mathbf{x}) + \sum_{i=1}^{N} \left\langle \partial_1 F^{h,N}(x_i, \mathbf{x}_{-i}), y_i - x_i \right\rangle,
\end{aligned}$$

thereby establishing the claim. □

We remark here that $F^{h,N}$ is convex in the limited sense established by the comparison result in Lemma 22, and this does not necessarily generalize to the entire domain Ω^N, due to which the function $F^{h,N}$ can be non-convex in general.

Estimate on the minimum value of $F^{h,N}$

From the l-smoothness of the functional F, we have:

$$F(\widehat{\mu}_{\mathbf{x}}^{h,N}) - F(\mu^*) \leq \frac{l}{2} W_2^2 \left(\widehat{\mu}_{\mathbf{x}}^{h,N}, \mu^* \right).$$

Moreover, from convexity of F, we get:

$$F(\widehat{\mu}_{\mathbf{x}}^{h,N}) - F(\mu^*) \geq 0.$$

We generalize from the above inequalities, with $m \geq 0$:

$$\frac{m}{2} W_2^2 \left(\widehat{\mu}_{\mathbf{x}}^{h,N}, \mu^* \right) \leq F(\widehat{\mu}_{\mathbf{x}}^{h,N}) - F(\mu^*) \leq \frac{l}{2} W_2^2 \left(\widehat{\mu}_{\mathbf{x}}^{h,N}, \mu^* \right).$$

In particular, if the functional F is strongly convex, we will have $0 < m < l$. Since $F(\widehat{\mu}_{\mathbf{x}}^{h,N}) = F^{h,N}(\mathbf{x})$, we therefore have the following estimate on the minimum value of $F^{h,N}$:

$$\frac{m}{2} \min_{\mathbf{x} \in \Omega^N} W_2^2 \left(\widehat{\mu}_{\mathbf{x}}^{h,N}, \mu^* \right) \leq \min_{\mathbf{x} \in \Omega^N} F^{h,N}(\mathbf{x}) - F(\mu^*) \leq \frac{l}{2} \min_{\mathbf{x} \in \Omega^N} W_2^2 \left(\widehat{\mu}_{\mathbf{x}}^{h,N}, \mu^* \right). \tag{2.10}$$

2.4.2 Multi-agent proximal descent

We formulate the proximal descent on the function $F^{h,N}$ as follows:

$$\mathbf{x}^+ = \arg\min_{\mathbf{z} \in \tilde{\Omega}^N} \frac{1}{2\tau} \|\mathbf{x} - \mathbf{z}\|^2 + F^{h,N}(\mathbf{z}), \tag{2.11}$$

where $\|\mathbf{x} - \mathbf{z}\|^2 = \sum_{i=1}^N |x_i - z_i|^2$. We now establish strong convexity of the proximal descent objective function in (2.11) through the following lemma, under α-smoothness of $F^{h,N}$ from Lemma 19:

Lemma 23. *The objective function in* (2.11) *is* $\left(\frac{1}{\tau} - \alpha\right)$*-strongly convex for* $\tau < \frac{1}{\alpha}$.

Proof. From Lemma 19 on α-smoothness of $F^{h,N}$, we have:

$$\left| \left\langle \nabla F^{h,N}(\mathbf{y}) - \nabla F^{h,N}(\mathbf{x}), \mathbf{y} - \mathbf{x} \right\rangle \right| \leq \alpha \|\mathbf{y} - \mathbf{x}\|^2.$$

With $G_{\mathbf{x}}^{h,N}(\mathbf{z}) = \frac{1}{2\tau}\|\mathbf{x}-\mathbf{z}\|^2 + F^{h,N}(\mathbf{z})$ being the objective function of (2.11), we have:

$$
\begin{aligned}
&\left\langle \nabla G_{\mathbf{x}}^{h,N}(\mathbf{z}_1) - G_{\mathbf{x}}^{h,N}(\mathbf{z}_2), \mathbf{z}_1 - \mathbf{z}_2 \right\rangle \\
&= \left\langle \frac{1}{\tau}(\mathbf{z}_1 - \mathbf{x}) + \nabla F^{h,N}(\mathbf{z}_1) - \frac{1}{\tau}(\mathbf{z}_2 - \mathbf{x}) - \nabla F^{h,N}(\mathbf{z}_2) \ , \ \mathbf{z}_1 - \mathbf{z}_2 \right\rangle \\
&= \left\langle \frac{1}{\tau}(\mathbf{z}_1 - \mathbf{z}_2) + \nabla F^{h,N}(\mathbf{z}_1) - \nabla F^{h,N}(\mathbf{z}_2) \ , \ \mathbf{z}_1 - \mathbf{z}_2 \right\rangle \\
&= \frac{1}{\tau}\|\mathbf{z}_1 - \mathbf{z}_2\|^2 + \left\langle \nabla F^{h,N}(\mathbf{z}_1) - \nabla F^{h,N}(\mathbf{z}_2) \ , \ \mathbf{z}_1 - \mathbf{z}_2 \right\rangle \\
&\geq \frac{1}{\tau}\|\mathbf{z}_1 - \mathbf{z}_2\|^2 - \alpha\|\mathbf{z}_1 - \mathbf{z}_2\|^2 \\
&= \left(\frac{1}{\tau} - \alpha\right)\|\mathbf{z}_1 - \mathbf{z}_2\|^2,
\end{aligned}
$$

thereby establishing the claim. □

Now, with $\mathbf{x}_{-i} = (x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_N) \in \tilde{\Omega}^{N-1}$, we can write:

$$
F^{h,N}(x_1, \ldots, x_N) = \frac{1}{N}\sum_{i=1}^{N} F^{h,N}(x_1, \ldots, x_N) = \frac{1}{N}\sum_{i=1}^{N} F^{h,N}(x_i, \mathbf{x}_{-i}).
$$

The proximal gradient descent (2.11) can be decomposed into the following agent-wise update scheme:

$$
x_i^+ = \arg\min_{z \in \tilde{\Omega}} \frac{1}{2\tau}|x_i - z|^2 + F^{h,N}(z, \mathbf{x}_{-i}^+).
$$

Note that the above scheme requires as argument $\mathbf{x}_{-i}^+$. From a computational perspective, to implement the above algorithm, every agent i at time instant k requires the positions of the other agents at time $k+1$, which poses a hurdle for implementation. Therefore, to avoid this problem we instead consider the following proximal descent scheme for every agent i:

$$
x_i^+ = \arg\min_{z \in \tilde{\Omega}} \frac{1}{2\tau}|x_i - z|^2 + F^{h,N}(z, \mathbf{x}_{-i}). \tag{2.12}
$$

It follows from Lemma 23 that the objective function in (2.12) is also strongly convex, and thereby has a unique minimizer. We now present the following result on the convergence of (2.12) to the local minimizers of $F^{h,N}$:

Theorem 4 (Convergence of (2.12) to critical points of $F^{h,N}$)**.** *For $\tau < \frac{2}{3\alpha}$, the sequence $\{\mathbf{x}(k)\}_{k\in\mathbb{N}}$ generated by the update scheme* (2.12) *converges to a critical point $\mathbf{x}^*$ of $F^{h,N}$ that is not a local maximizer, for all initial conditions $\mathbf{x}(0) \in \tilde{\Omega}^N$. Moreover, if the critical point $\mathbf{x}^* \in \Delta$, it is a local minimizer.*

Proof. We first consider the objective function in (2.12), $\eta_{x_i}(z) = \frac{1}{2\tau}|x_i - z|^2 + F^{h,N}(z, \mathbf{x}_{-i})$. The inner product of the gradient of η on the boundary $\partial\tilde{\Omega}$ with the outward normal $\tilde{\mathbf{n}}$ to $\partial\Omega$, is given by:

$$\begin{aligned}\nabla\eta(z)\cdot\tilde{\mathbf{n}}(z) &= \frac{1}{\tau}(z - x_i)\cdot\tilde{\mathbf{n}}(z) + \partial_1 F^{h,N}(z, \mathbf{x}_{-i})\cdot\tilde{\mathbf{n}}(z)\\ &= \frac{1}{\tau}(z - x_i)\cdot\tilde{\mathbf{n}}(z)\\ &\geq 0,\end{aligned}$$

with the inequality being strict when $x_i \notin \partial\tilde{\Omega}$. This implies that the $x_i^+ \in \partial\tilde{\Omega}$ cannot be a minimizer if $x_i \notin \partial\tilde{\Omega}$, and if $x_i \in \partial\tilde{\Omega}$, we will have $x_i^+ = x_i$. In both cases, we will then have that the minimizer x_i^+ is also a critical point of the function η. This allows us to express (2.12) equivalently by:

$$x_i^+ = x_i - \tau\partial_1 F^{h,N}(x_i^+, \mathbf{x}_{-i}). \tag{2.13}$$

We note that in the limit $\tau \to 0$, we get a gradient flow that can be shown to converge to a critical point of $F^{h,N}$. We therefore hope that this property is preserved over a neighborhood of $\tau = 0$. In what follows, we establish that this is indeed the case and provide a sufficient strict upper bound on τ for which the property is preserved.

From α-smoothness of $F^{h,N}$, we get:

$$\left| F(\mathbf{x}^+) - F(\mathbf{x}) - \sum_{i=1}^{N} \left\langle \partial_1 F^{h,N}(x_i, \mathbf{x}_{-i}), x_i^+ - x_i \right\rangle \right| \le \frac{\alpha}{2} \|\mathbf{x}^+ - \mathbf{x}\|^2.$$

We can rewrite the above as:

$$\begin{aligned} &\left| F(\mathbf{x}^+) - F(\mathbf{x}) - \sum_{i=1}^{N} \left\langle \partial_1 F^{h,N}(x_i^+, \mathbf{x}_{-i}), x_i^+ - x_i \right\rangle \right. \\ &\left. - \sum_{i=1}^{N} \left\langle \partial_1 F^{h,N}(x_i, \mathbf{x}_{-i}) - \partial_1 F^{h,N}(x_i^+, \mathbf{x}_{-i}), x_i^+ - x_i \right\rangle \right| \le \frac{\alpha}{2} \|\mathbf{x}^+ - \mathbf{x}\|^2. \end{aligned}$$

We now have $-\sum_{i=1}^{N} \left\langle \partial_1 F^{h,N}(x_i^+, \mathbf{x}_{-i}), x_i^+ - x_i \right\rangle = \frac{1}{\tau} \|\mathbf{x}^+ - \mathbf{x}\|^2$ and by α-smoothness of $F^{h,N}$ again that:

$$\left| \sum_{i=1}^{N} \left\langle \partial_1 F^{h,N}(x_i, \mathbf{x}_{-i}) - \partial_1 F^{h,N}(x_i^+, \mathbf{x}_{-i}), x_i^+ - x_i \right\rangle \right| \le \alpha \|\mathbf{x}^+ - \mathbf{x}\|^2.$$

From the above inequalities, we therefore get:

$$F^{h,N}(\mathbf{x}^+) \le F^{h,N}(\mathbf{x}) - \left(\frac{1}{\tau} - \frac{3\alpha}{2} \right) \|\mathbf{x}^+ - \mathbf{x}\|^2.$$

Thus, for $\tau < \frac{2}{3\alpha}$, when every agent follows the update (2.12), we get a descent in $F^{h,N}$, and $\mathbf{x}^+$ belongs to the $F^{h,N}$-sublevel set of $\mathbf{x}$. We can express the above inequality for any time instant $k \in \mathbb{N}$ as:

$$F(\mathbf{x}(k+1)) \le F(\mathbf{x}(k)) - \left(\frac{1}{\tau} - \frac{3\alpha}{2} \right) \|\mathbf{x}(k+1) - \mathbf{x}(k)\|^2.$$

Summing over $k = 0, \ldots, K-1$, we get:

$$F(\mathbf{x}(K)) \le F(\mathbf{x}(0)) - \left(\frac{1}{\tau} - \frac{3\alpha}{2} \right) \sum_{k=1}^{K} \|\mathbf{x}(k) - \mathbf{x}(k-1)\|^2,$$

and it follows that:

$$\sum_{k=1}^{K} \|\mathbf{x}(k) - \mathbf{x}(k-1)\|^2 \leq \left(\frac{1}{\frac{1}{\tau} - \frac{3\alpha}{2}} \right) (F(\mathbf{x}(0)) - F(\mathbf{x}(K))),$$

and since the sequence $\{\mathbf{x}(k)\}_{k\in\mathbb{N}}$ belongs to the $F^{h,N}$-sublevel set of $\mathbf{x}(0)$ (for all $\mathbf{x}(0) \in \tilde{\Omega}^N$), which is a subset of $\tilde{\Omega}^N$ (compact), it is precompact. By the boundedness above, in the limit $K \to \infty$, we get that $\lim_{K\to\infty} \|\mathbf{x}(K) - \mathbf{x}(K-1)\|^2 = 0$, which implies that the sequence is Cauchy. Moreover, since the sequence is contained in a compact $\tilde{\Omega}^N$, which is also complete, we get that $\{\mathbf{x}(k)\}_{k\in\mathbb{N}}$ is convergent. Let the limit $\lim_{k\to\infty} \mathbf{x}(k) = \mathbf{x}^*$, and from (2.13) we thereby get:

$$\partial_1 F^{h,N}(x_i^*, \mathbf{x}_{-i}^*) = 0, \quad \forall\, i \in \{1, \ldots, N\},$$

which implies that $\nabla F^{h,N}(\mathbf{x}^*) = 0$. Therefore, the sequence $\{\mathbf{x}(k)\}_{k\in\mathbb{N}}$ converges to a critical point of $F^{h,N}$ which cannot be a local maximizer since $\{F(\mathbf{x}(k))\}_{k\in\mathbb{N}}$ is decreasing and consequently every neighborhood of $\mathbf{x}^*$ contains atleast one point with a higher value of $F^{h,N}$.

Moreover, if $\mathbf{x}^* \in \Delta$, from Lemmas 22 and 20, we get that there exists an open ball $B(\mathbf{x}^*) \subset \Omega^N$ such that for all $\mathbf{x} \in B(\mathbf{x}^*)$, we have $F(\mathbf{x}) \geq F(\mathbf{x}^*)$, which implies that $\mathbf{x}^*$ must be a local minimizer. □

Theorem 4 establishes that the multi-agent proximal descent converges to critical points of the function $F^{h,N}$, which is a discretization of the functional F, and that those critical points are not local maximizers. This is a weaker result than Theorem 2, which established convergence of transport of measures by the scheme (2.7) to the global minimizer μ^* of F. The weakening of the guarantee is due to the discretization of F, involved in defining the multi-agent transport scheme. However, we can still hope to achieve convergence to the global minimizer in the limit $N \to \infty$, thereby guaranteeing best performance asymptotically. In what follows, we show that this is indeed the case and that we can retrieve the property of convergence to the global minimizer μ^* of F in the limit $N \to \infty$ of the multi-agent proximal descent.

$N \to \infty$ and continuous-time limit

We now begin with a macroscopic model of multi-agent transport in continuous-time under a vector field $\mathbf{v} \in L^\infty([0,T] \times Lip(\Omega)^d)$, where the agent dynamics are given by $\dot{x}_i(t) = \mathbf{v}(t, x_i(t))$ and $x_i(0) \sim_{i.i.d} \mu_0$, for $i \in \{1, \ldots, N\}$.

Proposition 1 (Continuity equation as the model of transport in the continuous-time and $N \to \infty$ limit)**.** *The sequence of solutions $\{\mathbf{x}^N = (x_1, \ldots, x_N)\}_{N \in \mathbb{N}}$ (with $x_i(0) \sim_{i.i.d} \mu_0$ for $i \in \mathbb{N}$) to the multi-agent transport by the vector field $\mathbf{v} \in L^\infty([0,T] \times Lip(\Omega)^d)$ converge to a solution μ (where $d\mu(t) = \rho(t)\,\mathrm{dvol}$ and $\mu(0) = \mu_0$) of the continuity equation:*

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{v}) = 0, \tag{2.14}$$

satisfied in the distributional sense.

Proof. Let the flow corresponding to the vector field $\mathbf{v} \in L^\infty([0,T] \times Lip(\Omega)^d)$ be given by:

$$\partial_t X^t(x) = \mathbf{v}(t, X^t(x)),$$

with $X^0(x) = x$, and let $\mu(t) = X^t_{\#}\mu_0$ be the pushforward of μ_0 by the flow at time $t \in \mathbb{R}_{\geq 0}$. Now, with $d\widehat{\mu}^N_{\mathbf{x}(t)}(z) = \frac{1}{N}\sum_{i=1}^N \delta_{x_i(t)}(z)\,\mathrm{dvol}(z)$ (δ here is the Dirac Delta function), we can write:

$$\begin{aligned} \frac{\partial \rho^{h,N}_{\mathbf{x}}}{\partial t}(t,x) &= -\frac{1}{N}\sum_{i=1}^N \nabla_x K(x - x_i(t), h) \cdot \mathbf{v}(t, x_i) \\ &= -\int_\Omega \mathbf{v}(t,z) \cdot \nabla_x K(x - z, h)\; d\widehat{\mu}^N_{\mathbf{x}(t)}(z). \end{aligned}$$

We note that, by the Glivenko-Cantelli Theorem, the measure $\widehat{\mu}^N_{\mathbf{x}(\mathbf{0})}$ converges uniformly almost surely to μ_0. This implies that $\widehat{\mu}^N_{\mathbf{x}(t)}$ converges uniformly almost surely to the pushforward

$\mu(t) = X^t_{\#}\mu_0$ at time t. We therefore have:

$$
\begin{aligned}
\lim_{h\to 0}\lim_{N\to\infty} \rho^{h,N}_{\mathbf{x}}(t,x) &= \lim_{h\to 0}\lim_{N\to\infty} \int_\Omega K(x-z,h)\; d\widehat{\mu}^N_{\mathbf{x}(t)}(z) \\
&= \lim_{h\to 0}\left(\lim_{N\to\infty}\int_\Omega K(x-z,h)\; d\widehat{\mu}^N_{\mathbf{x}(t)}(z)\right) \\
&=_{a.s.} \lim_{h\to 0}\int_\Omega K(x-z,h)\; d\mu(t,z) \\
&= \lim_{h\to 0}\int_\Omega \rho(t,z)K(x-z,h)\; \mathrm{dvol}(z) \\
&= \rho(t,x).
\end{aligned}
\tag{2.15}
$$

From the above, we get that for a smooth test function $\zeta \in C^\infty([0,T]\times\Omega)$ such that $\zeta(0) = 0 = \zeta(T)$, we have:

$$
\begin{aligned}
&\int_{[0,T]}\int_\Omega \frac{\partial\zeta}{\partial t}\rho\,\mathrm{dvol}\; dt =_{a.s.} \lim_{h\to 0}\lim_{N\to\infty}\int_{[0,T]}\int_\Omega \frac{\partial\zeta}{\partial t}\rho^{h,N}_{\mathbf{x}}\,\mathrm{dvol}\; dt \\
&= \lim_{h\to 0}\lim_{N\to\infty} -\int_{[0,T]}\int_\Omega \zeta\frac{\partial\rho^{h,N}_{\mathbf{x}}}{\partial t}\,\mathrm{dvol}\; dt \\
&= \lim_{h\to 0}\lim_{N\to\infty}\int_{[0,T]}\int_\Omega \zeta(t,x)\int_\Omega \mathbf{v}(t,z)\cdot\nabla_x K(x-z,h)\; d\widehat{\mu}^N_{\mathbf{x}(t)}(z)\,\mathrm{dvol}(x)\; dt \\
&= \lim_{h\to 0}\lim_{N\to\infty} -\int_{[0,T]}\int_\Omega\int_\Omega \rho^{h,N}_{\mathbf{x}}(t,z)\nabla_x\zeta(t,x)\cdot\mathbf{v}(t,z)K(x-z,h)\; \mathrm{dvol}(z)\; \mathrm{dvol}(x)\; dt \\
&= -\int_{[0,T]}\int_\Omega \nabla\zeta\cdot\mathbf{v}\rho\,\mathrm{dvol}\; dt.
\end{aligned}
\tag{2.16}
$$

The above is the distributional sense of the continuity equation (2.14). □

We now obtain the vector field corresponding to the multi-agent transport scheme (2.12) in the $N\to\infty$ and continuous-time limit. We first rewrite (2.12) as follows:

$$
\begin{aligned}
x^+ &= \arg\min_{z\in\tilde{\Omega}} \frac{1}{2\tau}|x-z|^2 + F^{h,N}(z,\xi) \\
x &\sim \frac{1}{N}\sum_{i=1}^N \delta_{x_i}, \quad \xi \sim \otimes^N_{\substack{i=1,\\ x_i\neq x}} \frac{1}{N}\delta_{x_i},
\end{aligned}
$$

where $\otimes_{\substack{i=1,\\ x_i\neq x}}^{N} \frac{1}{N}\delta_{x_i}$ is the product measure describing the independent coupling between the discrete measures $\frac{1}{N}\delta_{x_i}$. From the arguments in the proof of Theorem 4, the above update scheme can be expressed equivalently as:

$$x^{+} = x - \tau\partial_1 F^{h,N}(x^{+},\xi)$$
$$x \sim \frac{1}{N}\sum_{i=1}^{N}\delta_{x_i}, \quad \xi \sim \otimes_{\substack{i=1,\\ x_i\neq x}}^{N} \frac{1}{N}\delta_{x_i}.$$

For $x_i \sim_{i.i.d} \mu$, we know that $\lim_{N\to\infty}\frac{1}{N}\sum_{i=1}^{N}\delta_{x_i} = \mu$ uniformly, almost surely. Now, in the limit $N\to\infty$ and $h\to 0$, with sufficient regularity of the kernel K ensuring uniform integrability of $\{\nabla\varphi^{h,N}\}$, we will have:

$$\lim_{h\to 0}\lim_{N\to\infty}\partial_1 F^{h,N}(z,\xi) \;=\; \lim_{h\to 0}\lim_{N\to\infty}\int_{\Omega}\nabla\varphi^{h,N}d\mu_z^h \;=\; \nabla\varphi(z),$$

with $\xi \sim \otimes_{\substack{i=1,\\ x_i\neq x}}^{N} \frac{1}{N}\delta_{x_i}$ and $\varphi = \frac{\delta F}{\delta \nu}\mid_{\mu}$ above, and we therefore get:

$$x^{+} = x - \tau\nabla\varphi(x^{+}),$$
$$x \sim \mu.$$

or equivalently:

$$x^{+} = \arg\min_{z\in\Omega}\frac{1}{2\tau}|x-z|^2 + \varphi(z)$$
$$x \sim \mu. \tag{2.17}$$

We thereby retrieve (2.7) from (2.12) in the $N\to\infty$ and $h\to 0$ limit. We know from Theorem 2 that transport of a probability measure μ_0 by (2.7) (which is identical to (2.17)) is guaranteed to converge to the global minimizer μ^* of F.

Informally, we see that as $\tau \to 0$ in (2.17), we have $x^+ \to x$ and we let:

$$\mathbf{v}(x) = \lim_{\tau \to 0} \frac{x^+ - x}{\tau} = -\nabla \varphi(x).$$

We can therefore expect the solutions to (2.17) converge to the solution of the gradient flow under the vector field $\mathbf{v} = -\nabla \varphi$. We now show, in a weak sense, that the above reasoning holds.

Proposition 2 (Continuous-time limit of (2.17))**.** *For every decreasing sequence $\{\tau_n\}_{n \in \mathbb{N}}$ satisfying $\tau_0 < \frac{1}{l}$ and $\lim_{n\to\infty} \tau_n = 0$, the sequence of solutions $\{x^n\}_{n\in\mathbb{N}}$ to* (2.17) *(with $\tau = \tau_n$) contains a convergent subsequence, and the limit is a weak solution to the gradient flow given by:*

$$\partial_t X^t(x) = -\nabla \varphi_t(X^t(x)), \tag{2.18}$$

with $X^0(x) = x$, $\mu(t) = X^t_\# \mu_0$ and $\varphi_t = \frac{\delta F}{\delta \nu}\Big|_{\mu(t)}$.

Proof. We begin by noting from (2.17) that $\frac{1}{2\tau}|x^+ - x|^2 \leq \varphi(x) - \varphi(x^+)$ (where $\tau < \frac{1}{l}$, $x \sim \mu$ and $\varphi = \frac{\delta F}{\delta \nu}\Big|_{\mu}$). Now let $\{\tau_n\}_{n\in\mathbb{N}}$ be a decreasing sequence such that $\tau_0 < \frac{1}{l}$ and $\lim_{n\to\infty} \tau_n = 0$. Let $\{x^n\}$ be the sequence of solutions to (2.17) starting from the same initial condition x_0. We note that $x^n(k) \in \Omega$ for all $n, k \in \mathbb{N}$. We now define continuous curves $\bar{x}^n$, such that $\bar{x}^n(t) = \left(1 + \lfloor \frac{t}{\tau_n} \rfloor - t\right) x^n(\lfloor \frac{t}{\tau_n} \rfloor) + \left(t - \lfloor \frac{t}{\tau_n} \rfloor\right) x^n(\lfloor \frac{t}{\tau_n} \rfloor + 1)$. From compactness of Ω, we get that the sequence $\{\bar{x}^n\}$ is uniformly bounded. Moreover, we have for $0 \leq k' \leq k$ that:

$$\begin{aligned}
\left|x^n(k) - x^n(k')\right| &\leq \sum_{m=k'+1}^{k} |x^n(m) - x^n(m-1)| \\
&\leq \left(\sum_{m=k'+1}^{k} |x^n(m) - x^n(m-1)|^2\right)^{1/2} \left(k-k'\right)^{1/2} \\
&= \sqrt{2\tau_n}\left(\sum_{m=k'+1}^{k} \frac{1}{2\tau_n} |x^n(m) - x^n(m-1)|^2\right)^{1/2} \left(k-k'\right)^{1/2} \\
&\leq \sqrt{2\tau_0}\left(\sum_{m\in\mathbb{N}} \frac{1}{2\tau_n} |x^n(m) - x^n(m-1)|^2\right)^{1/2} \left(k-k'\right)^{1/2} \\
&\leq \sqrt{2\tau_0}\left(\varphi_0(x_0) - \lim_{m\to\infty} \varphi_m(x^n(m))\right)^{1/2} \left(k-k'\right)^{1/2},
\end{aligned}$$

where $\phi_m = \frac{\delta F}{\delta \nu}\Big|_{\mu(m)}$, $\mu(m) = T_{m-1\#}\ldots T_{0\#}\mu_0$ and $T_k = (id + \tau_n \nabla \varphi_{k-1})^{-1}$. From Theorem 2, it follows that $\lim_{m\to\infty} \varphi_m = \frac{\delta F}{\delta \mu}\Big|_{\mu^*} = C$, a constant function. We therefore have:

$$\left|x^n(k) - x^n(k')\right| \leq \sqrt{2\tau_0}\left(\varphi_0(x_0) - C\right)^{1/2} \left(k-k'\right)^{1/2}. \tag{2.19}$$

It now follows for $0 \leq t' \leq t$ that:

$$\begin{aligned}
\left|\bar{x}^n(t) - \bar{x}^n(t')\right| &= \left|\bar{x}^n(t) - x^n\left(\left\lfloor \frac{t}{\tau_n} \right\rfloor\right) + x^n\left(\left\lfloor \frac{t}{\tau_n} \right\rfloor\right) - x^n\left(\left\lfloor \frac{t'}{\tau_n} \right\rfloor + 1\right) + x^n\left(\left\lfloor \frac{t'}{\tau_n} \right\rfloor + 1\right) - \bar{x}^n(t')\right| \\
&\leq \left|\bar{x}^n(t) - x^n\left(\left\lfloor \frac{t}{\tau_n} \right\rfloor\right)\right| + \left|x^n\left(\left\lfloor \frac{t}{\tau_n} \right\rfloor\right) - x^n\left(\left\lfloor \frac{t'}{\tau_n} \right\rfloor + 1\right)\right| + \left|x^n\left(\left\lfloor \frac{t'}{\tau_n} \right\rfloor + 1\right) - \bar{x}^n(t')\right| \\
&\leq \left|\bar{x}^n(t) - x^n\left(\left\lfloor \frac{t}{\tau_n} \right\rfloor\right)\right| + \sum_{m=\left\lfloor \frac{t'}{\tau_n} \right\rfloor + 1}^{\left\lfloor \frac{t'}{\tau_n} \right\rfloor - 1} |x^n(m+1) - x^n(m)| + \left|x^n\left(\left\lfloor \frac{t'}{\tau_n} \right\rfloor + 1\right) - \bar{x}^n(t')\right| \\
&\leq \sqrt{2\tau_0}\left(\varphi_0(x_0) - C\right)^{1/2} \left(t-t'\right),
\end{aligned}$$

where the final inequality follows from the definition of $\bar{x}^n$ and (2.19). The above inequality holds for any $n \in \mathbb{N}$, and it thereby follows that the sequence $\{\bar{x}^n\}$ is equicontinuous. Therefore, from the Arzelá-Ascoli Theorem, we have that $\{\bar{x}^n\}$ contains a uniformly convergent subsequence, and let the limit be the curve $\{x(t)\}_{t\in\mathbb{R}_{\geq 0}}$. Moreover, by isolating the uniformly convergent subsequence and using a smooth test function $\zeta \in C^\infty([0,T])$, we have:

$$\begin{aligned}
\int_{[0,T]} \frac{d\zeta}{dt} x(t)\, dt = \lim_{n\to\infty} \int_{[0,T]} \frac{d\zeta}{dt} \bar{x}^n(t)\, dt &= \lim_{n\to\infty} \int_{[0,T]} \left(\frac{\zeta(t+\tau_n) - \zeta(t)}{\tau_n} \right) \bar{x}^n(t)\, dt \\
&= \lim_{n\to\infty} \int_{[\tau_n,T]} \zeta(t) \left(\frac{\bar{x}^n(t-\tau_n) - \bar{x}^n(t)}{\tau_n} \right)\, dt \\
&= \int_{[0,T]} \zeta(t) \nabla \varphi_t(x(t))\, dt,
\end{aligned}$$

where the final equality follows from (2.17). The above is the weak form of the gradient flow (2.18). □

We observe that the vector field $\mathbf{v} = -\nabla\varphi$ satisfies a zero-flux boundary condition $\mathbf{v}\cdot\mathbf{n} = \nabla\varphi\cdot\mathbf{n} = 0$ on $\partial\Omega$ owing to the definition of the functional F.

Definition 13 (Gradient flows in the space of probability measures)**.** *For a C^1 function $F : \mathscr{P}(\Omega) \to \mathbb{R}$, the transport by* (2.14) *(satisfied in the distributional sense* (2.16)*) with $\mathbf{v} = -\nabla\left(\frac{\delta F}{\delta\mu}\right)$ is called a gradient flow on F.*

The following theorem establishes the asymptotic stability of the gradient flow on the functional F, with convergence to $\mu^* \in \mathscr{P}(\Omega)$, the global minimizer of F as $t \to \infty$:

Theorem 5 (Asymptotic stability of gradient flow)**.** *The solutions to the gradient flow w.r.t F converge to μ^* in the limit $t \to \infty$.*

Proof. Let $\{\mu_t\}_{t\geq 0}$ be an orbit of the gradient flow w.r.t. F in $\mathscr{P}(\Omega)$. We have:

$$\begin{aligned}
\frac{d}{dt} F(\mu_t) &= \int_\Omega \left\langle \nabla\left(\frac{\delta F}{\delta\mu}\right), \mathbf{v} \right\rangle d\mu_t \\
&= -\int_\Omega \left| \nabla\left(\frac{\delta F}{\delta\mu}\right) \right|^2 d\mu_t \\
&\leq 0.
\end{aligned}$$

This implies that $F(\mu_t) \leq F(\mu_0)$ for all $t \geq 0$, and therefore $\{\mu_t\}_{t\geq 0}$ is contained in the sublevel set $\mathscr{S}(\mu_0) = \{\nu \in \mathscr{P}(\Omega) | F(\nu) \leq F(\mu_0)\}$. From Lemma 15, we have that $S(\mu_0)$ is compact in $(\mathscr{P}(\Omega), W_2)$, which implies that the orbit $\{\mu_t\}_{t\geq 0}$ is precompact. Moreover, the functional F is lower bounded in $S(\mu_0)$ by $F(\mu^*)$. By the LaSalle invariance principle for Banach spaces, in Lemma 6, we therefore have that the orbit converges in $(\mathscr{P}(\Omega), W_2)$ (also weakly, from Lemma 10) asymptotically to the largest invariant set contained in $\dot{F}^{-1}(0)$. We have:

$$\dot{F}^{-1}(0) = \left\{ \mu \in \mathscr{P}(\Omega) \middle| \nabla \left(\frac{\delta F}{\delta \mu} \right) = 0,\ \mu - \text{a.e. in } \Omega \right\},$$

which implies that the Fréchet derivative of F is zero in the set $\dot{F}^{-1}(0)$. This corresponds to the set of critical points of F and from the strict geodesic convexity of F, we therefore get that $\dot{F}^{-1}(0) = \{\mu^*\}$. □

2.5 Multi-agent coverage control

We now investigate the problem of multi-agent coverage control within the theoretical framework developed in this section. The multi-agent coverage control problem is characterized by the objective of deploying a group of agents across a spatial domain to maximize an appropriate notion of coverage, specified as a locational optimization problem given a target coverage profile [36]. This is accomplished by distributed algorithms that steer the agents towards the (often local) minima of the aggregate objective function of the locational optimization problem. This can alternatively be viewed as a problem of optimally quantizing an absolutely continuous target probability measure by a discrete probability measure, minimizing a quantization cost, where the discrete measure is supported on the set of agent positions. The quantization cost can be formulated as the optimal transport cost between the absolutely continuous probability measure and the discrete probability measure, and is called the semi-discrete optimal transport problem [26, 66].

The general theory developed in this paper allows for investigation of the behavior of the locational optimization algorithms in the limit $N \to \infty$, and the convergence to the global minimizer of the aggregate objective function. This is in addition to the macroscopic perspective of coverage control offered by analysis within this framework. The order of presentation in this section is reversed in comparison to the rest of the paper, in that we begin with the discretized perspective, i.e., widely used aggregate objective functions in locational optimization, and search for their functional counterparts in the space of probability measures from which they can be seen to be discretized. The same approach is taken for the locational optimization algorithms.

2.5.1 Aggregate objective functions

In what follows, we study a widely used aggregate objective function for locational optimization, the quantization energy, interpret it as the optimal transport cost between the target probability measure and a weighted discrete measure, and investigate its behavior in the limit $N \to \infty$. We then discuss the limitations of this function and seek to mitigate its limitations by considering an alternative function used in the context of area/weight-constrained locational optimization. This alternative objective function is again interpreted as an optimal transport cost and its behavior is investigated in the limit $N \to \infty$, and its convexity is established.

Optimal transport cost

We first define the cost of optimal transport between measures μ and ν with the unit cost of transport $c(x,y) = f(|x-y|)$, where $f : \mathbb{R} \to \mathbb{R}$ is a non-decreasing differentiable function with $f(0) = 0$, as:

$$C_f(\mu,\nu) = \inf_{\substack{T:\Omega\to\Omega \\ T_{\#}\mu=\nu}} \int_\Omega f(|x - T(x)|)\, d\mu(x). \tag{2.20}$$

Moreover, if the function f is strictly convex and $\mu \in \mathscr{P}^r(\Omega)$, it also follows that there exists a unique optimal transport map [106] minimizing $C_f(\mu,\nu)$ in (2.20). Assuming strict convexity of

f, we now establish the following result on the strict convexity of $C_f(\cdot, \mu^*)$:

Lemma 24 (Strict convexity of $C_f(\cdot, \mu^*)$)**.** *Fix $\mu^* \in \mathscr{P}(\Omega)$ (absolutely continuous) as the reference measure and let $\mu_0, \mu_1 \in \mathscr{P}(\Omega)$. Let $T_{\mu^* \to \mu_0}$ and $T_{\mu^* \to \mu_1}$ be optimal transport maps from μ^* to μ_0 and μ^* to μ_1 respectively, corresponding to the optimal transport cost C_f, and let $T_t = (1-t)T_{\mu^* \to \mu_0} + tT_{\mu^* \to \mu_1}$ for $t \in [0,1]$. For $\mu_t = T_{t\#}\mu^*$, we have:*

$$C_f(\mu_t, \mu^*) < (1-t)C_f(\mu_0, \mu^*) + tC_f(\mu_0, \mu^*).$$

Proof. We have:

$$\begin{aligned} C_f(\mu_t, \mu^*) \leq \int_\Omega f(|T_t(x) - x|) d\mu^*(x) &= \int_\Omega f\left(\left|(1-t)T_{\mu^* \to \mu_0}(x) + tT_{\mu^* \to \mu_1}(x) - x\right|\right) d\mu^*(x) \\ &= \int_\Omega f\left(\left|(1-t)\left[T_{\mu^* \to \mu_0}(x) - x\right] + t\left[T_{\mu^* \to \mu_1}(x) - x\right]\right|\right) d\mu^*(x) \\ &\leq \int_\Omega f\left((1-t)\left|T_{\mu^* \to \mu_0}(x) - x\right| + t\left|T_{\mu^* \to \mu_1}(x) - x\right|\right) d\mu^*(x), \end{aligned}$$

where the final inequality is a consequence of the fact that f is non-decreasing. Further, if f is strictly convex in Ω, we will have:

$$\begin{aligned} C_f(\mu_t, \mu^*) &< \int_\Omega \left[(1-t)f\left(\left|T_{\mu^* \to \mu_0}(x) - x\right|\right) + tf\left(\left|T_{\mu^* \to \mu_1}(x) - x\right|\right)\right] d\mu^*(x) \\ &= (1-t)\int_\Omega f\left(\left|T_{\mu^* \to \mu_0}(x) - x\right|\right) d\mu^*(x) + t\int_\Omega f\left(\left|T_{\mu^* \to \mu_1}(x) - x\right|\right) d\mu^*(x) \\ &= (1-t)C_f(\mu_0, \mu^*) + tC_f(\mu_0, \mu^*). \end{aligned}$$

□

The lemma above can be applied to the quadratic case $f(x) = x^2$, where we get the squared L^2-Wasserstein distance and convexity can similarly be shown. In the case of the squared L^2-Wasserstein distance, the interpolants essentially turn out to be generalized geodesics with μ^* as the reference measure, and we thereby get (generalized) geodesic convexity. This is noted in the following corollary:

Corollary 5 (Strict convexity of squared L^2-Wasserstein distance to μ^*)**.** *The squared L^2-Wasserstein distance functional $W_2^2(\cdot,\mu^*)$ is (generalized) geodesically strictly convex.*

We now establish the following result on the l-smoothness of $C_f(\cdot,\mu^*)$:

Lemma 25 (l-smoothness of $C_f(\cdot,\mu^*)$)**.** *If the function $h:\mathbb{R}^d\to\mathbb{R}$ such that $h(\mathbf{v})=f(|\mathbf{v}|)$ is l-smooth, then the functional $F(\mu)=C_f(\mu,\mu^*)$, with ξ_μ being the Fréchet derivative of F at $\mu\in\mathscr{P}^r(\Omega)$, satisfies:*

$$\left|\int_\Omega \left\langle \xi_{\mu_1}-\xi_{\mu_2}, T_{\mu_2\to\mu_1}-id\right\rangle d\mu_2\right| \le \int_\Omega \left|T_{\mu_2\to\mu_1}-id\right|^2 d\mu_2,$$

where $T_{\mu_2\to\mu_1}$ is the optimal transport map from μ_2 to μ_1 w.r.t. the cost C_f.

Proof. Let $\phi_\mu=\frac{\delta C_f(\mu,\mu^*)}{\delta\mu}$ the Kantorovich potential for the optimal transport from μ to μ^*. We now have the following relation [106]:

$$T_{\mu\to\mu^*}=id-(\nabla h)^{-1}(\nabla\phi_\mu).$$

With the above, and from the l-smoothness of h, we have:

$$\begin{aligned}\left|\int_\Omega \left\langle \xi_{\mu_1}-\xi_{\mu_2}, T_{\mu_2\to\mu_1}-id\right\rangle d\mu_2\right| &= \left|\int_\Omega \left\langle \nabla h\left(id-T_{\mu_1\to\mu^*}\right)-\nabla h\left(id-T_{\mu_2\to\mu^*}\right), T_{\mu_2\to\mu_1}-id\right\rangle d\mu_2\right|\\ &\le \int_\Omega \left|\left\langle \nabla h\left(id-T_{\mu_1\to\mu^*}\right)-\nabla h\left(id-T_{\mu_2\to\mu^*}\right), T_{\mu_2\to\mu_1}-id\right\rangle\right| d\mu_2\\ &\le l\int_\Omega \left|T_{\mu_2\to\mu_1}-id\right|^2 d\mu_2.\end{aligned}$$

□

The corollary below follows immediately from Lemma 25:

Corollary 6 (l-smoothness of squared L^2-Wasserstein distance to μ^*)**.** *The functional $\frac{1}{2}W_2^2(\cdot,\mu^*)$ is 1-smooth.*

Quantization energy

We now define an aggregate objective function for the locational optimization problem [36], also known as the quantization energy, as follows:

$$\mathscr{H}_f(\mathbf{x}) = \int_\Omega \min_{i\in\{1,\ldots,N\}} f(|x-x_i|)d\mu^*(x). \tag{2.21}$$

Now, the following lemma establishes the relationship between the aggregate objective function in (2.21) and the optimal transport cost defined in (2.20):

Lemma 26 (Optimal transport formulation of locational optimization objective)**.** *The aggregate objective function $\mathscr{H}$ for the locational optimization problem, as defined in* (2.21)*, satisfies:*

$$\mathscr{H}_f(\mathbf{x}) = \min_{\mathbf{w}\in\mathbb{R}^N_{\geq 0}} C_f\left(\sum_{i=1}^N w_i\delta_{x_i}\ ,\ \mu^*\right) = C_f\left(\sum_{i=1}^N \mu^*(\mathscr{V}_i)\delta_{x_i}\ ,\ \mu^*\right).$$

Proof. Now, let $\widehat{\mu}^N_{\mathbf{x},\mathbf{w}} = \sum_{i=1}^N w_i\delta_{x_i}$ be a weighted discrete probability measure corresponding to the set of points $\{x_i\}_{i=1}^N$ with corresponding weights $\{w_i\}_{i=1}^N$, such that $w_i \in [0,1]$ and $\sum_{i=1}^N w_i = 1$. The optimal transport cost between $\widehat{\mu}^N_{\mathbf{x},\mathbf{w}}$ and μ^* is given by:

$$C_f(\widehat{\mu}^N_{\mathbf{x},\mathbf{w}}, \mu^*) = \inf_{\substack{T:\Omega\to\Omega \\ T_\#\mu^*=\widehat{\mu}^N_{\mathbf{x}}}} \int_\Omega f(|x-T(x)|)\, d\mu^*(x),$$

where the infimum is over the set of maps T that pushforward μ^* to $\widehat{\mu}^N_{\mathbf{x},\mathbf{w}}$ (we note that since $\widehat{\mu}^N_{\mathbf{x},\mathbf{w}}$ has finite support, pushforward maps exist only from μ^* to $\widehat{\mu}^N_{\mathbf{x},\mathbf{w}}$ and not the other way around). The maps $T : \Omega \to \{x_i\}_{i=1}^N$ partition Ω into N regions $\{\mathscr{W}_i\}_{i=1}^N$ of mass $\mu^*(\mathscr{W}_i) = w_i$. Let T^* :

$\Omega \to \{x_i\}_{i=1}^N$ be the optimal transport map from μ^* to $\widehat{\mu}_{\mathbf{x},\mathbf{w}}^N$, which allows us to write:

$$\begin{aligned} C_f(\widehat{\mu}_{\mathbf{x},\mathbf{w}}^N, \mu^*) &= \inf_{\substack{T:\Omega\to\Omega \\ T_\#\mu^* = \widehat{\mu}_{\mathbf{x},\mathbf{w}}^N}} \int_\Omega f(|x - T(x)|)\, d\mu^*(x) \\ &= \int_\Omega f(|x - T^*(x)|)\, d\mu^*(x) \\ &\geq \int_\Omega \min_{i\in\{1,\ldots,N\}} f(|x - x_i|)\, d\mu^*(x). \end{aligned} \tag{2.22}$$

The above inequality is due to the fact that for a family $\mathscr{G}$ of bounded functions over Ω, we have $\inf_{g\in\mathscr{G}} \int_\Omega g(x)\,\mathrm{dvol}(x) \geq \int_\Omega \min_{g\in\mathscr{G}} g(x)\,\mathrm{dvol}(x)$. Since f is non-decreasing, we also have:

$$\int_\Omega \min_{i\in\{1,\ldots,N\}} f(|x - x_i|)\, d\mu^*(x) = \sum_{i=1}^N \int_{\mathscr{V}_i} f(|x - x_i|)\, d\mu^*(x),$$

where $\{\mathscr{V}_i\}_{i=1}^N$ is such that $\mathring{\mathscr{V}}_i \cap \mathring{\mathscr{V}}_j = \emptyset$ for $i \neq j$ and $\cup_{i=1}^N \mathscr{V}_i = \Omega$, is the Voronoi partition of Ω.

We now define a map $T_\mathscr{V} : \Omega \to \Omega$ such that $T_\mathscr{V}(x) = x_i$ for $x \in \mathscr{V}_i$, with $T_\mathscr{V}(\Omega) = \{x_1, \ldots, x_N\}$, and we get:

$$\begin{aligned} \int_\Omega f(|x - T_\mathscr{V}(x)|)\, d\mu^*(x) &= \sum_{i=1}^N \int_{\mathscr{V}_i} f(|x - x_i|)\, d\mu^*(x) \\ &= \int_\Omega \min_{i\in\{1,\ldots,N\}} f(|x - x_i|)\, d\mu^*(x). \end{aligned}$$

From (2.22) and the above, we therefore get:

$$\int_\Omega f(|x - T_\mathscr{V}(x)|)\, d\mu^*(x) \leq C_f(\widehat{\mu}_{\mathbf{x},\mathbf{w}}^N, \mu^*).$$

For the particular choice of the weights $w_i = \mu^*(\mathscr{V}_i)$ such that $\widehat{\mu}_{\mathbf{x},\mu^*(\mathscr{V})}^N = \sum_{i=1}^N \mu^*(\mathscr{V}_i)\delta_{x_i}$, we also get the reverse inequality:

$$\int_\Omega f(|x - T_\mathscr{V}(x)|)\, d\mu^*(x) \geq C_f(\widehat{\mu}_{\mathbf{x},\mu^*(\mathscr{V})}^N, \mu^*).$$

This is because:

$$C_f\left(\sum_{i=1}^{N}\mu^*(\mathcal{V}_i)\delta_{x_i},\mu^*\right) = \inf_{\substack{T:\Omega\to\Omega \\ T_{\#}\mu^*=\sum_{i=1}^{N}\mu^*(\mathcal{V}_i)\delta_{x_i}}} \int_\Omega f(|x-T(x)|)\,d\mu^*(x)$$
$$\leq \int_\Omega f(|x-T_{\mathcal{V}}(x)|)\,d\mu^*(x),$$

and we therefore get:

$$C_f\left(\sum_{i=1}^{N}\mu^*(\mathcal{V}_i)\delta_{x_i},\mu^*\right) = \int_\Omega \min_{i\in\{1,\dots,N\}} f(|x-x_i|)\,d\mu^*(x),$$

which establishes that:

$$\min_{\mathbf{w}\in\mathbb{R}^N_{\geq 0}} C_f\left(\sum_{i=1}^{N} w_i\delta_{x_i},\mu^*\right) = \int_\Omega \min_{i\in\{1,\dots,N\}} f(|x-x_i|)\,d\mu^*(x) = \mathcal{H}_f(\mathbf{x}),$$

with the minimizing weights $w_i^* = \mu^*(\mathcal{V}_i)$. □

Lemma 26 establishes the connection between the locational optimization problem and the problem of semi-discrete optimal transport. Moreover, we also observe the connection to quantization of probability measures, wherein the aggregate objective function is expressed as the optimal transport cost between the absolutely continuous target probability measure μ^* and a discrete probability measure. For this reason, the aggregate objective function $\mathcal{H}_f$ is also known as the quantization energy.

Corollary 7 (Aggregate objective function as L^2-Wasserstein distance)**.** *It follows from Lemma 26 that with a quadratic cost $f(x) = x^2$ (and the corresponding aggregate objective function $\mathcal{H}_2$), we have:*

$$\mathcal{H}_2(\mathbf{x}) = \min_{\mathbf{w}\in\mathbb{R}^N_{\geq 0}} W_2^2\left(\sum_{i=1}^{N} w_i\delta_{x_i},\mu^*\right),$$

where the minimizing weights $w_i^* = \mu^*(\mathcal{V}_i)$.

We now investigate the properties of the aggregate objective function $\mathscr{H}_f$ in the limit $N \to \infty$. Recall that we have:

$$\mathscr{H}_f(\mathbf{x}) = C_f\left(\sum_{i=1}^{N} \mu^*(\mathscr{V}_i)\delta_{x_i}, \mu^*\right).$$

Now let $x_i \sim_{i.i.d} \mu$, where $\mu \in \mathscr{P}(\Omega)$ is any absolutely continuous probability measure such that $\text{supp}(\mu) \supseteq \text{supp}(\mu^*)$. In the limit $N \to \infty$, the weighted measure $\sum_{i=1}^{N} \mu^*(\mathscr{V}_i)\delta_{x_i}$ converges weakly almost surely to μ^* (this can be seen by evaluating the expectation w.r.t. $\sum_{i=1}^{N} \mu^*(\mathscr{V}_i)\delta_{x_i}$ of any simple function, in the limit $N \to \infty$, along with the application of the Glivenko Cantelli Theorem). Therefore, by continuity of C_f, in the limit $N \to \infty$ the value of the aggregate objective function converges to zero almost surely, i.e., $\lim_{N\to\infty} \mathscr{H}_f(\mathbf{x}) = \lim_{N\to\infty} C_f\left(\sum_{i=1}^{N} \mu^*(\mathscr{V}_i)\delta_{x_i}, \mu^*\right) = C_f(\mu^*, \mu^*) = 0$ a.s. In other words, the value of the aggregate objective function $\mathscr{H}_f$ converges almost surely to zero in the limit $N \to \infty$ irrespective of the configuration of the points $\{x_i\}_{i=1}^{N}$. While the empirical measure $\frac{1}{N}\sum_{i=1}^{N} \delta_{x_i}$ corresponding to the points $\{x_i\}_{i=1}^{N}$ sampled i.i.d. from μ converges uniformly almost surely to μ (Glivenko-Cantelli theorem), the quantization energy $\mathscr{H}_f$, by converging to zero, does not reflect the discrepancy between the measures μ and μ^*. The quantization energy $\mathscr{H}_f$ therefore suffers from this deficiency as a candidate aggregate function for coverage control.

Constrained quantization energy

To mitigate the above deficiency, let us define another aggregate function $\bar{\mathscr{H}}_f$ as follows:

$$\bar{\mathscr{H}}_f(\mathbf{x}) = C_f\left(\frac{1}{N}\sum_{i=1}^{N} \delta_{x_i}\ ,\ \mu^*\right). \tag{2.23}$$

The function $\bar{\mathscr{H}}_f$ is again a commonly used aggregate objective function in the literature, used in the area (weight)-constrained coverage control problem (where the weights $w_i = 1/N$ are balanced in the case of (2.23)). This can be seen from the following:

$$\begin{aligned}
\bar{\mathscr{H}}_f(\mathbf{x}) &= C_f\left(\frac{1}{N}\sum_{i=1}^{N}\delta_{x_i}\ ,\ \mu^*\right) \\
&= \min_{\substack{T:\Omega\to\{x_i\}_{i=1}^N \\ T_\#\mu^*=\frac{1}{N}\sum_{i=1}^N\delta_{x_i}}} f(|x-T(x)|)d\mu^*(x) \\
&= \min_{T:\Omega\to\{x_i\}_{i=1}^N}\left\{f(|x-T(x)|)d\mu^*(x)\ \middle|\ \mu^*\left(T^{-1}(\{x_i\})\right)=\frac{1}{N},\ \forall\ i\right\}.
\end{aligned}$$

As was the case with $\mathscr{H}_f(\mathbf{x})$, the quantization energy, that it could be decomposed into the sum of energies evaluated over the Voronoi partition of Ω with $\{x_i\}_{i=1}^N$ as its generators (we recall this from the proof of Lemma 26), a similar statement could be made of $\bar{\mathscr{H}}_f$, except that this case involves a generalized Voronoi partition $\{\mathscr{W}_i\}_{i=1}^N$, with $\cup_{i=1}^N\mathscr{W}_i=\Omega$, $\mathring{\mathscr{W}}_i\cap\mathring{\mathscr{W}}_j=\emptyset$ for $i\neq j$, where:

$$\mathscr{W}_i=\left\{x\in\Omega\,\middle|\,f(|x-x_i|)-\omega_i\leq f(|x-x_j|)-\omega_j\right\},$$

where $\{\omega_1,\ldots,\omega_N\}$ are chosen such that $\mu^*(\mathscr{W}_i)=1/N$ for all $i\in\{1,\ldots,N\}$. We refer the reader to [33] for a detailed treatment of the computation of $(\omega_1,\ldots,\omega_N)$. We can now write:

$$\bar{\mathscr{H}}_f(\mathbf{x})=\sum_{i=1}^{N}\int_{\mathscr{W}_i}f(|x-x_i|)\,d\mu^*(x). \tag{2.24}$$

We again investigate the properties of $\bar{\mathscr{H}}_f$ in the limit $N\to\infty$. By letting $x_i\sim_{i.i.d}\mu$, where $\mu\in\mathscr{P}(\Omega)$ is any absolutely continuous probability measure, in the limit $N\to\infty$, we will have $\frac{1}{N}\sum_{i=1}^N\delta_{x_i}$ converging uniformly almost surely to μ, and by the continuity of C_f, we will have:

$$\lim_{N\to\infty}\bar{\mathscr{H}}_f(\mathbf{x})=C_f(\mu,\mu^*),\ \ \mu-a.s. \tag{2.25}$$

2.5.2 Generalized Lloyd proximal descent

We now turn our attention to the generalized Lloyd descent algorithm, which is a gradient descent algorithm on the aggregate objective function $\bar{\mathscr{H}}_f$ for the quadratic case $f(x) = x^2$. We begin with its formulation, introduce a proximal descent version and investigate its convergence to the global minimizer in the limit $N \to \infty$.

We now formulate the multi-agent proximal descent on the aggregate objective function $\bar{\mathscr{H}}_f$, with $f(x) = x^2$, as follows:

$$x_i^+ = \arg\min_{z\in\Omega} \frac{1}{2\tau}|x_i - z|^2 + \bar{\mathscr{H}}_f(z, \mathbf{x}_{-i}). \tag{2.26}$$

Theorem 6 (Convergence to centroidal generalized Voronoi partition)**.** *The Lloyd proximal descent* (2.26)*, with $f(x) = x^2$, converges to a local minimizer of $\bar{\mathscr{H}}_f$.*

Proof. The statement of the theorem follows from Corollary 6 and the application of Theorem 4. □

It is known that the generalized Lloyd descent algorithm results in convergence to centroidal generalized Voronoi partitions [33], in which case, the generators $\{x_1, \ldots, x_N\}$ of the generalized Voronoi partition are also the centroids of their respective generalized Voronoi cells. The centroidal generalized Voronoi partition is, however, not unique, and this relates to the fact that the convergence is to the local minimizers of $\bar{\mathscr{H}}_f$, which is typically nonconvex. From Section 2.4.2, we have that in the limit $N \to \infty$, the proximal descent scheme (2.26) converges to:

$$x^+ = \arg\min_{z\in\Omega} \frac{1}{2\tau}|x - z|^2 + \phi(z), \tag{2.27}$$

with $x \sim \mu$ and $\phi = \left.\frac{\delta W_2^2(\nu,\mu^*)}{\delta \nu}\right|_{\mu}$ is the Kantorovich potential for optimal transport from μ to μ^*.

Theorem 7 (Convergence of (2.27))**.** *The sequence $\{\mu_k\}_{k\in\mathbb{N}}$ obtained as the transport of measure $\mu_0 \in \mathscr{P}^r(\Omega)$ by* (2.27)*, with $x_0 \sim \mu_0$, converges weakly to μ^* as $k \to \infty$.*

Proof. The claim follows from Theorem 3, with $F(\nu) = W_2^2(\nu, \mu^*)$ and the strict (generalized) geodesic convexity of $W_2^2(\cdot, \mu^*)$. □

2.6 Summary

In this chapter, we presented a macroscopic description of the behavior of multi-agent gradient descent algorithms as transport in the space of probability measures, and developed a multiscale theory bridging the microscopic and macroscopic scales. The chapter also contributes a framework for developing iterative, gradient-based algorithms for multi-agent transport with provable convergence, based on descent schemes in the space of probability measures. Within the above framework, we investigated the behavior of coverage optimization algorithms, particularly the asymptotic convergence to the global minimizer, in the limit that the number of agents $N \to \infty$.

The material in this chapter is currently being prepared for submission as *Multiscale Analysis of Multi-Agent Transport by Gradient Descent*, V. Krishnan and S. Martínez. The dissertation author was the primary investigator and author of this paper.

Chapter 3

Multi-agent optimal transport

In this chapter, we propose and investigate an iterative scheme for large-scale optimal transport of multi-agent collectives based on a scalable, distributed online optimization. Working with a reduction of the Kantorovich duality for metric costs conformal to the Euclidean metric, we note that the Kantorovich potential is almost everywhere differentiable and obtain a bound on the norm of its gradient. We then obtain an iterative scheme for optimal transport of probability measures based on Kantorovich duality, showing it to be equivalent to optimal transport along geodesics. We propose a distributed primal-dual algorithm to be implemented online by the agents to obtain local estimates of the Kantorovich potential, which are then used as local objectives in a proximal algorithm for transport. In the continuous-time limit and as $N \to \infty$, we derive a PDE-based flow for optimal transport, and obtain convergence results for an online implementation of the transport. The material in this chapter contains a novel scalable, distributed algorithm for multi-agent optimal transport, addressing a longstanding concern in the research on multi-agent systems.

3.1 Bibliographical comments

The applications of optimal transport in image processing and various engineering domains has motivated a search for efficient computational methods for the optimal transport problem, and we refer the reader to [100] for a comprehensive account. Entropic regularization of the Kantorovich formulation has been an efficient tool for approximate computation of the optimal transport cost using the Sinkhorn algorithms [37], [39]. Data-driven approaches to the computation of the optimal transport cost between two distributions from their samples have been investigated in [82, 114], and with an eye towards large-scale problems in [61], [109], [93]. A related problem of computation of Wasserstein barycenters was addressed in [38]. Optimal transport from continuous to discrete probability distributions has been studied under the name of semi-discrete optimal transport, with connections to the problem of optimal quantization of probability measures, in [26]. While computational approaches to optimal transport often work with the static, Monge or Kantorovich formulations of the problem, investigations involving dynamical formulations was initiated by [20], where the authors recast the L^2 Monge-Kantorovich mass transfer problem in a fluid mechanics framework. This largely owes to notion of displacement interpolation originally introduced in [87]. The underlying rationale is that the optimal transport cost defines a metric in the space of probability measures, which allows for the interpretation of optimal transport between two probability measures as transport along distance-minimizing geodesics connecting them. [98] and [21] are other works in this vein. The problem of optimal transport was also explored from a stochastic control perspective in [91] and [30], where the latter further explored connections to Schrodinger bridges. However, there has remained a gap in this literature with regard to distributed computation of optimal transport, which arises as a rather stringent constraint in multi-agent transport scenarios. We note that despite the potential for the application of optimal transport ideas to the multi-agent setting, it has hitherto largely remained unsuccessful. The papers [57] and [15] represent attempts in this direction, while in the first paper the problem is formulated as one of optimal control, the second is placed in the framework of optimal transport. These works, however, present significant limitations either because they require centralized offline planning [57], or because of a need for costly computation and information exchange between agents [15]. This serves as a strong motivation for the distributed iterative scheme for optimal transport presented in this chapter.

Markov Chain Monte Carlo (MCMC) methods [9, 63, 103] present another framework for the problem of rearranging probability measures, and can be traced back to early works by Metropolis [90] and Hastings [67]. MCMC methods involve the construction of a Markov chain with the target probability measure as its equilibrium measure, and yield samples of the target measure as $t \to \infty$. From a computational perspective, MCMC methods allow for the agents to be transported independently of one another, which results in a fully decentralized implementation. However, MCMC methods are inefficient with respect to the cost of transport. On the other hand, an optimal transport-based approach suffers from the need for a centralized implementation, as the optimal transport plan is computed using global information of the intial and target probability measures. This further motivates our search for scalable, distributed iterative algorithms that occupy the middle ground. We attempt to improve the cost of transport by imposing more structure to the set of agents in the form of a nearest-neighbor network and using the information from the neighbors to compute the successive iterates. From a computational standpoint, such an approach would neither be decentralized to the point of complete independence between agents as in the case of MCMC, nor would it be centralized as is typical of conventional optimal transport-based methods.

3.2 On the Monge and Kantorovich formulations of optimal transport

We begin this section with an overview of the Monge and Kantorovich formulations of optimal transport, followed by preliminary results used later in the paper.

Let $\mu, \nu \in \mathscr{P}(\Omega)$ be absolutely continuous probability measures on Ω. Let $c : \Omega \times \Omega \to \mathbb{R}_{\geq 0}$ be such that for $x, y \in \Omega$, $c(x, y)$ is the unit cost of transport from x to y. We now make the following assumptions on the cost c:

Assumption 5. *The cost c is continuous and is a metric on Ω conformal to the Euclidean metric (with strictly positive conformal factor $\xi \in C^1(\Omega)$).*

In the Monge (deterministic) formulation, the optimal cost of transporting the probability measure μ onto ν is defined as the infimum of the transport cost over the set of all maps for which ν is obtained as the pushforward measure of μ, as given below:

$$C_M(\mu,\nu) = \inf_{\substack{T:\Omega\to\Omega \\ T_\#\mu=\nu}} \int_{\Omega\times\Omega} c(x,T(x))d\mu(x). \tag{3.1}$$

The Kantorovich formulation relaxes the above formulation by minimizing the transport cost over the set of joint probability measures $\Pi(\mu,\nu) \subset \mathscr{P}(\Omega\times\Omega)$, for which μ and ν are the respective marginals over Ω. The optimal transport cost from μ to ν is defined as follows:

$$C_K(\mu,\nu) = \inf_{\pi\in\Pi(\mu,\nu)} \int_{\Omega\times\Omega} c(x,y)\ d\pi(x,y). \tag{3.2}$$

We now present the following lemma on the existence of minimizers to the Monge and Kantorovich formulations and the equivalence between them. We refer the reader to [6] for proofs.

Lemma 27 (Existence of minimizers)**.** *Under Assumption 5, there exists a minimizer π^* to the Kantorovich problem. Moreover, if the measure μ is atomless (i.e., $\mu(\{x\}) = 0$ for all $x \in \Omega$), the Monge formulation has a minimizer T^* and it holds that $\pi^* = (id,T^*)_\#\mu$.*

Following Lemma 27, we denote by $C(\mu,\nu) = C_M(\mu,\nu) = C_K(\mu,\nu)$ the optimal transport cost from μ to ν. We now present the following key result that the optimal transport cost C defines a metric on the space of probability measures $\mathscr{P}(\Omega)$:

Lemma 28 (Corollary 3.2, 3.3 [106])**.** *Under Assumption 5, the optimal transport cost C : $\mathscr{P}(\Omega)\times\mathscr{P}(\Omega) \to \mathbb{R}_{\geq 0}$ defines a metric on $\mathscr{P}(\Omega)$.*

Kantorovich duality

The Kantorovich formulation (3.2) allows the following dual formulation [117]:

$$\begin{aligned} K(\mu,\nu) = \sup_{\phi\in L^1(\Omega);\psi\in L^1(\Omega)} & \int_\Omega \phi(x)d\mu(x) + \int_\Omega \psi(y)d\nu(y) \\ \text{s.t} \;\; & \phi(x)+\psi(y) \le c(x,y), \;\; \forall x,y\in\Omega. \end{aligned} \tag{3.3}$$

The maximizers of the above dual formulation are pairs of functions (ϕ,ψ), called Kantorovich potentials, which occur at the boundary of the inequality constraint, and satisfy the equations:

$$\phi(x) = \inf_{y\in\Omega}\left(c(x,y)-\psi(y)\right), \quad \psi(y) = \inf_{z\in\Omega}\left(c(z,y)-\phi(z)\right). \tag{3.4}$$

We refer to (ϕ,ψ) defined above as a c-conjugate pair, and write $\psi = \phi^c$ to denote that ψ is the conjugate of ϕ. We therefore have:

$$\phi(x) = \inf_{y\in\Omega}\left[c(x,y) - \inf_{z\in\Omega}\left(c(z,y)-\phi(z)\right)\right]. \tag{3.5}$$

The Kantorovich duality (3.3) can now be rewritten as:

$$K(\mu,\nu) = \sup_{\phi\in L^1(\Omega)} \int_\Omega \phi(x)d\mu(x) + \int_\Omega \phi^c(y)d\nu(y). \tag{3.6}$$

We recall the following lemma on the strong duality property of the Kantorovich formulation. We refer the reader to Theorem 5.10 in [117] for a detailed proof.

Lemma 29 (Theorem 5.10, [117])**.** *Strong duality holds for the Kantorovich formulation. In other words, the gap between the costs defined in the Kantorovich formulation* (3.2) *and its dual* (3.3) *is zero, i.e.,* $C(\mu,\nu) = K(\mu,\nu)$.

Under Assumption 5 on the transport cost function c, we can obtain a further reduction of the Kantorovich duality (3.6). The following key lemma allows for such a reduction:

Lemma 30. *Under Assumption 5 and from* (3.5)*, the conjugate of the Kantorovich potential satisfies* $\phi^c = -\phi$ *and* $|\phi(x) - \phi(y)| \leq c(x,y)$ *for all* $x, y \in \Omega$.

Proof. From (3.4), we have:

$$\begin{aligned}
\phi(x) &= \inf_{y\in\Omega}\Big(c(x,y) - \inf_{z\in\Omega}\left(c(z,y) - \phi(z)\right)\Big) = \inf_{y\in\Omega}\sup_{z\in\Omega}\Big(c(x,y) - c(z,y) + \phi(z)\Big) \\
&\geq \inf_{y\in\Omega}\Big(c(x,y) - c(z,y) + \phi(z)\Big) = \inf_{y\in\Omega}\Big(c(x,y) - c(z,y)\Big) + \phi(z) \\
&\geq -c(x,z) + \phi(z),
\end{aligned}$$

where we have used the fact that c is a metric to obtain the final inequality (for any y, we have $c(x,y) - c(z,y) = c(x,y) - c(y,z) \geq -c(x,z)$, which implies that $\inf_{y\in\Omega}\left(c(x,y) - c(z,y)\right) \geq -c(x,z)$). Moreover, since the above inequality holds for any $x, z \in \Omega$, we have $|\phi(x) - \phi(z)| \leq c(x,z)$.

Now, when $|\phi(x) - \phi(y)| \leq c(x,y)$, we have that $-\phi(x) \leq c(x,y) - \phi(y)$, which implies that $-\phi(x) \leq \inf_y\left(c(x,y) - \phi(y)\right) = \phi^c(x)$. Equivalently, we obtain the relation $\phi(x) \geq -\phi^c(x)$.

Similarly, from (3.4) $\phi^c(x) = \inf_y c(x,y) - \phi(y)$, we obtain $\phi^c(x) \leq c(x,y) - \phi(y)$. By setting $y = x$ in the above inequality, and using $c(x,x) = 0$ we get $\phi(x) \leq -\phi^c(x)$.

In all, we have that $\phi^c(x) = -\phi(x)$ when $|\phi(x) - \phi(y)| \leq c(x,y)$. □

Following Lemma 30, we can now reduce the Kantorovich duality (3.6) to obtain:

$$\begin{gathered}
K(\mu,\nu) = \sup_{\phi\in\mathscr{L}(\Omega)} \mathbb{E}_\mu[\phi] - \mathbb{E}_\nu[\phi], \\
\text{where } \ \mathscr{L}(\Omega) = \{\phi \in L^1(\Omega) : |\phi(x) - \phi(y)| \leq c(x,y), \ \forall\, x,y \in \Omega\}.
\end{gathered} \tag{3.7}$$

Remark 2. *We note from* (3.7) *that functions* $\phi \in \mathscr{L}(\Omega)$ *are Lipschitz continuous (since* c *is conformal to the Euclidean metric from Assumption 5, and* Ω *is compact). It then follows from Rademacher's theorem (in Lemma 3) that* ϕ *is differentiable* μ*-almost everywhere in* Ω*. Moreover, its (pointwise a.e.) derivative is equal to its weak derivative, and we interpret the* *derivative of the Kantorovich potential in the weak sense in the rest of the paper. Moreover, we have that the Kantorovich potential ϕ is differentiable at every $x \in \Omega$ that is not a fixed point of the optimal transport map T^* [106].*

Furthermore, we would like to obtain a bound on the gradient of functions in the set $\mathscr{L}(\Omega)$, with the added assumption that they are everywhere differentiable. To this end, we characterize the set $\mathscr{L}(\Omega) \cap \mathscr{C}^1(\Omega)$ through the following lemma:

Lemma 31. *Under Assumption 5, the set $\mathscr{L}(\Omega) \cap \mathscr{C}^1(\Omega)$ is given by:*

$$\mathscr{L}(\Omega) \cap \mathscr{C}^1(\Omega) = \left\{ \phi \in L^1(\Omega) \cap \mathscr{C}^1(\Omega) : |\nabla \phi| \leq \xi \text{ in } \Omega \right\}. \tag{3.8}$$

Proof. Let $\phi \in \mathscr{L}(\Omega) \cap \mathscr{C}^1(\Omega)$, and $x, y \in \mathring{\Omega}$ with $x \neq y$ such that the line segment joining x and y is contained in Ω. By the Mean Value Theorem and the definition of $\mathscr{L}(\Omega)$ in (3.7), for some $m \in [0,1]$, we get:

$$\frac{|\phi(y) - \phi(x)|}{|y-x|} = \frac{|\nabla \phi((1-m)x + my) \cdot (y-x)|}{|y-x|} \leq \frac{c(x,y)}{|y-x|}.$$

With $y = x + tv$, where $v \in \mathscr{T}_x\Omega$ (tangent space of Ω at $x \in \Omega$), in the limit $t \to 0$, we get:

$$\frac{|\nabla \phi(x) \cdot v|}{|v|} \leq \xi(x),$$

where the above inequality holds for all $v \in T_x\Omega$ and ξ is the conformal factor for the metric c w.r.t the Euclidean metric, which implies that $|\nabla \phi(x)| \leq \xi(x)$ for any $x \in \mathring{\Omega}$.

Now, to prove the converse, we suppose that $|\nabla \phi(x)| \leq \xi(x)$ for any $x \in \mathring{\Omega}$. For $x, y \in \mathring{\Omega}$ with $x \neq y$, along the geodesic γ (w.r.t the metric c) joining x and y, we have:

$$\begin{aligned} |\phi(y) - \phi(x)| = \left| \int_0^1 \nabla \phi(\gamma(t)) \cdot \dot{\gamma}(t)\, dt \right| &\leq \int_0^1 |\nabla \phi(\gamma(t))|\, |\dot{\gamma}(t)|\, dt \\ &\leq \int_0^1 \xi(\gamma(t))\, |\dot{\gamma}(t)|\, dt = c(x,y) \end{aligned}$$

□

We now define the restricted Kantorovich duality as follows:

$$K(\mu,\nu) = \sup_{\phi \in \mathscr{L}(\Omega)\cap\mathscr{C}^1(\Omega)} \mathbb{E}_\mu[\phi] - \mathbb{E}_\nu[\phi], \tag{3.9}$$

where it is restricted in the sense that the constraint set is $\mathscr{L}(\Omega)\cap\mathscr{C}^1(\Omega)$ as opposed to $\mathscr{L}(\Omega)$ as given in (3.7).

3.3 Iterative scheme for multi-stage optimal transport

In this section, we establish a framework for multi-stage optimal transport of probability measures.

Let $\mu_0 \in \mathscr{P}(\Omega)$ be a given initial probability measure and $\mu^* \in \mathscr{P}(\Omega)$ the target probability measure. Our objective is to optimally transport μ_0 onto μ^* by an iterative scheme. To this end, we begin by constructing a finite sequence $\{\mu_k\}_{k=1}^K$ such that $\mu_K = \mu^*$, and carrying out optimal transport in stages $\{\mu_{k-1} \to \mu_k\}_{k=1}^K$. The net cost of transport along the sequence would then be given by $\sum_{k=1}^K C(\mu_{k-1},\mu_k)$, the sum of the (optimal) stage costs. We now have the following lemma on the retrieval of the optimal transport cost:

Lemma 32. *Given atomless probability measures $\mu_0, \mu^* \in \mathscr{P}(\Omega)$, the cost of optimal transport from μ_0 to μ^* satisfies:*

$$C(\mu_0,\mu^*) = \min_{\substack{(\mu_1,\ldots,\mu_K)\\ \mu_k\in\mathscr{P}(\Omega)\\ \mu_K=\mu^*}} \sum_{k=1}^K C(\mu_{k-1},\mu_k) \tag{3.10}$$

Proof. We begin by noting that there clearly exists at least one minimizing sequence for the optimization problem (3.10) (the trivial sequence $\mu_k = \mu^*$ for all $k = 1,\ldots,T$, minimizes the cost).

From the Monge formulation (3.1) and Lemma 27, we have:

$$C(\mu_0, \mu^*) = \min_{\substack{T:\Omega\to\Omega \\ T_\#\mu_0=\mu^*}} \int_\Omega c(x, T(x))\, d\mu_0(x),$$

and let T^* be a minimizing map above. Let T_0 be the identity map on Ω and let $\{T_k\}_{k=1}^K$ be a sequence of maps on Ω such that $T_K \circ \ldots \circ T_0 = T^*$, with $\mu_k = (T_k \circ \ldots \circ T_0)_\# \mu_0 = T_{k\#} \ldots T_{0\#}\mu_0$. Since c is a metric, we have:

$$c(x, T^*(x)) \leq \sum_{k=1}^{K} c(T_{k-1} \circ \ldots \circ T_0(x), T_k \circ \ldots \circ T_0(x)).$$

It then follows that:

$$\begin{aligned}
C(\mu_0, \mu^*) &= \int_\Omega c(x, T^*(x))\, d\mu_0(x) \\
&\leq \int_\Omega \sum_{k=1}^{K} c(T_{k-1} \circ \ldots \circ T_0(x), T_k \circ \ldots \circ T_0(x)) d\mu_0(x) \\
&= \sum_{k=1}^{K} \int_\Omega c(T_{k-1} \circ \ldots \circ T_0(x), T_k \circ \ldots \circ T_0(x)) d\mu_0(x) \\
&= \sum_{k=1}^{K} \int_\Omega c(x, T_k(x)) d\mu_{k-1}(x) \\
&= \sum_{k=1}^{K} C(\mu_{k-1}, \mu_k).
\end{aligned}$$

We also have:

$$c(x, T^*(x)) = \min_{\substack{T_1,\ldots,T_K \\ T_k:\Omega\to\Omega \\ T_K\circ\ldots\circ T_0=T^*}} \sum_{k=1}^{K} c(T_{k-1} \circ \ldots \circ T_0(x), T_k \circ \ldots \circ T_0(x)),$$

where the minimum is attained when the point $T_k \circ \ldots \circ T_0(x)$ lies on the geodesic from $T_{k-1} \circ \ldots \circ T_0(x)$ to $T^*(x)$. This can be seen from the fact that for any $x_1, x_2 \in \Omega$, $z^* \in \arg\min_{z\in\Omega} c(x_1, z) +$

$c(z, x_2)$ lies on the geodesic from x_1 to x_2. Thus, we get:

$$C(\mu_0, \mu^*) = \min_{\substack{T_1,\ldots,T_K \\ T_k:\Omega\to\Omega \\ T_K\circ\ldots\circ T_0 = T^* \\ T_{k\#}\mu_{k-1}=\mu_k}} \sum_{k=1}^{K} C(\mu_{k-1}, \mu_k).$$

We further note that any minimizing sequence $\{\mu_k\}_{k=1}^K$ must be generated by a sequence of maps $\{T_k\}_{k=1}^K$ such that for any $x \in \Omega$, $T_k \circ \ldots \circ T_0(x)$ lies on the geodesic from $T_{k-1} \circ \ldots \circ T_0(x)$ to $T^*(x)$, which yields (3.10). □

From the set of minimizing sequences characterized by Lemma 32, we are interested in those sequences for which the individual stage costs are upper bounded by an $\varepsilon > 0$. We thereby consider the following optimization-based iterative scheme to generate a minimizing sequence:

$$\begin{aligned} \mu_{k+1} \in \arg\min_{\nu\in\mathscr{P}(\Omega)} \; & C(\mu_k, \nu) + C(\nu, \mu^*) \\ \text{s.t.} \quad & C(\mu_k, \nu) \leq \varepsilon, \end{aligned} \tag{3.11}$$

where the iterative scheme (3.11) additionally satisfies the constraint $\lim_{k\to\infty} \mu_k = \mu^*$. Now, let T_k^* be an optimal transport map from μ_k to μ^*. We now construct the following optimization-based iterative process:

$$\begin{aligned} x(k+1) \in \arg\min_{z\in\Omega} \; & c(x(k), z) + c(z, T_k^*(x(k))) \\ \text{s.t.} \quad & c(x(k), z) \leq \varepsilon, \end{aligned} \tag{3.12}$$

where $x(k+1)$ obtained from the above process lies on the geodesic connecting $x(k)$ and $T_k^*(x(k))$. We now have the following lemma on the connection between the process (3.12) and the iterative scheme (3.11):

Lemma 33. *The law of the process* (3.12)*, when $x(0) \sim \mu_0$, evolves according to* (3.11).

Proof. This result follows from the arguments in the proof of Lemma 32. □

Following Lemma 33, it is clear that if we can compute the optimal transport map T_k^*, then (3.12) defines an iterative scheme for multi-stage optimal transport from an initial μ_0 to μ^*. We achieve this equivalently using the Kantorovich duality via the following process:

$$x(k+1) \in \arg \min_{z \in B_\varepsilon^c(x(k))} c(x(k), z) + \phi_{\mu_k \to \mu^*}(z), \tag{3.13}$$

We recall that $B_\varepsilon^c(x(k))$ is the closed ε-ball with respect to the metric c, centered at $x(k)$. The following lemma establishes that the processes (3.12) and (3.13) are equivalent.

Lemma 34. *The processes* (3.12) *and* (3.13) *are equivalent. The equivalence is in the sense that the sets of minimizers in* (3.12) *and* (3.13) *are equal.*

Proof. We recall from (3.4) and Lemma 30 that for the transport $\mu_k \to \mu^*$, and for any $x \in \Omega$, we have:

$$\phi_{\mu_k \to \mu^*}(x) = \inf_{y \in \Omega} c(x, y) + \phi_{\mu_k \to \mu^*}(y). \tag{3.14}$$

Also, for any $x, y \in \Omega$, we have the inequality $\phi_{\mu_k \to \mu^*}(x) \leq c(x, y) + \phi_{\mu_k \to \mu^*}(y)$. This implies in particular that for any transport map T_k from μ_k to μ^*, we get $\phi_{\mu_k \to \mu^*}(x) \leq c(x, T_k(x)) + \phi_{\mu_k \to \mu^*}(T_k(x))$. It then follows that:

$$\begin{aligned} \int_\Omega \left(\phi_{\mu_k \to \mu^*}(x) - \phi_{\mu_k \to \mu^*}(T_k(x))\right) d\mu_k(x) &= \int_\Omega \phi_{\mu_k \to \mu^*} d\mu_k - \int_\Omega \phi_{\mu_k \to \mu^*} d\mu^* \\ &\leq \int_\Omega c(x, T_k(x)) d\mu_k(x). \end{aligned}$$

We see that the LHS is the optimal transport cost obtained from the Kantorovich dual formulation, while an infimum over the RHS w.r.t. T_k would again yield the optimal transport cost from the Monge formulation and an equality would then be attained. Therefore, we get that the equality is attained when $T_k = T_k^*$, the corresponding optimal transport map from μ_k to μ^*. Thus, we infer that $\phi_{\mu_k \to \mu^*}(x) = c(x, T_k^*(x)) + \phi_{\mu_k \to \mu^*}(T_k^*(x))$ μ_k-almost everywhere in Ω. Since $c(x, T_k^*(x)) =$

$c(x,z)+c(z,T_k^*(x))$ for any (and only) z on the geodesic from x to $T_k^*(x)$, we can write:

$$\phi_{\mu_k\to\mu^*}(x)-\phi_{\mu_k\to\mu^*}(z)+\phi_{\mu_k\to\mu^*}(z)-\phi_{\mu_k\to\mu^*}(T_k^*(x))=c(x,z)+c(z,T_k^*(x)),$$

which implies that:

$$\left[\phi_{\mu_k\to\mu^*}(x)-\phi_{\mu_k\to\mu^*}(z)-c(x,z)\right]+\left[\phi_{\mu_k\to\mu^*}(z)-\phi_{\mu_k\to\mu^*}(T_k^*(x))-c(z,T_k^*(x))\right]$$
$$=0.$$

Moreover, since the expressions on the LHS are each non-positive, and their sum is zero, we get that they are individually zero. In other words, for any (and only) z on the geodesic from x to $T_k^*(x)$ we get $\phi_{\mu_k\to\mu^*}(x)-\phi_{\mu_k\to\mu^*}(z)-c(x,z)=0$, and these $z\in\Omega$ are in fact minimizers in (3.14). Therefore, set of minimizers obtained from (3.13) is essentially the segment of the geodesic from $x(k)$ to $T_k^*(x(k))$ contained in the ball $B_\varepsilon^c(x(k))$ which is also the set of minimizers obtained from (3.12), establishing equivalence in this sense between the processes (3.12) and (3.13). □

3.4 Multi-agent optimal transport

Working within the framework established in Section 3.3, we develop in this section the algorithm for multi-agent optimal transport based on distributed online optimization.

Let $\{x_i(0)\}_{i=1}^N$ be the positions of the N agents, distributed independently and identically according to a probability measure μ_0. The idea is to transport the agents by the iterative scheme (3.13) to obtain $\{x_i(k)\}_{i=1}^N$ at any time k. Let $\widehat{\mu}_N(k)=\frac{1}{N}\sum_{i=1}^N\delta_{x_i(k)}$ be the empirical measure generated by the agents $\{x_i(k)\}_{i=1}^N$ at time k. To this end, we formulate a (finite) N-dimensional distributed optimization to be implemented by the agents to obtain local estimates of the Kantorovich potential. We approximate the true Kantorovich potential by a $\Phi^d:\mathbb{N}\times\Omega\to\mathbb{R}$ generated by an (finite) N-dimensional vector $\phi(k)=(\phi^1(k),\dots,\phi^N(k))\in\mathbb{R}^N$, such that $\Phi^d(k,x_i(k))=\phi^i(k)$ for $i\in\{1,\dots,N\}$ and $\Phi^d(k,x)$ for $x\in\Omega\setminus\{x_1(k),\dots,x_N(k)\}$ is defined by a suitable multivariate interpolation. In particular, let $\{\mathscr{V}_i(k)\}_{i=1}^N$ be the Voronoi partition of Ω generated by $\{x_1(k),\dots,x_N(k)\}$ w.r.t. the metric c, and $\Phi^d = \sum_{i=1}^N \phi^{\mathscr{V}_i(k)}$ (decomposed into a sum of N functions $\phi^{\mathscr{V}_i(k)}$ with supports $\mathscr{V}_i(k)$). We assume that at time k, the agents i, j corresponding to neighboring cells $\mathscr{V}_i(k)$ and $\mathscr{V}_j(k)$ are connected by an edge, which defines a connected graph $G(k) = \left(\{x_i(k)\}_{i=1}^N, E(k)\right)$ (where $E(k)$ is the edge set of the graph $G(k)$ at time k).

Dropping the index k (as is clear from context), the finite dimensional approximation of the Kantorovich duality (3.7) for the transport between $\widehat{\mu}_N$ and μ^*, restricted to the graph G, is given by:

$$\begin{aligned} \max_{(\phi^1,\dots,\phi^N)} \; & \sum_{i=1}^N \left(\frac{1}{N} \cdot \phi^i - \mathbb{E}_{\mu^*}[\phi^{\mathscr{V}_i}] \right) \\ \text{s.t.} \quad & |\phi^i - \phi^j| \le c(x_i, x_j), \quad \forall (i,j) \in E. \end{aligned} \tag{3.15}$$

We call (3.15) a restriction of (3.7) to the graph G because we only impose the constraint $|\phi^i - \phi^j| \le c(x_i, x_j)$ on neighbors i, j on the graph.

We solve the optimization problem (3.15) by a primal-dual algorithm, and its solution is used to update the agent positions by (3.13). We take Φ^d here to be a simple function, such that $\phi^{\mathscr{V}_i}(x) = \phi^i$ for $x \in \mathscr{V}_i$. The Lagrangian for the problem (3.15), with Φ^d a simple function and $c(x_i, x_j) = c_{ij}$, is given by:

$$L_d = \sum_{i=1}^N \phi^i \left(\frac{1}{N} - \mu^*(\mathscr{V}_i) \right) - \frac{1}{2} \sum_{i=1}^N \sum_{j \in \mathscr{N}_i} \lambda_{ij} \left(\left|\phi^i - \phi^j\right|^2 - c_{ij}^2 \right).$$

The primal-dual (primal-ascent, dual-descent) algorithm (with step size τ) is then given by:

$$\begin{aligned} \phi^i(l+1) &= \phi^i(l) - \tau \sum_{j \in \mathscr{N}_i} \lambda_{ij}(l) \left(\phi^i(l) - \phi^j(l) \right) + \left(\frac{1}{N} - \mu^*(\mathscr{V}_i) \right), \\ \lambda_{ij}(l+1) &= \max \left\{ 0, \lambda_{ij}(l) + \tau \left(\frac{1}{2} \left| \phi^i(l) - \phi^j(l) \right|^2 - c_{ij}^2 \right) \right\}, \quad \text{where } j \in \mathscr{N}_i. \end{aligned} \tag{3.16}$$

We note from the structure of the above algorithm that it renders itself to a distributed implementation by the agents, where agent i uses information from its neighbors $j \in \mathscr{N}_i$ to update ϕ^i and $\{\lambda_{ij}\}_{j\in\mathscr{N}_i}$. The primal algorithm is in fact a weighted Laplacian-based update.

At the end of every step $x_i(k) \mapsto x_i(k+1)$ from (3.13), the agent i assigns $\phi^i \leftarrow \Phi_k^d(x_i(k+1))$ as the initial condition for the primal algorithm (3.16) at the time step $k+1$ of the transport. Moreover, we are interested in an on-the fly implementation of the transport, in that the agents do not wait for convergence of the distributed primal-dual algorithm but carry out n iterations of it for every update step (3.13), as outlined formally in the algorithm below..

Algorithm 1. Multi-agent (on-the-fly) optimal transport

Input: Target measure μ^*, Transport cost $c(x,y)$, Bound on step size ε, Time step τ

For each agent i at time instant k of transport:

Obtain: Positions $x_j(k)$ of neighbors within communication/sensing radius r ($r \leq \text{diam}(\Omega)$, large enough to cover Voronoi neighbors)

Compute: Voronoi cell $\mathscr{V}_i(k)$, Mass of cell $\mu^*(\mathscr{V}_i(k))$, Voronoi neighbors $\mathscr{N}_i(k)$

Initialize: $\phi^i \leftarrow \Phi_{k-1}^d(x_i(k))$, $\lambda_{ij} \leftarrow \lambda_{ij}(k-1)$ (with $\Phi_0^d = 0$, $\lambda_{ij}(0) = 0$)

Implement n iterations of primal-dual algorithm (3.16) (synchronously, in communication with neighbors $j \in \mathscr{N}_i$) to obtain $\phi^i(k)$, $\lambda_{ij}(k)$

Communicate with neighbors $j \in \mathscr{N}_i$ to obtain $\phi^j(k)$, construct local estimate of Φ_k^d by multivariate interpolation

Implement transport step (3.13) with local estimate of Φ_k^d (which approximates $\phi_{\mu_k \to \mu^*}$)

3.5 Analysis of PDE model

We investigate the behavior of the multi-agent transport by the update scheme (3.13) by studying the candidate system of PDEs for the continuous time and $N \to \infty$ limit. The results contained in this section are summarized below:

1. The candidate PDE model for transport in the continuous-time and $N \to \infty$ limit of the transport scheme (3.13) is derived formally in Section 3.5.1. The transport is described by the continuity equation (2.14) with the transport vector field (3.17).

2. In Section 3.5.2, we first derive the candidate PDE model (3.20) for the primal-dual algorithm (3.16) in the continuous-time and $N \to \infty$ limit. We then establish analytically

that the solutions to (3.20) converge as $t \to \infty$ to the optimality condition of the Kantorovich duality, in Lemma 36.

3. Section 3.5.3 deals with the stability of the feedback interconnection between the transport PDE (continuity equation with the transport vector field (3.17)) and the primal-dual flow (3.20). Convergence of the probability density (as solutions to the transport PDE) to the target density in the limit $t \to \infty$, provided that the primal-dual flow is always at steady state, is first established in Theorem 8. On-the-fly implementation is considered next, and a convergence result is obtained under a second-order relaxation of the dual flow in Theorem 9. Although the primal-dual flow is asymptotically stable and the transport PDE under the action of the field (3.17) is asymptotically stable, the stability of the feedback interconnected system of PDEs, in general, does not follow. This motivates the second-order relaxation of the dual flow, and we are able to establish asymptotic stability of the feedback interconnected system through a backstepping control of the dual flow. We reserve the feedback interconnection of the transport PDE directly with the primal-dual flow (3.20) for investigated by numerical simulations in Section 4.5. Although we are only able to obtain analytical stability results for the feedback interconnection with a relaxed dual flow, we observe convergence in simulation of the original feedback interconnected system, which motivates us to conjecture that it is indeed stable.

3.5.1 Transport PDE

We recall that the continuous-time evolution of a probability density function is described by the continuity equation (2.14):

$$\partial_t \rho + \nabla \cdot (\rho \mathbf{v}) = 0,$$

where $\mathbf{v}$ is the underlying transport vector field. In what follows, we derive the transport vector field which is the candidate for the continuous-time limit of the update scheme (3.13). We assume that all the probability measures considered have the same support Ω.

Let $x \in \Omega$, $\mu \in \mathscr{P}(\Omega)$ and $x^+ \in \arg\min_{z \in B_\varepsilon^c(x)} c(x,z) + \phi_{\mu \to \mu^*}(z)$, where x^+ is the update by the scheme (3.13). It then follows that:

$$c(x,x^+) + \phi_{\mu \to \mu^*}(x^+) \leq \phi_{\mu \to \mu^*}(x),$$

where $x^+ \in B_\varepsilon^c(x)$. We interpret the iterative scheme as a discrete-time dynamical system with uniform timestep Δt between successive instants, and derive the continuous-time limit by letting $\Delta t = g(\varepsilon) \to 0$ (where $g : \mathbb{R} \to \mathbb{R}$ is a monotonically increasing function). We have that $\phi_{\mu \to \mu^*}$ is bounded and continuously differentiable, which implies that $\lim_{\varepsilon \to 0} x^+ = x$ and $\lim_{x^+ \to x} \nabla \phi_{\mu \to \mu^*}(x^+) = \nabla \phi_{\mu \to \mu^*}(x)$. Let $\mathbf{v}(x) = \lim_{\Delta t \to 0} \frac{1}{\Delta t}(x^+ - x)$ (we note that this limit indeed exists), and we have:

$$\lim_{\Delta t \to 0} \frac{1}{\Delta t} c(x,x^+) \leq \lim_{\Delta t \to 0} \frac{1}{\Delta t} \left(\phi_{\mu \to \mu^*}(x) - \phi_{\mu \to \mu^*}(x^+) \right),$$

and it follows that:

$$\xi(x)\, |\mathbf{v}(x)| \leq -\nabla \phi_{\mu \to \mu^*}(x) \cdot \mathbf{v}(x).$$

The above inequality is satisfied only if $\mathbf{v}(x) = -\alpha \nabla \phi_{\mu \to \mu^*}(x)$ when $\lambda \neq 0$ (the Lagrangian dual function corresponding to the constraint $|\nabla \phi_{\mu \to \mu^*}| \leq \xi$) and for any $\alpha \geq 0$. This follows from the fact that $|\nabla \phi_{\mu \to \mu^*}| = \xi$, as $\phi_{\mu \to \mu^*}$ is the solution to (3.9) for the transport from μ to μ^*, satisfies (3.8) and occurs at the boundary of the constraint. Therefore, as $\Delta t \to 0$, we have the candidate velocity field:

$$\mathbf{v} = -\alpha \nabla \phi_{\mu \to \mu^*}, \tag{3.17}$$

where α can be any non-negative function on Ω. The implementation of the transport with the vector field (3.17) requires the computation of the Kantorovich potential $\phi_{\mu_t \to \mu^*}$ at any time t. Thus, we set up a primal-dual flow to obtain the Kantorovich potential as the solution to (3.9) (to which (3.15) is seen as the discrete counterpart as noted earlier).

3.5.2 Primal-Dual flow

The Lagrangian functional corresponding to the restricted Kantorovich duality (3.9) (to which (3.15) is seen as the discrete counterpart as noted earlier) for the optimal transport from μ to μ^* is given by:

$$L(\phi,\lambda) = \int_\Omega \phi(\rho - \rho^*) - \frac{1}{2}\int_\Omega \lambda(|\nabla\phi|^2 - |\xi|^2), \tag{3.18}$$

where all the integrals are with respect to the Lebesgue measure, and $\lambda \geq 0$ is the Lagrange multiplier function for the constraint $|\nabla\phi| \leq \xi$ (which corresponds to the set $\mathscr{L}(\Omega)\cap\mathscr{C}^1(\Omega)$), as specified in (3.8), which we have rewritten here as $|\nabla\phi|^2 \leq |\xi|^2$.

Lemma 35 (Optimality conditions)**.** *The necessary and sufficient conditions for a feasible solution $\bar{\phi}$ of* (3.9) *to be optimal are:*

$$\begin{aligned} -\nabla\cdot\left(\bar{\lambda}\nabla\bar{\phi}\right) &= \rho - \rho^*, \quad (\textit{in } \Omega) \\ \bar{\lambda}\nabla\bar{\phi}\cdot\mathbf{n} &= 0, \quad (\textit{on } \partial\Omega) \\ \bar{\lambda} \geq 0, &\quad |\nabla\bar{\phi}| \leq \xi, \quad \textit{(Feasibility)} \\ \bar{\lambda}(|\nabla\bar{\phi}| - \xi) &= 0 \ \ a.e., \quad \textit{(Complementary slackness)} \end{aligned} \tag{3.19}$$

where $\bar{\lambda}$ is the optimal Lagrange multiplier function.

Proof. We consider the Lagrangian (3.18), for which the first variation with respect to a variation $\delta\phi$, is given by:

$$\left\langle \frac{\delta L}{\delta \phi}, \delta\phi \right\rangle = \int_\Omega (\rho - \rho^*)\delta\phi - \int_\Omega \lambda \nabla\phi \cdot \nabla\delta\phi$$
$$= \int_\Omega (\rho - \rho^*)\delta\phi + \int_\Omega \nabla\cdot(\lambda\nabla\phi)\delta\phi - \int_{\partial\Omega} \lambda\nabla\phi\cdot\mathbf{n}\delta\phi,$$

where we have used the divergence theorem to obtain the final equality. We have $\left\langle \frac{\delta L}{\delta \phi}, \delta\phi \right\rangle = 0$ for any variation $\delta\phi$ around the stationary point $(\bar{\phi}, \bar{\lambda})$. Therefore, we obtain $-\nabla\cdot\left(\bar{\lambda}\nabla\bar{\phi}\right) = \rho - \rho^*$ in Ω and $\bar{\lambda}\nabla\bar{\phi}\cdot\mathbf{n} = 0$ on $\partial\Omega$. Also, $\bar{\lambda} \geq 0$ is the feasibility condition for the Lagrange multiplier, $|\nabla\bar{\phi}| \leq \xi$ is the feasibility condition on ϕ and $\bar{\lambda}(|\nabla\bar{\phi}| - \xi) = 0$ is the complementary slackness condition. These correspond to the necessary KKT conditions, which for this problem (linear objective function and a convex constraint) are also the sufficient conditions for optimality. □

We now define a primal-dual flow to converge to the saddle point of the Lagrangian (3.18). For this, we henceforth consider the functions ϕ and λ to be additionally parametrized by time t. The primal-dual flow for the Lagrangian (3.18) is given by:

$$\begin{aligned} \partial_t\phi &= \nabla\cdot(\lambda\nabla\phi) + \rho - \rho^*, \\ \nabla\phi\cdot\mathbf{n} &= 0, \quad \text{on } \partial\Omega, \\ \partial_t\lambda &= \frac{1}{2}\left[|\nabla\phi|^2 - |\xi|^2\right]^+_\lambda, \\ \phi(0,x) &= \phi_0(x), \;\; \lambda(0,x) = \lambda_0(x), \end{aligned} \tag{3.20}$$

where $[f]^+_\lambda = \begin{cases} f & \text{if } \lambda > 0 \\ \max\{0, f\} & \text{if } \lambda = 0 \end{cases}$ is a projection operator.

We note that $\partial_t\phi = \frac{\delta L}{\delta\phi}$ and $\partial_t\lambda = \left[-\frac{\delta L}{\delta\lambda}\right]^+_\lambda$, and we have a gradient ascent on $L(\phi,\lambda)$ w.r.t. ϕ and a projected gradient descent on $L(\phi,\lambda)$ w.r.t. λ.

Remark 3 (On the connection between (3.20) and (3.16))**.** *The primal-dual algorithm* (3.16) *is the discretization of the primal-dual flow* (3.20) *over a graph G (as defined in the previous subsection) with a step size* τ*. The term* $-\sum_{j\in\mathcal{N}_i}\lambda_{ij}(l)\left(\phi^i(l)-\phi^j(l)\right)$ *in* (3.16) *is the action of the weighted Laplacian matrix (with weights* $\lambda_{ij}(l)$*) on* $\phi(l)$*, which is the discretization over the graph of the term* $\nabla\cdot(\lambda\nabla\phi)$ *in* (3.20)*.*

Remark 4 (Existence and Uniqueness of solutions to (3.20))**.** *We first note that* (3.20) *generates a strongly continuous semigroup of operators and we interpret any solution of* (3.20) *as generated by this operator semigroup. We now consider the evolution of the Lagrange multiplier function* λ*. Letting* $\lambda_0 \equiv 0$ *and* $h = \frac{1}{2}\left[|\nabla\phi|^2 - |\xi|^2\right]^+_\lambda$*, we note that at any* $x\in\Omega$*, we have* $\lambda(t,x) = \int_0^t h(\tau,x)d\tau$*. Thus, a unique solution* λ *exists if* $h(t,x) = \frac{1}{2}\left[|\nabla\phi|^2 - |\xi|^2\right]^+_\lambda$ *is integrable in time at every* $x\in\Omega$*, which depends on the regularity of the solution* ϕ*. However, we do not apriori characterize or establish the desired level of regularity of the solutions* ϕ*, but instead assume that* $\lambda \in L^\infty(0,T;L^\infty(\Omega))$ *for any given* $T>0$*. For any given* $T>0$*, under the assumptions that* $\lambda\in L^\infty(0,T;L^\infty(\Omega))$ *and* $\rho,\rho^*\in L^2(0,T;L^2(\Omega))$*, there exists a unique weak solution* $\phi\in L^2(0,T;H^1(\Omega))$ *to the primal-dual flow* (3.20) *(we recall that we impose the Neumann boundary condition as* $\nabla\phi\cdot\mathbf{n}=0$ *on* $\partial\Omega$*). The existence and uniqueness results follow by adapting the arguments presented in [55], Section 7.1 to the current problem (a homogenous second order parabolic PDE with a Neumann boundary condition). We note that the solution* ϕ *completely determines* λ*. To guarantee that* $\lambda\in L^\infty(0,T;L^\infty(\Omega))$ *is consistent with the solution* ϕ*, it may be necessary to add further regularity assumptions on* ρ,ρ^**. However, further investigation into the regularity of solutions of the primal-dual flow is beyond the scope of this present work.*

Assumption 6 (Well-posedness of primal-dual flow)**.** *We assume that* (3.20) *is well-posed, with solution* (ϕ,λ) *such that* $\phi\in L^\infty(0,\infty;H^1(\Omega))$ *and the Lagrange multiplier function* $\lambda \in L^\infty(0,\infty;L^\infty(\Omega))$ *and is precompact in* $L^2(\Omega)$*.*

The following lemma establishes the convergence of solutions of (3.20) to the optimality conditions (3.19):

Lemma 36 (Convergence of primal-dual flow)**.** *Solutions* (ϕ_t, λ_t) *to the primal-dual flow* (3.20)*, under Assumption 6 on the well-posedness of the primal-dual flow, converge to an optimizer* $(\tilde{\phi}, \tilde{\lambda})$ *given in* (3.19) *in the* L^2 *norm as* $t \to \infty$*, for any fixed* $\rho, \rho^* \in L^2(\Omega)$.

Proof. Let $(\tilde{\phi}, \tilde{\lambda})$ be an optimizer of (3.19) and let:

$$V(\phi,\lambda) = \frac{1}{2}\int_\Omega |\phi - \bar{\phi}|^2 \mathrm{dvol} + \frac{1}{2}\int_\Omega |\lambda - \bar{\lambda}|^2 \mathrm{dvol}.$$

Clearly, $V(\phi,\lambda) \geq 0$ for all $\phi, \lambda \in L^2(\Omega)$. The time-derivative of V along the solutions of the primal-dual flow (3.20) is given by:

$$\begin{aligned}
\dot{V} &= \left\langle \frac{\delta L}{\delta \phi}, \phi - \bar{\phi} \right\rangle + \left\langle \left[-\frac{\delta L}{\delta \lambda}\right]^+_\lambda, \lambda - \bar{\lambda} \right\rangle \\
&= \left\langle \frac{\delta L}{\delta \phi}, \phi - \bar{\phi} \right\rangle - \left\langle \frac{\delta L}{\delta \lambda}, \lambda - \bar{\lambda} \right\rangle + \left\langle \frac{\delta L}{\delta \lambda} + \left[-\frac{\delta L}{\delta \lambda}\right]^+_\lambda, \lambda - \bar{\lambda} \right\rangle.
\end{aligned}$$

Since L is concave in ϕ and convex in λ, we get:

$$\begin{aligned}
\dot{V} &\leq L(\phi,\lambda) - L(\bar{\phi},\lambda) + L(\phi,\bar{\lambda}) - L(\phi,\lambda) + \left\langle \frac{\delta L}{\delta \lambda} + \left[-\frac{\delta L}{\delta \lambda}\right]^+_\lambda, \lambda - \bar{\lambda} \right\rangle \\
&= L(\bar{\phi},\bar{\lambda}) - L(\bar{\phi},\lambda) + L(\phi,\bar{\lambda}) - L(\bar{\phi},\bar{\lambda}) + \left\langle \frac{\delta L}{\delta \lambda} + \left[-\frac{\delta L}{\delta \lambda}\right]^+_\lambda, \lambda - \bar{\lambda} \right\rangle.
\end{aligned}$$

We have that $L(\bar{\phi},\bar{\lambda}) - L(\bar{\phi},\lambda) \leq 0$ and $L(\phi,\bar{\lambda}) - L(\bar{\phi},\bar{\lambda}) \leq 0$ (recall that $(\bar{\phi},\bar{\lambda})$ is a saddle point of L). Moreover, by definition, when $\lambda(t,x) > 0$, we have $\left[-\frac{\delta L}{\delta \lambda}\right]^+_\lambda = -\frac{\delta L}{\delta \lambda}$ at (t,x), and when $\lambda(t,x) = 0$ (which implies that $\lambda - \bar{\lambda} \leq 0$), we have $\left[-\frac{\delta L}{\delta \lambda}\right]^+_\lambda \geq -\frac{\delta L}{\delta \lambda}$ at (t,x). This implies that $\left\langle \frac{\delta L}{\delta \lambda} + \left[-\frac{\delta L}{\delta \lambda}\right]^+_\lambda, \lambda - \bar{\lambda} \right\rangle \leq 0$ at any (t,x). We therefore can say that $\dot{V} \leq 0$. Moreover, by Assumption 6, it holds that the orbit ϕ is bounded in $H^1(\Omega)$ which, by Lemma 5, is compactly embedded in $L^2(\Omega)$. It then follows that the orbit is precompact in $L^2(\Omega)$. Moreover, by Assumption 6, we have that λ is precompact in $L^2(\Omega)$. We get that $\dot{V} = 0$ only at an optimizer $(\tilde{\phi}, \tilde{\lambda})$, which implies that the flow converges asymptotically to a $(\tilde{\phi}, \tilde{\lambda})$. □

3.5.3 Convergence of PDE-based transport

Following the outline from earlier in the section, we now investigate the convergence properties of transport by the vector field:

$$\mathbf{v} = -\alpha\nabla\phi_{\mu\to\mu^*}.$$

However, as discussed earlier, the scenario of particular interest to us is that of an on-the-fly implementation of the transport, where we do not wait for the convergence of the primal-dual flow to its steady state to obtain $\phi_{\mu\to\mu^*}$. This results in a coupling between the transport PDE and the primal-dual flow, and we investigate the convergence of solutions of this system of PDEs later in this section.

Lemma 37. *The transport* (2.14) *by the vector field* (3.17) *is a gradient flow, in the sense of Definition 13, on the optimal transport cost* $C(\cdot,\mu^*) : \mathscr{P}(\Omega) \to \mathbb{R}_{\geq 0}$.

Proof. From Lemma 29, we have the strong duality $K(\mu,\mu^*) = C(\mu,\mu^*)$. The Kantorovich potential $\phi_{\mu\to\mu^*}$ is such that $\nabla\phi_{\mu\to\mu^*} = \nabla\left(\frac{\delta K}{\delta\mu}\right)$ (since $\phi_{\mu\to\mu^*} = \frac{\delta K}{\delta\mu}$, and we refer the reader to Chapter 7 in [106] for a proof). Therefore, the transport vector field $\mathbf{v} = -\alpha\nabla\phi_{\mu\to\mu^*}$ yields a gradient flow on the optimal transport cost $C(\mu,\mu^*)$. □

Remark 5 (Existence and uniqueness of solutions to the transport PDE)**.** *We refer the reader to [7] for a detailed treatment of existence and uniqueness results for the continuity equation, for transport vector fields with Sobolev regularity. We make the necessary well-posedness assumption for our purposes.*

Assumption 7 (Well-posedness of gradient flow on optimal transport cost)**.** *We assume that the desired distribution* μ^* *is absolutely continuous (with density function* ρ^* *in* $H^1(\Omega)$*) with* $\mathrm{supp}(\mu^*) = \Omega$*. Further, we assume that* (2.14) *is well-posed for the gradient flow on the optimal transport cost, with solution* $\rho \in L^\infty(0,\infty;H^1(\Omega))$.

Theorem 8. *Under Assumption 7 on the well-posedness of the gradient flow on the optimal transport cost and for absolutely continuous initial distributions* μ_0 *with* $\text{supp}(\mu_0) = \Omega$*, the solutions* ρ *to the transport* (2.14) *by the vector field* $\mathbf{v} = -\frac{\lambda_{\mu\to\mu^*}}{\rho}\nabla\phi_{\mu\to\mu^*}$ *(where* $\phi_{\mu\to\mu^*}$ *and* $\lambda_{\mu\to\mu^*}$ *are the Kantorovich potential and the optimal Lagrange multiplier function for the transport* $\mu \to \mu^*$*) converge exponentially to* ρ^* *in the* L^2 *norm as* $t \to \infty$*.*

Proof. From the optimality conditions (3.19), we have that $\nabla\cdot(\lambda_{\mu\to\mu^*}\nabla\phi_{\mu\to\mu^*}) = \rho^* - \rho$, which implies that $\partial_t\rho = -\nabla\cdot(\rho\mathbf{v}) = \nabla\cdot(\lambda_{\mu\to\mu^*}\nabla\phi_{\mu\to\mu^*}) = \rho^* - \rho$ when $\rho > 0$. Moreover, we have that ρ_0 and ρ^* are strictly positive in Ω. Therefore, for any $t \in [0,\infty]$ and $x \in \mathring{\Omega}$, we have $\rho(t,x) > 0$. Consequently, since $\rho(t,x) > 0$, the transport vector field $\mathbf{v} = -\frac{\lambda_{\mu\to\mu^*}}{\rho}\nabla\phi_{\mu\to\mu^*}$ is well-defined on Ω. Let $V : L^2(\Omega) \to \mathbb{R}_{\geq 0}$ be defined by $V(\rho) = \frac{1}{2}\int_\Omega |\rho - \rho^*|^2 \text{ dvol}$, where ρ is the density function of the absolutely continuous probability measure μ. The time derivative $\dot{V}$, under the transport (2.14) by $\mathbf{v} = -\frac{\lambda_{\mu\to\mu^*}}{\rho}\nabla\phi_{\mu\to\mu^*}$ is given by:

$$\begin{aligned}\dot{V} &= \int_\Omega (\rho - \rho^*)\partial_t\rho = -\int_\Omega (\rho - \rho^*)\nabla\cdot(\rho\mathbf{v}) \\ &= \int_\Omega (\rho - \rho^*)\nabla\cdot(\lambda_{\mu\to\mu^*}\nabla\phi_{\mu\to\mu^*}).\end{aligned}$$

Further, from (3.19), we get:

$$\dot{V} = -\int_\Omega |\rho - \rho^*|^2 = -2V,$$

which implies that V is a Lyapunov functional for the transport by the vector field $\mathbf{v} = -\frac{\lambda_{\mu\to\mu^*}}{\rho}\nabla\phi_{\mu\to\mu^*}$. Moreover, by Assumption 7, we have that the solution ρ is bounded in $H^1(\Omega)$, which by the Rellich-Kondrachov theorem 5 is compactly contained in $L^2(\Omega)$. We then infer that the solution ρ to the transport (2.14) by the vector field $\mathbf{v} = -\frac{\lambda_{\mu\to\mu^*}}{\rho}\nabla\phi_{\mu\to\mu^*}$ is precompact, and therefore by the invariance principle in Lemma 6, converges to ρ^* in the L^2-norm in the limit $t \to \infty$, i.e. $\lim_{t\to\infty}\|\rho - \rho^*\|_{L^2} = 0$. Moreover, since we have $\dot{V} = -2V$, we note that the convergence is exponential. □

Remark 6 (Adaptation to tracking of time-varying target distributions)**.** *The exponential convergence result in the above theorem permits adaptation of the transport scheme to multi-agent tracking scenarios involving target distributions that evolve on a much slower timescale.*

We now present an on-the-fly implementation of the transport, where we do not wait for the primal-dual flow to reach steady state, but instead set the transport vector field as $\mathbf{v} = -\frac{\lambda}{\rho}\nabla\phi$, where ϕ and λ are supplied by (3.20). This results in a coupling between the transport PDE (2.14) and the primal-dual flow (3.20), and we investigate the behavior of the transport in simulation in Section 4.5.

We now establish the convergence of the on-the-fly transport under the primal flow and a fixed dual function $\lambda > 0$, which we define as follows:

$$\begin{aligned} \partial_t \phi &= \nabla \cdot (\lambda \nabla \phi) + \rho - \rho^*, \\ \nabla \phi \cdot \mathbf{n} &= 0, \quad \text{on } \partial\Omega, \\ \lambda &= \lambda(x) > 0. \end{aligned} \tag{3.21}$$

We note that transport under the relaxed primal-dual flow differs from the transport under (3.20) only in that the Lagrange multiplier function λ that weights the primal flow is fixed and does not vary in time.

Remark 7 (Existence and uniqueness of solutions to on-the-fly transport)**.** *We note that* (3.21) *generates a strongly continuous semigroup of operators and we interpret any solution of* (3.21) *as generated by this operator semigroup. We recall from Remark 4 that a unique weak solution $\phi \in L^2(0,T;H^1(\Omega))$ to the primal flow exists if $\rho, \rho^* \in L^2(0,T;H^1(\Omega))$ and $\lambda \in L^\infty(\Omega)$.*

Assumption 8. *We assume that the desired distribution μ^* is absolutely continuous (with density function ρ^* in $H^1(\Omega)$) and supported on Ω. Further, we assume that the primal flow* (3.21) *and the transport* (2.14) *are well-posed, with solutions ϕ and ρ such that $\phi \in L^\infty(0,\infty;H^1(\Omega))$, and strictly positive $\rho \in L^\infty(0,\infty;H^1(\Omega))$.*

Theorem 9 (Convergence of on-the-fly transport)**.** *Under Assumption 8, the solutions ρ to* (2.14) *with the transport vector field $\mathbf{v} = -\frac{\lambda}{\rho}\nabla\phi$, with ϕ from* (3.21)*, converge in the L^2-norm to ρ^* as $t \to \infty$, while the solutions to the primal flow* (3.21) *converge to the optimality condition* (3.19) *corresponding to $\rho = \rho^*$.*

Proof. We first note that since $\rho > 0$ from Assumption 8, the transport vector field $\mathbf{v} = -\frac{\lambda}{\rho}\nabla\phi$ is well-defined on Ω. We now consider the following Lyapunov functional:

$$E = \frac{1}{2}\int_\Omega \lambda|\nabla\phi|^2 + \frac{1}{2}\int_\Omega |\rho - \rho^*|^2,$$

where all the integrals are with respect to dvol. The time derivative of E under the flow (3.21) and $\mathbf{v} = -\frac{\lambda}{\rho}\nabla\phi$ is given by:

$$\dot{E} = \int_\Omega \lambda\nabla\phi\cdot\nabla\partial_t\phi + \int_\Omega (\rho-\rho^*)\partial_t\rho.$$

Applying the divergence theorem and using the boundary condition for the first term, and the continuity equation (2.14) for the second, we obtain:

$$\begin{aligned}\dot{E} &= -\int_\Omega \nabla\cdot(\lambda\nabla\phi)\partial_t\phi - \int_\Omega(\rho-\rho^*)\nabla\cdot(\rho\mathbf{v})\\ &= -\int_\Omega |\nabla\cdot(\lambda\nabla\phi)|^2 - \int_\Omega \nabla\cdot(\lambda\nabla\phi)(\rho-\rho^*)\\ &\quad + \int_\Omega(\rho-\rho^*)\nabla\cdot(\lambda\nabla\phi)\\ &= -\int_\Omega |\nabla\cdot(\lambda\nabla\phi)|^2.\end{aligned}$$

By Assumption 8 we have that the orbits ϕ and ρ are bounded in $H^1(\Omega)$ and by Lemma 5 (Rellich-Kondrachov theorem) we have that the orbits are precompact in $L^2(\Omega)$. Now, from the invariance principle in Lemma 6, we infer that the orbits of the system converge to the largest invariant set in $\dot{E}^{-1}(0)$. We have that $\dot{E} = 0$ implies $\|\nabla\cdot(\lambda\nabla\phi)\|_{L^2(\Omega)} = 0$, from which it follows by substitution in (3.21) that the transport (2.14) with $\mathbf{v} = -\frac{\lambda}{\rho}\nabla\phi$ yields $\rho \to_{L^2} \rho^*$, while ϕ

converges to the optimality conditions corresponding to $\rho = \rho^*$. $\square$

This leads us to the following algorithm for on-the-fly multi-agent transport under fixed, positive dual weighting:

Algorithm 2. Multi-agent (on-the-fly) optimal transport with fixed (dual) weighting

---

Input: Target measure μ^*, Weights (dual variable) λ_{ij}, Bound on step size ε, Time step τ
For each agent i **at time instant** k **of transport:**

1: Obtain: Positions $x_j(k)$ of neighbors within communication/sensing radius r ($r \leq \text{diam}(\Omega)$, large enough to cover Voronoi neighbors)
2: Compute: Voronoi cell $\mathscr{V}_i(k)$, Mass of cell $\mu^*(\mathscr{V}_i(k))$, Voronoi neighbors $\mathscr{N}_i(k)$
3: Initialize: $\phi^i \leftarrow \Phi^d_{k-1}(x_i(k))$ (with $\Phi^d_0 = 0$)
4: Implement n iterations of primal algorithm (3.21) (synchronously, in communication with neighbors $j \in \mathscr{N}_i$) to obtain $\phi^i(k)$
5: Communicate with neighbors $j \in \mathscr{N}_i$ to obtain $\phi^j(k)$, construct local estimate of Φ^d_k by multivariate interpolation
6: Implement transport step (3.13) with local estimate of Φ^d_k (which approximates $\phi_{\mu_k \to \mu^*}$)

3.6 Simulation studies and discussion

In this section, we present simulation results for multi-agent optimal transport in $\mathbb{R}^2$, based on the the iterative multi-stage transport scheme (3.13) (with c being the Euclidean metric and $\varepsilon = 0.02$), where the local estimates of the Kantorovich potential are computed by the distributed online algorithm (3.16) with a step size $\tau = 1$. We also present simulation results for the PDE-based transport (2.14) under the primal-dual flow (3.20).

We considered a bivariate Gaussian distribution with covariance $\Sigma = \begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix}$ and mean randomly chosen in $[0,1]^2$ as the target probability measure, and $N = 30$ agents for the transport. Figure 3.1 shows the agents along with the corresponding Voronoi partition of the domain, at three different stages (time instants $k = 0, 5, 10$) during the course of their transport. We observe that the agents are transported towards the target probability measure and that a quantization of the target measure is obtained. This is clarified further in Figure 3.2, as described below.

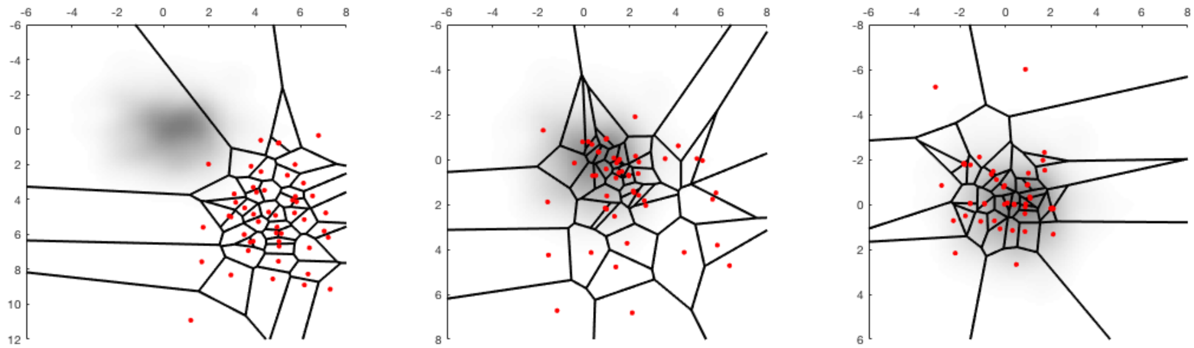

Figure 3.1. Positions of agents along with the Voronoi partition generated by them at three different stages (time instants $k = 0, 5, 10$) of transport by the iterative scheme (3.13) with local estimates of Kantorovich potential supplied by (3.16). Target probability measure shown in grayscale with a darker shade indicating a region of higher target density. The plots show convergence in time of the agents to full coverage of the target coverage profile (represented by the target probability distribution).

As we had noted in the previous section, there exists a fundamental trade-off between optimality and an on-the-fly implementation of the distributed optimal transport. We sought to investigate the extent of this trade-off in simulation by running multiple iterations n of the primal-dual algorithm (3.16) for every iteration of the transport (3.13). The underlying rationale is that the distributed computation is many times faster than the transport. Figure 3.2 shows the rate of convergence (w.r.t. the variance in target mass $\mu^*(\mathcal{V}_i)$ across the partition) for various values of n. Figure 3.3 is a plot of the net cost of transport w.r.t. n.

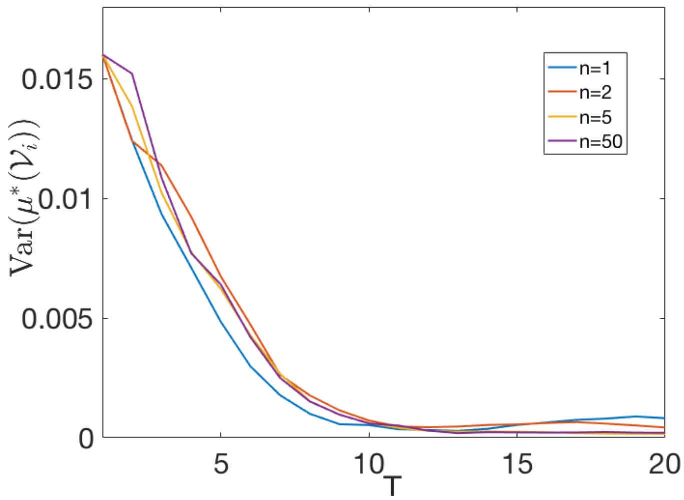

Figure 3.2. Variance in target mass $\mu^*(\mathscr{V}_i)$ across the partition vs time for various iteration steps n of the primal-dual algorithm (3.16) for every step of the transport (3.13).

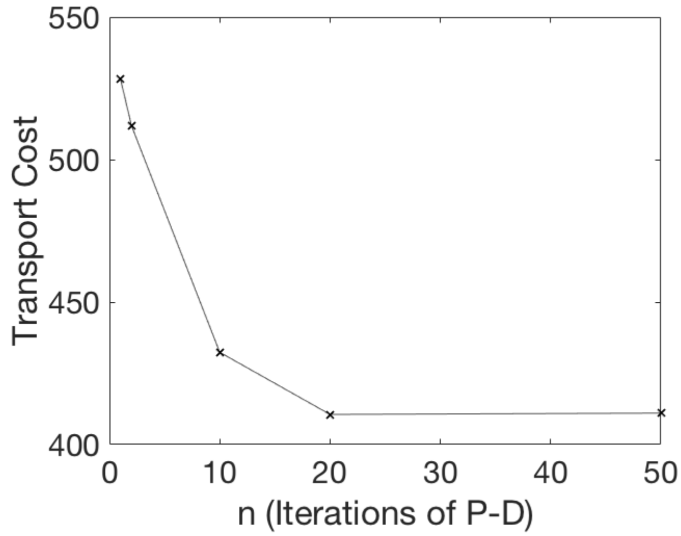

Figure 3.3. Net cost of transport for various iteration steps n of the primal-dual algorithm (3.16) for every step of the transport (3.13).

Figure 3.4 shows the evolution of the distribution of the agents over time. The grayscale images show the distribution of the agents in the domain, with darker shades representing higher density of agents at any given location. The domain is a 50×50 grid, and the PDE (3.20) was discretized over the grid. The initial distribution value was randomly generated (a random number was generated by the *rand* function in MATLAB for each cell of the grid and then normalized to obtain the probability distribution over the grid). The target density was defined by a grayscale image, as seen in the final subfigure in Figure 3.4. The cost of transport was chosen to be $c_{ij} = 1$ between neighboring cells i and j in the grid.

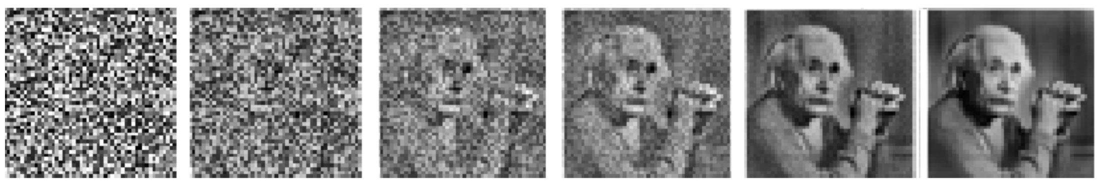

Figure 3.4. Distribution at various stages of the PDE-based transport (2.14) under the primal-dual flow (3.20). The figure shows convergence in time of the distribution to the target distribution represented by the final image.

We observe convergence of the on-the-fly transport under the primal-dual flow (3.20). Although we have established convergence of the transport analytically only under a primal flow with a fixed dual function, we conjecture that an on-the-fly transport under the primal-dual flow (3.20) also possesses the asymptotic stability property.

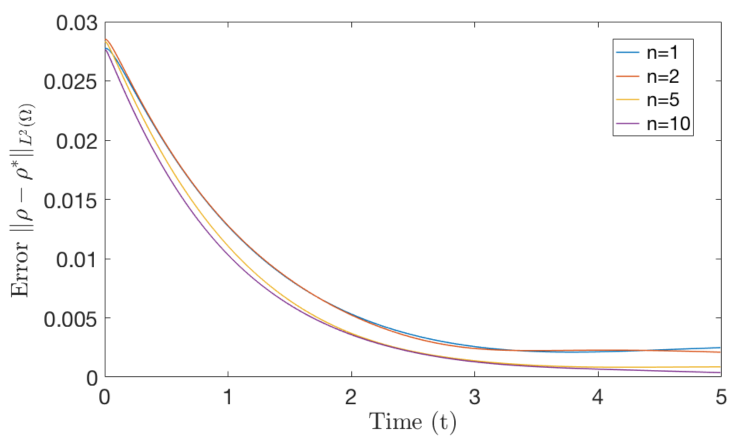

Figure 3.5. Density error $\|\rho - \rho^*\|_{L^2(\Omega)}$ vs time for various multiples n of the time scale of primal-dual flow (3.20) w.r.t the time scale of transport (2.14). The plot shows the rate of convergence to the optimal transport gradient flow (represented by $n = 10$).

Figure 3.5 is the plot of the L^2-density error $e(t) = \|\rho - \rho^*\|_{L^2(\Omega)}$ as a function of time, for various iteration steps of the primal-dual flow (to converge to the optimal gradient flow velocity) for every iteration of the transport PDE. We notice a significant improvement in the tracking performance (as measured by $e(t)$) within a few iterations of the primal-dual flow per iteration of the transport PDE, and the convergence to true optimal transport (in the sense of decay rate of the error $e(t)$) is obtained with approximately an order ($n \approx 10^1$) of magnitude time scale separation between computation and transport.

3.7 Summary

In this chapter, we proposed a scalable, distributed iterative proximal point algorithm for large-scale optimal transport of multi-agent collectives. We obtained a dynamical formulation of optimal transport of agents, for metric transport costs that are conformal to the Euclidean distance. We proposed a distributed primal-dual algorithm to be implemented by the agents to obtain local estimates of the Kantorovich potential, which are then used as local objectives in a proximal point algorithm for transport. We studied the behavior of the transport in simulation and presented an analysis of the candidate PDE model for the continuous time and $N \to \infty$ limit, establishing asymptotic stability of the transport. We explored in simulation the suboptimality of the on-the-fly implementation.

The material in this chapter, in full, has been submitted for publication to the SIAM Journal on Control and Optimization and is under review. It may appear as *Distributed Online Optimization for Multi-Agent Optimal Transport*, V. Krishnan and S. Martínez. A preliminary version of the work appeared in the proceedings of the IEEE Conference on Decision and Control, Miami Beach, USA, December 2018, as *Distributed optimal transport for the deployment of swarms*, V. Krishnan and S. Martínez. The dissertation author was the primary investigator and author of these papers.

Chapter 4

Self-organizing multi-agent transport

In this chapter, we adopt a viewpoint outlined in [13], wherein we make an amorphous medium abstraction of the swarm, which is essentially a manifold with an agent located at each point. We then model the system using PDEs and design distributed control laws for them. An important component of this paper is the Laplacian-based distributed algorithm which we call pseudo-localization algorithm, which the agents implement to localize themselves in a new coordinate frame. The convergence properties of the graph Laplacian to the manifold Laplacian have been studied in [19], which find useful applications in this paper.

The main contribution of this paper is the development of distributed control laws for the index- and position-free density control of swarms to achieve general 1D and a large class of 2D density profiles. In very large swarms with thousands of agents, particularly those deployed indoors or at smaller scales, presupposing the availability of position information or pre-assignment of indices to individual agents would be a strong assumption. In this paper, in addition to not making the above assumptions, the agents are only capable of measuring the local density, and in the 2D case, the density gradient and the normal direction to the boundary.

Under these assumptions, we present distributed pseudo-localization algorithms for one and two dimensions that agents implement to compute their position identifiers. Since every agent occupies a unique spatial position, we are able to rigorously characterize the resulting position assignment as a one-to-one correspondence between the set of spatial coordinates and the set of position identifiers, which corresponds to a diffeomorphism of the continuum domain. Based on this assignment, we then design control strategies for self-organization in one and two dimensions under the assumption that the motion control of agents is noiseless. The extension to the 2D case leads to new difficulties related to the control of the swarm boundaries. To address these, we implement a variant of the 1D pseudo-localization algorithm at the boundary during an initialization phase.

4.1 Bibliographical comments

In the literature, Markov chain-based methods have been widely used in addressing some of the key theoretical problems pertaining to swarm self-organization. By means of it, the swarm configuration is described through the partitioning the spatial domain in a finite number of larger size disjoint subregions, on which a probability distribution is defined. Then, the self-organization problem is reduced to the design of the transition matrix governing the evolution of this probability density function to ensure its convergence to a desired profile. A recent approach to density control using Markov chains is presented in [41], which includes additional conflict-avoidance constraints. In this setting every agent is able to determine the bin to which it belongs at every instant of time, which essentially means that individual agents have self-localization capabilities. Also, the dimensional transition matrix is synthesized in a central way at every instant of time by solving a convex optimization problem. In [14], the authors make use of inhomogeneous Markov chains to minimize the number of transitions to achieve a swarm formation. In this approach, the algorithm necessitates the estimation of the current swarm distribution, and computes the transition Markov matrices for each agent, at each instant of time. The fact that every agent needs to have an estimate of the global state (swarm distribution) at every time may not be desirable or feasible. The localization of each agent still remains a main assumption. Under similar conditions, one can find the manuscripts [12] and [29], which describe probabilistic swarm guidance algorithms. In [22], the authors present an approach to task allocation for a homogeneous swarm of robots. This is a Markov-chain based approach, where the goal is to converge to the desired population distribution over the set of tasks.

Many works in the literature use PDE-based methods to model swarm behaviour, where control action is applied along the boundary of the swarm. Previous works on PDE-based methods with boundary control include [60], where the authors present an algorithm for the deployment of agents onto families of planar curves. Here, the swarm collective dynamics are modeled by the reaction-advection-diffusion PDE and the particular family of curves to which the swarm is controlled to is parametrized by the continuous agent identity in the interval of unit length. An extension of this work to deployment on a family of 2D surfaces in 3D space can be found in [101]. The problem of planning and task allocation is addressed in the framework of advection-diffusion-reaction PDEs in [49]. In [59] and [57], the authors present an optimal control problem formulation for swarm systems, where microscopic control laws are derived from the optimal macroscopic description using a potential function approach. The problem of position-free extremum-seeking of an external scalar signal using a swarm of autonomous vehicles, inspired by bacterial chemotaxis, has been studied in [89].

4.2 Problem description and conceptual approach

In this section, we provide a high-level description of the proposed problem and explain the conceptual idea behind our approach. The technical details can be found in the following sections.

The problem at hand is to ultimately design a distributed control law for a swarm to converge to a desired configuration. Here, a swarm configuration is a density function ρ of the multi-agent system and the objective is that agents reconfigure themselves into a desired known density ρ^*. To do this, an agent at position x is able to measure the current local density value, $\rho(t,x)$; however, its position x within the swarm is unknown. Thus, given ρ^*, an agent at x cannot directly compute $\rho^*(x)$ nor a feedback law based on $\rho - \rho^*$. To solve this problem, we devise a mechanism that allows agents to determine their coordinates in a distributed way in an equivalent coordinate system.

Note that, given a diffeomorphism Θ^* from the spatial domain of the swarm onto the unit interval or disk (i.e. a coordinate transformation), we can equivalently provide the agents with a transformed density function p^*, such that $p^* = \rho^* \circ (\Theta^*)^{-1}$. In this way, instead of ρ^* the agents are given p^*, but still do not have access to Θ^*. The pseudo-localization algorithm is a mechanism that agents employ to progressively compute an appropriate (configuration-dependent) diffeomorphism by local interactions.

In 1D, the pseudo-localization algorithm is a continuous-time PDE system in a new variable or pseudo-coordinate X which plays the role of an "approximate x coordinate" that agents can use to know where they are. The input to this system is the current density value ρ, see Figure 4.1 for an illustration, and the objective is that X converges to a ρ-dependent diffeomorphism. On the other hand, the variable X and the function p^* are used to define the control input of another PDE system in the density ρ. In this way, we have a feedback interconnection of two systems, one in X and one in ρ, with the goal to achieve $X \to \Theta^*$ (the pseudo-coordinate X converges to a true coordinate given by Θ^*) and $\rho \to \rho^*$.

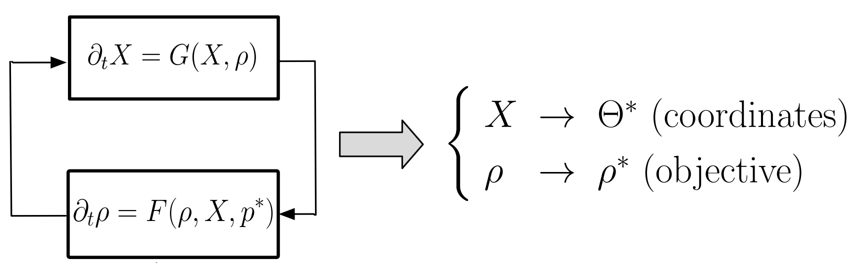

Figure 4.1. Feedback interconnection of pseudo-localization system in X and system in ρ in the 1D case. The function p^* is an equivalent density objective provided to agents in terms of a diffeomorphism Θ^*. The variables X play the role of coordinates and eventually converge to the true coordinates given by Θ^*. Agents use p^* and X to compute the control in the equation ρ. In turn, agents move and this will require a re-computation of coordinates or update in X. The control strategy in the 2D case (stages 2 and 3) can be interpreted similarly.

As for the control design methodology, we follow a constructive, Lyapunov-based approach to designing distributed control laws for the swarm dynamics modeled by PDEs. For this, we define appropriate non-negative energy functionals that encode the objective and choose control laws that keep the time derivative of the energy functional non-positive. This, along with well-known results on the precompactness of solutions as in Lemma 5, the Rellich Kondrachov compactness theorem, allows us to apply the LaSalle Invariance Principle in Lemma 6 and other technical arguments to establish the convergence results that we seek.

In the 1D case, we can identify a set of diffeomorphisms Θ associated with any ρ that eventually converge to Θ^*, and simultaneously control boundary agents into a desired final domain (the support of ρ^*). These are given by the cumulative distribution function associated with the density function; see Section 4.3.1. The 2D case is more complex, and analogous results could not be derived in their full generality. Unlike the 1D case, estimating the cumulative distribution is not straightforward in the 2D case. Instead, we set out to find diffeomorphisms as the result of a distributed algorithm. Given that the discretization of heat flow naturally leads to distributed algorithms, we investigate under what conditions this is the case via harmonic map theory. On the control side, there also are additional difficulties, and because of this, we simplify the control strategy into three stages. In the first stage, the boundary agents are re-positioned onto the boundary of the desired domain while containing the others in the interior. Once this is achieved, the second and third stages can be seen again as the interconnection of two systems in pseudo-coordinates $R = (X, Y)$ (instead of X) and ρ, analogously to Figure 4.1. However, we apply a two time-scale separation for analysis by which coordinates are computed in a fast-time scale and reconfiguration is done in a slow-time scale, which allows for a sequential analysis of the two stages. We then study the robustness of this approach.

4.3 Self-organization in one dimension

In this section, we present our proposed pseudo-localization algorithm and the distributed control law for the 1D self-organization problem.

For each $t \in \mathbb{R}_{\geq 0}$, let $M(t) = (0, L(t)) \subset \mathbb{R}$ be the interval (with boundary $\{0, L(t)\}$) in which the agents are distributed in 1D, and let $\rho : \mathbb{R} \times \mathbb{R} \to \mathbb{R}_{\geq 0}$ be the normalized density function supported on $\bar{M}(t)$, for all $t \geq 0$ (with $\rho(t,x) > 0$, $\forall x \in \bar{M}(t)$), describing the swarm on that interval. Without loss of generality, we place the origin at the leftmost agent of the swarm. We also assume that the leftmost and the rightmost agents, l and r, are aware that they are at the boundary. Let $\rho^* : \bar{M}^* = [0, L^*] \to \mathbb{R}_{>0}$ be the desired normalized density function.

Since a direct feedback control law can not be implemented by agents because they do not have access to their positions, we introduce an equivalent representation of the density ρ^*, p^*, depending on a particular diffeomorphism Θ^*. First, define $\Theta^* : \bar{M}^* \to [0,1]$ such that $\Theta^*(x) = \int_0^x \rho^*(\bar{x}) d\bar{x}$ and $\Theta^*(L^*) = 1$.

Now, let $p^* : [0,1] \to \mathbb{R}_{>0}$, and $\theta^* \in \Theta^*(\bar{M}^*) = [0,1]$, be such that:

$$p^*(\theta^*) = \rho^*((\Theta^*)^{-1}(\theta^*)) = \rho^*(x).$$

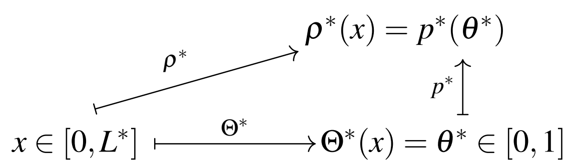

The function p^*, which represents the desired density function mapped onto the unit interval $[0,1]$, is computed offline and is broadcasted to the agents prior to the beginning of the self-organization process. We use p^* to derive the distributed control law which the agents implement. We assume that p^* is a Lipschitz function in the sequel.

Assumption 9 (Uniform boundedness of density function)**.** *We assume that the density function and its derivative are uniformly bounded in its support, that is, for $\rho(t,\cdot)$ and $\partial_x \rho(t,\cdot)$ there exist uniform lower bounds d_l, D_l and uniform upper bounds d_u, D_u (where $0 < d_l \leq d_u < \infty$ and $0 < D_l \leq D_u < \infty$) (that is, $d_l \leq \rho(t,x) \leq d_u$ for all $t \in \mathbb{R}_{\geq 0}$ and $x \in [0, L(t)]$ and $D_l \leq \partial_x \rho(t,x) \leq D_u$ for all $t \in \mathbb{R}_{\geq 0}$ and $x \in (0, L(t))$).*

4.3.1 Pseudo-localization algorithm in one dimension

We first consider the static case, that is, the design of the pseudo-localization dynamics on X of the upper block in Figure 4.1, when the agents and ρ are stationary. We define $\Theta : \bar{M} = [0,L] \to [0,1]$ as:

$$\Theta(x) = \int_0^x \rho(\bar{x}) d\bar{x}, \tag{4.1}$$

such that $\Theta(L) = 1$. In other words, Θ is the cumulative distribution function (CDF) associated with ρ. (Note that the domains are static and hence the argument t has been dropped, which will be reintroduced later.)

Lemma 38 (The CDF diffeomorphism)**.** *Given* $\rho : \bar{M} \to \mathbb{R}_{>0}$*, a* C^1 *function, the mapping* $\Theta : \bar{M} \to [0,1]$ *as defined above, is a diffeomorphism and* $\Theta(\bar{M}) = [0,1]$.

Proof. Since $\rho(x) > 0$, $\forall x \in \bar{M}$, it follows that Θ is a strictly increasing function of x, and is therefore a one-to-one correspondence on $\bar{M}$. Moreover, Θ is atleast C^1 and has a differentiable inverse, which implies it is a diffeomorphism. Finally, since $\Theta(L) = 1$, we have $\Theta(\bar{M}) = [0,1]$. □

Our goal here is to set up a partial differential equation with appropriate boundary conditions that yield the diffeomorphism Θ as its asymptotically stable steady-state solution. We begin by setting up the pseudo-localization dynamics for a stationary swarm (for which the spatial domain M and the density function ρ are fixed). Let $X : \mathbb{R} \times \bar{M} \to \mathbb{R}$ be such that $(t,x) \mapsto X(t,x) \in \mathbb{R}$, with:

$$\begin{aligned}
&\partial_t X = \frac{1}{\rho} \partial_x \left(\frac{\partial_x X}{\rho} \right), \\
&X(t,0) = \alpha(t), \quad X(t,L) = \beta(t), \quad X(0,x) = X_0(x), \\
&\dot{\alpha}(t) = -\alpha(t), \quad \dot{\beta}(t) = 1 - \beta(t),
\end{aligned} \tag{4.2}$$

where $\alpha : \mathbb{R} \to \mathbb{R}$ is a control input at the boundary $x = 0$ and $\beta : \mathbb{R} \to \mathbb{R}$ is a control input at the boundary $x = L$. From (4.1), we observe that $\partial_x \left(\frac{\partial_x \Theta}{\rho} \right) = 0$. Letting $w = X - \Theta$ denote the error, we obtain:

$$
\begin{aligned}
&\partial_t w = \frac{1}{\rho} \partial_x \left(\frac{\partial_x w}{\rho} \right), \\
&\frac{d}{dt} w(t,0) = -w(t,0), \quad \frac{d}{dt} w(t,L) = -w(t,L), \quad w(0,x) = X_0(x) - \Theta(x).
\end{aligned}
\tag{4.3}
$$

Assumption 10 (Well-posedness of the pseudo-localization dynamics)**.** *We assume that the pseudo-localization dynamics* (4.2) *(and* (4.3)*) is well-posed, that the solution is sufficiently smooth (at least $\mathscr{C}^2$ in the spatial variable, even as $t \to \infty$) and belongs to $H^1(M)$.*

Lemma 39 (Pointwise convergence to diffeomorphism)**.** *Under Assumption 10, on the well-posedness of the pseudo-localization dynamics, and Assumption 9 on the boundedness of ρ, the solutions to PDE* (4.2) *converge pointwise to the CDF diffeomorphism Θ defined in* (4.1)*, as $t \to \infty$, for all C^2 initial conditions X_0.*

In this case, the swarm is stationary, which implies that the distribution ρ is fixed (and so is its support $\bar{M}$), and the uniform boundedness assumption 9 simply becomes a boundedness assumption.

Proof. We prove that the solutions to the PDE (4.2) converge pointwise to the diffeomorphism Θ by showing that $w \to 0$, as $t \to \infty$, pointwise for (4.3). For this, we consider a functional V, given by (integrations are with respect to the Lebesgue measure):

$$
V = \frac{1}{2} \int_M \rho |w|^2 + \frac{1}{2} \int_M \frac{1}{\rho} |\partial_x w|^2.
$$

The time derivative $\dot{V}$ is given by:

$$
\dot{V} = \int_M \rho w (\partial_t w) + \int_M \frac{1}{\rho} (\partial_x w)(\partial_t \partial_x w).
$$

Here, replace $\partial_t w$ in the first integral with the dynamics in (4.3), and then use $\partial_t\partial_x = \partial_x\partial_t$ in the second integral together with the Divergence Theorem in Lemma 1. We obtain:

$$\begin{aligned}\dot{V} &= \int_M w\partial_x\left(\frac{\partial_x w}{\rho}\right) - \int_M \partial_x\left(\frac{\partial_x w}{\rho}\right)\partial_t w + \frac{\partial_x w}{\rho}\partial_t w\bigg|_L - \frac{\partial_x w}{\rho}\partial_t w\bigg|_0 \\ &= -\int_M \frac{1}{\rho}\,|\partial_x w|^2 - \int_M \frac{1}{\rho}\left|\partial_x\left(\frac{\partial_x w}{\rho}\right)\right|^2 + \frac{w+\partial_t w}{\rho}\partial_x w\bigg|_L - \frac{w+\partial_t w}{\rho}\partial_x w\bigg|_0.\end{aligned}$$

(After the second equal sign, apply again the Divergence Theorem on the first integral of the previous line, and replace $\partial_t w$ from (4.3).) Substituting from (4.3), we have:

$$\dot{V} = -\int_M \frac{1}{\rho}\,|\partial_x w|^2 - \int_M \frac{1}{\rho}\left|\partial_x\left(\frac{\partial_x w}{\rho}\right)\right|^2.$$

Clearly, $\dot{V} \le 0$, and $w(t,\cdot) \in H^1(M)$, for all t. Moreover, since $V(t) \le V(0)$ and since ρ is uniformly bounded according to Assumption 9, we have that $w(t,\cdot)$ is bounded in $H^1(M)$. Moreover, by the Rellich-Kondrachov Theorem of Lemma 5, $H^1(M)$ is compactly contained in $L^2(M)$. Then it follows that the solutions $w(t,\cdot)$ are precompact. Thus, by the LaSalle Invariance Principle of Lemma 6, the solution to (4.3) converges in L^2-norm to the largest invariant subset of $\dot{V}^{-1}(0)$. Note that $\dot{V} = 0$ implies $\int_M \frac{1}{\rho}|\partial_x w|^2 = 0$. Thus, $\lim_{t\to\infty}\int_M \frac{1}{\rho}|\partial_x w|^2 = 0$. Since ρ is bounded ($\sup\rho < \infty$), we have $\lim_{t\to\infty}\frac{1}{\sup\rho}\int_M |\partial_x w|^2 \le \lim_{t\to\infty}\int_M \frac{1}{\rho}|\partial_x w|^2 = 0$, which implies $\lim_{t\to\infty}\int_M |\partial_x w|^2 = \lim_{t\to\infty}\|\partial_x w\|^2_{L^2(M)} = 0$. Now, $\lim_{t\to\infty}|w(t,x)| = \lim_{t\to\infty}|w(t,0) + \int_0^x \partial_x w(t,\cdot)| \le \lim_{t\to\infty}|w(t,0)| + \int_0^x|\partial_x w(t,\cdot)| \le \lim_{t\to\infty}|w(t,0)| + \sqrt{L(t)}\|\partial_x w(t,\cdot)\|_{L^2(M)} = 0$ (since $\lim_{t\to\infty} w(t,0) = 0$ and $\lim_{t\to\infty}\|\partial_x w(t,\cdot)\|_{L^2(M)} = 0$). Thus, $\lim_{t\to\infty} w(t,x) = 0$, for all $x \in M$. Therefore, the solutions to (4.3) converge to $w \equiv 0$ pointwise, as $t \to \infty$, from any smooth initial $w_0 = X_0 - \Theta$. □

We now have that the solution to the pseudo-localization dynamics converges to the diffeomorphism Θ in the stationary case. For the dynamic case, we modify (4.2) to account for agent motion. Let $X : \mathbb{R}\times\mathbb{R} \to \mathbb{R}$ be supported on $\bar{M}(t) = [0, L(t)]$ for all $t \ge 0$. Using the relation $\frac{dX}{dt} = \partial_t X + v\partial_x X$, where v is the velocity field on the spatial domain, we consider:

$$
\begin{aligned}
&\partial_t X = \frac{1}{\rho}\partial_x\left(\frac{\partial_x X}{\rho}\right) - v\partial_x X, \\
&X(t,0) = 0, \quad X(t,L(t)) = \beta(t), \quad X(0,x) = X_0(x).
\end{aligned} \tag{4.4}
$$

In the dynamic case, and w.l.o.g. we have set $\alpha(t) = 0$ for all $t \geq 0$, for simplicity. We will use the above PDE system in the design of the distributed motion control law, redesigning the boundary control β to achieve convergence of the entire system. We now discretize (4.4) to obtain a distributed pseudo-localization algorithm. Let $X_i(t) = X(t,x_i)$, where $x_i \in \bar{M}(t)$ is the position of the i^{th} agent. We identify the agent i with its desired coordinate in the unit interval at time t, i.e., $\Theta(t,x) = \theta \in [0,1]$, where $\Theta(t,x) = \int_0^x \rho(t,\bar{x})d\bar{x}$ from (4.1), which now shows the time dependency of ρ. In this way, $\rho(t,x) = \partial_x\Theta(t,x)$. It follows that $\partial_x(\cdot) = \partial_\theta(\cdot)\partial_x\theta = \partial_\theta(\cdot)\rho$. Therefore, $\frac{1}{\rho}\partial_x(\cdot) = \partial_\theta(\cdot)$. From (4.4), we have:

$$
\frac{dX}{dt} = \partial_t X + v\partial_x X = \frac{1}{\rho}\partial_x\left(\frac{\partial_x X}{\rho}\right) = \partial_\theta\left(\partial_\theta X\right) = \frac{\partial^2 X}{\partial\theta^2}. \tag{4.5}
$$

Now, we discretize (4.5) with the consistent finite differences $\frac{dX}{dt} \approx \frac{X_i(t+1)-X_i(t)}{\Delta t}$ and $\frac{\partial^2 X}{\partial\theta^2} \approx \frac{X_{i+1}-2X_i+X_{i-1}}{(\Delta\theta)^2}$ (that is, we have that $\lim_{\Delta t\to 0}\frac{X_i(t+1)-X_i(t)}{\Delta t} = \frac{dX}{dt}$ and that $\lim_{\Delta\theta\to 0}\frac{X_{i+1}-2X_i+X_{i-1}}{(\Delta\theta)^2} = \frac{\partial^2 X}{\partial\theta^2}$). Now, with the choice $3\Delta t = (\Delta\theta)^2$, and from (4.4), we obtain for $i \in \mathscr{S}\setminus\{l,r\}$:

$$
\begin{aligned}
&X_i(t+1) = \frac{1}{3}\left(X_{i-1}(t) + X_i(t) + X_{i+1}(t)\right), \\
&X_l(t) = 0, \quad X_r(t) = \beta(t), \quad X_i(0) = X_{0i}.
\end{aligned} \tag{4.6}
$$

Equation (4.6) is the discrete pseudo-localization algorithm to be implemented synchronously by the agents in the swarm, starting from any initial condition X_0. The leftmost agent holds its value at zero while the rightmost agent implements the boundary control β. In the following section we analyze its behavior together with that of the dynamics on ρ.

4.3.2 Distributed density control law and analysis

In this subsection, we propose a distributed feedback control law to achieve $\rho \to \rho^*$ and $w \to 0$, as $t \to \infty$, through a distributed control input v and a boundary control β. We refer the reader to [81] for an overview of Lyapunov-based methods for stability analysis of PDE systems.

From (2.14) and (4.4), we have the dynamics:

$$\begin{aligned} &\partial_t \rho = -\partial_x(\rho v), \\ &\partial_t X = \frac{1}{\rho}\partial_x\left(\frac{\partial_x X}{\rho}\right) - v\partial_x X, \\ &X(t,0) = 0, \quad X(t,L(t)) = \beta(t), \quad X(0,x) = X_0(x). \end{aligned} \tag{4.7}$$

This realizes the feedback interconnection of Figure 4.1.

Assumption 11 (Well-posedness of the full PDE system)**.** *We assume that* (4.7) *is well posed, and that the solutions* $(\rho(t,\cdot),X(t,\cdot))$ *are sufficiently smooth (both in t and* $x \in [0,L(t)]$*), satisfy Assumption 9 on the uniform boundedness of* ρ *and* $\partial_x\rho$*, and are bounded in the Sobolev space* $H^1((0,1/d_l))$.

We also assume that the agent at position x at time t is able to measure $\rho(t,x)$. However, the agents in the swarm do not have access to their positions, and therefore cannot access $\rho^*(x)$, which could be used to construct a feedback law. To circumvent this problem, we propose a scheme in which the agents use the position identifier or pseudo-localization variable X to compute $p^* \circ X(t,x)$, using this as their dynamic set-point. The idea is to then design a distributed control law and a boundary control law such that $\rho \to p^* \circ X$ and $X \to \Theta^*$, as $t \to \infty$, to obtain $\rho \to p^* \circ \Theta^* = \rho^*$. Recall that the function p^* is computed offline and is broadcasted to the agents prior to the beginning of the self-organization process, and that p^* is assumed to be a Lipschitz function. Consider the distributed control law, defined as follows for all time t:

$$v(t,0) = 0, \quad \partial_x v = (\rho - p^* \circ X) - \frac{\partial_X p^*}{\rho(\rho + p^* \circ X)}\partial_x\left(\frac{\partial_x X}{\rho}\right), \tag{4.8}$$

together with the boundary control law:

$$X(t,0)=0, \quad \beta_t = k\left(2-\beta(t)-\frac{X_x}{\rho}\Big|_{L(t)}\right). \tag{4.9}$$

We remark again that the agents implementing the control laws (4.8) and (4.9) do not require position information, because for the agent at position x at time t, $\rho(t,x)$ is a measurement, $X(t,x)$ is the pseudo-localization variable, through which $p^*\circ X(t,x)$ can be computed.

Theorem 10 (Convergence of solutions)**.** *Under the well-posedness Assumption 11, the solutions* $(\rho(t,\cdot),X(t,\cdot))$ *to* (4.7)*, under the control laws* (4.8) *and* (4.9)*, converge to* (ρ^*,Θ^*)*,* $\rho\to\rho^*$ *in* L^2-*norm and* $X\to\Theta^*$ *pointwise as* $t\to\infty$*, from any smooth initial condition* (ρ_0, X_0).

Proof. Consider the candidate control Lyapunov functional V:

$$V=\frac{1}{2}\int_0^{L(t)}|\rho-p^*\circ X|^2dx+\frac{1}{2}\int_0^{L(t)}\rho|w|^2dx+\frac{1}{2}|w(L(t))|^2.$$

Taking the time derivative of V along the dynamics (4.7), using Lemma 2 on the Leibniz integral rule, and applying Corollary 1 on the derivative of energy functionals, we obtain:

$$\begin{aligned}\dot V=&\int_0^{L(t)}(\rho-p^*\circ X)\left(\frac{d\rho}{dt}-\frac{d(p^*\circ X)}{dt}\right)dx+\frac{1}{2}\int_0^{L(t)}|\rho-p^*\circ X|^2\partial_x v\,dx\\&+\int_0^{L(t)}\rho w\partial_t w\,dx+\frac{1}{2}\int_0^{L(t)}(\partial_t\rho)|w|^2\,dx+\frac{1}{2}\rho|w|^2v\Big|_0^{L(t)}+w(L)\frac{dw(L(t))}{dt}.\end{aligned}$$

Now, $\frac{d\rho}{dt}=\partial_t\rho+v\partial_x\rho=-\rho\partial_x v$ (since $\partial_t\rho=-\partial_x(\rho v)$, from (4.7)), and $\partial_t w=\frac{1}{\rho}\partial_x\left(\frac{\partial_x w}{\rho}\right)-$

$v\partial_x w$. Thus, we obtain:

$$\begin{aligned}\dot{V} =& \int_0^{L(t)} (\rho - p^* \circ X)\left[-\rho\partial_x v - \partial_X p^* \frac{1}{\rho}\partial_x\left(\frac{\partial_x X}{\rho}\right)\right] dx + \frac{1}{2}\int_0^{L(t)} |\rho - p^* \circ X|^2 \partial_x v \, dx \\ &+ \int_0^{L(t)} w\partial_x\left(\frac{\partial_x w}{\rho}\right) dx - \int_0^{L(t)} \rho v w \partial_x w \, dx - \frac{1}{2}\int_0^{L(t)} \partial_x(\rho v)|w|^2 \, dx + \frac{1}{2}\rho|w|^2 v\Big|_0^{L(t)} \\ &+ w(L)\frac{dw(L(t))}{dt}.\end{aligned}$$

Now, using the above equation, applying the Divergence theorem (1.1) (integration by parts) and rearranging the terms, we obtain:

$$\begin{aligned}\dot{V} =& -\frac{1}{2}\int_0^{L(t)} (\rho - p^* \circ X)\left[(\rho + p^* \circ X)(\partial_x v) + \frac{\partial_X p^*}{\rho}\partial_x\left(\frac{\partial_x X}{\rho}\right)\right] dx \\ &+ \frac{w\partial_x w}{\rho}\Big|_0^{L(t)} - \int_0^{L(t)} \frac{|\partial_x w|^2}{\rho} \, dx - \int_0^{L(t)} \rho v w \partial_x w \, dx - \frac{1}{2}\rho v |w|^2\Big|_0^{L(t)} \\ &+ \int_0^{L(t)} \rho v w \partial_x w \, dx + \frac{1}{2}\rho|w|^2 v\Big|_0^{L(t)} + w(L)\frac{dw(L(t))}{dt}.\end{aligned}$$

Since $w(0) = 0$, the above equation reduces to:

$$\begin{aligned}\dot{V} =& -\frac{1}{2}\int_0^{L(t)} (\rho - p^* \circ X)\left[(\rho + p^* \circ X)(\partial_x v) + \frac{\partial_X p^*}{\rho}\partial_x\left(\frac{\partial_x X}{\rho}\right)\right] dx \\ &- \int_0^{L(t)} \frac{|\partial_x w|^2}{\rho} \, dx + w(L(t))\left(\frac{d}{dt}w(L(t)) + \frac{\partial_x w}{\rho}\right).\end{aligned}$$

From (4.8) and (4.9), we have $\partial_x v = (\rho - p^* \circ X) - \frac{\partial_X p^*}{\rho(\rho + p^* \circ X)}\partial_x\left(\frac{\partial_x X}{\rho}\right)$, and

$$\frac{dw}{dt}\Big|_{L(t)} = -\left(\frac{\partial_x w}{\rho} + kw\right)\Big|_{L(t)},$$

and we obtain:

$$\dot{V} = -\frac{1}{2}\int_0^{L(t)} (\rho + p^* \circ X)|\rho - p^* \circ X|^2 dx - \int_0^{L(t)} \frac{|\partial_x w|^2}{\rho} \, dx - k|w(L(t))|^2. \tag{4.10}$$

Clearly, $\dot{V} \leq 0$, and $\rho(t,\cdot), w(t,.) \in H^1((0,1/d_l))$, for all t. By Lemma 5, the Rellich-Kondrachov Compactness Theorem, the space $H^1((0,1/d_l))$ is compactly contained in $L^2((0,1/d_l))$, and the bounded solutions (by Assumption 11) in $H^1((0,1/d_l))$ are then precompact in $L^2((0,1/d_l))$. Moreover, the set of (ρ, X) satisfying Assumption 11 is dense in $L^2((0,1/d_l))$. Then, by the LaSalle Invariance Principle, Lemma 6, we have that the solutions to (4.7) converge in the L^2-norm to the largest invariant subset of $\dot{V}^{-1}(0)$. This implies that:

$$\lim_{t\to\infty} \|\rho(t,\cdot) - p^* \circ X(t,\cdot)\|_{L^2((0,L(t)))} = 0,$$
$$\lim_{t\to\infty} \|\frac{\partial_x w}{\rho}\|_{L^2((0,L(t)),\rho)} = 0, \quad \lim_{t\to\infty} w(t,L(t)) = 0.$$

Thus, we have:

$$\lim_{t\to\infty} \left\|\frac{\partial_x w}{\rho}\right\|_{L^2((0,L(t)),\rho)} = 0 \quad \Rightarrow \lim_{t\to\infty} \|\partial_x w\|_{L^2((0,L(t)))} = 0.$$

Using the Poincaré-Wirtinger inequality, Lemma 4, again, we note that this implies $\lim_{t\to\infty} \|w - \int_0^{L(t)} w\|_{L^2((0,L(t)))} = 0$. We have $\lim_{t\to\infty} |\int_0^{L(t)} w| = |\int_0^{L(t)} \int_0^x \partial_x w| \leq L(t)^{3/2} \|\partial_x w\|_{L^2((0,L(t)))} = 0$, which implies that $\lim_{t\to\infty} \int_0^{L(t)} w = 0$ and therefore $\lim_{t\to\infty} \|w\|_{L^2((0,L(t)))} = 0$. Thus, we get $\lim_{t\to\infty} \|w(t,\cdot)\|_{H^1((0,L(t)))} = 0$, or in other words, $w \to_{H^1} 0$. Now, $\lim_{t\to\infty} |w(t,x)| = \lim_{t\to\infty} |w(t,0) + \int_0^x \partial_x w(t,\cdot)| \leq \lim_{t\to\infty} |w(t,0)| + \int_0^x |\partial_x w(t,\cdot)| \leq \lim_{t\to\infty} |w(t,0)| + \sqrt{L(t)} \|w(t,\cdot)\|_{H^1((0,L(t)))} = 0$, which implies that $w \to 0$ pointwise. Given that $w = X - \Theta$, we have $\lim_{t\to\infty} X(t,\cdot) - \Theta(t,\cdot) = 0$. Let $\lim_{t\to\infty} L(t) = L$ and $\lim_{t\to\infty} \Theta(t,\cdot) = \bar{\Theta}(\cdot)$, which implies that $X \to \bar{\Theta}$ pointwise.

From the above, we have $\lim_{t\to\infty} \|\rho(t,\cdot) - p^* \circ \bar{\Theta}\|_{L^2((0,L(t)))} = \lim_{t\to\infty} \|\rho(t,\cdot) - p^* \circ X(t,\cdot) + p^* \circ X(t,\cdot) - p^* \circ \bar{\Theta}\|_{L^2((0,L(t)))} \leq \lim_{t\to\infty} \|\rho(t,\cdot) - p^* \circ X(t,\cdot)\|_{L^2((0,L(t)))} + \|p^* \circ X(t,\cdot) - p^* \circ \bar{\Theta}\|_{L^2((0,L(t)))} = 0$ (this follows from the assumption that p^* is Lipschitz, since $\|p^* \circ X - p^* \circ \bar{\Theta}\|_{L^2} \leq c\|X - \bar{\Theta}\|_{L^2}$ for some Lipschitz constant c). Thus, we have $\rho \to_{L^2} p^* \circ \bar{\Theta}$.

Now, we are interested in the limit density function $\bar{\rho} = p^* \circ \bar{\Theta}$, and by the definition of $\bar{\Theta}$

we have $\bar{\Theta}(x) = \int_0^x \bar{\rho}$. We now prove that this limit $(\bar{\rho}, \bar{\Theta})$ is unique, and that $(\bar{\rho}, \bar{\Theta}) = (\rho^*, \Theta^*)$. From the definition of $\bar{\Theta}$, we get $\frac{d\bar{\Theta}}{dx}(x) = \bar{\rho}(x) = p^*(\bar{\Theta}(x)) > 0$, $\forall \bar{\Theta}(x) \in [0,1]$. We therefore have:

$$x = \int_0^{\bar{\Theta}(x)} (p^*(\theta))^{-1} \, d\theta.$$

Recall from the definition of p^* and (4.1) that $p^* \circ \Theta^*(x) = \rho^*(x)$, and $\frac{d}{dx}\Theta^*(x) = \rho^*(x) = p^* \circ \Theta^*(x)$, which implies that $\frac{d\Theta^*}{dx} = p^*(\theta^*) > 0$, where $\theta^* = \Theta^*(x)$. Therefore:

$$x = \int_0^{\Theta^*(x)} (p^*(\theta))^{-1} \, d\theta.$$

From the above two equations, we get:

$$\int_0^{\bar{\Theta}(x)} (p^*(\theta))^{-1} \, d\theta = \int_0^{\Theta^*(x)} (p^*(\theta))^{-1} \, d\theta,$$

for all x, and since p^* is strictly positive, it implies that $\bar{\Theta} = \Theta^*$, and we obtain $\bar{\rho} = p^* \circ \bar{\Theta} = p^* \circ \Theta^* = \rho^*$. And we know that $\rho \to_{L^2} p^* \circ \bar{\Theta} = p^* \circ \Theta^* = \rho^*$. In other words, ρ converges to ρ^* in the L^2 norm. □

Physical interpretation of the density control law

For a physical interpretation of the control law, we first rewrite some of the terms in a suitable form. From (4.7), we know that:

$$\frac{1}{\rho}\partial_x\left(\frac{\partial_x X}{\rho}\right) = \frac{\partial X}{\partial t} + v\partial_x X = \frac{dX}{dt}.$$

The second term in the expression for $\partial_x v$ in the law (4.8) can thus be rewritten as:

$$\frac{\partial_X p^*}{\rho(\rho + p^* \circ X)}\partial_x\left(\frac{\partial_x X}{\rho}\right) = \frac{1}{(\rho + p^* \circ X)}\,\partial_X p^* \frac{dX}{dt} = \frac{1}{(\rho + p^* \circ X)}\frac{dp^*}{dt}.$$

Now, from above and (4.8), we obtain:

$$v(t,x) = \int_0^x (\rho - p^* \circ X) - \int_0^x \frac{1}{(\rho + p^* \circ X)} \frac{dp^*}{dt}. \tag{4.11}$$

Equation (4.11) gives the velocity of the agent at x at time t. Now, to interpret it, we first consider the case where the pseudo-localization error is zero, that is, when $X = \Theta^*$. This would imply that $p^* \circ X = p^* \circ \Theta^* = \rho^*$, $\frac{dX}{dt} = \frac{d\Theta^*}{dt} = 0$, and we obtain:

$$v(t,x) = \int_0^x (\rho - \rho^*). \tag{4.12}$$

The term $\int_0^x (\rho - \rho^*) = \int_0^x \rho - \int_0^x \rho^*$ is the difference between the number of agents in the interval $[0,x]$ and the desired number of agents in $[0,x]$. If the term is positive, it implies that there are more than the desired number of agents in $[0,x]$ and the control law essentially exerts a pressure on the agent to move right thereby trying to reduce the concentration of agents in the interval $[0,x]$, and, vice versa, when the term is negative. This eventually accomplishes the desired distribution of agents over a given interval. This would be the physical interpretation of the control law for the case where the pseudo-localization error is zero (that is, the agents have full information of their positions).

However, in the transient case when the agents do not possess full information of their positions and are implementing the pseudo-localization algorithm for that purpose, the control law requires a correction term that accounts for the fact that the transient pseudo coordinates $X(t,x)$ cannot be completely relied upon. This is what the second term $\int_0^x \frac{1}{(\rho+p^* \circ X)} \frac{dp^*}{dt}$ in (4.11) corrects for. When this term is positive, that is, $\int_0^x \frac{1}{(\rho+p^* \circ X)} \frac{dp^*}{dt} > 0$, it roughly implies that the "estimate" of the desired number of agents in the interval $[0,x]$ is increasing (indicating that an increase in the concentration of agents in $[0,x]$ is desirable), and the term essentially reduces the "rightward pressure" on the agent (note that this term will have a negative contribution to the velocity (4.11)).

4.3.3 Discrete implementation

In this section, we present a scheme to compute p^* (the transformed desired density profile) and a consistent discretization scheme for the distributed control law. We follow that up with a discussion on the convergence of the discretized system and a pseudo-code for the implementation.

On the computation of p^*

We now provide a method for computing p^* from a given ρ^* via interpolation. Let the desired domain $M^* = [0, L^*]$ be discretized uniformly to obtain $M_d^* = \{0 = x_1, \ldots, x_m = L^*\}$ such that $x_j - x_{j-1} = h$ (constant step-size). Note that m is the number of interpolation points, not equal to the number of agents. The desired density $\rho^* : [0, L^*] \to \mathbb{R}_{>0}$ is known, and we compute the value of ρ^* on M_d^* to get $\rho^*(x_1, \ldots, x_m) = (\rho_1^*, \ldots, \rho_m^*)$. We also have $\Theta^*(x) = \int_0^x \rho^* d\mu$, for all $x \in [0, L^*]$. Now, computing the integral with respect to the Dirac measure for the set M_d^*, we obtain $\Theta_d^*(x_1, \ldots, x_m) = (\theta_1^*, \ldots, \theta_m^*)$, where $\theta_1^* = 0$ and $\theta_k^* = \frac{1}{2}\sum_{j=1}^{k}(\rho_{j-1}^* + \rho_j^*)h$, for $k = 2, \ldots, m$ (note that $0 = \theta_1^* \le \theta_2^* \le \ldots \le \theta_m^* \le 1$ and $\lim_{h\to 0}\theta_m^* = \Theta^*(L^*) = 1$). Now, the value of the function p^* at any $X \in [0,1]$ can be now obtained from the relation $p^*(\theta_k^*) = \rho_k^*$, for $k = 1, \ldots, m$, by an appropriate interpolation.

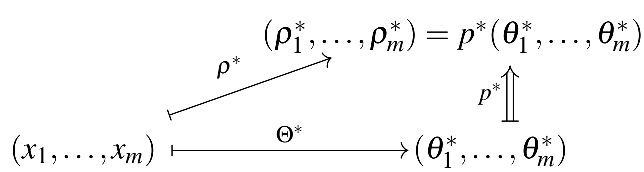

Discrete control law

A discretized pseudo-localization algorithm is given by (4.6). We now discretize (4.8) to obtain an implementable control law for a finite number of agents $i \in \mathscr{S}$, and a numerical simulation of this law is later presented in Section 4.5.

Let $i \in \mathscr{S} \setminus \{l, r\}$. First note that $\partial_x v = (\partial_\theta v)\Big|_{\theta=\Theta(x)}(\partial_x \Theta) = (\partial_\theta v)\Big|_{\theta=\Theta(x)}\rho$ (where $v \equiv v(\Theta(x))$). Using a consistent backward differencing approximation, and recalling that $\Delta\theta = \varepsilon$, we can write:

$$(\partial_x v)_i \approx \rho_i \frac{v_i - v_{i-1}}{\Delta\theta} = \rho_i \frac{v_i - v_{i-1}}{\varepsilon}, \quad i \in \mathscr{S}$$

where ρ_i is agent i's density measurement.

From Section 4.3.1, recall the consistent finite-difference approximation:

$$\frac{1}{\rho}\partial_x\left(\frac{\partial_x X}{\rho}\right)_i \approx \frac{1}{\varepsilon^2}(X_{i-1} - 2X_i + X_{i+1}).$$

With $\kappa = \frac{1}{2\varepsilon}$, from (4.8) and the above equation, we obtain the law for agent i as:

$$v_i = v_{i-1} + \frac{\rho_i - p^*(X_i)}{2\kappa\rho_i} - \frac{2\kappa}{\rho_i(\rho_i + p^*(X_i))} \cdot \frac{p^*(X_{i+1}) - p^*(X_{i-1})}{X_{i+1} - X_{i-1}} \cdot (X_{i-1} - 2X_i + X_{i+1}) \tag{4.13}$$

with $v_l = 0$. The computation in v can be implemented by propagating from the leftmost agent to the rightmost agent along a line graph $\mathscr{G}_{line}$ (with message receipt acknowledgment). Note that this propagation can alternatively be formulated by each agent averaging appropriate variables with left and right neighbors, which will result in a process similar to a finite-time consensus algorithm. Now, the boundary control (4.9) is discretized (with $\partial_t\beta \approx \frac{\beta(t+1)-\beta(t)}{\Delta t}$), with the choice $k = \frac{1}{\varepsilon}$ to:

$$\beta(t+1) = \beta(t) + k\Delta t(2 - \beta(t) - 2\kappa(\beta(t) - X_{r-1}(t))) = \frac{4-2\varepsilon}{3}\beta(t) + \frac{1}{3}X_{r-1}(t) \tag{4.14}$$

On the convergence of the discrete system

The discretized pseudo-localization algorithm (4.6) with the boundary control law (4.9), can be rewritten as:

$$X(t+1) = X(t) - \frac{1}{3}LX(t) + u(t), \tag{4.15}$$

where $X(t) = (X_l(t), \ldots, X_r(t))$, L is the Laplacian of the line graph $\mathscr{G}_{line}$ and the input $u(t) = \left(0, \ldots, 0, \frac{\varepsilon}{3}(2 - \beta(t))\right)$. This discretized system is stable and we thereby have that the discretized pseudo-localization algorithm is consistent and stable. Thus, by the Lax Equivalence Theorem [112], the solution of (4.15) converges to the solution of (4.4) with the boundary control (4.9) as $N \to \infty$. Due to the nonlinear nature of the discrete implementation of the equation in ρ, we are only certain that we have a consistent discrete implementation in this case (no similar convergence theorem exists for discrete approximations of nonlinear PDEs.)

Algorithm 3. Self-organization algorithm for 1D environments

```
Input: ρ*, K (number of iterations), Δt (time step)
Requires:
    Offline computation of p* as outlined in Section 4.3.3
    Initialization X_i(0) = X_{0i}, v_i = 0
    Leftmost and rightmost agents, l, r, resp., are aware they are at boundary
for k := 1 to K do
    if i = l then
        agent l holds onto X_l(k) = 0 and v_l(k) = 0
    else if agent i ∈ {l+1, ..., r−1} then
        agent i receives X_{i−1}(k) and X_{i+1}(k) from its left and right neighbors
        agent i implements the update (4.6)
    else if i = r then
        agent r receives X_{r−1}(k) from its left neighbor
        agent r implements the update (4.14)
    for i := l to r do
        agent i computes velocity v_i from (4.13)
    agent i moves to x_i(k+1) = x_i(k) + v_i(k)Δt
```

4.4 Self-organization in two dimensions

In this section, we present the two-dimensional self-organization problem. Although our approach to the 2D problem is fundamentally similar to the 1D case, we encounter a problem in the two-dimensional case that did not require consideration in one dimension, and it is the need to control the shape of the spatial domain in which the agents are distributed. We overcome this problem by controlling the shape of the domain with the agents on the boundary, while controlling the density function of the agents in the interior.

Let $M : \mathbb{R} \rightrightarrows \mathbb{R}^2$ be a smooth one-parameter family of bounded open subsets of $\mathbb{R}^2$, such that $\bar{M}(t)$ is the spatial domain in which the agents are distributed at time $t \geq 0$. Let $\rho : \mathbb{R} \times \mathbb{R}^2 \to \mathbb{R}_{\geq 0}$ be the spatial density function with support $\bar{M}(t)$ for all $t \geq 0$; that is, $\rho(t,x) > 0$, $\forall x \in \bar{M}(t)$, and $t \geq 0$. Without loss of generality, we shift the origin to a point on the boundary of the family of domains, such that $(0,0) \in \partial M(t)$, for all t. Let $\rho^* : M^* \to \mathbb{R}_{>0}$ be the desired density function, where M^* is the target spatial domain. From here on, we view $\bar{M}$ as a one-parameter family of compact 2-submanifolds with boundary of $\mathbb{R}^2$. Just as in the 1D case, the agents do no have access to their positions but know the true x- and y-directions.

In what follows we present our strategy to solve this problem, which we divide into three stages for simplicity of presentation and analysis. In the first stage, the agents converge to the target spatial domain M^* with the boundary agents controlling the shape of the domain. In stage two, the agents implement the pseudo-localization algorithm to compute the coordinate transformation. In the third stage, the boundary agents remain stationary and the agents in the interior converge to the desired density function. This simplification is performed under the assumption that, once the agents have localized themselves at a given time, they can accurately update this information by integrating their (noiseless) velocity inputs. Noisy measurements would require that these phases are rerun with some frequency; e.g. using fast and slow time scales as described in Section 4.2.

4.4.1 Pseudo-localization algorithm for boundary agents

To begin with, we propose a pseudo-localization algorithm for the boundary agents which allows for their control in the first stage. To do this, we assume that the agents have a boundary detection capability (can approximate the normal to the boundary), the ability to communicate with neighbors immediately on either side along the boundary curve, and can measure the density of boundary agents.

Let $M_0 \subset \mathbb{R}^2$ be a compact 2-manifold with boundary ∂M_0 and let $(0,0) \in \partial M_0$. To localize themselves, the agents on ∂M_0 implement the distributed 1D pseudo-localization algorithm presented in Section 4.3.1. This yields a parametrization of the boundary $\Gamma : \partial M_0 \to [0,1)$, with $\Gamma(0,0) = 0$, such that the closed curve which is the boundary ∂M_0 is identified with the interval $[0,1)$. We have that, for $\gamma \in [0,1)$, $\Gamma^{-1}(\gamma) \in \partial M_0$. For $\gamma \in [0,1)$, let $s(\gamma)$ be the arc length of the curve ∂M_0 from the origin, such that $s(0) = 0$ and $\lim_{\gamma \to 1} s(\gamma) = l$. We assume that the boundary agents have access to the unit outward normal $\mathbf{n}(\gamma)$ to the boundary, and thus the unit tangent $\mathbf{s}(\gamma)$.

Let $q : [0,l) \to \mathbb{R}_{>0}$ denote the normalized density of agents on the boundary, such that we have $\int_0^l q(s)ds = 1$. Now the 1D pseudo-localization algorithm of Section 4.3.1 serves to provide a 2D boundary pseudo-localization as follows. Note that $\frac{ds}{d\gamma} = \frac{1}{q(\gamma)}$, and $(dx,dy) = \mathbf{s}ds$, which implies $(dx,dy) = \frac{1}{q(\gamma)}\mathbf{s}(\gamma)d\gamma$. Therefore, we get the position of the boundary agent at γ, $(x(\gamma),y(\gamma))$, as $(x(\gamma),y(\gamma)) = \int_0^\gamma \frac{1}{q(\bar{\gamma})}\mathbf{s}(\bar{\gamma})d\bar{\gamma}$, and the arc-length $s(\gamma) = \int_0^\gamma \frac{1}{q(\bar{\gamma})}d\bar{\gamma}$, which is discretized by a consistent scheme to obtain:

$$(x_i,y_i) = \frac{1}{2}\Delta\gamma \sum_{k=0}^{i-1}\left(\frac{\mathbf{s}_k}{q_k} + \frac{\mathbf{s}_{k+1}}{q_{k+1}}\right), \qquad \text{for } i \in \partial M_0, \tag{4.16}$$

and we recall that the agents have access to q and $\mathbf{s}$. The computation of (x_i,y_i) can be implemented by propagating from the agent with $\gamma_i = 0$ along the boundary agents in the direction as $\gamma_i \to 1$, along a line graph $\mathscr{G}_{\text{line}}$ (with message receipt acknowledgment). Note that this propagation can alternatively be formulated by each agent averaging appropriate variables with left and right neighbors, which will result in a process similar to a finite-time consensus algorithm.

This way, the boundary agents are localized at time $t = 0$, and they update their position estimates using their velocities, for $t \geq 0$.

4.4.2 Pseudo-localization algorithm in two dimensions

In this subsection, we present the pseudo-localization algorithm for the agents in the interior of the spatial domain. We first describe the idea of the coordinate transformation (diffeomorphism) we employ and construct a PDE that converges asymptotically to this diffeomorphism. We then discretize the PDE to obtain the distributed pseudo-localization algorithm.

The main idea is to employ harmonic maps to construct a coordinate transformation or diffeomorphism from the spatial domain of the swarm onto the unit disk. We begin the construction with the static case, where the agents are stationary. Let $M \subseteq \mathbb{R}^2$ be a compact, static 2-manifold with boundary and $N = \{(x,y) \in \mathbb{R}^2 \,|\, (x-1)^2 + y^2 \leq 1\}$ be the unit disk. The manifolds M and N are both equipped with a Euclidean metric $g = h = \delta$.

First, we define a mapping for the boundary of M. Let $\Gamma : \partial M \to [0,1)$ be a parametrization of the boundary of M, as outlined in Section 4.4.1. Let $\xi : \bar{M} \to N$ be any diffeomorphism that takes the following form on the boundary of M:

$$\xi(\Gamma^{-1}(\gamma)) = (1-\cos(2\pi\gamma), \sin(2\pi\gamma)), \qquad \gamma \in [0,1), \tag{4.17}$$

and we know that $\Gamma^{-1}[0,1) = \partial M$.

Now, from Lemma 7, on harmonic diffeomorphisms, there is a unique harmonic diffeomorphism, $\Psi : M \to N$, such that $\Psi = \xi$ on ∂M. We know that, by definition, the mapping $\Psi = (\psi_1, \psi_2)$ satisfies:

$$\begin{cases} \Delta\psi_1 = 0, \\ \Delta\psi_2 = 0, \end{cases} \quad \text{for } \mathbf{r} \in \mathring{M}, \tag{4.18}$$
$$\Psi = \xi, \quad \text{on } \partial M,$$

where Δ is the Laplace operator. Let Ψ^* be the corresponding map from the target domain M^* to the unit disk N. Now, we define a function $p^* : N \to \mathbb{R}_{>0}$ by $p^* = \rho^* \circ (\Psi^*)^{-1}$, the image of the desired spatial density distribution on the unit disk, which is computed offline and is broadcasted to the agents prior to the beginning of the self-organization process. We later use p^* to derive the distributed control law which the agents implement.

$$
\begin{array}{ccc}
 & & \rho^*(\mathbf{r}) = p^*(\Psi^*(\mathbf{r})) \\
 & \overset{\rho^*}{\nearrow} & \uparrow p^* \\
\mathbf{r} \in M^* & \xmapsto{\Psi^*} & \Psi^*(\mathbf{r}) \in N
\end{array}
$$

We now construct a PDE that asymptotically converges to the harmonic diffeomorphism, which we then discretize to obtain a distributed pseudo-localization algorithm. We use the heat flow equation as the basis to define the pseudo-localization algorithm, which yields a harmonic map as its asymptotically stable steady-state solution. We begin by setting up the system for a stationary swarm, for which the spatial domain is fixed.

Let $M \subset \mathbb{R}^2$ be a compact 2-manifold with boundary, N be the unit disk of $\mathbb{R}^2$, and $\mathbf{R} = (X,Y) : M \to N$. The heat flow equation is given by:

$$
\begin{cases} \partial_t X = \Delta X, \\ \partial_t Y = \Delta Y, \end{cases} \quad \text{for } \mathbf{r} \in \mathring{M}, \qquad \mathbf{R} = \xi, \quad \text{on } \partial M. \tag{4.19}
$$

The heat flow equation has been studied extensively in the literature. For well-known existence and uniqueness results, we refer the reader to [48].

Lemma 40. *[Pointwise convergence of the heat flow equation to a harmonic diffeomorphism] The solutions of the heat flow equation* (4.19) *converge pointwise to the harmonic map satisfying* (4.18)*, exponentially as $t \to \infty$, from any smooth initial $\mathbf{R}_0 \in H^1(M) \times H^1(M)$.*

Proof. Let Ψ be the solution to (4.18), which is a harmonic map by definition. Let $\tilde{\mathbf{R}} = \mathbf{R} - \Psi$ be the error where $\mathbf{R} = (X,Y)$ is the solution to (4.19). Subtracting (4.18) from (4.19), we obtain:

$$\begin{cases} \partial_t X = \Delta X, \\ \partial_t Y = \Delta Y, \end{cases} \quad \text{for } \mathbf{r} \in \mathring{M}, \qquad \tilde{\mathbf{R}} = 0, \quad \text{on } \partial M. \tag{4.20}$$

The Laplace operator Δ with the Dirichlet boundary condition in (4.20) is self-adjoint and has an infinite sequence of eigenvalues $0 < \lambda_1 < \lambda_2 < \ldots$, with the corresponding eigenfunctions $\{\phi_i\}_{i=1}^{\infty}$ forming an orthonormal basis of $L^2(M)$ (where $\phi_i \in L^2(M)$ and $\Delta\phi_i = \lambda_i\phi_i$ for all i, with $\phi_i = 0$ on the boundary) [55]. Let the initial condition be $\tilde{X}_0 = \sum_{i=1}^{\infty} a_i\phi_i$ and $\tilde{Y}_0 = \sum_{i=1}^{\infty} b_i\phi_i$ (where a_i and b_i are constants for all i). The solution to (4.20) is then given by $\tilde{X}(t,\mathbf{r}) = \sum_{i=1}^{\infty} a_i e^{-\lambda_i t}\phi_i(\mathbf{r})$ and $\tilde{Y}(t,\mathbf{r}) = \sum_{i=1}^{\infty} b_i e^{-\lambda_i t}\phi_i(\mathbf{r})$. Since $\lambda_i > 0$, for all i, we obtain $\lim_{t\to\infty}\tilde{X}(t,\mathbf{r}) = 0$ and $\lim_{t\to\infty}\tilde{Y}(t,\mathbf{r}) = 0$, for all $\mathbf{r} \in \bar{M}$. Therefore, $\lim_{t\to\infty}\mathbf{R}(t,\mathbf{r}) = \Psi(\mathbf{r})$, for all $\mathbf{r} \in \bar{M}$, and the convergence is exponential. □

We now have a PDE that converges to the diffeomorphism given by (4.18) for the stationary case (agents in the swarm are at rest). For the dynamic case, and to describe the algorithm while the agents are in motion, we modify (4.19) as follows. Let $\mathbf{R} = (X,Y) : \mathbb{R} \times \mathbb{R}^2 \to \mathbb{R}$. We are only interested in the restriction to $M(t)$, $\mathbf{R}|_{M(t)}$, at any time t, so we drop the restriction and just identify $\mathbf{R} \equiv \mathbf{R}_{|_{M(t)}}$. Using the relation $\frac{dX}{dt} = \partial_t X + \nabla X \cdot \mathbf{v}$, where $\mathbf{v}$ is a velocity field, we obtain:

$$\begin{cases} \partial_t X = \Delta X - \nabla X \cdot \mathbf{v}, \\ \partial_t Y = \Delta Y - \nabla Y \cdot \mathbf{v}, \end{cases} \quad \text{for } \mathbf{r} \in \mathring{M}(t), \qquad \mathbf{R} = \xi, \quad \text{on } \partial M(t). \tag{4.21}$$

We now discretize (4.21) to derive the distributed pseudo-localization algorithm. Now, we have $\rho : \mathbb{R} \times \mathbb{R}^2 \to \mathbb{R}_{\geq 0}$ with support $M(t)$, the density function of the swarm on the domain $M(t)$.

We view the swarm as a discrete approximation of the domain $M(t)$ with density ρ, and the PDE (4.21) as approximated by a distributed algorithm implemented by the swarm.

Here, we propose a candidate distributed algorithm, which would yield the heat flow equation via a functional approximation. Our candidate algorithm is a time-varying weighted Laplacian-based distributed algorithm, owing to the connection between the graph Laplacian and the manifold Laplacian [19]:

$$X_i(t+1) = X_i(t) + \sum_{j \in \mathcal{N}_i(t)} w_{ij}(t)(X_j(t) - X_i(t)), \tag{4.22}$$

and a similar equation for Y. We show how to derive next the values for the weights $w_{ij}(t) \in \mathbb{R}$, for all t. First, the set of neighbors, $j \in \mathcal{N}_i(t)$, of i at time t, are the spatial neighbors of i in $M(t)$, that is, $\mathcal{N}_i(t) = \{j \in \mathcal{S} \mid \|\mathbf{r}_j(t) - \mathbf{r}_i(t)\| \leq \varepsilon\} \equiv B_\varepsilon(\mathbf{r}_i(t))$. Using $X_i(t+1) - X_i(t) = \frac{dX}{dt}\delta t$, for a small δt, we make use of a functional approximation of (4.22):

$$\frac{dX}{dt}\delta t = \int_{B_\varepsilon(\mathbf{r}_i(t))} w(t, \mathbf{r}_i, \mathbf{s})(X(t,\mathbf{s}) - X(t,\mathbf{r}_i))\, \rho(t,\mathbf{s}) d\mu, \tag{4.23}$$

where $d\nu = \rho\, d\mu$ is a density-dependent measure on the manifold, and the weighting function w satisfies $w(t, \mathbf{r}_i(t), \mathbf{r}_j(t)) = w_{ij}(t)$, for all $i, j \in \mathcal{S}$. We note that the summation term in (4.22) is a special form of the integral in (4.23) with a Dirac measure $d\nu$ supported on the set $\{\mathbf{r}_1(t), \ldots, \mathbf{r}_N(t)\}$ at time t. Now, with the choice $w(t, \mathbf{r}_i, \mathbf{s}) = \frac{1}{\int_{B_\varepsilon(\mathbf{s}(t))} \rho(t, \bar{\mathbf{s}}) d\mu}$ and for very small ε (making $\mathcal{O}(\varepsilon^3)$ terms negligible), (4.23) reduces to:

$$\frac{dX}{dt}\delta t = a\Delta X,$$

where $a = \frac{1}{4\varepsilon}\int_{B_\varepsilon(\mathbf{r}_i(t))}(\mathbf{s} - \mathbf{r}_i(t)) \cdot (\mathbf{s} - \mathbf{r}_i(t)) d\mu$ is a constant. Now, with the choice $\delta t = a$, we obtain:

$$\frac{dX}{dt} = \frac{\partial X}{\partial t} + \mathbf{v} \cdot \nabla X = \Delta X,$$

which is the PDE (4.21). Let $d(t, \mathbf{r}_i(t)) = \int_{B_\varepsilon(\mathbf{r}_i(t))} \rho(t, \mathbf{s}) d\mu$ and $d_i(t) = |\mathscr{N}_i(t)|$, for $i \in \mathscr{S}$. Substituting $w_{ij}(t) = w(t, \mathbf{r}_i(t), \mathbf{r}_j(t)) = \frac{1}{\int_{B_\varepsilon(\mathbf{r}_j(t))} \rho(t, \bar{\mathbf{s}}) d\mu} = \frac{1}{d(t, \mathbf{r}_j(t))} \approx \frac{1}{d_j(t)}$, in (4.22), we get the distributed pseudo-localization algorithm for the agents in the interior of the swarm to be:

$$\begin{aligned} X_i(t+1) &= X_i(t) + \sum_{j \in \mathscr{N}_i(t)} \frac{1}{d_j(t)} (X_j(t) - X_i(t)), \\ Y_i(t+1) &= Y_i(t) + \sum_{j \in \mathscr{N}_i(t)} \frac{1}{d_j(t)} (Y_j(t) - Y_i(t)). \end{aligned} \tag{4.24}$$

For the agents on the boundary $\partial M(t)$, we have:

$$\mathbf{R}_i = (X_i, Y_i) = \xi_i,$$

where $\xi_i = \xi(\mathbf{r}_i(t))$, for $\mathbf{r}_i(t) \in \partial M(t)$. Note that the discretization scheme is consistent, in that as the number of agents $N \to \infty$, the discrete equation (4.24) converges to the PDE (4.21). In this way, from (4.24), the pseudo-localization algorithm is a Laplacian-based distributed algorithm, with a time-varying weighted graph Laplacian.

4.4.3 Distributed density control law and analysis

In this section, we derive the distributed feedback control law to converge to the desired density function over the target domain in the two-dimensional case. The swarm dynamics are given by:

$$\begin{aligned} \partial_t \rho &= -\nabla \cdot (\rho \mathbf{v}), \quad \text{for } \mathbf{r} \in \mathring{M}(t), \\ \partial_t \mathbf{r} &= \mathbf{v}, \quad \text{on } \partial M(t). \end{aligned} \tag{4.25}$$

Assumption 12 (Well-posedness of the PDE system)**.** *We assume that* (4.25) *is well-posed, and that its solution* $\rho(t,\cdot)$ *is sufficiently smooth and is bounded in the Sobolev space* $H^1(\cup_t M(t))$*, the components of the velocity field* $\mathbf{v}$ *are bounded in the Sobolev space* $H^1(\cup_t M(t))$ *and of the parametrized velocity on the boundary are bounded in the Sobolev space* $H^1((0,1))$.

In what follows, we describe the control strategy based on three different stages.

Stage 1

In this stage, the objective is for the swarm to converge to the target spatial domain M^*.

Let $\mathbf{r}^* : [0,1] \to \partial M^*$ be the closed curve describing the desired boundary. Let $\mathbf{e}(\gamma) = \mathbf{r}(\gamma) - \mathbf{r}^*(\gamma)$ be the position error of agent γ on the boundary, where $\mathbf{r}(\gamma)$ is the actual position of agent γ computed as presented in Section 4.4.1. We define a distributed control law for swarm motion as follows:

$$\begin{cases} \mathbf{v} = -\frac{\nabla\rho}{\rho}, & \text{for } \mathbf{r} \in \mathring{M}(t), \\ \partial_t \mathbf{v} = -\mathbf{e} - \mathbf{v}, & \text{on } \partial M(t). \end{cases} \tag{4.26}$$

Theorem 11 (Convergence to the desired spatial domain)**.** *Under the well-posedness Assumption 12, the domain* $M(t)$ *of the system* (4.25)*, with the distributed control law* (4.26) *converges to the target spatial domain* M^* *as* $t \to \infty$*, from any initial domain* M_0 *with smooth boundary.*

Proof. We consider an energy functional E given by:

$$E = \frac{1}{2}\int_{\partial M(t)} |\mathbf{e}|^2 + \frac{1}{2}\int_{\partial M(t)} |\mathbf{v}|^2.$$

Its time derivative, $\dot{E}$, using (4.26), is given by:

$$\dot{E} = \int_{\partial M(t)} \mathbf{e}\cdot\mathbf{v} + \int_{\partial M(t)} \mathbf{v}\cdot\partial_t\mathbf{v} = \int_{\partial M(t)} (\mathbf{e}+\mathbf{v})\cdot\partial_t\mathbf{v} = -\int_{\partial M(t)} |\mathbf{v}|^2.$$

Clearly, $\dot{E} \leq 0$, and considering a parametrization of $\partial M(t)$ by the interval $[0,1)$, we have

$\mathbf{v}(t,\cdot) \in H^1((0,1))$ and bounded. By Lemma 5, the Rellich-Kondrachov Compactness theorem, $H^1((0,1))$ is compactly contained in $L^2((0,1))$ (and we also have that $H^1((0,1))$ is dense in $L^2((0,1))$). Thus, by the LaSalle Invariance Principle, Lemma 6, we have that the solutions to (4.25) with the control law (4.26) converge in the L^2-norm to the largest invariant subset of $\dot{E}^{-1}(0)$, which satisfies:

$$\lim_{t\to\infty} |||\mathbf{v}|||_{L^2(\partial M(t))} = 0, \quad \lim_{t\to\infty} \partial_t |||\mathbf{v}|||_{L^2(\partial M(t))} = \lim_{t\to\infty} \int_{\partial M(t)} \mathbf{v}\cdot\partial_t \mathbf{v} = 0.$$

The set $\dot{E}^{-1}(0)$ is characterized by the first equality above and the second equality is further satisfied by the invariant subset of $\dot{E}^{-1}(0)$. We know from (4.26) that $\partial_t \mathbf{v} = -\mathbf{e} - \mathbf{v}$ on $\partial M(t)$, which upon multiplying on both sides by $\mathbf{v}$, integrating over $\partial M(t)$ and applying the previous equality on the integral of $\mathbf{v}\cdot\partial_t\mathbf{v}$, yields $\lim_{t\to\infty}\int_{\partial M(t)} \mathbf{e}\cdot\mathbf{v} = 0$. Now, we have $|\partial_t\mathbf{v}|^2 = |\mathbf{e}|^2 + |\mathbf{v}|^2 + 2\mathbf{e}\cdot\mathbf{v}$, which on integrating over $\partial M(t)$ yields $\lim_{t\to\infty} |||\partial_t\mathbf{v}|||_{L^2(\partial M(t))} = \lim_{t\to\infty} |||\mathbf{e}|||_{L^2(\partial M(t))}$. By multiplying $\partial_t\mathbf{v} = -\mathbf{e} - \mathbf{v}$ on both sides by $\partial_t\mathbf{v}$, integrating over $\partial M(t)$, and using the Cauchy-Schwarz inequality, we obtain:

$$\begin{aligned}\lim_{t\to\infty} |||\partial_t\mathbf{v}|||^2_{L^2(\partial M(t))} &= \lim_{t\to\infty} -\int_{\partial M(t)} \mathbf{e}\cdot\partial_t\mathbf{v} \leq \lim_{t\to\infty}\int_{\partial M(t)} |\mathbf{e}||\partial_t\mathbf{v}| \\ &\leq \lim_{t\to\infty} |||\mathbf{e}|||_{L^2(\partial M(t))} |||\partial_t\mathbf{v}|||_{L^2(\partial M(t))} = \lim_{t\to\infty} |||\partial_t\mathbf{v}|||^2_{L^2(\partial M(t))}\end{aligned}$$

In this way, the Cauchy-Schwarz inequality becomes an equality, which implies that $\lim_{t\to\infty}\int_{\partial M(t)} [|\mathbf{e}||\partial_t\mathbf{v}| - (-\mathbf{e})\cdot\partial_t\mathbf{v}] = 0$ (since the integrand is non-negative and its integral is zero, it is zero almost everywhere), thus $\lim_{t\to\infty}\partial_t\mathbf{v} = -\lim_{t\to\infty}\mathbf{e}$ almost everywhere (a.e.) on the boundary, and, in turn, implies that $\lim_{t\to\infty}\mathbf{v} = 0$ a.e. on the boundary (since $\partial_t\mathbf{v} = -\mathbf{e} - \mathbf{v}$ and $\lim_{t\to\infty}\partial_t\mathbf{v} = -\lim_{t\to\infty}\mathbf{e}$). From here, and owing to the Invariance Principle, we have $\lim_{t\to\infty}\partial_t\mathbf{v} = 0 = \lim_{t\to\infty}\mathbf{e}$ a.e. on the boundary. Thus, we have that $\lim_{t\to\infty} M(t) = M^*$. □

Stage 2

Here, the agents in the swarm implement the pseudo-localization algorithm presented in Section 4.4.2. Since the agents are distributed across the target spatial domain M^*, implementing the pseudo-localization algorithm yields the coordinate transformation Ψ^* characteristic of the domain M^*. We therefore have $\partial_t \Psi^* = 0$, which implies that $\frac{d\Psi^*}{dt} = \partial_t \Psi^* + \nabla(\Psi^*)\mathbf{v} = \nabla(\Psi^*)\mathbf{v}$, which will be used in Stage 3.

Stage 3

In this stage, the boundary agents of the swarm remain stationary and interior agents converge to the desired density function.

Consider the distributed control law, defined as follows for all time t:

$$\begin{cases} \frac{d\mathbf{v}}{dt} = -\rho\nabla(\rho - p^* \circ \Psi^*) + (\mathbf{v}\cdot\nabla)\mathbf{v} + \Delta\mathbf{v} - \mathbf{v}, & \text{for } \mathbf{r} \in \mathring{M}^*, \\ \mathbf{v} = 0, & \text{on } \partial M^*, \end{cases} \tag{4.27}$$

where $\frac{d\mathbf{v}}{dt}$ at $\mathbf{r} \in M$ is the acceleration of the agent at $\mathbf{r}$, the control input. Using the relation $\frac{d}{dt} = \partial_t + \mathbf{v}\cdot\nabla$, it follows from (4.27) that $\partial_t \mathbf{v} = -\rho\nabla(\rho - p^* \circ \Psi^*) + \Delta\mathbf{v} - \mathbf{v}$.

Theorem 12 (Convergence to the desired density)**.** *The solutions $\rho(t,\cdot)$ to* (4.25) *for the fixed domain M^*, under the distributed control law* (4.27) *and the well-posedness Assumption 12, converge to the desired density distribution ρ^* in the L^2-norm as $t \to \infty$.*

Proof. We consider an energy functional E given by:

$$E = \frac{1}{2}\int_{M^*} |\rho - p^* \circ \Psi^*|^2 + \frac{1}{2}\int_{M^*} |\mathbf{v}|^2.$$

Using Corollary 1, to compute the derivative of energy functionals, we obtain $\dot{E}$ (letting $\bar{\nabla} =$

$(\partial_X, \partial_Y))$ as follows:

$$\begin{aligned}
\dot{E} &= \int_{M^*} (\rho - p^* \circ \Psi^*) \left(\frac{d\rho}{dt} - \frac{d(p^* \circ \Psi^*)}{dt} \right) + \frac{1}{2} \int_{M^*} |\rho - p^* \circ \Psi^*|^2 \nabla \cdot \mathbf{v} + \int_{M^*} \mathbf{v} \cdot \partial_t \mathbf{v} \\
&= -\int_{M^*} (\rho - p^* \circ \Psi^*) \left(\rho \nabla \cdot \mathbf{v} + \bar{\nabla} p^* \cdot \frac{d\Psi^*}{dt} \right) + \frac{1}{2} \int_{M^*} |\rho - p^* \circ \Psi^*|^2 \nabla \cdot \mathbf{v} + \int_{M^*} \mathbf{v} \cdot \partial_t \mathbf{v} \\
&= -\frac{1}{2} \int_{M^*} (\rho^2 - (p^* \circ \Psi^*)^2) \nabla \cdot \mathbf{v} - \int_{M^*} (\rho - p^* \circ \Psi^*) \bar{\nabla} p^* \cdot \frac{d\Psi^*}{dt} + \int_{M^*} \mathbf{v} \cdot \partial_t \mathbf{v},
\end{aligned}$$

where, to obtain the third equality, we expand the square $|\rho - p^* \circ \Psi^*|^2$ in the second integral of the second equality. Since $\mathbf{v} = 0$ on ∂M^* and from Section 4.4.3, we have $\frac{d\Psi^*}{dt} = \nabla(\Psi^*)\mathbf{v}$, we obtain:

$$\dot{E} = \frac{1}{2} \int_{M^*} \nabla(\rho^2 - (p^* \circ \Psi^*)^2) \cdot \mathbf{v} - \int_{M^*} (\rho - p^* \circ \Psi^*) \bar{\nabla} p^* \cdot (\nabla \Psi^* \mathbf{v}) + \int_{M^*} \mathbf{v} \cdot \partial_t \mathbf{v}.$$

We have $\bar{\nabla} p^* \nabla \Psi^* = \nabla(p^* \circ \Psi^*)$, and $\nabla(\rho^2 - (p^* \circ \Psi^*)^2) = (\rho - p^* \circ \Psi^*) \nabla(\rho + p^* \circ \Psi^*) + (\rho + p^* \circ \Psi^*) \nabla(\rho - p^* \circ \Psi^*)$. Thus, we get:

$$\begin{aligned}
\dot{E} = \frac{1}{2} \int_{M^*} (\rho + p^* \circ \Psi^*) \nabla(\rho - p^* \circ \Psi^*) \cdot \mathbf{v} + \frac{1}{2} \int_{M^*} (\rho - p^* \circ \Psi^*) \nabla(\rho + p^* \circ \Psi^*) \cdot \mathbf{v} \\
- \int_{M^*} (\rho - p^* \circ \Psi^*) \nabla(p^* \circ \Psi^*) \cdot \mathbf{v} + \int_{M^*} \mathbf{v} \cdot \partial_t \mathbf{v}.
\end{aligned}$$

We therefore get:

$$\dot{E} = \int_{M^*} \rho \nabla(\rho - p^* \circ \Psi^*) \cdot \mathbf{v} + \int_{M^*} \mathbf{v} \cdot \partial_t \mathbf{v} = \int_{M^*} \mathbf{v} \cdot (\rho \nabla(\rho - p^* \circ \Psi^*) + \partial_t \mathbf{v}).$$

From (4.27), we have $\partial_t \mathbf{v} = -\rho \nabla(\rho - p^* \circ \Psi^*) + \Delta \mathbf{v} - \mathbf{v}$, and we obtain:

$$\dot{E} = -\int_{M^*} |\mathbf{v}|^2 - \int_{M^*} |\nabla \mathbf{v}_x|^2 - \int_{M^*} |\nabla \mathbf{v}_y|^2.$$

Clearly, $\dot{E} \leq 0$, with $\rho(t,.), \mathbf{v} \in H^1(M^*)$ and bounded (by Assumption 12). By Lemma 5, the Rellich-Kondrachov Compactness theorem, $H^1(M^*)$ is compactly contained in $L^2(M^*)$ (and we also know that the set of all $(\rho, \mathbf{v})$ satisfying Assumption 12 is dense in $L^2(M^*)$). Thus, by the Invariance Principle, Lemma 6, we have that the solution to (4.25) converges in the L^2-norm to the largest invariant subset of $\dot{E}^{-1}(0)$, which satisfies:

$$|||\mathbf{v}|||_{H^1(M^*)} = 0, \quad \frac{1}{2}\partial_t |||\mathbf{v}|||^2_{L^2(M^*)} = \int_{M^*} \mathbf{v} \cdot \partial_t \mathbf{v} = 0. \tag{4.28}$$

The set $\dot{E}^{-1}(0)$ is characterized by the first equality above and the second equality is further satisfied by the invariant subset of $\dot{E}^{-1}(0)$. We know from (4.27) that

$$\partial_t \mathbf{v} = -\rho \nabla(\rho - p^* \circ \Psi^*) + \Delta \mathbf{v} - \mathbf{v}, \tag{4.29}$$

which substituted in (4.28) yields $\int_{M^*} \rho \mathbf{v} \cdot \nabla(\rho - p^* \circ \Psi^*) = 0$. Now, from (4.29), we obtain $|||\partial_t \mathbf{v}|||^2_{L^2(M^*)} = \int_{M^*} |\rho \nabla(\rho - p^* \circ \Psi^*)|^2 + \int_{M^*} |\mathbf{v}|^2 + 2\int_{M^*} \rho \mathbf{v} \cdot \nabla(\rho - p^* \circ \Psi^*) = \int_{M^*} |\rho \nabla(\rho - p^* \circ \Psi^*)|^2$; that is, $|||\partial_t \mathbf{v}|||_{L^2(M^*)} = |||\rho \nabla(\rho - p^* \circ \Psi^*)|||_{L^2(M^*)}$. By multiplying (4.29) by $\partial_t \mathbf{v}$ on both sides and applying the Cauchy-Schwarz inequality, we can also get that $|||\partial_t \mathbf{v}|||^2_{L^2(M^*)} = -\int_{M^*} \rho \partial_t \mathbf{v} \cdot \nabla(\rho - p^* \circ \Psi^*) \leq \int_{M^*} |\partial_t \mathbf{v}||\rho \nabla(\rho - p^* \circ \Psi^*)| \leq |||\partial_t \mathbf{v}|||_{L^2(M^*)} ||\rho \nabla(\rho - p^* \circ \Psi^*)|||_{L^2(M^*)} = |||\partial_t \mathbf{v}|||^2_{L^2(M^*)}$. Thus, the Cauchy-Schwarz inequality is in fact an equality, which implies that $\partial_t \mathbf{v} = -\rho \nabla(\rho - p^* \circ \Psi^*)$ almost everywhere in M^*, which, from (4.29) implies in turn that $\mathbf{v} = 0$ a.e. in M^*. It thus follows that $\partial_t \mathbf{v} = 0$ and $\nabla(\rho - p^* \circ \Psi^*) = 0$ a.e in M^*, and therefore $\rho - p^* \circ \Psi^*$ is constant a.e. in M^*. Using the Poincare-Wirtinger inequality, Lemma 4, we obtain that $\|(\rho - p^* \circ \Psi^*) - (\rho - p^* \circ \Psi^*)_{M^*}\| \leq C\|\nabla(\rho - p^* \circ \Psi^*)\| = 0$, where $(\rho - p^* \circ \Psi^*)_{M^*} = \frac{1}{|M^*|}\int_{M^*}(\rho - p^* \circ \Psi^*)$. Since $\int_{M^*} \rho = \int_N p^* = \int_{M^*} p^* \circ \Psi^* = 1$, we have that $(\rho - p^* \circ \Psi^*)_{M^*} = 0$, and therefore $\|\rho - p^* \circ \Psi^*\|_{L^2(M^*)} = 0$. □

Robustness of the distributed control law

The self-organization algorithm in 2D has been divided into three stages, where asymptotic convergence is achieved in each stage (with exponential convergence in the second stage). We now present a robustness result for convergence in Stage 3 under incomplete convergence in the preceding stages.

Lemma 41 (Robustness of the control law)**.** *For every $\delta > 0$, there exist $T_1, T_2 < \infty$ such that when Stages 1 and 2 are terminated at $t_1 > T_1$ and $t_2 > T_2$ respectively, we have that $\lim_{t\to\infty} \|\rho(t,\cdot) - \rho^*\|_{L^2(M(t_1))} < \delta$.*

Proof. In Stage 1, it follows from Theorem 11 on the convergence to the desired spatial domain that $\lim_{t\to\infty} M(t) = M^*$. Then for every $\varepsilon_1 > 0$, we have $T_1 < \infty$, such that $d_H(M(t), M^*) < \varepsilon_1$ for all $t > T_1$, where d_H is the Hausdorff distance between two sets. (Note that any appropriate notion of distance can alternatively be used here.) Let Stage 1 be terminated at $t_1 > T_1$, which implies that the swarm is distributed across the domain $M(t_1)$. In Stage 2, it follows from Lemma 40 on the convergence of the heat flow equation to the harmonic map, that for a domain $M(t_1)$, we have that $\lim_{t\to\infty} \mathbf{R}(t,\cdot) = \Psi_{M(t_1)}$ pointwise, where $\Psi_{M(t_1)}$ is the harmonic map from $M(t_1)$ to N (the unit disk). Then, for every $\varepsilon_2 > 0$, we have a $T_2 < \infty$, such that $\|\mathbf{R}(t,\cdot) - \Psi_{M(t_1)}\|_\infty < \varepsilon_2$ for all $t > T_2$. Let Stage 2 be terminated at $t_2 > T_2$, which implies that the map from the spatial domain to the disk is $\mathbf{R}(t_2,\cdot)$. In Stage 3, it follows from the arguments in the proof of Theorem 12 (on the convergence to the desired density function) that $\lim_{t\to\infty} \rho(t,\cdot) = p^* \circ \mathbf{R}(t_2,\cdot)$ a.e. in $M(t_1)$ if the map at the end of Stage 2 is $\mathbf{R}(t_2,\cdot)$. We characterize the error as $\lim_{t\to\infty} \|\rho - \rho^*\|_{L^2(M(t_1))} = \|p^* \circ \mathbf{R}(t_2,\cdot) - p^* \circ \Psi^*\|_{L^2(M(t_1))} = \|p^* \circ \mathbf{R}(t_2,\cdot) - p^* \circ \Psi_{M(t_1)} + p^* \circ \Psi_{M(t_1)} - p^* \circ \Psi^*\|_{L^2(M(t_1))} \leq \|p^* \circ \mathbf{R}(t_2,\cdot) - p^* \circ \Psi_{M(t_1)}\|_{L^2(M(t_1))} + \|p^* \circ \Psi_{M(t_1)} - p^* \circ \Psi^*\|_{L^2(M(t_1))}$. Recall that $\|\mathbf{R}(t_2,\cdot) - \Psi_{M(t_1)}\|_\infty < \varepsilon_2$, and since p^* is Lipschitz, we can get the bound $\|p^* \circ \mathbf{R}(t_2) - p^* \circ \Psi_{M(t_1)}\|_{L^2(M(t_1))} < \delta_1 = c\varepsilon_2$ (where c is the Lipschitz constant times the area of $M(t_1)$). The harmonic map also depends continuously on its domain [68], which yields the bound $\|\Psi_{M(t_1)} - \Psi^*\|_\infty < \varepsilon_3$, since $d_H(M(t_1), M^*) < \varepsilon_1$. Thus, we get another bound $\|p^* \circ \Psi_{M(t_1)} - p^* \circ \Psi^*\|_{L^2(M(t_1))} < \delta_2 = c\varepsilon_3$, and that $\|\rho - \rho^*\|_{L^2(M(t_1))} < \delta_1 + \delta_2 = \delta$. Therefore, going backwards, for all $\delta > 0$, we can find T_1 and T_2 such that the density error is bounded by δ, when the Stages 1 and 2 are terminated at $t_1 > T_1$ and $t_2 > T_2$ respectively. □

4.4.4 Discrete implementation

In this section, we present consistent schemes for discrete implementation of the distributed control laws (4.26) and (4.29), where the key aspect is the computation of spatial gradients (of ρ in Stage 1, and of ρ, Ψ^* and the components of velocity $\mathbf{v}$ in Stage 3). The network graph underlying the swarm is a random geometric graph, where the nodes are distributed according to the density function over the spatial domain. According to this, every agent communicates with other agents within a disk of given radius (say r) determined by the hardware capabilities, which reduces to the graph having an edge between two nodes if and only if the nodes are separated by a distance less than r. We recall the earlier stated assumption that the agents know the true x- and y-directions.

On the computation of p^*

We first begin with an approach to compute offline the map p^* via interpolation. Let the desired domain $M^* \in \mathbb{R}^2$ be discretized into a uniform grid to obtain $M_d^* = \{\mathbf{r}_1, \ldots, \mathbf{r}_m\}$ (the centers of finite elements, where $\mathbf{r}_k = (x_k, y_k)$). The desired density $\rho^* : M^* \to \mathbb{R}_{>0}$ is known, and we compute the value of ρ^* on M_d^* to get $\rho^*(\mathbf{r}_1, \ldots, \mathbf{r}_m) = (\rho_1^*, \ldots, \rho_m^*)$. We also have $\Psi^*(x,y) = (X^*, Y^*) \in N$, for all $(x,y) \in M^*$. Now, computing the integral with respect to the Dirac measure for the set M_d^*, we obtain $\Psi^*(\mathbf{r}_1, \ldots, \mathbf{r}_m) = (\Psi_1^*, \ldots, \Psi_m^*)$. The value of the function p^* at any $(X,Y) \in N$ can be obtained from the relation $p^*(\Psi_1^*, \ldots, \Psi_m^*) = \rho^*(\mathbf{r}_1, \ldots, \mathbf{r}_m)$ for $k = 1, \ldots, m$ by an appropriate interpolation.

$$\begin{array}{ccc} & & (\rho_1^*, \ldots, \rho_m^*) = p^*(\Psi_1^*, \ldots, \Psi_m^*) \\ & \overset{\rho^*}{\nearrow} & \Uparrow p^* \\ (\mathbf{r}_1, \ldots, \mathbf{r}_m) & \xmapsto{\quad \Psi^* \quad} & (\Psi_1^*, \ldots, \Psi_m^*) \end{array}$$

Commutative diagram

Discrete control law

As stated earlier, for the discrete implementation of the distributed control laws (4.26) and (4.29), the key aspect is the computation of spatial gradients (of ρ in Stage 1, and of ρ, Ψ^* and the components of velocity $\mathbf{v}$ in Stage 3). In the subsequent sections we present two alternative, consistent schemes for computing the spatial gradient (of any smooth function, with the above being the ones of interest), one using the Jacobian of the harmonic map and the other without it.

Computing the Jacobian of the harmonic map

Let $J(\mathbf{r}) = \nabla\Psi(\mathbf{r})$ be the (non-singular) Jacobian of the harmonic diffeomorphism $\Psi : M \to N$. When the steady-state is reached in the pseudo-localization algorithm (4.24) (i.e., $X_i(t+1) = X_i(t) = \psi_1^i$ and $Y_i(t+1) = Y_i(t) = \psi_2^i$), we have, $\forall\, i \in \mathscr{S}$:

$$\sum_{j\in\mathscr{N}_i} \frac{1}{d_j}(\psi_1^j - \psi_1^i) = 0, \qquad \sum_{j\in\mathscr{N}_i} \frac{1}{d_j}(\psi_2^j - \psi_2^i) = 0,$$

where i is the index of the agent located at $\mathbf{r} \in M$ and $\mathscr{N}_i$ is the set of agents in a disk-shaped neighborhood $B_\varepsilon(\mathbf{r})$ of area ε centered at $\mathbf{r}$. Rewriting the above, we get, $\forall\, i \in \mathscr{S}$:

$$\psi_1^i = \frac{\sum_{j\in\mathscr{N}_i} \frac{1}{d_j}\psi_1^j}{\sum_{j\in\mathscr{N}_i} \frac{1}{d_j}}, \qquad \psi_2^i = \frac{\sum_{j\in\mathscr{N}_i} \frac{1}{d_j}\psi_2^j}{\sum_{j\in\mathscr{N}_i} \frac{1}{d_j}}. \tag{4.30}$$

We assume that the agents have the capability in their hardware to perturb the disk of communication $B_\varepsilon(\mathbf{r})$ (by moving an antenna, for instance). The Jacobian $J = \nabla\Psi$, where $\Psi = (\psi_1, \psi_2)$ is computed through perturbations to $\mathscr{N}_i$ (i.e., the neighborhood $B_\varepsilon(\mathbf{r})$) and using consistent discrete approximations:

$$\partial_x\psi_1 \approx \frac{\psi_1(\mathbf{r}+\delta x\mathbf{e}_1) - \psi_1(\mathbf{r})}{\delta x}, \qquad \partial_y\psi_1 \approx \frac{\psi_1(\mathbf{r}+\delta y\mathbf{e}_2) - \psi_1(\mathbf{r})}{\delta y},$$

and similarly for ψ_2. Now, $\psi_1(\mathbf{r}+\delta x\mathbf{e}_1)$ is computed as in (4.30) for $\mathcal{N}_i^{\delta x}$, the set of agents in $B_\varepsilon(\mathbf{r}+\delta x\mathbf{e}_1)$ and $\psi_1(\mathbf{r}+\delta y\mathbf{e}_2)$ from $B_\varepsilon(\mathbf{r}+\delta y\mathbf{e}_2)$.

Computing the spatial gradient of a smooth function using the Jacobian of Ψ

Let $\nabla = (\partial_x, \partial_y)$ and $\bar{\nabla} = \left(\partial_{\psi_1}, \partial_{\psi_2}\right)$, where $\Psi = (\psi_1, \psi_2)$. We have $\partial_x = (\partial_x\psi_1)\partial_{\psi_1} + (\partial_x\psi_2)\partial_{\psi_2}$ and $\partial_y = (\partial_y\psi_1)\partial_{\psi_1} + (\partial_y\psi_2)\partial_{\psi_2}$. Therefore, $\nabla = J^\top\bar{\nabla}$. For a smooth function $f : M \to \mathbb{R}$, we have, $\nabla f = J^\top\bar{\nabla}f$, and the agents can numerically compute $\bar{\nabla}$ by:

$$\left(\frac{\partial f}{\partial \psi_1}\right)_i \approx \frac{1}{|\mathcal{N}_i|}\sum_{j\in\mathcal{N}_i}\frac{f_j - f_i}{\psi_1^j - \psi_1^i}, \qquad \left(\frac{\partial f}{\partial \psi_2}\right)_i \approx \frac{1}{|\mathcal{N}_i|}\sum_{j\in\mathcal{N}_i}\frac{f_j - f_i}{\psi_2^j - \psi_2^i},$$

where i is the index of the agent located at $\mathbf{r} \in M$ and $\mathcal{N}_i$ is the set of agents in a ball $B_\varepsilon(\mathbf{r})$.

Computing the spatial gradient of a smooth function without the Jacobian of Ψ

In the absence of a Jacobian estimate, we use the following alternative method for computing an approximate spatial gradient estimate of a smooth function. This is used in Stage 1 of the self-organization process.

Let $\bar{f}(\mathbf{r})$ be the mean value of f over a ball $B_\varepsilon(\mathbf{r})$:

$$\bar{f}(\mathbf{r}) = \frac{1}{\varepsilon}\int_{B_\varepsilon(\mathbf{r})} f\,d\mu \approx \frac{1}{|\mathcal{N}_i|}\sum_{j\in\mathcal{N}_i} f_j.$$

We have:

$$\begin{aligned}\frac{1}{\varepsilon}\frac{\partial \bar{f}}{\partial x} &\approx \frac{1}{\varepsilon}\frac{\bar{f}(\mathbf{r}+\delta x\mathbf{e}_1) - \bar{f}(x)}{\delta x} = \frac{1}{\varepsilon}\frac{\int_{B_\varepsilon(\mathbf{r}+\delta x\mathbf{e}_1)} f\,d\mu - \int_{B_\varepsilon(\mathbf{r})} f\,d\mu}{\delta x} \\ &= \frac{1}{\varepsilon}\int_{B_\varepsilon(\mathbf{r})}\frac{(f(\mathbf{r}+\delta x\mathbf{e}_1) - f(\mathbf{r}))}{\delta x}d\mu \approx \frac{1}{\varepsilon}\int_{B_\varepsilon(\mathbf{r})}\frac{\partial f}{\partial x}d\mu = \overline{\left(\frac{\partial f}{\partial x}\right)}.\end{aligned}$$

Similarly,

$$\frac{1}{\varepsilon}\frac{\partial \bar{f}}{\partial y} \approx \frac{1}{\varepsilon}\frac{\bar{f}(\mathbf{r}+\delta y\mathbf{e}_2) - \bar{f}(x)}{\delta y} \approx \overline{\left(\frac{\partial f}{\partial y}\right)}.$$

In all, for any scalar function f, each agent can use the approximation:

$$(\nabla f)_i \approx \left(\overline{\left(\frac{\partial f}{\partial x}\right)}, \overline{\left(\frac{\partial f}{\partial y}\right)}\right) = \frac{1}{\varepsilon}\left(\frac{\partial \bar{f}}{\partial x}, \frac{\partial \bar{f}}{\partial y}\right), \tag{4.31}$$

to estimate of the gradient ∇f.

On the convergence of the discrete system

We have noted earlier that the pseudo-localization algorithm (4.24) satisfies the consistency condition in that as $N \to \infty$, Equation (4.24) converges to the PDE (4.21). The pseudo-localization algorithm is also essentially a weighted Laplacian-based distributed algorithm that is stable. Thus, by the Lax Equivalence theorem [112], the solution of (4.24) converges to the solution of (4.21) as $N \to \infty$. However, for the distributed control laws in Stages 1-3, we are only able to provide consistent discretization schemes. The dynamics of the swarm (4.25) with the control laws (4.26) and (4.27) are nonlinear for which is no equivalent convergence theorem. Further analysis to determine convergence is required, which falls out the scope of this present work.

4.5 Numerical simulations

In this section, we present numerical simulations of swarm self-organization, that is, of the control laws presented in Sections 4.3.2 and of Section 4.4.3.

4.5.1 Self-organization in one dimension

In the simulation of the 1D case, we consider a swarm of $N = 10000$ agents, the desired density function is given by $\rho^*(x) = a\sin(x) + b$, where $a = 1 - \frac{\pi}{2N}$ and $b = \frac{1}{N}$, $x \in \left[0, \frac{\pi}{2}\right]$. We

Algorithm 4. Self-organization algorithm for 2D environments

Input: M^*, ρ^* and k_1, k_2, K (number of iterations for each stage), Δt (time step)
Requires:
  Offline computation of p^* similar to the outline in Section 4.3.3
  Boundary agents are aware of being at boundary or interior of domain, can
    communicate with others along the boundary, can approximate the normal
    to the boundary, and can measure density of boundary agents,
  Agents have knowledge of a common orientation of a reference frame
Initialize: $\mathbf{r}_i$ (Agent positions), $\mathbf{v}_i = 0$ (Agent velocities)
Boundary agents localize as outlined in Section 4.4.1
Stage 1**:**
for $k := 1$ to k_1 **do**
  if agent i is at the interior of domain **then**
    compute $\mathbf{v}_i(k) = -\frac{(\nabla\rho)_i}{\rho_i}(k)$ from (4.26)
    move $\mathbf{r}_i(k+1) = \mathbf{r}_i(k) + \mathbf{v}_i(k)\Delta t$
  else if agent i is at the boundary of domain **then**
    compute $\mathbf{v}_i(k+1) = \mathbf{v}_i(k) - (\mathbf{r}_i(k) - \mathbf{r}_i^*(k) + \mathbf{v}_i(k))\Delta t$ from (4.26), and move $\mathbf{r}_i(k+1) = \mathbf{r}_i(k) + \mathbf{v}_i(k)\Delta t$
End Stage 1
Stage 2**:**
Boundary agents map themselves onto unit circle according to (4.17)
for $k := 1$ to k_2 **do**
  for agent i in the interior **do**
    compute $X_i(k+1)$, $Y_i(k+1)$ according to (4.24)
Stage 3**:**
for $k := 1$ to K **do**
  for agent i in the interior **do**
    compute $\mathbf{v}_i(k+1) = \mathbf{v}_i(k) + (-\rho_i(k)(\nabla(\rho - p^* \circ \Psi^*))_i(k) + (\mathbf{v}_i(k)\cdot\nabla)\mathbf{v}_i(k) - \mathbf{v}_i(k))\Delta t$ from (4.27) , with $(\nabla(\rho - p^* \circ \Psi^*))_i(k)$ as in (4.31)
    update $\mathbf{r}_i(k+1) = \mathbf{r}_i(k) + \mathbf{v}_i(k)\Delta t$

use a kernel-based method to approximate the continuous density function, which is given by:

$$\rho(t,\mathbf{r}) = \sum_{i\in\mathscr{S}} K\left(\frac{\|\mathbf{r}-\mathbf{r}_i(t)\|}{d}\right), \quad K(x) = \begin{cases} \frac{c_d}{d^n}, & \text{for } 0 \le x < 1, \\ 0, & \text{for } x \ge 1, \end{cases}$$

is a flat kernel and $c_d \in \mathbb{R}_{>0}$ is a constant [31]. We discretize the spatial domain with $\Delta x = 0.001$ units, and use an adaptive time step. The self-organization begins from an arbitrary initial density distribution. Figure 4.2 shows the initial density distribution, an intermediate distribution and the final distribution. We observe that there is convergence to the desired density function, even with noisy density measurements.

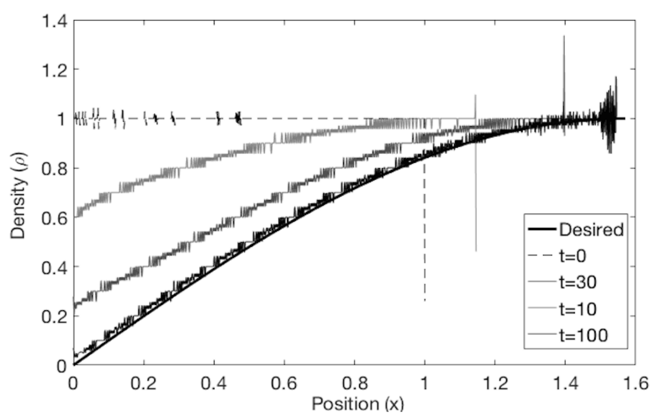

Figure 4.2. Density $\rho(x)$ plotted against position x at different instants of time.

4.5.2 Self-organization in two dimensions

In the simulation of the 2D case, we first present in Figure 4.3 the evolution of the boundary of the swarm in Stage 1, where the swarm converges to the target spatial domain M^* from an initial spatial domain. The target spatial domain, a circle of radius 0.5 units, given by $M^* = \{(x,y) \in \mathbb{R}^2 \,|\, (x-0.6)^2 + y^2 \le 0.25\}$, with the desired density function ρ^* given by $\rho^*(x,y) = \frac{1}{\left((x-0.4)^2+y^2\right)^{0.3}}$.

We present in Figures 4.4 and 4.5 the result of implementation of the pseudo-localization algorithm with the steady state distributions of $\Psi^* = (\psi_1^*, \psi_2^*)$ respectively. We note that the

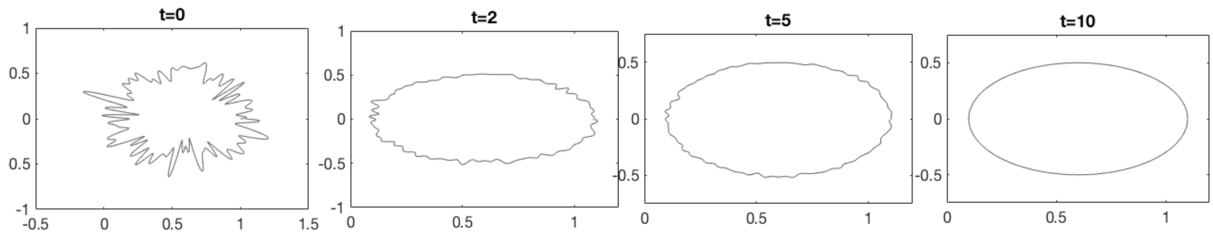

Figure 4.3. Evolution of the swarm boundary in Stage 1.

steady state distribution Ψ^* as a function of the spatial coordinates (x, y) in this case is linear.

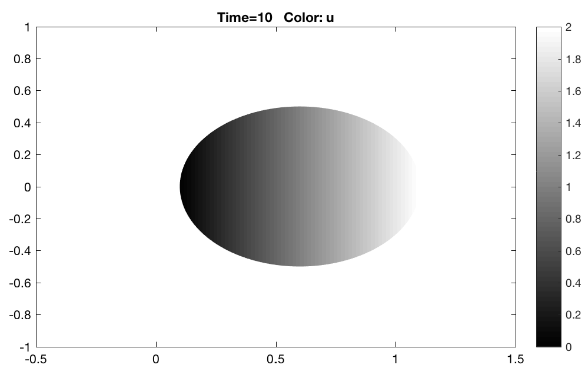

Figure 4.4. Steady-state distribution of ψ_1^*.

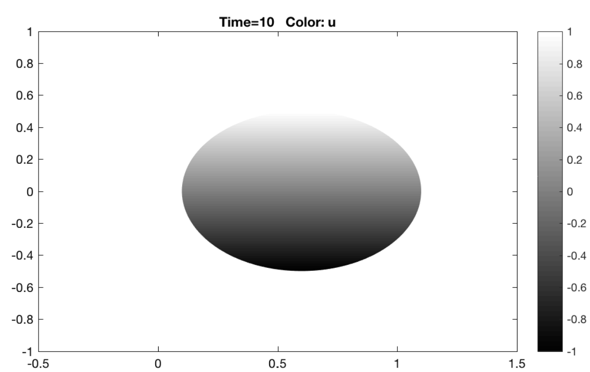

Figure 4.5. Steady-state distribution of ψ_2^*.

Next, we focus on Stage 3 of the self-organization process, where the agents already distributed over the target spatial domain, converge to the desired density function. The initial density function of the swarm is uniform, and the distributed control law of Stage 3 in Section 4.4.3

is implemented. Figure 4.6 shows the density function at a few intermediate time instants of implementation and figure 4.7 shows the spatial density error plot, where $e(\rho) = \int_{M^*} |\rho - \rho^*|^2$ is the spatial density error. The results show convergence as desired.

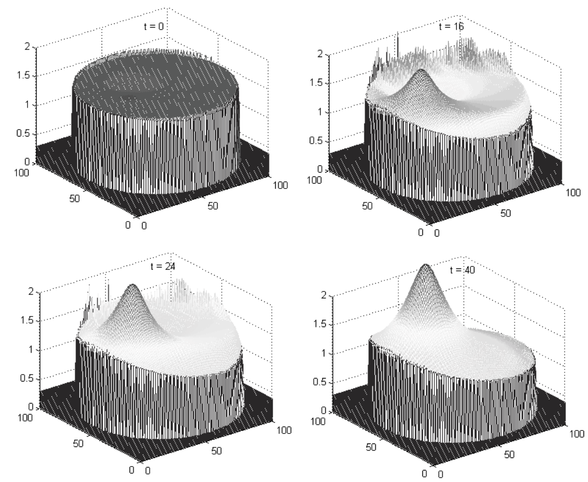

Figure 4.6. Evolution of density function in Stage 3.

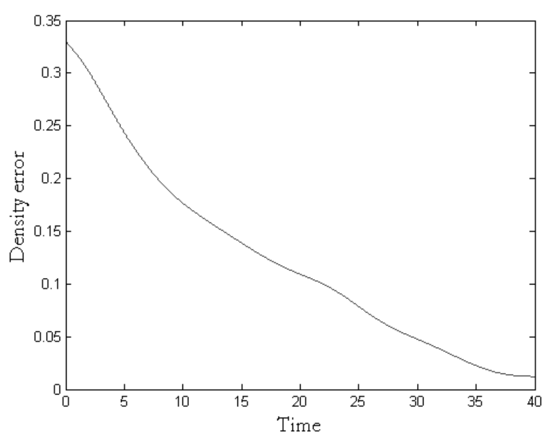

Figure 4.7. Spatial density error $e(\rho) = \int_{M^*} |\rho - \rho^*|^2$ vs time,

4.6 Summary

In this chapter, we considered the problem of self-organization in multi-agent swarms in one and two dimensions. The primary contribution of the work is the analysis and design of position and index-free distributed control laws for swarm self-organization, aided by a distributed pseudo-localization algorithm for the assignment of agent identifiers.

The material in this chapter, in full, is a reprint of the material as it appears in the publication *Distributed Control for Spatial Self-Organization of Multi-Agent Swarms*, V. Krishnan and S. Martínez, SIAM Journal on Control and Optimization, 56(5), pp. 3642–3667, 2018. A preliminary version of the work appeared in the proceedings of the International Symposium on Mathematical Theory of Networks and Systems, Minneapolis, USA, July 2016 as *Self-Organization in Multi-Agent Swarms via Distributed Computation of Diffeomorphisms*, V. Krishnan and S. Martínez. The dissertation author was the primary investigator and author of these papers.

Chapter 5

State estimation for tracking and navigation

In this chapter, we begin with the well-studied notion of strong local observability of nonlinear, discrete-time systems and investigate its relationship to the optimization-based state estimation problem. To handle uncertain initial conditions and the possible non-uniqueness of solutions to the estimation problem, we adopt a generalized problem formulation over the space of probability measures over the state space. More precisely, we define the MHE as a proximal gradient descent in the space of probability measures, with a non-convex, time-varying cost function. This distributional setting serves as a unifying framework for moving-horizon estimation and allows us to develop different classes of moving-horizon estimators by simply varying the metric used to define the proximal operator, and to obtain implementable filters by Monte Carlo methods. We then consider the Wasserstein metric and the KL-divergence, which yield the more familiar MHE and a particle filter, respectively. Following this, we present an analysis of the convergence and robustness properties of these estimators in the distributional setting, under assumptions of strong local observability. Further, we modify the distributional optimization problem via an entropic regularization to derive conditions that guarantee a desired level of differential privacy for these filters.

5.1 Bibliographical comments

The origins of MHE can be traced back to the limited memory optimal filters introduced in [76]. In principle, its optimization-based formulation enables it to handle nonlinearities and state constraints much more effectively than other known methods. This, coupled with the adoption of increasingly powerful, inexpensive computing platforms has brought new impetus to the adoption of moving-horizon estimation in various data-driven applications. Theoretical investigations on MHE have broadly been directed at their asymptotic stability [3, 102, 121] and robustness [73, 77, 94] properties. These properties have primarily been built upon underlying assumptions of input/output-to-state (IOSS) stability, which is adopted as the notion of detectability, wherein the norm of the state is bounded given the sequences of inputs and outputs. However, alternative foundations for the stability results in other classical notions of observability, such as strong observability [95], have remained unexplored. The connection between nonlinear observability theory and estimation problems runs deep, see [83] and more recently [115], and it is worthwhile to explore this connection in the context of optimization-based estimation methods such as moving-horizon estimation.

Another important consideration in the MHE problem is the cost of computation. The problem formulation more commonly involves solving an optimization problem at every time instant, with the state estimate and disturbances as decision variables in the optimization, where the dimension of the problem scales with the size of the horizon. This approach, in general, tends to be computationally intensive, which poses a hurdle for implementation in real-time. This has motivated the search for fast MHE that implement one or more iterations of the optimization at every time instant. Recently, in [4], [5], the authors develop such a method for noiseless systems and provide theoretical guarantees on convergence. However, these works assume the convexity of the cost function, which is restrictive for general nonlinear systems, and not well connected to notions of observability.

The problem of state estimation is fundamentally about dealing with uncertainty, manifested as uncertainty in the initial conditions and/or in the evolution of the system in the presence of unknown disturbances. This is appropriately formulated in the space of probability measures over the state space of the system. Advances in gradient flows in the space of probability measures [8], [107], and the corresponding discrete-time movement-minimizing schemes [99] present powerful theoretical tools that can be applied to recursive optimization-based estimation methods such as moving-horizon estimation, and can serve as a unifying framework for their design and analysis.

In many applications, the measurement data is acquired from particular individuals or users, which introduces new ethical concerns about data collection and manipulation, highlighting an increasing need for data privacy. Such is the case in home monitoring and traffic estimation (with vehicle GPS data) applications, to name a few. Differential privacy [46] has emerged over the past decade as a benchmark in data privacy. The typical setting assumes independence between the records in static databases; however, basic existing mechanisms fail to provide guarantees when correlations exist between the records in the database. This is the case when data is employed by a state estimation process whose output is then released: there is a dynamic system from which a time series of sensor measurements is obtained, and the measurement data and the released estimates are correlated.

In [43, 44], the authors generalize the definition of differential privacy to include general notions of distance between datasets and design differentially private mechanisms for Bayesian inference. In [85, 113], the authors investigate privacy-preserving mechanisms for the case where correlations exist between database records. Privacy-preserving mechanisms for functions and functional data were investigated in [65]. The work [97] studies the problem of differentially-private state estimation, introducing the formal notion of differential privacy into the framework of Kalman filter design for dynamic systems. The authors of [56] consider the problem of optimal state estimation for linear discrete-time systems with measurements corrupted by Laplacian noise. A finite-dimensional distributed convex optimization is considered in [96], where differential privacy is achieved by perturbation of the objective function. We refer the reader to [34] for a broad overview of the systems and control-theoretic perspective on differential privacy.

5.2 Observability, Estimation and Differential Privacy

The notion of observability used in this paper is intricately related to solutions of inverse problems, with an associated notion of well-posedness that is introduced below:

Definition 14 (Well posedness [78]). *Let $\mathscr{X}$ and $\mathscr{Y}$ be normed spaces, and $P : \mathscr{X} \to \mathscr{Y}$ a mapping. The equation $P(x) = y$ is called well-posed if:*

1. *Existence: For every $y \in \mathscr{Y}$, there is (at least one) $x \in \mathscr{X}$ such that $P(x) = y$.*
2. *Uniqueness: For every $y \in \mathscr{Y}$, there is at most $x \in \mathscr{X}$ such that $P(x) = y$.*
3. *Stability: The solution x depends continuously on y, that is, for any sequence $\{x_i\} \subset \mathscr{X}$ such that $P(x_i) \to P(x)$, it follows that $x_i \to x$.*

We now introduce the notion of lower semicontinuity of set-valued maps, which underlies some of the results on optimization-based state estimation in this paper.

Definition 15 (Lower semicontinuity of set-valued maps). *A point-to-set mapping $H : \mathscr{Z} \subset \mathbb{R} \rightrightarrows \mathbb{R}^d$ is lower semicontinuous at a point $\alpha \in \mathscr{Z}$ if for any $x \in H(\alpha)$ and sequences $\{\alpha_i\} \subseteq \mathscr{Z}$, $\{x_i\} \subseteq \mathbb{R}^d$ with $\{\alpha_i\} \to \alpha$, $\{x_i\} \to x$ such that $x_i \in H(\alpha_i)$ for all i, it holds that $x \in H(\alpha)$. If H is lower semicontinuous at every $\alpha \in \mathscr{Z}$, then H is said to be lower semicontinuous on $\mathscr{Z}$.*

In this paper, we consider systems of the form:

$$\Omega : \begin{cases} x_{k+1} = f(x_k, w_k), \\ y_k = h(x_k) + v_k, \end{cases} \tag{5.1}$$

where $f : \mathbb{X} \times \mathbb{W} \to \mathbb{X}$ and $h : \mathbb{X} \to \mathbb{Y}$, $w_k \in \mathbb{W}$ is the process noise, $v_k \in \mathbb{V}$ is the measurement noise at time instant k, and $\mathbb{X} \subset \mathbb{R}^{d_X}$, $\mathbb{Y} \subset \mathbb{R}^{d_Y}$, $\mathbb{W} \subset \mathbb{R}^{d_W}$, and $\mathbb{V} \subset \mathbb{R}^{d_V}$.

Assumption 13 (Lipschitz continuity)**.** *The functions f and h are Lipschitz continuous, with $\|f(x_1,w_1)-f(x_2,w_2)\| \leq c_f^{(1)}\|x_1-x_2\|+c_f^{(2)}\|w_1-w_2\|$ and $\|h(x_1)-h(x_2)\| \leq c_h\|x_1-x_2\|$.*

Assumption 14 (Noise characteristics)**.** *The noise sequences $\{w_k\}_{k\in\mathbb{N}}$ and $\{v_k\}_{k\in\mathbb{N}}$ are i.i.d samples from distributions ω and ν (with supports in $\mathbb{W}$ and $\mathbb{V}$). The sets $\mathbb{W}$ and $\mathbb{V}$ are bounded, with $|w_k| \leq W$ and $|v_k| \leq V$. Moreover, we assume that $\mathbb{E}_\omega[w_k] = 0$ and $\mathbb{E}_\nu[v_k] = 0$.*

We also introduce the following autonomous system corresponding to (5.1):

$$\Sigma : \begin{cases} x_{k+1} = f(x_k,0) = f_0(x_k), \\ y_k = h(x_k). \end{cases} \tag{5.2}$$

With a slight abuse of notation, for any $x \in \mathbb{X}$, we let $\Sigma_T(x) = \left(h(x), h\circ f_0(x), \ldots, h\circ f_0^T(x)\right)$, the sequence of outputs over a horizon of length $T+1$ for the system (5.2) from the state $x \in \mathbb{X}$. Similarly, for the system (5.1), we let:

$$\Omega(x,\mathbf{w}_{i:j}) = (h(x), h\circ f(x,w_i), \ldots, h\circ f(\ldots f(f(x,w_i),w_{i+1}),\ldots,w_j),$$

for some sequence of process noise samples $\{w_k\}$, where $\mathbf{w}_{i:j} = (w_i,\ldots,w_j)$.

The theoretical results in the moving-horizon estimation literature have largely been derived in the setting of input/output-to-state (IOSS) stability, as in [73, 77, 102] to name a few, which is a notion of norm-observability, see [71], wherein the norm of the state is bounded using the sequences of inputs and outputs. However, there are other classical notions of observability based on the notion of distinguishability, which generalize the approach taken to linear systems. For a detailed treatment, we refer the reader to [95] and [2]. In this paper, we explore the connection between the classical notion of strong local observability and moving-horizon estimation.

We now introduce the notion of strong local observability used in this paper:

Definition 16 (Strong local observability)**.** *The system* Σ *defined in* (5.2) *is called strongly locally observable if there exists a* $T_0 \in \mathbb{N}$ *such that for any given* $\mathbf{y}_T = \Sigma_T(x) \in \mathbb{Y}^{T+1}$ *and* $T \geq T_0$*, we have that* ${\Sigma_T}^{-1}(\mathbf{y}_T)$ *is a set of isolated points, and, in addition,* ${\Sigma_{T_1}}^{-1}(\mathbf{y}_1) = {\Sigma_{T_2}}^{-1}(\mathbf{y}_2)$*, for all* $\mathbf{y}_1 = \Sigma_{T_1}(x)$ *and* $\mathbf{y}_2 = \Sigma_{T_2}(x)$*, and* $T_1, T_2 \geq T_0$*. We call* T_0 *the minimum horizon length of* Σ*.*

The above definition is equivalent to the definitions contained in [2, 95], which has been restated it in a manner suitable for the optimization-based estimation framework considered here.

For systems with process noise, of the form Ω in (5.1), we introduce the notion of almost sure strong local observability.

Definition 17 (Almost sure strong local observability)**.** *The system* Ω *defined in* (5.1) *is called almost surely strongly locally observable if there exists a* $T^w \in \mathbb{N}$ *such that, given a process noise sequence* $\mathbf{w}_{0:T-1} \in \mathbb{W}^T$*, for* $T \geq T^w$*, any* $\mathbf{y}_{0:T} = \Omega_{\mathbf{w}_{0:T-1}}(x) \in \mathbb{Y}^{T+1}$*, and* $T \geq T^w$*, we have that* $\Omega^{-1}_{\mathbf{w}_{0:T-1}}(\mathbf{y}_{0:T})$ *is a set of isolated points almost surely. More precisely, the set of noise sequences* $\mathbf{w}_{0:T-1}$ *for which* $\Omega^{-1}_{\mathbf{w}_{0:T-1}}(\mathbf{y}_{0:T})$ *is not a set of isolated points, is of measure zero. Moreover, we call* T^w *the minimum horizon length of* Ω*.*

We now present a fundamental result that characterizes strong local observability via a rank condition.

Lemma 42 (Observability rank condition [95])**.** *The system* Σ *is locally strongly observable with minimum horizon length* T_0 *if and only if* $Rank(\nabla\Sigma_T(x)) = dim(\mathbb{X})$ *for all* $T \geq T_0$ *and* $x \in \mathbb{X}$*. The system* Ω *is almost surely locally strongly observable with minimum horizon length* T^w *if and only if* $Rank(\nabla\Omega_{\mathbf{w}_{0:T-1}}(x)) = dim(\mathbb{X})$ *almost surely for all* $T \geq T^w$ *and* $x \in \mathbb{X}$*.*

We now present an example to illustrate these concepts.

Example 1. *Consider a system with the state space* $\mathbb{X} = (0, \infty)$*, with* $x_{k+1} = f_0(x_k)$ *and* $y_k = h(x_k)$*,* *such that:*

$$f_0(x) = \begin{cases} 3x, & \text{for } x \in (0, a\pi - \varepsilon], \\ \gamma(x) & \text{for } x \in (a\pi - \varepsilon, a\pi + \varepsilon], \\ 2x + a\pi, & \text{for } x \in (a\pi + \varepsilon, \infty), \end{cases}$$

for some $a \in \mathbb{N}$*,* ε *small and a smooth function* γ *such that* $\gamma(a\pi - \varepsilon) = 3(a\pi - \varepsilon)$ *and* $\gamma(a\pi + \varepsilon) = 2(a\pi + \varepsilon) + a\pi$*. Moreover, let the output* $h(x) = \sin x$*. We note that* $\nabla h(x) = \cos x$ *which implies that* $\nabla h((2m+1)\pi/2) = 0$ *for all* $m \in \mathbb{N}$*. Applying Lemma 42 for this system, we can infer that for* $a = 2$*, we get that the minimum horizon length* $T_0 = 3$*. This is because the system becomes strongly locally observable at* $x = \pi/2$ *only over a horizon of length* $T_0 = 3$*, that is* $\nabla \Sigma_k(\pi/2) = \mathbf{0}_{k+1}$ *for* $k \in \{0, 1, 2\}$*. This is a case of a one-dimensional system which is strongly locally observable with a minimum horizon of length* $T_0 = 3$*. With larger values of a, the minimum horizon length is further increased.*

We make the following assumption in the rest of the paper:

Assumption 15 (Strong local observability)**.**

1. *The system* Σ *in* (5.2) *is strongly locally observable with minimum horizon length* T_0*.*
2. *The system* Ω *in* (5.1) *is almost surely strongly locally observable with minimum horizon length* T^w*.*

The KL-divergence from μ_1 to μ_2 is given by:

$$\begin{aligned} D_{\mathrm{KL}}(\mu_1 || \mu_2) &= \int_{\mathscr{X}} \log\left(\frac{d\mu_1(x)}{d\mu_2(x)}\right) d\mu_1(x) \\ &= \int_{\mathscr{X}} \rho_1(x) \log\left(\frac{\rho_1(x)}{\rho_2(x)}\right) \mathrm{dvol}(x). \end{aligned}$$

The max-divergence between μ_1 and μ_2 is defined as:

$$D_{\max}(\mu_1, \mu_2) = \sup_{x \in \mathscr{X}} \left| \log\left(\frac{\rho_1(x)}{\rho_2(x)}\right)\right|.$$

We refer the reader to [62] for a detailed overview of the relations between the various metrics and divergences in probability spaces.

We define an estimator $\mathscr{E} : \mathscr{Y} \to \mathscr{P}(\mathscr{X})$ as a function that accepts as input data y from the metric space $\mathscr{Y}$ and releases as output $\mathscr{E}[y]$, a probability measure over the space $\mathscr{X}$.

Definition 18 (Differential privacy)**.** *Given δ, an estimator $\mathscr{E}$ is ε-differentially private if for any two δ-adjacent measurements $y_1, y_2 \in \mathscr{Y}$ (that is $d_{\mathscr{Y}}(y_1, y_2) \leq \delta$), and any measurable $A \subseteq \mathscr{X}$, we have $\mathscr{E}[y_1](A) \leq e^{\varepsilon}\mathscr{E}[y_2](A)$.*

Note that the condition $d_{\mathscr{Y}}(y_1^m, y_2^m) \leq \delta$ is a generalization of the notion of adjacency to arbitrary metric spaces that we adopt in this paper. We now have the following lemma on the connection between the notions of differential privacy and max-divergence introduced above:

Lemma 43 (Differential privacy and max-divergence)**.** *An estimator $\mathscr{E}$ is ε-differentially private iif $D_{\max}(\mathscr{E}[y_1], \mathscr{E}[y_2]) \leq \varepsilon$ for any $y_1, y_2 \in \mathscr{Y}$ with $d_{\mathscr{Y}}(y_1, y_2) \leq \delta$.*

Proof. Clearly, if for any $y_1, y_2 \in \mathscr{Y}$ with $d_{\mathscr{Y}}(y_1, y_2) \leq \delta$, we have $D_{\max}(\mathscr{E}[y_1], \mathscr{E}[y_2]) \leq \varepsilon$, then:

$$\begin{aligned} \varepsilon \geq D_{\max}(\mathscr{E}[y_1], \mathscr{E}[y_2]) &= \sup_{x \in \mathscr{X}} \left| \log\left(\frac{\rho_1(x)}{\rho_2(x)}\right)\right| \\ &\geq \left| \log\left(\frac{\rho_1(x)}{\rho_2(x)}\right)\right|. \end{aligned}$$

This implies that for any $x \in \mathscr{X}$, we have $\rho_1(x) \leq e^{\varepsilon}\rho_2(x)$, from which differential privacy follows. Now, for any $A \subseteq \mathscr{X}$, we have $\mathscr{E}[y_1^m](A) = \int_{\mathscr{A}} \rho_1(x)\,\mathrm{dvol} \leq \int_{\mathscr{A}} e^{\varepsilon}\rho_2(x)\,\mathrm{dvol} = e^{\varepsilon}\int_{\mathscr{A}} \rho_2(x)\,\mathrm{dvol} = e^{\varepsilon}\mathscr{E}[y_2^m](A)$, which implies that $\mathscr{E}$ is ε-differentially private. It is easy to verify the forward implication holds. □

Thus, ε-differential privacy essentially imposes an upper bound on the sensitivity of the estimate generated by $\mathscr{E}$ (in the sense of the max-divergence $D_{\max}$), to the measurement.

5.3 Optimization-based state estimation

We now begin by addressing the state estimation problem for the autonomous system Σ, and develop a recursive moving-horizon estimator for it.

5.3.1 Full-Information Estimation (FIE)

Let $\{y_k\}_{k\in\{0\}\cup\mathbb{N}}$ be a sequence of measurements generated by the system Σ. Let $\{0,\ldots,T\}$ be a time horizon such that $T \geq T_0$, the minimum horizon length of the system Σ, and denote $\mathbf{y}_{0:T} = (y_0,\ldots,y_T)$. The problem of estimation essentially aims at characterizing ${\Sigma_T}^{-1}(\mathbf{y}_{0:T})$, which is an inverse problem, and optimal estimation formulates this problem as an optimization. Assumptions 13, and 15, on Lipschitz continuity and strong local observability, respectively, ensure that the inverse problem is locally well-posed as in Definition 14.

To formulate the inverse problem as an optimization, consider a convex function $J_T(\mathbf{y}_{0:T},\cdot) : \mathbb{Y}^{T+1} \to \mathbb{R}_{\geq 0}$ such that $J_T(\mathbf{y}_{0:T},\xi) = 0$ if and only if $\xi = \mathbf{y}_{0:T}$. Moreover, let $\lim_{T\to\infty} J_T(\mathbf{y}_{0:T},\Sigma_T(x)) = \infty$ if $x \notin {\Sigma_T}^{-1}(\mathbf{y}_{0:T})$ for $T \geq T_0$. Now, the problem of interest becomes:

$$x_0 \in \arg\min_{x\in\mathbb{X}} J_T(\mathbf{y}_{0:T},\Sigma_T(x)). \tag{5.3}$$

In the above, $\mathbf{y}_{0:T}$ is the data in the estimation problem, which is given. Since the objective is to solve the original inverse problem, and we would like to use gradient descent-based methods, we would like for every local minimizer of $J_T(\mathbf{y}_{0:T},\Sigma_T(x))$ to belong to the set ${\Sigma_T}^{-1}(\mathbf{y}_{0:T})$, or, in other words, that every local minimizer is also global. We therefore make the following additional assumption on the system Σ and the choice of J_T. For a conciseness of notation, in the following assumption and lemma, we let $J_T(\cdot) = J_T(\mathbf{y}_{0:T},\cdot)$, suppressing the data $\mathbf{y}_{0:T}$ in the notation where useful, and is understood from context.

Assumption 16 (Lower semicontinuity of sublevel sets)**.** *We assume that, for all $T \geq T_0$, the convex function $J_T : \mathbb{Y}^{T+1} \to \mathbb{R}$ is such that the set-valued map $\mathscr{S}_{\mathbb{X}}(\alpha) = {\Sigma_T}^{-1}\left(\mathscr{S}^{J_T}_{\mathbb{Y}^{T+1}}(\alpha) \cap \Sigma_T(\mathbb{X})\right)$ is lower semicontinuous, where $\mathscr{S}^{J_T}_{\mathbb{Y}^{T+1}}(\alpha) = \{\xi \in \mathbb{Y}^{T+1} | J_T(\xi) \leq \alpha\}$.*

The above assumption ensures that the function $J_T(\mathbf{y}_{0:T}, \Sigma_T(\cdot))$ satisfies the condition for the local minimizers to be global (Theorem 1 from [122]). The following lemma provides a sufficient condition for it to hold.

Lemma 44 (Second-order sufficient condition for lower semicontinuity)**.** *Assumption 16 holds if for any $x \in \mathbb{X}$ such that $\nabla\left(J_T(\mathbf{y}_{0:T}, \Sigma_T(x))\right) = 0$ we have $J_T(\mathbf{y}_{0:T}, \Sigma_T(x)) = 0$, or the following condition holds when $J_T(\mathbf{y}_{0:T}, \Sigma_T(x)) \neq 0$ for any $v \in \mathbb{R}^{d_X}$, $v \neq 0$:*

$$\frac{\left\langle \nabla^2 \Sigma_T[v,v](x), \nabla J_T \Big|_{\Sigma_T(x)} \right\rangle}{\|\nabla \Sigma_T[v]\|^2} \leq -\lambda_{\max}\left(\operatorname{Hess} J_T \Big|_{\Sigma_T(x)}\right),$$

where $\operatorname{Hess} J_T$ *is the Hessian of* J_T*.*

The final inequality in Lemma 44 merely states that those critical points at which the cost function does not reach the global minimum value are local maximizers.

We are now ready to present the following theorem that establishes the equivalence between the inverse problem of characterizing the set ${\Sigma_T}^{-1}(\mathbf{y}_{0:T})$ and the optimization (5.3).

Theorem 13 (Inverse as minimizer)**.** *Under Assumptions 15 and 16, for any $T \geq T_0$, it holds that $z \in {\Sigma_T}^{-1}(\mathbf{y}_{0:T})$ if and only if z is a minimizer of $J_T(\mathbf{y}_{0:T}, \Sigma_T(\cdot))$.*

Proof. If $z \in {\Sigma_T}^{-1}(\mathbf{y}_{0,T})$, we have that $h \circ f_0^k(z) = y_k$ for all $k \in \{0, \ldots, T\}$. It now follows that $J_T(\mathbf{y}_{0:T}, \Sigma_T(z)) = 0$. Since, $J_T(\mathbf{y}_{0:T}, \Sigma_T(z)) \geq 0$ by definition, we infer that z is a global minimizer of $J_T(\mathbf{y}_{0:T}, \Sigma_T(\cdot))$.

Suppose that z is a local minimizer of $J_T(\mathbf{y}_{0:T}, \Sigma_T(\cdot))$. By Assumption 16 and Theorem 1 in [122], we get that the local minima of $J_T(\mathbf{y}_{0:T}, \Sigma_T(\cdot))$ are also global, which implies that $J_T(\mathbf{y}_{0:T}, \Sigma_T(z)) = 0$, and therefore $\Sigma_T(z) = \mathbf{y}_{0:T}$. $\square$

Theorem 13 suggests that the state estimates for the system Σ can be obtained by minimizing $J_T(\mathbf{y}_{0:T}, \Sigma_T(\cdot))$ over a horizon of length $T \geq T_0$. This is also called the full information estimation (FIE) problem in the optimal state estimation literature [73, 102], as it works with the entire sequence of output measurements over the horizon $\{0, \ldots, T\}$.

Now, from Assumption 15 and Theorem 13, we have that ${\Sigma_T}^{-1}(\mathbf{y}_{0:T})$ is a set of isolated points which are minimizers of $J_T(\mathbf{y}_{0:T}, \Sigma_T(\cdot))$. It then follows that ${\Sigma_T}^{-1}(\mathbf{y}_{0:T})$ is the set of stable fixed points of the negative gradient vector field of $J_T(\mathbf{y}_{0:T}, \Sigma_T(\cdot))$. We let $\mathscr{C}_0$ be the basin of attraction of this set. Moreover, we note that $f^k({\Sigma_T}^{-1}(\mathbf{y}_{0:T}))$ is the set of stable fixed points of the negative gradient vector field of $J_T\left(\mathbf{y}_{k:k+T}, f^k \circ \Sigma_T(\cdot)\right)$, and we let $\mathscr{C}_k$ be the basin of attraction of ${\Sigma_T}^{-1}(\mathbf{y}_{k:k+T})$. We have used above the fact that ${\Sigma_T}^{-1}(\mathbf{y}_{k:k+T}) = f_0^k({\Sigma_T}^{-1}(\mathbf{y}_{0:T}))$, which follows from the definition of strong local observability.

We now lift the FIE problem (5.3) to the space of probability measures over $\mathbb{X}$, as a minimization in expectation of the estimation objective function:

$$\mu_0 \in \arg\min_{\mu \in \mathscr{P}(\mathbb{X})} \mathbb{E}_\mu \left[J_T(\mathbf{y}_{0:T}, \Sigma_T(\cdot)) \right]. \tag{5.4}$$

The above formulation allows us to capture information about the (probably many) optimal estimates through a probability measure μ_0, and help encode distributional constraints, which will be considered in a forthcoming publication.

In the following, we develop recursive moving-horizon estimators that generate sequences $\{\mu_k\}_{k \in \mathbb{N}}$ of probability measures in $\mathscr{P}(\mathbb{X})$ as estimates. We then obtain practically implementable estimators using Monte Carlo methods to sample from the measures μ_k.

5.3.2 Moving-Horizon Estimation (MHE)

In the previous section, we presented a formulation of the full information estimation (FIE) problem for the autonomous system Σ, which uses the entire measurement sequence over a horizon of length $T \geq T_0$. However, the minimum horizon length T_0 may be large, which would make the estimation computationally intensive. Moreover, we would like to progressively assimilate the incoming measurements online. We therefore adopt a moving-horizon estimation method which, at any time instant $k+N$, uses the output measurements from the horizon $\{k+1,\dots,k+N\}$ (of length $N < T_0$), and the state estimate at the time instant $k-1$, to obtain the state estimate at instant k, recursively.

We let $G_k^N(z) = J_{N-1}\left(\mathbf{y}_{k+1:k+N}, \Sigma_N(z)\right)$ be the objective function over the horizon $\{k+1,\dots,k+N\}$, at the time instant $k+N$, where $\mathbf{y}_{k+1:k+N} = (y_{k+1},\dots,y_{k+N})$.

Assumption 17 (Moving-horizon cost)**.** *We make the following assumptions on the cost function G_k^N:*

1. *the cost G_k^N is l-smooth,*
2. *it holds that $|G_{k+1}^N(f_0(z)) - G_k^N(z)| \leq L\|\nabla G_k^N(z)\|^2$,*
3. *the previous constants are such that $lL \leq \frac{1}{2}$,*
4. *for any two δ-adjacent measurements $\mathbf{y}, \tilde{\mathbf{y}} \in \mathbb{Y}^{T+1}$, such that $\|\mathbf{y}-\tilde{\mathbf{y}}\| \leq \delta$ and with corresponding costs G_k^N and $\widetilde{G}_k^N$, for $k \in \{0,\dots,T\}$ and $N \leq T-k$, we have $\|\nabla(G_k^N - \widetilde{G}_k^N)(x)\| \leq l\delta$ for all $x \in \mathbb{X}$.*

We now formulate the general moving-horizon estimation method as follows:

$$\begin{aligned} &\mu_k \in \arg\min_{\mu \in \mathscr{P}(\mathbb{X})} D(\mu, f_{0\#}\mu_{k-1}) + \eta \mathbb{E}_\mu\left[G_k^N\right], \\ &\text{given } \ \mu_0 \in \mathscr{P}(\mathbb{X}), \end{aligned} \tag{5.5}$$

where $D : \mathscr{P}(\mathbb{X}) \times \mathscr{P}(\mathbb{X}) \to \mathbb{R}_{\geq 0}$ is a placeholder for a metric, divergence or transport cost on $\mathscr{P}(\mathbb{X})$. We obtain implementable observers from the above formulation by sampling from the measures, by Monte Carlo methods. As discussed in the ensuing sections, using the 2-Wasserstein distance W_2 yields the more familiar MHE formulation, whereas with the KL-divergence we obtain a moving-horizon particle filter. Hence, this formulation is proposed as a distributional unifying framework for moving-horizon estimation, where different estimators are generated by different choices of D.

We now introduce the following asymptotic stability notion for estimators that will be used in investigating the properties of the estimators we design.

Definition 19 (Asymptotic stability of state estimator)**.** *We call an estimator of the form* (5.5) *an asymptotically stable observer for the system* Σ *if the sequence of estimates* $\{\mu_k\}_{k\in\mathbb{N}}$ *is such that* $\lim_{k\to\infty} \mu_k({\Sigma_T}^{-1}(\mathbf{y}_{k:k+T})) = 1$ *for* $T \geq T_0$.

5.4 A W_2-Moving-Horizon Estimator

In this section, we derive a moving-horizon estimator, which we refer to as the W_2-MHE, to generate a sequence of probability distributions $\{\mu_k\}_{k\in\mathbb{N}}$. This is based on the one-step minimization scheme of [107] in $\mathscr{P}(\mathbb{X})$ w.r.t. the Wasserstein metric W_2, which we extend to the moving-horizon setting. For every $k > 0$, consider:

$$\begin{aligned} &\mu_k \in \arg\min_{\mu\in\mathscr{P}(\mathbb{X})} \frac{1}{2} W_2^2(\mu, f_{0\#}\mu_{k-1}) + \eta \mathbb{E}_\mu \left[G_k^N\right], \\ &\text{given } \ \mu_0 \in \mathscr{P}(\mathbb{X}). \end{aligned} \tag{5.6}$$

We let $\mathscr{K}_k$ be the support of μ_k, with $\mathscr{K}_0 \subseteq \mathscr{C}_0$, where $\mathscr{C}_0$ is as defined earlier in Section 5.3.1.

5.4.1 Sample update scheme for W_2-MHE

We now derive a sample update scheme for W_2-MHE, which also yields an implementable filter for the W_2-MHE formulation.

We note that any local minimizer μ_k of (5.6) is a critical point of the objective functional, and, therefore, it satisfies:

$$\begin{aligned} c &= \frac{\delta}{\delta\mu}\left[\left(\frac{1}{2}W_2^2(\mu, f_{0\#}\mu_{k-1}) + \eta\mathbb{E}_\mu\left[G_k^N\right]\right)\right]\Bigg|_{\mu=\mu_k} \\ &= \phi_k + \eta G_k^N, \end{aligned}$$

where ϕ_k is the Kantorovich potential [107] associated with the transport from μ_k to $f_{0\#}\mu_{k-1}$, and c is a constant (from the constraint $\int_{\mathbb{X}} d\mu(x) = 1$, for $\mu \in \mathscr{P}(\mathbb{X})$, due to which the first variation is defined up to an additive constant). From the above equation, we now obtain:

$$\nabla\phi_k(x) + \eta\nabla G_k^N(x) = 0.$$

The gradient of the Kantorovich potential ϕ_k defines the deterministic optimal transport map T_k (note that this notation is not to be confused with that of the time horizon T) w.r.t. the W_2-distance from μ_k to $f_{0\#}\mu_{k-1}$, which determines $\nabla\phi_k(x) = x - T_k^{-1}(x)$ (where $\mu_k = T_{k\#}f_{0\#}\mu_{k-1}$). We therefore get:

$$x = T_k^{-1}(x) - \eta\nabla G_k^N(x). \tag{5.7}$$

The above equation allows us to design an implementable filter for the W_2-MHE (5.6). We let $z_k \sim \mu_k$, that is, $z_k \in \mathscr{K}_k$ is sampled from the distribution μ_k. From (5.7), it holds that $z_k = T_k^{-1}(z_k) - \eta\nabla G_k^N(z_k)$. Since $(T_k^{-1})_\#\mu_k = f_{0\#}\mu_{k-1}$, we let $T_k^{-1}(z_k) = f_0(z_{k-1})$, a sample of the distribution $f_{0\#}\mu_{k-1}$, and we obtain the following recursive estimator:

$$z_k = f_0(z_{k-1}) - \eta\nabla G_k^N(z_k), \quad k > 0. \tag{5.8}$$

We now note that the estimate z_k in (5.8) corresponds to a critical point of the following minimizing movement scheme:

$$
\begin{aligned}
& z_k \in \arg\min_z \frac{1}{2}\|z - f_0(z_{k-1})\|^2 + \eta G_k^N(z), \quad k > 0, \\
& z_0 \sim \mu_0 \in \mathscr{P}(\mathbb{X}).
\end{aligned}
\tag{5.9}
$$

Lemma 45 (Strong convexity)**.** *For $\eta < l^{-1}$, the objective function in* (5.9) *is strongly convex, and therefore* $\mathrm{prox}_{\eta G_k^N}(f_0(x))$ *is a singleton for any* $x \in \mathbb{X}$.

Proof. Let $\Theta(z) = \frac{1}{2}\|z - f_0(\tilde{z})\|^2 + \eta G_k^N(z)$. We have $\nabla\Theta(z_1) - \nabla\Theta(z_2) = z_1 - z_2 + \eta\left(\nabla G_k^N(z_1) - \nabla G_k^N(z_2)\right)$. It now follows that $\langle \nabla\Theta(z_1) - \nabla\Theta(z_2), z_1 - z_2\rangle = \|z_1 - z_2\|^2 + \eta\left\langle \nabla G_k^N(z_1) - \nabla G_k^N(z_2), z_1 - z_2\right\rangle$. From Assumption 17-(1), on the moving-horizon cost, we now get $\langle \nabla\Theta(z_1) - \nabla\Theta(z_2), z_1 - z_2\rangle \geq (1 - \eta l)\|z_1 - z_2\|^2$, and since $\eta l < 1$, we infer that Θ is strongly convex, and therefore has a unique minimizer. Thus, $\mathrm{prox}_{\eta G_k^N}(f_0(\tilde{z})) = \arg\min_z \Theta(z)$ is a singleton. □

We note that the minimization (5.9) defines a proximal mapping w.r.t. the Euclidean metric, which we represent in a compact form using the proximal operator as:

$$
\begin{aligned}
& z_k = \mathrm{prox}_{\eta G_k^N}(f_0(z_{k-1})), \quad k > 0, \\
& z_0 \sim \mu_0 \in \mathscr{P}(\mathbb{X}),
\end{aligned}
\tag{5.10}
$$

where $\mathrm{supp}(\mu_0) = \mathscr{K}_0 \subseteq \mathscr{C}_0$.

5.4.2 Asymptotic stability of W_2-MHE

We present the asymptotic stability result for W_2-MHE in this section, before which we introduce the following assumption on positive invariance of the discrete-time dynamics defined by the map $\mathrm{prox}_{\eta G_k^N} \circ f$.

Assumption 18 (Positive invariance)**.** *We assume that there exists $\alpha > (1 - \sqrt{1 - 2lL})l^{-1}$ such that for all $\eta \in (0, \alpha)$, we have* $\mathrm{prox}_{\eta G_k^N}(f(\mathscr{C}_{k-1})) \subseteq \mathscr{C}_k$.

The above assumption ensures that under the discrete-time dynamics defined by the map $\text{prox}_{\eta G_k^N} \circ f$, any sequence starting in the basin of attraction $\mathscr{C}_0$ of ${\Sigma_T}^{-1}(\mathbf{y}_{0:T})$ remains within the basins of attraction $\mathscr{C}_k$ of ${\Sigma_T}^{-1}(\mathbf{y}_{k:k+T})$ at the subsequent instants of time $k \in \mathbb{N}$.

We are now ready to present the asymptotic stability result for W_2-MHE:

Theorem 14 (Asymptotic stability of W_2-MHE)**.** *The estimator* (5.6)*, under Assumptions 15 to 18, with a constant step size* $\eta \in \left(\frac{1-\sqrt{1-2lL}}{l}, \min\left\{\alpha, \frac{1}{l}\right\}\right)$*, is an asymptotically stable observer for the system* Σ*.*

Proof. By Assumption 17-(1), on the moving-horizon cost, and Lemma 13, we have:

$$
\begin{aligned}
&|G_k^N(f_0(z_{k-1})) - G_k^N(z_k) - \langle \nabla G_k^N(z_k), f_0(z_{k-1}) - z_k \rangle| \\
&\leq \frac{l}{2} \|f_0(z_{k-1}) - z_k\|^2.
\end{aligned}
$$

Substituting from (5.8) into the above, we get:

$$
\begin{aligned}
&|G_k^N(f_0(z_{k-1})) - G_k^N(z_k) - \eta \|\nabla G_k^N(z_k)\|^2| \\
&\leq \eta^2 \frac{l}{2} \|\nabla G_k^N(z_k)\|^2.
\end{aligned}
$$

It now follows that:

$$
G_k^N(z_k) \leq G_k^N(f_0(z_{k-1})) - \eta \left(1 - \frac{l}{2}\eta\right) \|\nabla G_k^N(z_k)\|^2.
$$

From Assumption 17-(2), on the moving-horizon cost, we have:

$$
\begin{aligned}
G_k^N(z_k) \leq G_{k-1}^N(z_{k-1}) &+ L\|\nabla G_{k-1}^N(z_{k-1})\|^2 \\
&- \eta \left(1 - \frac{l}{2}\eta\right) \|\nabla G_k^N(z_k)\|^2.
\end{aligned}
$$

Summing the above inequality from $k=1$ to K, we get:

$$\eta\left(1-\frac{l}{2}\eta\right)\sum_{k=1}^{K}\|\nabla G_k^N(z_k)\|^2 - L\sum_{k=1}^{K}\|\nabla G_{k-1}^N(z_{k-1})\|^2$$
$$\leq G_0^N(z_0) - G_K^N(z_K).$$

From here, we obtain:

$$\left[\eta\left(1-\frac{l}{2}\eta\right)-L\right]\sum_{k=1}^{K}\|\nabla G_k^N(z_k)\|^2$$
$$\leq G_0^N(z_0) - G_K^N(z_K) + L\|\nabla G_0^N(z_0)\|^2$$
$$\leq G_0^N(z_0) + L\|\nabla G_0^N(z_0)\|^2.$$

Since $\eta \in \left(\frac{1-\sqrt{1-2lL}}{l}, \frac{1}{l}\right)$, we have that $\eta\left(1-\frac{l}{2}\eta\right)-L>0$ and therefore, taking limits in the previous inequality, we deduce that the series is summable. The latter implies that $\lim_{k\to\infty}\nabla G_k^N(z_k)=0$, and from (5.8), we have that $\lim_{k\to\infty}\|z_k - f(z_{k-1})\|=0$.

It now follows, by definition, from the above that:

$$\lim_{k\to\infty}\nabla G_k^{T+1}(z_k) = \lim_{k\to\infty}\nabla\left(J_T\left(\mathbf{y}_{k:k+T}, \Sigma_T(z_k)\right)\right)=0,$$

over a horizon of length $T+1$ (with $T \geq T_0$). We now have that the initial condition $z_0 \in \mathscr{K}_0 \subseteq \mathscr{C}_0$ and Assumption 18 ensure that $z_k \in \mathscr{C}_k$, the basin of attraction of $f^k\left({\Sigma_T}^{-1}(\mathbf{y}_{0:T})\right)$ and from the fact that $\lim_{k\to\infty}\nabla\left(J_T\left(\mathbf{y}_{k:k+T}, \Sigma_T(z_k)\right)\right)=0$, we infer that $\{z_k\}$ converges to the local minima of $J_T\left(\mathbf{y}_{k:k+T}, \Sigma_T(\cdot)\right)$. By Theorem 13, it now follows that $\{z_k\}$ converges to the set ${\Sigma_T}^{-1}(\mathbf{y}_{k:k+T})$. Therefore $\lim_{k\to\infty} d(z_k, {\Sigma_T}^{-1}(\mathbf{y}_{k:k+T}))=0$.

Moreover, since $\lim_{k\to\infty} d(z_k, {\Sigma_T}^{-1}(\mathbf{y}_{k:k+T})) = 0$ for all $z_0 \in \mathscr{K}_0$, it follows that $\lim_{k\to\infty}\mathscr{K}_k = {\Sigma_T}^{-1}(\mathbf{y}_{k:k+T})$. We know that $\text{supp}(\mu_k)=\mathscr{K}_k$, and therefore we get that $\lim_{k\to\infty}\mu_k\left({\Sigma_T}^{-1}(\mathbf{y}_{k:k+T})\right)=1$. □

5.4.3 Robustness of W_2-MHE

We now characterize the performance of the estimator (5.6) on the system Ω in (5.1). Since the true process and measurement noise sequences remain unknown, we are interested in the robustness properties of the estimator (5.11), in the form of an upper bound by the norms of the disturbance sequences on the estimation error.

We begin by constructing a reference estimator that recursively generates the estimate sequence, given the true disturbance sequences $\{w_k\}_{k\in\mathbb{N}}$ and $\{v_k\}_{k\in\mathbb{N}}$, as follows:

$$\begin{aligned} &\bar{\mu}_k \in \arg\min_{\mu\in\mathscr{P}(\mathbb{X})} \frac{1}{2}W_2^2(\mu, f_{0\#}\bar{\mu}_{k-1}) + \eta\mathbb{E}_\mu\left[\bar{G}_k^N\right], \\ &\text{given } \; \bar{\mu}_0 \in \mathscr{P}(\mathbb{X}). \end{aligned} \tag{5.11}$$

where, we employ for conciseness $\mathbf{w} \equiv \mathbf{w}_{k:k+N-1} = (w_k,\dots,w_{k+N-1})$ and $\mathbf{v} \equiv \mathbf{v}_{k+1:k+N} = (v_{k+1},\dots,v_{k+N})$, so that $\bar{G}_k^N(z) \equiv \bar{G}_k^N(z,\mathbf{w},\mathbf{v}) = J_{N-1}\left(\mathbf{y}_{k+1:k+N}, \Omega_{\mathbf{w}_{k:k+N-1}}(z) + \mathbf{v}_{k+1:k+N}\right)$. Note that $G_k^N = \bar{G}_k^N\big|_{\mathbf{w}=0,\mathbf{v}=0}$. We let $\bar{\mathscr{K}_k}$ be the support of $\bar{\mu}_k$, with $\bar{\mathscr{K}_0} \subseteq \bar{\mathscr{C}_0}$, where the definition of $\bar{\mathscr{C}_k}$ is similar to that of $\mathscr{C}_k$ but taking the noise $\{w_k\}$ and $\{v_k\}$ into account.

Assumption 19 (l-Smoothness w.r.t. disturbances)**.** *We assume that* $\|\nabla G_k^N(z) - \nabla\bar{G}_k^N(z)\| \leq l_w\|(\mathbf{w}_{k:k+N-1}, \mathbf{v}_{k+1:k+N})\|$ *for all* $z \in \mathbb{X}$.

Following the proof of Theorem 14, under the same set of underlying assumptions, we infer that the reference estimator (5.11) is almost surely an asymptotically stable observer for the system Ω, given a particular realization of the disturbances $\{w_k\}_{k\in\mathbb{N}}$ and $\{v_k\}_{k\in\mathbb{N}}$.

We now present the following theorem on the robustness of the estimator (5.6), characterized by a bound on the error in the estimates generated by (5.6) with respect to the estimates generated by the reference estimator (5.11):

Theorem 15 (Robustness of W_2-MHE)**.** *Under Assumptions 13, 15, 17, and 19, given the estimate sequences* $\{\mu_k\}_{k\in\mathbb{N}}$ *generated by* (5.6) *and* $\{\bar{\mu}_k\}_{k\in\mathbb{N}}$ *generated by the reference estimator* (5.11)*,* *with* $\mu_0 = \bar{\mu}_0$*, we have* $W_2(\mu_k, \bar{\mu}_k) \leq \frac{c_f^{(2)}}{c_f^{(1)}} W C_k + \frac{\eta l_w \sqrt{N}}{c_f^{(1)}}(W+V)C_k$*, for all* $k \in \mathbb{N}$*, where* $C_k = \sum_{\ell=1}^{k} (\frac{c_f^{(1)}}{1-\eta l})^{\ell}$.

Proof. The estimator (5.11) yields the following reference recursive scheme:

$$\bar{z}_k = f(\bar{z}_{k-1}, w_{k-1}) - \eta \nabla \bar{G}_k^N(\bar{z}_k), \tag{5.12}$$

where the above is derived similarly to the noiseless case. Let $\{z_k\}_{k\in\mathbb{N}}$ and $\{\bar{z}_k\}_{k\in\mathbb{N}}$ be the estimate sequences generated by (5.8) and (5.12) respectively, with $z_0 = \bar{z}_0$, for which we have:

$$\begin{aligned}
&\|z_k - \bar{z}_k\| \\
&= \|f_0(z_{k-1}) - f(\bar{z}_{k-1}, w_{k-1}) - \eta \nabla G_k^N(z_k) + \eta \nabla \bar{G}_k^N(\bar{z}_k)\| \\
&= \|f_0(z_{k-1}) - f_0(\bar{z}_{k-1}) + f_0(\bar{z}_{k-1}) - f(\bar{z}_{k-1}, w_{k-1}) \\
&- \eta \nabla G_k^N(z_k) + \eta \nabla G_k^N(\bar{z}_k) - \eta \nabla G_k^N(\bar{z}_k) + \eta \nabla \bar{G}_k^N(\bar{z}_k)\| \\
&\leq c_f^{(1)} \|z_{k-1} - \bar{z}_{k-1}\| + c_f^{(2)} \|w_{k-1}\| + \eta l \|z_k - \bar{z}_k\| \\
&\quad + l_w \eta \|(\mathbf{w}_{k:k+N-1}, \mathbf{v}_{k+1:k+N})\|,
\end{aligned}$$

where the final inequality follows from Assumptions 13, 17, and 19, on the several Lipschitz properties of f the gradient of G_k^N, and $\bar{G}_k^N$, respectively. Further, since $\eta l < 1$, we obtain from the above that:

$$
\begin{aligned}
&\|z_k - \bar{z}_k\| \\
&\leq \left(\frac{1}{1-\eta l}\right)\Big(c_f^{(1)}\|z_{k-1} - \bar{z}_{k-1}\| + c_f^{(2)}\|w_{k-1}\| \\
&\qquad\qquad + \eta l_w \|(\mathbf{w}_{k:k+N-1}, \mathbf{v}_{k+1:k+N})\|\Big) \\
&\leq \left(\frac{c_f^{(1)}}{1-\eta l}\right)^k \|z_0 - \bar{z}_0\| + \frac{c_f^{(2)}}{c_f^{(1)}} \sum_{\ell=1}^{k} \left(\frac{c_f^{(1)}}{1-\eta l}\right)^\ell \|w_{k-\ell}\| \\
&\quad + \frac{\eta l_w}{c_f^{(1)}} \sum_{\ell=1}^{k} \left(\frac{c_f^{(1)}}{1-\eta l}\right)^\ell \|(\mathbf{w}_{k-\ell+1:k-\ell+N}, \mathbf{v}_{k-\ell+2:k-\ell+N+1})\| \\
&\leq \frac{c_f^{(2)}}{c_f^{(1)}} W C_k + \frac{\eta l_w \sqrt{N}}{c_f^{(1)}} (W+V) C_k.
\end{aligned}
$$

We note that if $\frac{c_f^{(1)}}{1-\eta l} < 1$, we have that $\lim_{k\to\infty} C_k = \frac{c_f^{(1)}}{1-\eta l - c_f^{(1)}}$ is finite, and therefore, $\|z_k - \bar{z}_k\|$ is bounded as $k \to \infty$. We note here that even when $z_0 \neq \bar{z}_0$, the effect of this initial discrepancy vanishes as $k \to \infty$.

Now, let $T_k : \mathscr{K}_k \to \bar{\mathscr{K}_k}$ be a map such that for sequences $\{z_k\}$ and $\{\bar{z}_k\}$ generated by (5.8) and (5.12) respectively, with $z_0 = \bar{z}_0$, we have $T_k(z_k) = \bar{z}_k$. It then follows that $T_{k\#}\mu_k = \bar{\mu}_k$. Now, from the above, and by definition of the 2-Wasserstein distance, we have:

$$
\begin{aligned}
W_2(\mu_k, \bar{\mu}_k) &\leq \left(\int_{z\in\mathscr{K}_k} \|z - T_k(z)\|^2 d\mu_k(z)\right)^{\frac{1}{2}} \\
&\leq \left(\int_{z\in\mathscr{K}_k} \left|\frac{c_f^{(2)}}{c_f^{(1)}} W C_k + \frac{\eta l_w \sqrt{N}}{c_f^{(1)}} (W+V) C_k\right|^2 d\mu_k(z)\right)^{\frac{1}{2}} \\
&\leq \frac{c_f^{(2)}}{c_f^{(1)}} W C_k + \frac{\eta l_w \sqrt{N}}{c_f^{(1)}} (W+V) C_k.
\end{aligned}
$$

□

5.5 A KL-Moving-Horizon Estimator

In this section, we derive a moving-horizon estimator, which we refer to as KL-MHE, to generate a sequence of probability distributions $\{\mu_k\}_{k\in\mathbb{N}}$. Using the KL-divergence D_{KL} as the choice of divergence in the moving-horizon formulation (5.5), we obtain:

$$\begin{aligned} &\mu_k \in \arg\min_{\mu\in\mathscr{P}(\mathbb{X})} D_{\text{KL}}(\mu \| f_{0\#}\mu_{k-1}) + \eta \mathbb{E}_\mu \left[G_k^N\right], \\ &\text{given } \ \mu_0 \in \mathscr{P}(\mathbb{X}). \end{aligned} \tag{5.13}$$

We note that any local minimizer μ_k of (5.13) is a critical point of the objective functional, and, therefore, it satisfies:

$$c = \frac{\delta}{\delta\mu}\left[D_{\text{KL}}(\mu \| f_{0\#}\mu_{k-1}) + \eta \mathbb{E}_\mu \left[G_k^N\right]\right]\Big|_{\mu=\mu_k},$$

where c is a constant (from the constraint $\int_{\mathbb{X}} d\mu(x) = 1$, for $\mu \in \mathscr{P}(\mathbb{X})$, due to which the first variation is defined up to an additive constant). From the above, we get:

$$c = \log\left(\frac{\rho_k}{f_{0\#}\rho_{k-1}}\right)(x) + \eta G_k^N(x),$$

where for any $\ell \in \{0,1,\ldots\}$, ρ_ℓ is the density function corresponding to the measure μ_ℓ. Therefore, the corresponding recursive update scheme for the density function is given by:

$$\rho_k(x) = c_k\left(f_{0\#}\rho_{k-1}(x)\right)\exp\left(-\eta G_k^N(x)\right), \tag{5.14}$$

where c_k is the normalization constant. We note that the above is a particle filter formulation, with the horizon cost G_k^N defining the weighting function. Implementable filters are obtained by a Sequential Monte Carlo method, see [45]. We now present the asymptotic stability result for KL-MHE:

Theorem 16 (Asymptotic stability of KL-MHE)**.** *The estimator* (5.13)*, under Assumptions 13 to 16, is an asymptotically stable observer for the system* Σ*.*

Proof. We know that for any map $\mathscr{T}$ and measure μ, we have that $d\mathscr{T}_{\#}\mu(x) = d\mu\left(\mathscr{T}^{-1}(x)\right)$. It then follows from (5.14) that:

$$\rho_k(x) = c_k \rho_{k-1}(f_0^{-1}(x)) \exp\left(-\eta G_k^N(x)\right).$$

We now rewrite the above as:

$$\rho_k(f_0(x)) = c_k \rho_{k-1}(x) \exp\left(-\eta G_k^N(f_0(x))\right).$$

Repeating the above process k times, we obtain:

$$\rho_k(f_0^k(x)) = C_k \rho_0(x) \exp\left(-\eta \sum_{\ell=1}^{k} G_\ell^N(f_0^\ell(x))\right),$$

where $C_k = c_k c_{k-1} \ldots c_1$ is the normalization constant. If $x \notin {\Sigma_T}^{-1}(\mathbf{y}_{0:T})$, we have $\lim_{k\to\infty} \rho_k(f_0^k(x)) = 0$, since $\sum_{\ell=1}^{k} G_\ell^N(f_0^\ell(x)) \to \infty$ as $k \to \infty$ for all $x \notin {\Sigma_T}^{-1}(\mathbf{y}_{0:T})$ (by definition of the cost function, the sum diverges over an infinitely long horizon). Thus, we get:

$$\lim_{k\to\infty} \mu_k\left(f_0^k\left({\Sigma_T}^{-1}(\mathbf{y}_{0:T})\right)\right) = \lim_{k\to\infty} \mu_k\left({\Sigma_T}^{-1}(\mathbf{y}_{k:k+T})\right) = 1.$$

□

5.6 Differential privacy

In this section, we discuss the mechanism for encoding the desired level of differential privacy in moving-horizon estimators. We then apply this mechanism to the two estimators presented in the previous sections, the W_2-MHE and KL-MHE. We conclude the section with a discussion on differential privacy of the estimators over a time horizon.

Given the framework (5.5), we encode differential privacy by an entropic regularization of the estimation objective function, as follows:

$$\begin{aligned}\mu_k \in \arg\min_{\mu\in\mathscr{P}(\mathbb{X})} \Big[& s_k D(\mu, f_{0\#}\mu_{k-1}) + s_k \eta \mathbb{E}_\mu \left[G_k^N\right] \\ & -(1-s_k) S^{\mathscr{K}_k}(\mu)\Big], \\ \text{given } \ \mu_0 \in \mathscr{P}(\mathbb{X}), & \end{aligned} \tag{5.15}$$

where $s_k \in [0,1]$ is a tunable time-dependent parameter and $\mathscr{K}_k$ is the support of $f_{0\#}\mu_{k-1}$ (with $\mathscr{K}_0$ being the support of μ_0). Moreover, $S^A(\mu) = \int_A \rho \log(\rho)\,\mathrm{dvol}$, where $A \subset \mathbb{X}$ and $d\mu = \rho\,\mathrm{dvol}$. We note that when $s_k = 1$, the above formulation reduces to (5.5) and when $s_k = 0$, it is equivalent to an entropy maximization problem, yielding a uniform distribution over the set $f_0(\mathscr{K}_{k-1})$ as the solution. Clearly, the uniform distribution is insensitive to the measurements, and therefore offers maximum privacy, while being of no value to the estimation objective. The ensuing analysis in this section is directed at determining upper bounds on the parameter sequence $\{s_k\}_{k\in\mathbb{N}}$ such that the MHE offers ε-differential privacy. We rewrite the optimization problem (5.15) for $s_k \in (0,1]$ as follows:

$$\begin{aligned}\mu_k \in \arg\min_{\mu\in\mathscr{P}(\mathbb{X})} \Big[& D(\mu, f_{0\#}\mu_{k-1}) + \eta \mathbb{E}_\mu \left[G_k^N\right] \\ & -\left(\frac{1-s_k}{s_k}\right) S^{\mathscr{K}_k}(\mu)\Big], \\ \text{given } \ \mu_0 \in \mathscr{P}(\mathbb{X}), & \end{aligned} \tag{5.16}$$

Let $\mathbf{y}, \tilde{\mathbf{y}} \in \mathbb{Y}^{T+N+1}$ be two δ-adjacent measurement sequences as in Definition 18, over a horizon $\{0,\ldots,T+N\}$, such that $\|\mathbf{y}-\tilde{\mathbf{y}}\| \le \delta$ and let $\{\mu_k\}_{k\in\mathbb{N}}$ and $\{\tilde{\mu}_k\}_{k\in\mathbb{N}}$ be the sequences of estimates derived from (5.16). In the following, we determine conditions on $\{s_k\}_{k\in\mathbb{N}}$ that guarantee differential privacy for each of the estimators derived in previous sections.

5.6.1 Differentially private W_2-MHE

We now design a differentially private W_2-moving-horizon estimator. We begin by considering:

$$\begin{aligned} \mu_k \in \arg\min_{\mu \in \mathscr{P}(\mathbb{X})} \Bigg[& \frac{1}{2} W_2^2(\mu, f_{0\#}\mu_{k-1}) + \eta \mathbb{E}_\mu \left[G_k^N \right] \\ & - \left(\frac{1-s_k}{s_k} \right) S^{\mathscr{K}_k}(\mu) \Bigg], \\ \text{given } \ \mu_0 \in \mathscr{P}(\mathbb{X}), & \end{aligned} \tag{5.17}$$

for $s_k \in (0,1]$.

The following theorem provides a sufficient upper bound on s_T such that the entropy-regularized W_2-MHE in (5.17) is ε_T-differentially private at a time instant T.

Theorem 17 (Sensitivity of W_2-MHE)**.** *Given two δ-adjacent measurement sequences $\mathbf{y}, \tilde{\mathbf{y}} \in \mathbb{Y}^{T+N+1}$, under Assumption 17, we have that the estimates generated by* (5.17) *satisfy $D_{\max}(\mu_T, \tilde{\mu}_T) \leq \varepsilon_T$ if $s_T \leq \varepsilon_T \left(\varepsilon_T + c_f^T diam(\mathscr{K}_0) \left(\eta l \delta + c_f^T diam(\mathscr{K}_0) q(\delta) \right) \right)^{-1}$, where $q : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is a class-$\mathscr{K}$ function that satisfies $q(0) = 0$.*

Proof. Let G_k^N and $\widetilde{G}_k^N$ be the estimation objective functions at time instant k, corresponding to the measurement sequences $\mathbf{y}$ and $\widetilde{\mathbf{y}}$ respectively, and let μ_k and $\widetilde{\mu}_k$ be the respective estimated probability measures, with $\rho_k, \widetilde{\rho}_k$ the corresponding density functions. From (5.17), we get that for all $k \in \{0, \dots, T\}$, μ_k, being the local minimizer is also a critical point of the objective functional. We therefore obtain:

$$\phi_k(x) + G_k^N(x) + \left(\frac{1-s_k}{s_k} \right) \log(\rho_k(x)) = c,$$

where ϕ_k is the Kantorovich potential associated with the transport from μ_k to $f_{0\#}\mu_{k-1}$ and c is a constant. It now follows that:

$$\nabla\phi_k(x)+\nabla G_k^N(x)+\left(\frac{1-s_k}{s_k}\right)\nabla\log(\rho_k)(x)=0.$$

Similarly, we have:

$$\nabla\widetilde{\phi}_k(x)+\nabla\widetilde{G}_k^N(x)+\left(\frac{1-s_k}{s_k}\right)\nabla\log(\widetilde{\rho}_k)(x)=0.$$

Taking the difference between the above two equations:

$$\begin{aligned}\nabla\left[\log\left(\frac{\rho_k}{\widetilde{\rho}_k}\right)\right](x)=-\left(\frac{s_k}{1-s_k}\right)\Big[\nabla(\phi_k-\widetilde{\phi}_k)(x)\\ +\nabla(G_k^N-\widetilde{G}_k^N)(x)\Big].\end{aligned}$$

We have that $\nabla\phi_k(x)=x-T_k^{-1}(x)$, where $\mu_k=T_{k\#}(f_{0\#}\mu_{k-1})$. This implies that $\nabla(\phi_k-\widetilde{\phi}_k)(x)=-(T_k^{-1}(x)-\widetilde{T}_k^{-1}(x))$. However, $T_k^{-1}(x),\widetilde{T}_k^{-1}(x)\in f_0(\mathscr{K}_{k-1})=f_0^k(\mathscr{K}_0)$, and therefore $\|\nabla(\phi_k-\widetilde{\phi}_k)(x)\|\leq c_f^k\text{diam}(\mathscr{K}_0)q(\delta)$, for all $x\in f_0^k(\mathscr{K}_0)$ and some class-$\mathscr{K}$ function q. We let q characterize the dependence of ϕ on the measurement sequence, and we get that $\|\nabla(\phi_k-\widetilde{\phi}_k)(x)\|=0$ for all $x\in\mathbb{X}$, when $\delta=0$. Moreover, by Assumption 17, we get $\|\nabla(G_k^N-\widetilde{G}_k^N)(x)\|\leq l\delta$. Therefore, we obtain:

$$\left\|\nabla\left[\log\left(\frac{\rho_k}{\widetilde{\rho}_k}\right)\right]\right\|\leq\left(\frac{s_k}{1-s_k}\right)\left(c_f^k\text{diam}(\mathscr{K}_0)q(\delta)+l\delta\right). \tag{5.18}$$

We also have that for any $x\in f_0^k(\mathscr{K}_0)$:

$$\begin{aligned}\log\left(\frac{\rho_k}{\widetilde{\rho}_k}\right)(x)=&\log\left(\frac{\rho_k}{\widetilde{\rho}_k}\right)(\bar{x})\\ &+\int_0^1\nabla\left[\log\left(\frac{\rho_k}{\widetilde{\rho}_k}\right)\right](\gamma(t))\cdot\dot{\gamma}(t)dt,\end{aligned} \tag{5.19}$$

where $\gamma(0)=\bar{x}$ and $\gamma(1)=x$. Since ρ_k and $\widetilde{\rho}_k$ are continuous, with $\int_{f_0^k(\mathscr{K}_0)}(\rho_k-\widetilde{\rho}_k)=0$ (since

$\int_{f_0^k(\mathscr{K}_0)} \rho_k = \int_{f_0^k(\mathscr{K}_0)} \widetilde{\rho}_k = 1$), there exists an $\bar{x} \in f_0^k(\mathscr{K}_0)$ such that $\rho_k(\bar{x}) = \widetilde{\rho}_k(\bar{x})$, which implies that $\log\left(\frac{\rho_k}{\widetilde{\rho}_k}\right)(\bar{x}) = 0$. From (5.18) and (5.19), for a straight line segment γ, we therefore obtain:

$$\left|\log\left(\frac{\rho_k}{\widetilde{\rho}_k}\right)(x)\right| \leq \left(\frac{s_k}{1-s_k}\right)\left(c_f^k \text{diam}(\mathscr{K}_0) q(\delta) + l\delta\right) \times$$
$$c_f^k \text{diam}(\mathscr{K}_0),$$

where we have used the fact that $\int_0^1 |\dot{\gamma}(t)| dt = \|x - \bar{x}\| \leq \text{diam}(f_0^k(\mathscr{K}_0)) \leq c_f^k \text{diam}(\mathscr{K}_0)$. Thus, for $k = T$, we let:

$$\left|\log\left(\frac{\rho_T}{\widetilde{\rho}_T}\right)(x)\right| \leq \left(\frac{s_T}{1-s_T}\right)\left(c_f^T \text{diam}(\mathscr{K}_0) q(\delta) + l\delta\right) \times$$
$$c_f^T \text{diam}(\mathscr{K}_0)$$
$$\leq \varepsilon_T,$$

from which we obtain that:

$$s_T \leq \frac{\varepsilon_T}{\left(\varepsilon_T + c_f^T \text{diam}(\mathscr{K}_0)\left(\eta l \delta + c_f^T \text{diam}(\mathscr{K}_0) q(\delta)\right)\right)},$$

and since $\left|\log\left(\frac{\rho_T}{\widetilde{\rho}_T}\right)(x)\right| \leq \varepsilon_T$ for all $x \in f_0^T(\mathscr{K}_0)$, we have that $\sup_{x \in f_0^T(\mathscr{K}_0)} \left|\log\left(\frac{\rho_T}{\widetilde{\rho}_T}\right)\right| = D_{\max}(\mu_T, \widetilde{\mu}_T) \leq \varepsilon_T$. □

As noted earlier, Theorem 17 provides a sufficient upper bound on s_T for differential privacy of the estimate at T. The goal, however, is to guarantee the desired level of differential privacy over a time horizon $\{0, \ldots, T\}$. The key issue here is that the recursive update scheme of the estimator introduces a dependence between the estimates at different time instants. This essentially means that imposing an upper bound on sensitivity for the marginal distributions μ_k individually, without regard to the dependence between these distributions, may not be sufficient. Therefore, to guarantee the desired level of differential privacy over the time horizon, we must impose an upper bound on the sensitivity of the joint distribution $\sigma \in \mathscr{P}(\mathbb{X}^{T+1})$, where the estimates μ_k are the marginals of σ over $\mathbb{X}$.

The following theorem provides a sufficient upper bound on $\{s_k\}_{k=1}^{T}$ such that the entropy-regularized W_2-MHE in (5.17) is ε-differentially private over a time horizon $\{0,\ldots,T\}$.

Theorem 18 (Differentially private W_2-MHE)**.** *Given two δ-adjacent measurement sequences $\mathbf{y},\widetilde{\mathbf{y}} \in \mathbb{Y}^{T+N+1}$, under Assumption 17, we have that the estimates generated by* (5.17) *satisfy $D_{\max}(\sigma,\widetilde{\sigma}) \leq \varepsilon$ if $\sum_{k=1}^{T}\left(\frac{s_k}{1-s_k}\right)c_f^k \leq \frac{\varepsilon}{l\delta diam(\mathscr{K}_0)}$.*

Proof. Let G_k^N and $\widetilde{G}_k^N$ be the estimation objective functions at time instant k, corresponding to the measurement sequences $\mathbf{y}$ and $\widetilde{\mathbf{y}}$ respectively, and let σ and $\widetilde{\sigma}$ be the respective joint probability measures over the horizon $\{0,\ldots,T\}$. With a slight abuse of notation, we allow σ and $\widetilde{\sigma}$ to also denote the joint density function. We now have:

$$
\begin{aligned}
\sigma(x_0,x_1,\ldots,x_T) &= \rho_0(x_0)\sigma(x_1,\ldots,x_T|x_0) \\
&= \rho_0(x_0)\rho_1(x_1|x_0)\rho_2(x_2|x_1)\ldots\rho_T(x_T|x_{T-1}),
\end{aligned}
$$

where $\rho_k(x_k|x_{k-1})$ is the marginal density at x_k at time instant k, given that the distribution at time instant $k-1$ is concentrated at x_{k-1}. Moreover, we note that the W_2-MHE (5.17) yields a Markov process, which allows us to express $\rho_k(x_k|x_{k-1},\ldots,x_0) = \rho_k(x_k|x_{k-1})$. Now, $\rho_k(x_k|x_{k-1})$ is the density corresponding to the measure obtained by the following:

$$
\begin{aligned}
\mu_k \in \arg\min_{\mu\in\mathscr{P}(\mathbb{X})} \Bigg[&\frac{1}{2}W_2^2(\mu,\partial_{f_0(x_{k-1})}) + \eta\mathbb{E}_{\mu}\left[G_k^N\right] \\
&- \left(\frac{1-s_k}{s_k}\right) S^{\mathscr{K}_k}(\mu)\Bigg],
\end{aligned}
$$

where ∂_ξ is the Dirac measure concentrated at ξ. From the above, we get that for all $k \in \{0,\ldots,T\}$, μ_k, being the local minimizer is also a critical point of the objective functional.

Applying similar steps to those in the proof of Theorem 17, we obtain:

$$\left|\log\left(\frac{\rho_k}{\widetilde{\rho}_k}\right)(x|x_{k-1})\right| \leq \left(\frac{s_k}{1-s_k}\right) l\delta c_f^k \mathrm{diam}(\mathscr{K}_0).$$

Now, we have:

$$\begin{aligned}\left|\log\left(\frac{\sigma}{\widetilde{\sigma}}\right)(x_0,\ldots,x_T)\right| &\leq \sum_{k=1}^{T}\left|\log\left(\frac{\rho}{\widetilde{\rho}}\right)(x_k|x_{k-1})\right| \\ &\leq \sum_{k=1}^{T}\left(\frac{s_k}{1-s_k}\right) l\delta c_f^k \mathrm{diam}(\mathscr{K}_0).\end{aligned}$$

By taking

$$l\delta \mathrm{diam}(\mathscr{K}_0)\sum_{k=1}^{T}\left(\frac{s_k}{1-s_k}\right) c_f^k \leq \varepsilon,$$

we obtain the following inequality:

$$\sum_{k=1}^{T}\left(\frac{s_k}{1-s_k}\right) c_f^k \leq \frac{\varepsilon}{l\delta \mathrm{diam}(\mathscr{K}_0)},$$

and that $D_{\max}(\sigma,\widetilde{\sigma}) \leq \varepsilon$. □

We note that for a given ε, the upper bound on the sequence $\{s_k\}$ decreases with δ. In other words, guaranteeing ε-differential privacy w.r.t. measurement sequences that are farther apart requires the addition of more noise and a greater loss in estimation accuracy. This is because the weighting on the entropic regularization term in the estimation objective increases when s_k is reduced. The same is the case when ε is reduced for a given δ, which corresponds to a more stringent privacy requirement.

5.6.2 Differentially private KL-MHE

We now design a differentially private KL-moving-horizon estimator. We begin by considering the entropy-regularized KL-MHE formulation, given by:

$$\begin{aligned} \mu_k \in \arg\min_{\mu\in\mathscr{P}(\mathbb{X})} & \Big[D_{\mathrm{KL}}(\mu \| f_{0\#}\mu_{k-1}) + \eta\mathbb{E}_\mu\left[G_k^N\right] \\ & - \left(\frac{1-s_k}{s_k}\right) S^{\mathscr{K}_k}(\mu)\Big], \\ \text{given } & \mu_0 \in \mathscr{P}(\mathbb{X}), \end{aligned} \tag{5.20}$$

for $s_k \in (0,1]$. The corresponding recursive update scheme for (5.20) is given by:

$$\rho_k(x) = c_k\left(f_{0\#}\rho_{k-1}(x)\right)^{s_k} e^{-\eta s_k G_k^N(x)}, \tag{5.21}$$

which will be derived in the proof of Theorem 19 below.

The following theorem provides a sufficient upper bound on s_k such that the entropy-regularized KL-MHE in (5.20) is ε_T-differentially private at a time instant T, while ignoring the correlations between the estimates μ_k across time.

Theorem 19 (Sensitivity of KL-MHE)**.** *Given two δ-adjacent measurement sequences $\mathbf{y}, \widetilde{\mathbf{y}} \in \mathbb{Y}^{T+N+1}$, under Assumption 17, we have that the estimates generated by* (5.20) *satisfy* $D_{\max}(\mu_T, \widetilde{\mu}_T) \leq \varepsilon_T$ *if* $\sum_{k=1}^{T}\left(\prod_{i=k}^{T} s_i\right) \leq \varepsilon_T\left(2\eta\max_{k\in\{0,\dots,T\}}\left(\alpha_k + lc_f^k\delta diam(\mathscr{K}_0)\right)\right)^{-1}$, *where* $\alpha_k = \min_{\xi\in f_0^k(\mathscr{K}_0)}\left|\left(G_k^N - \widetilde{G}_k^N\right)(\xi)\right|$.

Proof. Let G_k^N and $\widetilde{G}_k^N$ be the estimation objective functions at time instant k, corresponding to the measurement sequences $\mathbf{y}$ and $\widetilde{\mathbf{y}}$ respectively, and let μ_k and $\widetilde{\mu}_k$ be the respective estimated probability measures, with $\rho_k, \widetilde{\rho}_k$ the corresponding density functions. From (5.20), we get that for all $k \in \{0,\dots,T\}$, μ_k, being the local minimizer is also a critical point of the objective functional. We therefore obtain:

$$\frac{\delta}{\delta\mu}\left[D_{\mathrm{KL}}(\mu\|f_{0\#}\mu_{k-1})+\eta\mathbb{E}_{\mu}\left[G_k^N\right]-\left(\frac{1-s_k}{s_k}\right)S^{\mathscr{K}_k}(\mu)\right]\Bigg|_{\mu_k}$$
$$=\bar{c}_k,$$

from which we derive that:

$$\log\left(\frac{\rho_k}{f_{0\#}\rho_{k-1}}\right)(x)+\eta G_k^N(x)+\left(\frac{1-s_k}{s_k}\right)\log\rho_k(x)=\bar{c}_k.$$

The above equation can be rewritten as follows:

$$\begin{aligned}\rho_k(x)&=c_k\left(f_{0\#}\rho_{k-1}(x)\right)^{s_k}e^{-\eta s_k G_k^N(x)}\\&=c_k\left(\rho_{k-1}(f_0^{-1}(x))\right)^{s_k}e^{-\eta s_k G_k^N(x)},\end{aligned}$$

where c_k is the normalization constant. We therefore obtain:

$$\rho_k(f_0(x))=c_k\left(\rho_{k-1}(x)\right)^{s_k}e^{-\eta s_k G_k^N(f_0(x))}.$$

Expanding the above, we get:

$$\rho_T(f_0^T(x))=C_T\left(\rho_0(x)\right)^{\Pi_{k=1}^T s_k}e^{-\eta\Sigma_{k=1}^T\left(\Pi_{i=k}^T s_i\right)G_k^N(f_0^k(x))},$$

where $C_T=c_1c_2\dots c_T$. Similarly, we have:

$$\widetilde{\rho}_T(f_0^T(x))=\widetilde{C}_T\left(\widetilde{\rho}_0(x)\right)^{\Pi_{k=1}^T s_k}e^{-\eta\Sigma_{k=1}^T\left(\Pi_{i=k}^T s_i\right)\widetilde{G}_k^N(f_0^k(x))},$$

where $\widetilde{C}_T=\widetilde{c}_1\widetilde{c}_2\dots\widetilde{c}_T$ and $\rho_0=\widetilde{\rho}_0$, as we assume that the estimator starts with the same initial μ_0.

From the above two equations, we obtain:

$$\log\left(\frac{\rho_T}{\widetilde{\rho}_T}\right)(f_0^T(x)) = \log\left(\frac{C_T}{\widetilde{C}_T}\right) - \eta\sum_{k=1}^{T}\left(\prod_{i=k}^{T}s_i\right)\left(G_k^N - \widetilde{G}_k^N\right)(f_0^k(x)).$$

The max-divergence between μ_T and $\widetilde{\mu}_T$ can be upper bounded now by:

$$\begin{aligned}
D_{\max}(\mu_T, \widetilde{\mu}_T) &= \sup_{x\in\mathscr{K}_0}\left|\log\left(\frac{\rho_T}{\widetilde{\rho}_T}\right)(f_0^T(x))\right| \\
&\leq \left|\log\left(\frac{C_T}{\widetilde{C}_T}\right)\right| + \sup_{x\in\mathscr{K}_0}\eta\sum_{k=1}^{T}\left(\prod_{i=k}^{T}s_i\right)\times \\
&\qquad \left|\left(G_k^N - \widetilde{G}_k^N\right)(f_0^k(x))\right| \\
&\leq 2\sup_{x\in\mathscr{K}_0}\eta\sum_{k=1}^{T}\left(\prod_{i=k}^{T}s_i\right)\left|\left(G_k^N - \widetilde{G}_k^N\right)(f_0^k(x))\right|,
\end{aligned}$$

where the final inequality is due to the following (note that we use the fact that $\rho = \tilde{\rho}$, as mentioned earlier):

$$\begin{aligned}
&\left|\log\left(\frac{C_T}{\widetilde{C}_T}\right)\right| \\
&= \left|\log\left(\frac{\int_{x\in\mathscr{K}_0}(\rho_0(x))^{\prod_{k=1}^T s_k} e^{-\eta\sum_{k=1}^T\left(\prod_{i=k}^T s_i\right)G_k^N(f_0^k(x))}}{\int_{x\in\mathscr{K}_0}(\widetilde{\rho}_0(x))^{\prod_{k=1}^T s_k} e^{-\eta\sum_{k=1}^T\left(\prod_{i=k}^T s_i\right)\widetilde{G}_k^N(f_0^k(x))}}\right)\right| \\
&\leq \sup_{x\in\mathscr{K}_0}\left|\log\left(\frac{e^{-\eta\sum_{k=1}^T\left(\prod_{i=k}^T s_i\right)G_k^N(f_0^k(x))}}{e^{-\eta\sum_{k=1}^T\left(\prod_{i=k}^T s_i\right)\widetilde{G}_k^N(f_0^k(x))}}\right)\right| \\
&\leq \sup_{x\in\mathscr{K}_0}\eta\sum_{k=1}^{T}\left(\prod_{i=k}^{T}s_i\right)\left|\left(G_k^N - \widetilde{G}_k^N\right)(f_0^k(x))\right|.
\end{aligned}$$

We now have, for all $k \in \{1, \dots, T\}$:

$$\left(G_k^N - \widetilde{G}_k^N\right)(f_0^k(x)) = \left(G_k^N - \widetilde{G}_k^N\right)(\xi_k) + \int_0^1 \nabla \left(G_k^N - \widetilde{G}_k^N\right)(\gamma_k(t)) \cdot \dot{\gamma}_k(t) dt,$$

where $\gamma_k(0) = \xi_k$ and $\gamma_k(1) = f_0^k(x)$. From Assumption 17, we have $\left\| \nabla \left(G_k^N - \widetilde{G}_k^N\right)(\xi) \right\| \le l\delta$. Moreover, let $\xi_k \in f_0^k(\mathscr{K}_0)$ such that $\left| \left(G_k^N - \widetilde{G}_k^N\right)(\xi_k) \right| = \min_{f_0^k(\mathscr{K}_0)} \left| \left(G_k^N - \widetilde{G}_k^N\right) \right| = \alpha_k$, and we obtain:

$$\left| \left(G_k^N - \widetilde{G}_k^N\right)(f_0^k(x)) \right| \le \alpha_k + l\delta \operatorname{diam}(f_0^k(\mathscr{K}_0)) \le \alpha_k + lc_f^k \delta \operatorname{diam}(\mathscr{K}_0).$$

This yields the following inequality:

$$\begin{aligned} & 2 \sup_{x \in \mathscr{K}_0} \eta \sum_{k=1}^T \left(\prod_{i=k}^T s_i \right) \left| \left(G_k^N - \widetilde{G}_k^N\right)(f_0^k(x)) \right| \\ & \le 2\eta \sum_{k=1}^T \left(\prod_{i=k}^T s_i \right) \left(\alpha_k + lc_f^k \delta \operatorname{diam}(\mathscr{K}_0) \right) \\ & \le 2\eta \max_k \left(\alpha_k + lc_f^k \delta \operatorname{diam}(\mathscr{K}_0) \right) \sum_{k=1}^T \left(\prod_{i=k}^T s_i \right). \end{aligned}$$

We now let:

$$2\eta \max_k \left(\alpha_k + lc_f^k \delta \operatorname{diam}(\mathscr{K}_0) \right) \sum_{k=1}^T \left(\prod_{i=k}^T s_i \right) \le \varepsilon_T,$$

which yields the bound

$$\sum_{k=1}^T \left(\prod_{i=k}^T s_i \right) \le \frac{\varepsilon_T}{2\eta \max_k \left(\alpha_k + lc_f^k \delta \operatorname{diam}(\mathscr{K}_0) \right)},$$

and we get $D_{\max}(\mu_T, \widetilde{\mu}_T) \leq \varepsilon_T$. $\square$

We note here that, in practice, with the choice of a sufficiently large domain $\mathscr{K}_0$, we can ensure that $\alpha_k = \min_{\xi \in f_0^k(\mathscr{K}_0)} \left| \left(G_k^N - \widetilde{G}_k^N \right)(\xi) \right| = 0$ for all $k \in \{0, \ldots, T\}$. This is owing to the fact that for a large enough $\mathscr{K}_0$, we will have $\min_{\xi \in f_0^k(\mathscr{K}_0)} \left(G_k^N - \widetilde{G}_k^N \right)(\xi) \leq 0 \leq \max_{\xi \in f_0^k(\mathscr{K}_0)} \left(G_k^N - \widetilde{G}_k^N \right)(\xi)$. Moreover, since the function $G_k^N - \widetilde{G}_k^N$ is continuous, there must therefore exist a point ξ^* such that $\left(G_k^N - \widetilde{G}_k^N \right)(\xi^*) = 0$.

As with the W_2-MHE, we now characterize the differential privacy of the KL-MHE over a horizon $\{0, \ldots, T\}$. We recall that the KL-MHE yields a sequence of distributions $\{\mu_k\}_{k=0}^T$ over the time horizon. Differential privacy over the horizon requires an upper bound on the sensitivity of the joint distribution σ over the horizon, where μ_k is the marginal of σ at the time instant k. As before, with a slight abuse of notation, letting σ also denote the joint density function, we have:

$$
\begin{aligned}
\sigma(x_0, x_1, \ldots, x_T) &= \rho_0(x_0)\sigma(x_1, \ldots, x_T | x_0) \\
&= \rho_0(x_0)\rho_1(x_1|x_0)\rho_2(x_2|x_1)\ldots\rho_T(x_T|x_{T-1}).
\end{aligned}
$$

From the above, we infer that to estimate the sensitivity of the joint density function, we must estimate the sensitivity of the conditionals $\rho_k(x_k|x_{k-1})$. The conditional $\rho_k(x_k|x_{k-1})$ at any time instant k, is obtained from the coupling between the marginal distributions μ_k and μ_{k-1}.

We now obtain an upper bound for the case where the marginals μ_k are independently coupled. In other words, we suppose that:

$$
\begin{aligned}
\sigma(x_0, x_1, \ldots, x_T) &= \rho_0(x_0)\sigma(x_1, \ldots, x_T | x_0) \\
&= \rho_0(x_0)\rho_1(x_1)\rho_2(x_2)\ldots\rho_T(x_T).
\end{aligned}
\tag{5.22}
$$

Theorem 20 (Differentially private KL-MHE)**.** *Given two δ-adjacent measurement sequences $\mathbf{y}, \widetilde{\mathbf{y}} \in \mathbb{Y}^{T+N+1}$, under Assumption 17 and the independent coupling* (5.22)*, we* *have that the estimates generated by* (5.20) *satisfy* $D_{\max}(\sigma, \widetilde{\sigma}) \leq \varepsilon$ *if* $\sum_{k=1}^{T} \sum_{l=1}^{k} \left(\prod_{i=l}^{k} s_i \right) \leq \varepsilon \left(2\eta \max_k \left(\alpha_k + l c_f^k \delta \mathit{diam}(\mathscr{K}_0) \right) \right)^{-1}$*, where* $\alpha_k = \min_{\xi \in f_0^k(\mathscr{K}_0)} \left| \left(G_k^N - \widetilde{G}_k^N \right)(\xi) \right|$.

Proof. Let G_k^N and $\widetilde{G}_k^N$ be the estimation objective functions at time instant k, corresponding to the measurement sequences $\mathbf{y}$ and $\widetilde{\mathbf{y}}$ respectively, and let σ and $\widetilde{\sigma}$ be the respective joint probability measures over the horizon $\{0, \ldots, T\}$. With a slight abuse of notation, we allow σ and $\widetilde{\sigma}$ to also denote the joint density function. From (5.22), we get:

$$\log \left(\frac{\sigma}{\widetilde{\sigma}} \right)(x_0, \ldots, x_T) = \sum_{k=1}^{T} \log \left(\frac{\rho_k}{\widetilde{\rho}_k} \right)(x_k),$$

which implies that:

$$D_{\max}(\sigma, \widetilde{\sigma}) \leq \sum_{k=1}^{T} D_{\max}(\mu_k, \widetilde{\mu}_k).$$

From the proof of Theorem 19 on the sensitivity of KL-MHE, we further get:

$$\begin{aligned} D_{\max}(\sigma, \widetilde{\sigma}) &\leq \sum_{k=1}^{T} D_{\max}(\mu_k, \widetilde{\mu}_k) \\ &\leq 2\eta \max_k \left(\alpha_k + l c_f^k \delta \mathrm{diam}(\mathscr{K}_0) \right) \sum_{k=1}^{T} \sum_{l=1}^{k} \left(\prod_{i=l}^{k} s_i \right). \end{aligned}$$

Therefore, it holds that $D_{\max}(\sigma, \widetilde{\sigma}) \leq \varepsilon$ if:

$$\sum_{k=1}^{T} \sum_{l=1}^{k} \left(\prod_{i=l}^{k} s_i \right) \leq \frac{\varepsilon}{2\eta \max_k \left(\alpha_k + l c_f^k \delta \mathrm{diam}(\mathscr{K}_0) \right)}.$$

□

5.7 Simulation results

In this section, we present results from numerical simulations of the estimators studied in this paper. The simulations were performed in MATLAB (version R2017a) on a 2.5 GHz Intel Core i5 processor.

We considered the following nonlinear discrete-time system:

$$\begin{aligned}
x_1(k+1) &= x_1(k) + \tau x_2(k), \\
x_2(k+1) &= x_2(k) - \tau \frac{x_1(k)}{1+|x_1(k)|^2+|x_2(k)|^2} + w_k, \\
y(k) &= x_1(k) + v_k,
\end{aligned}$$

with $\tau = 0.1$, w_k and v_k are i.i.d disturbances, sampled uniformly from the intervals $[-0.1, 0.1]$ and $[-0.15, 0.15]$ respectively.

We first present the simulation results for W_2-MHE. We ran 30 trials of the estimator (5.9) on the same measurement sequence, with randomly generated initial conditions and over a time horizon of length $T = 100$. The length of the moving-horizon was chosen to be $N = 10$. Figure 5.1 contains the plots of the mean of the estimates along with the true states. The root mean squared error (RMSE) for the mean state estimate sequences were found to be ${z_1}^{\text{RMSE}} = 0.0856$ and ${z_2}^{\text{RMSE}} = 0.0846$ for the estimates of x_1 and x_2, respectively. The average time for computing the state estimate through the minimization (5.9) using the *fminunc* function in MATLAB was observed to be $t_{\text{comp}} = 0.012 \pm 0.02s$.

We then implemented the estimator (5.13) with 30 samples, over a time horizon of length $T = 100$. The length of the moving-horizon was chosen to be $N = 10$. Figure 5.2 contains the plots of the mean of the estimates along with the true states. The root mean squared error (RMSE) for the mean state estimate sequences were found to be ${z_1}^{\text{RMSE}} = 0.1073$ and ${z_2}^{\text{RMSE}} = 0.1144$ for the estimates of x_1 and x_2, respectively. The average run-time for the minimization (5.13) by a resampling method was observed to be $t_{\text{comp}} = (4.8 \pm 0.4) \times 10^{-4}s$.

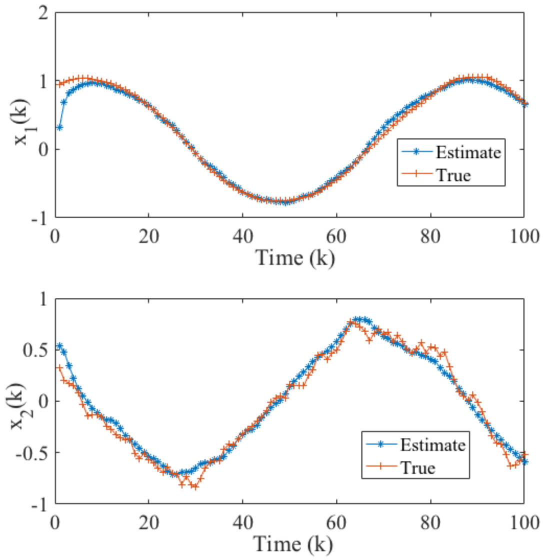

Figure 5.1. Mean state estimates from 30 trials of $W2$-MHE

In simulation, with 30 samples, we find that the W_2-MHE performs better with respect to the root mean squared error, while the KL-MHE is much faster. The performance of the KL-MHE is determined by the richness of the sample set and effectiveness of the resampling procedure, choices that depend on context and experience. In this manuscript, we did not attempt to investigate improvements in performance with respect to these choices. The performance of W_2-MHE does not necessarily improve with the richness of the sample set, but for systems for which ${\Sigma_T}^{-1}(\mathbf{y}_{0:T})$ is not a singleton, a richer sample set allows for a more complete characterization of the set of feasible estimates.

Figure 5.3 illustrates the typical trade-off between accuracy and privacy in moving-horizon estimation. We considered constant weights $s_k = s$ for the entropic regularization terms in (5.17) and (5.20). The values of s were chosen such that they satisfied the bounds specified in Theorems 18 and 20 for ε-differential privacy of the estimators over the horizon. In Figure 5.3, we plot the RMSE (for the estimates of the state x_1) for W_2-MHE, averaged over the 30 samples, specifying the accuracy, for different values of ε, the privacy parameter. We recall that a higher value of ε indicates a less stringent privacy requirement. We notice that the the accuracy of the

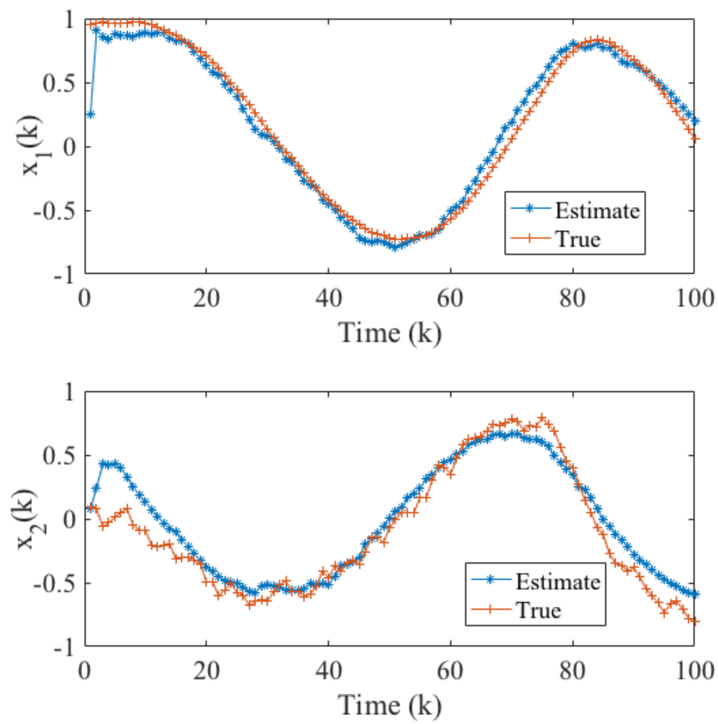

Figure 5.2. Mean state estimates from KL-MHE with 30 samples

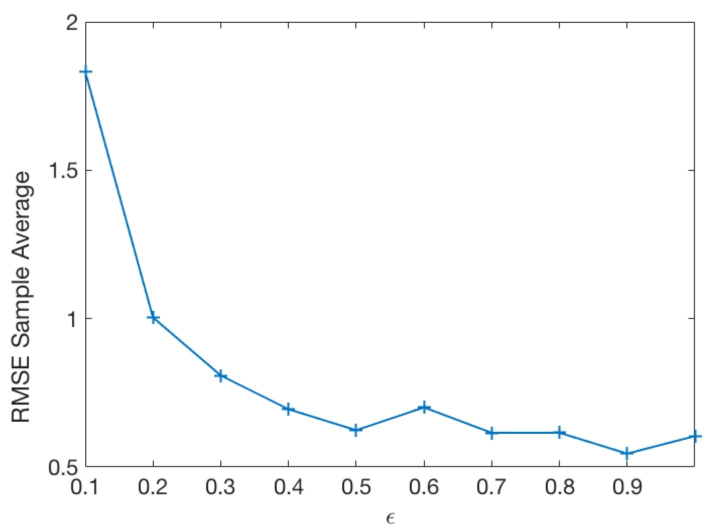

Figure 5.3. RMSE in estimates of state x_1 for W_2-MHE, averaged over 30 samples for different values of ε

estimators improves with an increase in the privacy parameter.

5.8 Summary

In this chapter, we laid out a unifying distributional framework for moving-horizon estimation. We clearly established the connection between the classical notion of strong local observability and the stability of moving-horizon estimation, for nonlinear discrete-time systems. We then proposed a differentially private mechanism based on entropic regularization and derived conditions under which ε-differential privacy is guaranteed at any given time instant and over time horizons.

The material in this chapter, in full, has been submitted for publication to the IEEE Transactions on Automatic Control and is under review. It may appear as *A Distributional Framework for Moving Horizon Estimation: Stability and Privacy Considerations*, V. Krishnan and S. Martínez. A preliminary version of the work appeared in the proceedings of the American Control Conference, Philadelphia, USA, July 2019 as *On Observability and Stability of Moving-Horizon Estimation in a Distributional Framework*, V. Krishnan and S. Martínez. The dissertation author was the primary investigator and author of these papers.

Chapter 6

Robustness of multi-agent networks

In this chapter, we study a critical node set identification problem for large-scale spatial networks with an associated weight-balanced Laplacian matrix. By considering a graph embedding technique, we reduce the problem to spatial networks with uniformly distributed nodes and nearest-neighbors communication topologies. Then we consider a special case of a hole-placement problem, which consists of identifying the optimal location of the center of a ball in the domain that minimizes the smallest positive eigenvalue of the Laplace operator for the residual domain. With the help of the Min-max theorem, we formulate our objective as an infinite-dimensional, non-convex and nested optimization problem. This limits our goal at the outset to achieving convergence to a local optimum. Since the solution is hard to obtain analytically, we develop an algorithmic approach to such problem. First, we consider the inner optimization or eigenvalue problem, whose KKT points include the eigenvalues of the Laplace operator. We then provide a closed-form expression for the projected gradient flow in a Banach space for this problem that converges to the set of KKT points. Exploiting further the special properties of these dynamics, we prove that the only locally asymptotically stable equilibrium point for the dynamics is the second eigenfunction of the Laplace operator. Moreover, since the other KKT points are saddle points that are non-degenerate, we infer almost global asymptotic stability of the second eigenfunction. Building on these results, we then design a novel hole-placement dynamics for the nested-optimization problem, and prove its local asymptotic stability to strict local minima. Finally, we provide a characterization of critical balls in the interior of the domain, and study the limiting case when its radius approaches zero. We conclude that the location of such critical nodes is at the nodal set of the second eigenfunction of the Laplace operator, which has an intuitive geometric interpretation in some cases.

6.1 Bibliographical comments

We first review some works that present combinatorial approaches to the problem of critical node identification. In [1, 10, 111, 116], the authors investigate the problem of identifying nodes whose deletion minimizes some network connectivity metric. An alternative approach to improving network robustness involves incorporating redundancy in the network by adding nodes and links, also called network augmentation [54]. In [47], the authors study the problem of network design as a function of the comparative costs of augmentation and defense against attack/failure.

The approximation of large networks by weighted graphs over a continuum set of infinite cardinality appears in previous literature. In this way, in [86] large networks are approximated by the so-called graphons, which result from the limit of convergent sequences of large dense graphs. Extending this idea to spatial networks, where the nodes are embedded in a domain $\Omega \in \mathbb{R}^N$, the nodes can be thought to be indexed by their positions $x \in \Omega$, and interactions restricted between the nearest spatial neighbors. Combining these notions in the context of network consensus dynamics, the object of interest is the continuum counterpart of the graph Laplacian, the Laplace operator on the domain. Theoretical results concerning the convergence of the graph Laplacian to the Laplace operator can be found in [17] and [18], which motivates the approach adopted in this paper.

There have been severals attempts to investigate problems linking the shape of a domain with the sequence of eigenvalues of the Laplace operator, for various boundary conditions, although those related to the critical subset identification are fewer in number. The work [69] contains an overview of the literature on extremum problems for eigenvalues of elliptic (e.g. Laplace) operators. In [79], the authors consider the problem of placing small holes in a domain to optimize the smallest Neumann eigenvalue of the Laplace operator (but with Dirichlet boundary condition on the hole).

6.2 Problem Formulation

We begin this section with the necessary background for setting up the critical node identification problem addressed in this paper. We begin by explaining how we employ a graph embedding along with a continuum approximation to go from the graph Laplacian to the Laplace operator on the domain. Using the Min-max theorem, we are then able to characterize the second eigenvalue of the Laplace operator corresponding to the algebraic connectivity of the graph. We finally point out to a connection to agreement algorithms in networked systems.

Let $G = (V, E)$ be a weight-balanced directed graph such that $|V| = n$, and w_{ij} be the edge weight corresponding to $(i, j) \in E$. A map $\mathbf{x} : V \to \Omega \subset \mathbb{R}^N$, is called a graph embedding ($N \ll n$ and Ω bounded), if $x_i = \mathbf{x}(i) \in \mathbb{R}^N$ is the (spatial) position assigned to node $i \in V$, and the map $\mathbf{x}$ preserves some proximity measure on the graph G. There exists a vast literature on graph embeddings [64, 110], of which we adopt the notion of the structure-preserving embedding. Starting with the unweighted, undirected graph corresponding to G (where the weighted directed edges in G are replaced by unweighted undirected edges), a structure preserving embedding can be constructed such that any node j which is a neighbor of i in the graph G is within a ball of radius h centered at at x_i in the embedding. Once the graph is embedded in $\Omega \subset \mathbb{R}^N$, we view the nodes V as having been sampled from an underlying distribution $\mu \in \mathscr{P}(\Omega)$ (with density function ρ, such that $d\mu = \rho\text{dvol}$). It is always possible to obtain the weighted adjacency matrix $W = [w_{ij}]$ of the digraph G as the discretization of a smooth weight function $\mathscr{W} : \Omega \times \Omega \to \mathbb{R}_{\geq 0}$, such that $w_{ij} = \mathscr{W}(x_i, x_j)$. The weight function $\mathscr{W}$ encodes the weights and directionality of the edges, and since the number of nodes V is finite, such a smooth weight function always exists.

Let $\varphi : \Omega \to \mathbb{R}$ be a real-valued function on Ω and $\phi^d : V \to \mathbb{R}$ such that $\phi_i^d = \phi^d(i) = \varphi(x_i)$. We define the $\mathscr{W}$-weighted average variation in φ around a point $x \in \Omega$, averaged over a ball $B_h(x)$ of radius $h > 0$ and centered at x as follows:

$$\frac{1}{\mu(B_h(x))} \int_{B_h(x)} \mathscr{W}(x,y)(\varphi(y) - \varphi(x)) d\mu(y).$$

We see next that the weighted Laplace operator on Ω can be obtained as the limit of a $\mathscr{W}$-weighted average variation as $h \to 0$. We first let $w(x) = \mathscr{W}(x,x)$ and $\nabla w(x) = \frac{1}{2}(\partial_1 \mathscr{W} + \partial_2 \mathscr{W})(x,x)$, and we obtain the following by means of a Taylor expansion:

$$\begin{aligned} \lim_{h \to 0} \frac{c}{h^2} \frac{1}{\mu(B_h(x))} \int_{B_h(x)} \mathscr{W}(x,y)(\varphi(y) - \varphi(x)) d\mu(y) \\ = \frac{1}{\rho} \nabla \cdot (w \rho \nabla \varphi), \end{aligned}$$

where c is a constant. The graph Laplacian matrix $L(G)$ corresponding to G can now be viewed as the discretization of the (negative) w-weighted Laplace operator $-\frac{1}{\rho} \nabla \cdot (w \rho \nabla)$. Alternatively, the w-weighted Laplace operator can be viewed as an approximation of $L(G)$, with closer approximations obtained as $n = |V| \to \infty$ and $h \to 0$.

In addition, approximating the Laplacian matrix $L(G)$ by the Laplace operator on Ω requires the specification of a boundary condition. This condition is obtained by observing that $\mathbf{1}_n \in \mathrm{Null}(L^\top(G))$, that is, $\left\langle \mathbf{1}, L(G)\phi^d \right\rangle = \mathbf{1}_n^\top L(G) \phi^d = 0$ for any ϕ^d. In the continuous setting, this translates into the Neumann boundary condition $\nabla \varphi \cdot \mathbf{n} = 0$ on $\partial \Omega$. This can be seen from an application of the Divergence theorem, that is, $\left\langle 1, \frac{1}{\rho} \nabla \cdot (w \rho \nabla \varphi) \right\rangle = \int_\Omega \frac{1}{\rho} \nabla \cdot (w \rho \nabla \varphi) d\mu = \int_{\partial \Omega} w \rho \nabla \varphi \cdot \mathbf{n} \, dS = 0$ (if $\nabla \varphi \cdot \mathbf{n} = 0$). Thus, the Neumann boundary condition is imposed as the natural boundary condition here.

Remark 8 (Problem reduction to uniformly spatially embedded graphs)**.** *Based on the previous considerations, and without loss of generality, in the following we focus on networks that are spatially embedded in an open bounded domain Ω according to a uniform distribution* *(the distribution μ is uniform above) and such that the underlying graph is undirected and unweighted. Note that the following derivations are analogous for the case of a non-uniform μ and weight-balanced directed graph: all results carry through by keeping the weights w and ρ in the weighted Laplace operator.*

The Laplace operator Δ with the Neumann boundary condition, has an infinite sequence of eigenvalues $0 = \lambda_1 \leq \lambda_2 \leq \ldots \leq \lambda_m \leq \ldots$, whose corresponding eigenfunctions $\{\psi_i\}_{i=1}^{\infty}$ form an orthonormal basis for $L^2(\Omega)$, [55]. Using the Min-max theorem [55] for the operator Δ, one can determine:

$$\lambda_2(\Omega) = \inf_{\psi \in \{\psi_1\}^{\perp}} \frac{\langle \psi, \Delta\psi \rangle_{L^2(\Omega)}}{\langle \psi, \psi \rangle_{L^2(\Omega)}}, \tag{6.1}$$

where $\{\psi_1\}^{\perp} = \{\psi \in H^1(\Omega) \,|\, \psi \neq 0, \ \int_{\Omega} \psi_1 \psi \, d\nu = 0\}$, and ψ_1 is constant, the eigenfunction corresponding to $\lambda_1 = 0$. This implies $\{\psi_1\}^{\perp} = \{\psi \in H^1(\Omega) \,|\, \int_{\Omega} \psi \, d\nu = 0\}$. Thus, using the Divergence theorem, applying the Neumann boundary condition, and normalizing the functions, we obtain an equivalent reformulation of (6.1) as:

$$\lambda_2(\Omega) = \inf_{\substack{\psi \in H^1(\Omega), \\ \int_{\Omega} \psi d\nu = 0, \\ \int_{\Omega} |\psi|^2 d\nu = 1}} \int_{\Omega} |\nabla\psi|^2 d\nu. \tag{6.2}$$

Remark 9 (Connection to agreement algorithms)**.** *The second eigenvalue is also of relevance to Laplacian-based agreement/consensus algorithms in networked systems, as it governs the convergence rate of these algorithms.*

We now define the notion of criticality adopted in this manuscript. We define critical nodes as those nodes in the graph whose removal results in the maximum deterioration in algebraic connectivity for the residual network, making them the most crucial nodes to be protected.

More precisely, this amounts to identifying a set $K^* \subset \Omega$ of given measure $|K^*| = c > 0$

such that $\lambda_2(\Omega \setminus K^*)$ is an infimum. The problem of identifying the critical nodes, K^*, can be formulated as:

$$K^* \in \arg \inf_{\substack{K \subset \Omega, \\ |K| = c}} \inf_{\substack{\psi \in H^1(\Omega \setminus K), \\ \int_{\Omega \setminus K} \psi dv = 0, \\ \int_{\Omega \setminus K} |\psi|^2 dv = 1}} \int_{\Omega \setminus K} |\nabla \psi|^2 dv.$$

We restrict the search to a class of subsets $K = B_r(x) = \{y \in \Omega \,|\, |y - x| < r\} \subset \Omega$, open balls of radius r (such that $|B_r(x)| = c$). This reduces the search space to $\tilde{\Omega}_r = \{x \in \Omega \,|\, \mathrm{dist}(x, \partial\Omega) > r\}$, and the problem is reformulated as:

$$x^* \in \arg \inf_{x \in \tilde{\Omega}_r} \inf_{\substack{\psi \in H^1(\Omega \setminus B_r(x)), \\ \int_{\Omega \setminus B_r(x)} \psi dv = 0, \\ \int_{\Omega \setminus B_r(x)} |\psi|^2 dv = 1}} \int_{\Omega \setminus B_r(x)} |\nabla \psi|^2 dv. \tag{6.3}$$

which we refer to as the hole-placement problem in the sequel.

Remark 10 (Generalization using multiple balls)**.** *We note that any compact subset $K \subset \Omega$ can be covered by a finite number m of open balls of a given radius r, and with arbitrary precision (as $r \to 0$ and $m \to \infty$). Given a finite collection $\{B_r(x_i)\}_{i=1}^m$ of open balls, we can then formulate the above optimization w.r.t. $(x_1, \ldots, x_m)$, the positions of the m open balls. For simplicity, we just focus on the one-ball case.*

6.3 Functional optimization to determine the most critical nodes

Here, we present our main results and algorithms to determine the most critical nodes in the network, in a functional optimization framework. To do this, we begin with the eigenvalue problem (6.2) (which is the inner optimization problem in (6.3)) for D, a fixed domain, and design a projected gradient flow to converge to a local minimizer of the problem. This algorithm will help us build subsequently the dynamics that can be employed to solve the full hole placement problem (6.3) in an algorithmic manner. The analysis of the projected gradient flow will also be instrumental in evaluating the properties of the second dynamics.

6.3.1 Projected gradient flow to determine $\lambda_2(\Omega)$

In what follows, we study the eigenvalue problem (6.2), characterize its critical points, construct and analyze a novel projected gradient flow to converge to the infimum. We write the optimization problem (for the smallest positive eigenvalue of the Laplace operator on a domain D with a C^1, Lipschitz boundary) as:

$$\begin{aligned}
&\inf_{\psi\in H^1(D)} \int_D |\nabla\psi|^2, \\
&\text{s.t} \quad \int_D |\psi|^2 = 1, \quad \int_D \psi = 0, \\
&\nabla\psi\cdot\mathbf{n} = 0 \text{ on } \partial D.
\end{aligned}$$

Let $\mathscr{S}_D = \{\psi \in H^1(D) \,|\, \int_D |\psi|^2 = 1, \int_D \psi = 0, \nabla\psi\cdot\mathbf{n} = 0 \text{ on } \partial D\}$ and $J(\psi) = \int_D |\nabla\psi|^2$. We can now express the above problem as $\inf_{\psi\in\mathscr{S}_D} J(\psi)$.

Lemma 46 (Minimizer of $J(\psi)$)**.** *The eigenfunctions of $\Delta(D)$ are the critical points of the functional $J(\psi)$, and the second eigenfunction ψ_2 of $\Delta(D)$ is the only minimizer of the functional $J(\psi)$ in $\mathscr{S}_D$. Moreover, the critical points of $J(\psi)$ are non-degenerate, i.e., the Hessian of $J(\psi)$ is non-singular at the critical points.*

Remark 11. *The content of this Lemma follows from the Min-max theorem [55], and we provide below an alternative proof. We explicitly compute the analytical expression for the Hessian of the objective function $J(\psi)$ in the proof of Lemma 46, which allows us to infer the non-degeneracy of the saddle points of $J(\psi)$ which is useful in establishing almost-global convergence of the projected gradient flow that follows.*

Proof. The first variation of the Lagrangian $L(\psi,\lambda,\chi) = J(\psi) + \lambda\left(1 - \int_D |\psi|^2\right) + \chi\int_D \psi$, at a critical point ψ^* is zero (where $\int_D |\psi|^2 = 1$ and $\int_D \psi = 0$ are the constraints, as $\psi \in \mathscr{S}_D$

and the Neumann boundary condition is assumed implicitly.) Thus, for any $\delta\psi \in T_{\psi^*}\mathscr{S}_D$ the tangent space of $\mathscr{S}_D$ at ψ^*, we have $\left\langle \frac{\delta L}{\delta \psi}, \delta\psi \right\rangle (\psi^*, \lambda^*, \chi^*) = 2\int_D \nabla\psi^* \cdot \nabla(\delta\psi) - 2\lambda^* \int_D \psi^* \delta\psi + \chi^* \int_D \delta\psi = -2\int_D (\Delta\psi^* + \lambda^*\psi^* - \frac{1}{2}\chi^*)\ \delta\psi = 0$, for any $\delta\psi$ (note that the Neumann boundary condition was used in obtaining the equation.) Additionally, we also have $\left\langle \frac{\partial L}{\partial \lambda}, \delta\lambda \right\rangle (\psi^*, \lambda^*, \chi^*) = 1 - \int_D |\psi^*|^2 = 0$, and $\left\langle \frac{\partial L}{\partial \chi}, \delta\chi \right\rangle (\psi^*, \lambda^*, \chi^*) = \int_D \psi^* = 0$. Thus, the critical points of the objective functional $\psi^* \in \mathscr{S}_D$ are characterized by:

$$\Delta\psi^* + \lambda^*\psi^* - \frac{1}{2}\chi^* = 0.$$

Integrating the previous equation over D and using the Neumann boundary condition, we obtain $\chi^* = 0$. Therefore, the critical points ψ^* satisfy:

$$\Delta\psi^* + \lambda^*\psi^* = 0. \tag{6.4}$$

Let $\psi(x, \varepsilon, \eta)$, $x \in D$, be a smooth two-parameter family of functions in $\mathscr{S}_D$ with $\int_D \psi(x, \varepsilon, \eta) = 0$ for all ε and η. The first variation of J at $\varepsilon = 0$, $\eta = 0$ is given by:

$$\left.\frac{\delta J}{\delta \varepsilon}\right|_{\substack{\varepsilon=0,\\ \eta=0}} (\psi) = 2\int_D \nabla\psi \cdot \partial_\varepsilon \nabla\psi = 2\int_D \nabla\psi \cdot \nabla(\partial_\varepsilon \psi).$$

We let $\partial_\varepsilon \psi|_{\varepsilon=0,\eta=0} = X$ and $\partial_\eta \psi|_{\varepsilon=0,\eta=0} = Y$. The second variation of J at $\varepsilon = 0$, $\eta = 0$ is given by:

$$\begin{aligned}
\frac{\delta^2 J}{\delta\eta\delta\varepsilon}(X,Y) &= 2\int_D \nabla(\partial_\eta \psi) \cdot \nabla(\partial_\varepsilon \psi) + 2\int_D \nabla\psi \cdot \nabla(\partial_{\eta\varepsilon}\psi) \\
&= 2\int_D \nabla(\partial_\eta \psi) \cdot \nabla(\partial_\varepsilon \psi) - 2\int_D \Delta\psi(\partial_{\eta\varepsilon}\psi) \\
&= 2\int_D \nabla X \cdot \nabla Y - 2\int_D \Delta\psi(\partial_{\eta\varepsilon}\psi).
\end{aligned}$$

Evaluating the second variation at a critical point $\psi(x,0,0) = \psi^*$, and from (6.4), we obtain:

$$\frac{\delta^2 J}{\delta\eta\delta\varepsilon}(X,Y) = 2\int_D \nabla X \cdot \nabla Y + 2\lambda^* \int_D \psi^*(\partial_{\eta\varepsilon}\psi^*). \tag{6.5}$$

Since $\psi(x,\varepsilon,\eta)$ is a smooth two-parameter family of functions in $\mathscr{S}_D$, we have $\int_D |\psi(x,\varepsilon,\eta)|^2 = 1$ for all ε,η, which implies that $\int_D \psi(\partial_\varepsilon\psi) = 0$ and $\int_D \partial_\eta\psi\partial_\varepsilon\psi + \int_D \psi(\partial_{\eta\varepsilon}\psi) = \int_D XY + \int_D \psi(\partial_{\eta\varepsilon}\psi) = 0$. Substituting in (6.5), we obtain:

$$\frac{\delta^2 J}{\delta\eta\delta\varepsilon}(X,Y) = 2\int_D \nabla X \cdot \nabla Y - 2\lambda^* \int_D XY.$$

In particular, for $X \neq 0$, this implies:

$$\begin{aligned}\frac{\delta^2 J}{\delta\eta\delta\varepsilon}(X,X) &= 2\int_D |\nabla X|^2 - 2\lambda^* \int_D |X|^2 \\ &= 2\left(\int_D |X|^2\right)\left(\frac{\int_D |\nabla X|^2}{\int_D |X|^2} - \lambda^*\right).\end{aligned} \tag{6.6}$$

We also have that $\int_D \psi(x,\varepsilon,\eta) = 0$, which leads to $\int_D \partial_\varepsilon\psi = \int_D X = 0$. From (6.2), we have that $\inf_{\int_D X=0} \frac{\int_D |\nabla X|^2}{\int_D |X|^2} = \lambda_2$, which implies that if $\lambda^* > \lambda_2$ in (6.6), by the definition of infimum, there exists an X such that $\frac{\delta^2 J}{\delta\eta\delta\varepsilon}\Big|_{\varepsilon=0,\eta=0}(X,X) < 0$. Therefore, the only critical point for which $\frac{\delta^2 J}{\delta\eta\delta\varepsilon}\Big|_{\varepsilon=0,\eta=0}(X,X) \geq 0$ is the second eigenfunction $\psi^* = \psi_2$. Note that, for this case, $\frac{\delta^2 J}{\delta\eta\delta\varepsilon}\Big|_{\varepsilon=0,\eta=0}(X,X) = 0$ if and only if $X = k\psi_2$. Since $\int_D \psi_2\, X = 0$, it must be that $k = 0$, and therefore $X = 0$. Thus, for all $X \neq 0$, $\frac{\delta^2 J}{\delta\eta\delta\varepsilon}\Big|_{\varepsilon=0,\eta=0}(X,X) > 0$ at $\psi^* = \psi_2$. Therefore, the second eigenfunction ψ_2 is the only minimizer of the functional $J(\psi)$ in $\mathscr{S}_D$.

It further follows from the above argument that the Hessian $\frac{\delta^2 J}{\delta\eta\delta\varepsilon}\Big|_{\varepsilon=0,\eta=0}$ is non-degenerate (or non-singular) at the critical points of $J(\psi)$, that is, $\frac{\delta^2 J}{\delta\eta\delta\varepsilon}\Big|_{\varepsilon=0,\eta=0}(X,X) = 0$ at the critical points of $J(\psi)$ if and only if $X = 0$. □

We now provide a novel closed-form expression for a projected gradient flow to converge to the minimum value of $J(\psi)$ in $\mathscr{S}_D$. For smooth one-parameter families of functions $\{\psi(t,x)\}_{t\in\mathbb{R}_{\geq 0}}$ (with $x \in D$), the derivative of the objective functional J is given by:

$$\frac{d}{dt}[J(\psi(t))] = 2\int_D \nabla\psi \cdot \nabla(\partial_t \psi) = -2\int_D \partial_t \psi(\Delta\psi).$$

We obtain a gradient flow by setting $\partial_t \psi = \Delta\psi$. We project this flow onto the tangent space of the set $\mathscr{S}_D$. For $\psi \in \mathscr{S}_D$, we require that $\langle \psi, \partial_t \psi \rangle = 0$ and $\int_D \partial_t \psi = 0$, which are satisfied if (this will be shown in Proposition 3):

$$\partial_t \psi = \Delta\psi - \frac{\langle \Delta\psi, \psi \rangle}{\|\psi\|^2}\psi = \Delta\psi - \langle \Delta\psi, \psi \rangle\, \psi,$$

since $\|\psi\| = 1$ for $\psi \in \mathscr{S}_D$. Further, using $J(\psi) = -\langle \Delta\psi, \psi \rangle$, we get the projected gradient flow:

$$\partial_t \psi = \Delta\psi + J(\psi)\psi. \tag{6.7}$$

The equilibria ψ^* of (6.7) satisfy $\Delta\psi^* + J(\psi^*)\psi^* = 0$ and the Neumann boundary condition $\nabla\psi^* = 0$ on ∂D. Clearly, $J(\psi^*)$ is an eigenvalue, and so let $\lambda^* = J(\psi^*)$. It is also clear that the equilibria of the projected gradient flow are also the critical points of the functional J over the set $\mathscr{S}_D$.

Proposition 3 (Convergence of gradient flow)**.** *The set $\mathscr{S}_D$ is invariant with respect to the flow* (6.7)*, and the solutions to* (6.7) *in $\mathscr{S}_D$ converge in an L^2 sense to the set of equilibria of* (6.7)*. Moreover, the only locally asymptotically stable equilibrium in $\mathscr{S}_D$ for* (6.7) *is the second eigenfunction ψ_2.*

Proof of Proposition 3. Recall that $\mathscr{S}_D = \{\psi \in H^1(D) \mid \int_D |\psi|^2 = 1, \int_D \psi = 0\}$. Therefore, for a smooth one-parameter family $\{\psi(t,x)\}_{t\in\mathbb{R}_{\geq 0}}$, (with $x \in D$) to be in $\mathscr{S}_D$, we need to prove that

$\int_D \psi \, \partial_t \psi = 0$ and $\int_D \partial_t \psi = 0$, assuming that the initial condition is in $\mathscr{S}_D$. (Note that it will later be shown that $\frac{d}{dt}\|\nabla\psi\| \leq 0$, thus $\psi(t,\cdot) \in H^1(D)$ for all $t \geq 0$ if $\psi(0,\cdot) \in \mathscr{S}_D$).

From Equation (6.7), we have $\int_D \psi \;\; \partial_t \psi = \int_D \psi(\Delta\psi + J(\psi)\psi)$. Using the Divergence theorem and the Neumann boundary condition on $\partial\Omega$, we get $\int_D \psi \, \partial_t \psi = -\int_D |\nabla\psi|^2 + J(\psi)\int_D |\psi|^2 = 0$ (since $J(\psi) = \int_D |\nabla\psi|^2$ and $\int_D |\psi|^2 = 1$).

We also have $\int_D \partial_t \psi = \int_D \Delta\psi + J(\psi)\int_D \psi = \int_D \nabla\psi \cdot \mathbf{n} + J(\psi)\int_D \psi = 0$ because of the Neumann boundary condition, $\nabla\psi \cdot \mathbf{n} = 0$ on ∂D, and $\int_D \psi = 0$.

Let $\psi(t,x)$ be a solution of (6.7) in $\mathscr{S}_D$, with $t \in \mathbb{R}_{\geq 0}$, $x \in D$, such that $\psi(0,x) \in \mathscr{S}_D$. We also have $\int_D |\psi|^2 = 1$, for all $t \geq 0$. Thus, $J(\psi) = \int_D |\nabla\psi|^2 = \frac{\int_D |\nabla\psi|^2}{\int_D |\psi|^2}$. The time derivative of J is given by:

$$
\begin{aligned}
\frac{d}{dt}J &= \frac{2}{\int_D |\psi|^2}\int_D \nabla\psi \cdot \nabla\partial_t\psi - 2\frac{\int_D |\nabla\psi|^2}{(\int_D |\psi|^2)^2}\int_D \psi\partial_t\psi \\
&= -2\int_D \Delta\psi \; \partial_t\psi - 2J(\psi)\int_D \psi\partial_t\psi \\
&= -2\int_D (\Delta\psi + J(\psi)\psi)\partial_t\psi \\
&= -2\int_D |\Delta\psi + J(\psi)\psi|^2 \leq 0.
\end{aligned}
$$

We have that $J \geq 0$ and $\frac{d}{dt}J \leq 0$. We also have $\mathscr{S}_D \subset H^1(D)$, D a bounded, open subset of $\mathbb{R}^N$ with ∂D being C^1. Thus by the Rellich-Kondrachov Compactness Theorem [55], we get that the orbit ψ is precompact in $L^2(D)$. Therefore, by the LaSalle invariance principle for infinite dimensional spaces [70], the solutions converge in an L^2 sense to largest invariant set contained in $\{\psi^* \in \mathscr{S}_D \,|\, \Delta\psi^* + J(\psi^*)\psi^* = 0\}$, the set of equilibria of (6.7).

In what follows we use the shorthand $\partial_t\psi = F(\psi)$, where $F(\psi^*) = 0$, for the dynamics (6.7). We consider perturbations $\delta\psi \in \mathscr{T}_D$ along the tangent space of $\mathscr{S}_D$ at ψ^* (also note that ψ^* is an eigenfunction). Thus $\int_D \delta\psi = 0$ and $\int_D \psi^* \; \delta\psi = 0$. We have:

$$
F(\psi^* + \delta\psi) = \Delta(\psi^* + \delta\psi) + J(\psi^* + \delta\psi)(\psi^* + \delta\psi).
$$

Since ψ^* is a critical point of $J(\psi)$ it holds that $J(\psi^* + \delta\psi) = J(\psi^*) + \mathscr{O}(\|\delta\psi\|^2) = \lambda^* + \mathscr{O}(\|\delta\psi\|^2)$. Thus, up to first-order we have that:

$$\begin{aligned} F(\psi^* + \delta\psi) &= \Delta(\psi^* + \delta\psi) + J(\psi^* + \delta\psi)(\psi^* + \delta\psi) \\ &= -\lambda^*\psi^* + \Delta(\delta\psi) + \lambda^*\psi^* + \lambda^*\delta\psi \\ &= \Delta(\delta\psi) + \lambda^*\delta\psi. \end{aligned}$$

Therefore, we have $\partial_t(\delta\psi) = \Delta(\delta\psi) + \lambda^*\delta\psi$. Expressing $\delta\psi(t) = \sum_{i=2}^{\infty}\alpha_i(t)\psi_i$, where ψ_i are the eigenfunctions which form an orthonormal basis for $\mathscr{T}_D$, we have that:

$$\begin{aligned} \partial_t(\delta\psi) = \sum_{i=2}^{\infty}\frac{d}{dt}\alpha_i(t)\psi_i &= \Delta(\delta\psi) + \lambda^*\delta\psi \\ &= \sum_{i=2}^{\infty}\alpha_i(t)(-\lambda_i + \lambda^*)\psi_i, \end{aligned}$$

which implies that $\delta\psi(t) = \sum_{i=2}^{\infty} e^{(\lambda^*-\lambda_i)t}\alpha_i(0)\psi_i$. (Note that, from orthogonality, the previous equality leads to $\frac{d}{dt}\alpha_i(t) = \alpha_i(t)(-\lambda_i + \lambda^*)$, for each i.) We claim that the latter converges to $\delta\psi = 0$ for all initial conditions $\delta\psi(0) \in \mathscr{T}_D$ at ψ^* if and only if $\lambda^* = \lambda_2$ (correspondingly, $\psi^* = \psi_2$). To see this, first observe that, if $\lambda^* = \lambda_2$ (correspondingly, $\psi^* = \psi_2$), we have $\int_D \psi_2\ \delta\psi(0) = 0$ (since $\delta\psi \in \mathscr{T}_D$ at $\psi^* = \psi_2$), which implies that $\alpha_2(0) = \alpha_2(t) = 0$. Hence $\delta\psi(t) = \sum_{i=3}^{\infty} e^{\lambda_2-\lambda_i}\alpha_i(0)\psi_i$ and the exponent $\lambda_2 - \lambda_i < 0$ for all $i \geq 3$. Conversely, if $\delta\psi(t) = \sum_{i=2}^{\infty} e^{(\lambda^*-\lambda_i)t}\alpha_i(0)\psi_i$ converges to $\delta\psi = 0$ for all initial conditions $\delta\psi(0) \in \mathscr{T}_D$ at ψ^*, and $\psi^* = \psi_i$ for some $i \in \{2,3,\ldots\}$. We have that $\alpha_i(0) = \alpha_i(t) = 0$ (from orthogonality), and that $\delta\psi(t) = \sum_{j=2,j\neq i}^{\infty} e^{(\lambda_i-\lambda_j)t}\alpha_j(0)\psi_j$, which converges to $\delta\psi = 0$ only if $i = 2$. Therefore, the second eigenfunction ψ_2 is the only locally asymptotically stable equilibrium in $\mathscr{S}_D$ for the projected gradient flow. □

Remark 12 (Implication of Proposition 3)**.** *Proposition 3 states that we have global convergence to the set of isolated equilibria of the gradient flow* (6.7) *and that only the second eigenfunction ψ_2 is locally asymptotically stable among the set of isolated equilibria. Moreover, as seen in the* *proof of Lemma 46, we have that the other equilibria are saddle points of $J(\psi)$ and are non-degenerate (the Hessian of J at these saddle points are non-singular). From this we deduce almost global asymptotic stability of the second eigenfuction ψ_2 for the flow* (6.7) *, and we therefore have convergence from almost all initial conditions, see [92] for an overview of this property.*

6.3.2 Design of hole-placement dynamics

We now consider the full optimization problem (6.3), which can be expressed as:

$$x^* \in \arg \inf_{x \in \tilde{\Omega}_r} \lambda_2(\Omega \setminus B_r(x))$$

Assumption 20 (Simplicity of the second eigenvalue)**.** *We assume that the second eigenvalue $\lambda_2(\Omega \setminus B_r(x))$ is simple for any $x \in \tilde{\Omega}$.*

Remark 13 (Relaxing Assumption 20)**.** *The assumption that the eigenvalue λ_2 is simple is ensures differentiability of $\lambda_2(\Omega \setminus B_r(x))$ w.r.t. x. The eigenvalues of $\Delta(\Omega \setminus B_r(x))$ exist as branches $x \mapsto \lambda(\Omega \setminus B_r(x))$, which can then be ordered as $\lambda_1 \leq \lambda_2 \leq \ldots$ for any given x. The branches $x \mapsto \lambda(\Omega \setminus B_r(x))$ of eigenvalues are differentiable w.r.t. x (more generally w.r.t. the perturbation of domains with Lipschitz boundaries [69]). The case of a non-simple eigenvalue λ_2 occurs when multiple branches intersect, for some x, at which point the ordering of the branches may change and we lose differentiability of λ_2. This situation can however be mitigated by considering the subdifferential of λ_2 in place of the gradient of λ_2. The dynamics presented later in the paper can be modified in this sense, and the analysis would require further investigation on the regularity/lower-semicontinuity properties of these subdifferentials. We nevertheless avoid this problem through Assumption 20, which we leave as future work.*

The following lemma allows for a characterization of the critical points of the functional λ_2 in the interior of the domain.

Lemma 47 (Characterization of critical ball)**.** *The first-order condition for a critical point x^* of the functional λ_2 in the interior of the domain is given by:*

$$\lambda_2^* \left(\int_{\partial B_r(x^*)} |\psi_2^*|^2 \mathbf{n} \right) = \int_{\partial B_r(x^*)} |\nabla \psi_2^*|^2 \mathbf{n}, \tag{6.8}$$

where (λ_2^*, ψ_2^*) is the second eigenpair such that $\lambda_2^* \stackrel{\triangle}{=} \lambda_2(\Omega \setminus B_r(x^*))$.

Proof. Let $x(\varepsilon)$ for $\varepsilon \in \mathbb{R}$ be a smooth curve contained in $\tilde{\Omega}_r$. Let ψ_2^ε be the second eigenfunction of the Laplace operator with Neumann boundary condition in the domain $\Omega \setminus B_r(x(\varepsilon))$. Thus, we have $\lambda_2^\varepsilon = \int_{\Omega_\varepsilon} |\nabla \psi_2^\varepsilon|^2$, where $\Omega_\varepsilon = \Omega \setminus B_r(x(\varepsilon))$ and $\|\psi_2^\varepsilon\|_{\Omega_\varepsilon = 1}$. The derivative $\frac{d\lambda_2^\varepsilon}{d\varepsilon}$ is given by:

$$\frac{d\lambda_2^\varepsilon}{d\varepsilon} = \frac{d}{d\varepsilon} \int_{\Omega_\varepsilon} |\nabla \psi_2^\varepsilon|^2 = 2 \int_{\Omega_\varepsilon} \nabla \psi_2^\varepsilon \cdot \nabla \left(\frac{\partial \psi_2^\varepsilon}{\partial \varepsilon} \right) + \int_{\partial \Omega_\varepsilon} |\nabla \psi_2^\varepsilon|^2 \mathbf{v} \cdot \mathbf{n}, \tag{6.9}$$

where $\mathbf{v} = \frac{dx(\varepsilon)}{d\varepsilon}$, is constant on $\partial B_r(x_\varepsilon)$. Equation (6.9) becomes:

$$\begin{aligned}
\frac{d\lambda_2^\varepsilon}{d\varepsilon} &= 2 \int_{\Omega_\varepsilon} \nabla \psi_2^\varepsilon \cdot \nabla \left(\frac{\partial \psi_2^\varepsilon}{\partial \varepsilon} \right) + \mathbf{v} \cdot \left(\int_{\partial B_r(x_\varepsilon)} |\nabla \psi_2^\varepsilon|^2 \mathbf{n} \right) \\
&= -2 \int_{\Omega_\varepsilon} \frac{\partial \psi_2^\varepsilon}{\partial \varepsilon} \Delta \psi_2^\varepsilon + \mathbf{v} \cdot \left(\int_{\partial B_r(x_\varepsilon)} |\nabla \psi_2^\varepsilon|^2 \mathbf{n} \right) \\
&= 2 \int_{\Omega_\varepsilon} \lambda_2^\varepsilon \psi_2^\varepsilon \frac{\partial \psi_2^\varepsilon}{\partial \varepsilon} + \mathbf{v} \cdot \left(\int_{\partial B_r(x_\varepsilon)} |\nabla \psi_2^\varepsilon|^2 \mathbf{n} \right) \\
&= \lambda_2^\varepsilon \frac{d}{d\varepsilon} \left(\int_{\Omega_\varepsilon} |\psi_2^\varepsilon|^2 \right) - \lambda_2^\varepsilon \mathbf{v} \cdot \left(\int_{\partial B_r(x_\varepsilon)} |\psi_2^\varepsilon|^2 \mathbf{n} \right) \\
&\quad + \mathbf{v} \cdot \left(\int_{\partial B_r(x_\varepsilon)} |\nabla \psi_2^\varepsilon|^2 \mathbf{n} \right) \\
&= -\lambda_2^\varepsilon \mathbf{v} \cdot \left(\int_{\partial B_r(x_\varepsilon)} |\psi_2^\varepsilon|^2 \mathbf{n} \right) + \mathbf{v} \cdot \left(\int_{\partial B_r(x_\varepsilon)} |\nabla \psi_2^\varepsilon|^2 \mathbf{n} \right),
\end{aligned} \tag{6.10}$$

since $\int_{\Omega_\varepsilon} |\psi_2^\varepsilon|^2 = 1$ for all $\varepsilon \in \mathbb{R}$, which implies that $\frac{d}{d\varepsilon} \left(\int_{\Omega_\varepsilon} |\psi_2^\varepsilon|^2 \right) = 0$. Let $x(0) = x^* \in \tilde{\Omega}$ be a critical point of $\lambda_2(x)$, such that $\lambda_2(x^*) = \lambda_2^*$, with ψ_2^* being the second eigenfunction. Thus we have $\frac{d\lambda_2^\varepsilon}{d\varepsilon}\Big|_{\varepsilon=0} = 0$ for all $\mathbf{v}$, which implies that:

$$\lambda_2^* \left(\int_{\partial B_r(x^*)} |\psi_2^*|^2 \mathbf{n} \right) = \int_{\partial B_r(x^*)} |\nabla \psi_2^*|^2 \mathbf{n}.$$

This is the first-order condition for critical points of λ_2 in the interior of the domain. □

We now construct the gradient dynamics to converge to a critical point of λ_2 in the interior of the domain. Note that the function $\lambda_2(\Omega \setminus B_r(x))$ is not known explicitly for a general domain $\Omega \setminus B_r(x)$. We reformulate the optimization problem (6.3) as:

$$x^* = \arg_1 \inf_{(x,\psi) \in \tilde{\Omega} \times \Psi(x)} \int_{\Omega \setminus B_r(x)} |\nabla \psi|^2 d\nu, \tag{6.11}$$

where the set $\Psi(x)$ is defined as:

$$\Psi(x) = \left\{ \psi \in H^1(\Omega \setminus B_r(x)) \,\middle|\, \int_{\Omega \setminus B_r(x)} \psi = 0, \int_{\Omega \setminus B_r(x)} |\psi|^2 = 1 \right\}, \tag{6.12}$$

where $\arg_1$ indicates the first argument x in (x, ψ). We also define the set $\Psi = \cup_{x \in \tilde{\Omega}_r} \Psi(x)$. We recall that $\tilde{\Omega}_r = \{x \in \Omega \,|\, \mathrm{dist}(x, \partial\Omega) > r\}$. Now let $\{x(t)\}_{t \in \mathbb{R}_{\geq 0}}$ be a smooth curve in $\tilde{\Omega}_r$ and $\{\psi(t,y)\}_{t \in \mathbb{R}_{\geq 0}}$ (with $y \in \Omega \setminus B_r(x(t))$,) a smooth one-parameter family of functions on $\Omega \setminus B_r(x(t))$. Also, let $\tilde{\mathbf{n}}(x)$ be the normal to the boundary $\partial\tilde{\Omega}_r$ at $x \in \partial\tilde{\Omega}_r$. We now consider the following hole-placement dynamics for our nested optimization problem:

$$\begin{aligned}
&\frac{dx}{dt} = \begin{cases} \mathbf{v}_{int}, \; x \in \mathrm{int}\, \tilde{\Omega}_r \\ \mathbf{v}_{int} - (\mathbf{v}_{int} \cdot \tilde{\mathbf{n}})\tilde{\mathbf{n}}, \; x \in \partial\tilde{\Omega}_r \end{cases} \\
&\mathbf{v}_{int} = -\int_{\partial B_r(x)} |\nabla \psi|^2 \mathbf{n} + J(\psi) \int_{\partial B_r(x)} |\psi|^2 \mathbf{n}, \\
&\partial_t \psi = \Delta\psi + J(\psi)\psi + a\psi + b, \\
&\nabla\psi \cdot \mathbf{n} = 0, \qquad \text{on } \partial\Omega \cup \partial B_r(x),
\end{aligned} \tag{6.13}$$

where $a = -\frac{1}{2}\mathbf{v}\cdot\left(\int_{\partial B_r(x)} |\psi|^2\,\mathbf{n}\right)$ and $b = -\frac{1}{|\Omega|-c}\mathbf{v}\cdot\left(\int_{\partial B_r(x)} \psi\,\mathbf{n}\right)$, with $c = |B_r(x)|$, for all $x \in \tilde{\Omega}_r$.

Theorem 21 (Convergence of the hole placement dynamics)**.** *The set* Ψ *in* (6.12) *is invariant with respect to the dynamics* (6.13)*. The solutions to the dynamics* (6.13) *converge to a critical point of the objective functional* λ_2 *in* (6.11)*. A critical point of* λ_2 *is locally asymptotically stable with respect to the dynamics* (6.13) *only if it is a strict local minimum.*

Proof. Let $\{x(t), \psi(t,y)\}_{t\in\mathbb{R}_{\geq 0}}$ (with $y \in \Omega \setminus B_r(x(t))$,) be a one-parameter family of functions that is a solution to the dynamics (6.13), and let $\psi(0,\cdot) \in \Psi(x(0))$. To prove the invariance of $\Psi(x(t))$, we need to show that $\frac{d}{dt}\left(\int_{\Omega\setminus B_r(x(t))} |\psi|^2\right) = 0$ and $\frac{d}{dt}\left(\int_{\Omega\setminus B_r(x(t))} \psi\right) = 0$ (Note that it will later be shown that $\frac{d}{dt}\|\nabla\psi\| \leq 0$, thus $\psi(t,\cdot) \in H^1(\Omega)$ for all $t \geq 0$ if $\psi(0,\cdot) \in \Psi$). From (6.13), we have (with $\Omega(t) = \Omega \setminus B_r(x(t))$):

$$
\begin{aligned}
\frac{d}{dt}\left(\int_{\Omega(t)} |\psi|^2\right) &= 2\int_{\Omega(t)} \psi\,\partial_t\psi + \mathbf{v}\cdot\int_{\partial B_r(x(t))} |\psi|^2\mathbf{n} \\
&= 2\int_{\Omega(t)} \psi\,\Delta\psi + 2J(\psi)\int_{\Omega(t)} |\psi|^2 + 2a(t)\times \\
&\quad \int_{\Omega(t)} |\psi|^2 + 2b(t)\int_{\Omega(t)} \psi + \mathbf{v}\cdot\int_{\partial B_r(x(t))} |\psi|^2\mathbf{n} \\
&= -2\int_{\Omega(t)} |\nabla\psi|^2 + 2J(\psi) + 2a(t) \\
&\quad + \mathbf{v}\cdot\int_{\partial B_r(x(t))} |\psi|^2\mathbf{n} \\
&= 0,
\end{aligned}
$$

because $J(\psi) = \int_{\Omega(t)} |\nabla\psi|^2$, $\int_{\Omega(t)} |\psi|^2 = 1$ and $\int_{\Omega(t)} \psi = 0$ (since $\psi(t,\cdot) \in \Psi(x(t))$.) We also have:

$$\frac{d}{dt}\left(\int_{\Omega(t)}\psi\right) = \int_{\Omega(t)}\partial_t\psi + \mathbf{v}\cdot\int_{\partial B_r(x(t))}\psi\,\mathbf{n}$$
$$= \int_{\Omega(t)}\Delta\psi + J(\psi)\int_{\Omega(t)}\psi$$
$$+ a(t)\int_{\Omega(t)}\psi + b(|\Omega| - c)$$
$$+ \mathbf{v}\cdot\int_{\partial B_r(x(t))}\psi\,\mathbf{n}$$
$$= 0.$$

Since we also have that $\psi(0,\cdot)\in\Psi(x(0))$, we conclude that the set Ψ is invariant with respect to the dynamics (6.13).

Let $\{x(t),\psi(t,y)\}_{t\in\mathbb{R}_{\geq 0}}$ (with $y\in\Omega\setminus B_r(x(t))$), be a one-parameter family of functions that is a solution to the dynamics (6.13), and let $\psi(t,\cdot)\in\Psi(x(t))$ for all $t\in\mathbb{R}_{\geq 0}$ (this assumption is justified by the invariance of Ψ). We have $J(\psi)=\int_{\Omega(t)}|\nabla\psi|^2 = \frac{\int_{\Omega(t)}|\nabla\psi|^2}{\int_{\Omega(t)}|\psi|^2}\geq 0$ for $\psi(t,\cdot)\in\Psi(x(t))$ (since $\int_{\Omega(t)}|\psi|^2=1$). Now:

$$\frac{d}{dt}J = 2\int_{\Omega(t)}\nabla\psi\cdot\nabla\partial_t\psi + \mathbf{v}\cdot\int_{\partial B_r(x(t))}|\nabla\psi|^2\mathbf{n}$$
$$- 2J(\psi)\int_{\Omega(t)}\psi\,\partial_t\psi - J(\psi)\mathbf{v}\cdot\int_{\partial B_r(x(t))}|\psi|^2\mathbf{n}$$
$$= -2\int_{\Omega(t)}|\Delta\psi + J(\psi)\psi|^2 - \mathbf{v}\cdot\mathbf{v}_{int} \qquad \leq 0,$$

where we have used (6.13) to obtain the second equality. By the Rellich-Kondrachov Compactness Theorem [55], we see that the orbit ψ is precompact in $L^2(\Omega)$. Thus, by the invariance principle [70], the solutions $\{x(t),\psi(t,y)\}_{t\in\mathbb{R}_{\geq 0}}$ (with $y\in\Omega\setminus B_r(x(t))$), converge to x^*,ψ^* (the convergence $\psi(t,\cdot)\to\psi^*$, is in the sense of L^2) such that $\mathbf{v}=0$ and $\Delta\psi^*+J(\psi^*)\psi^*=0$. We already have that the only asymptotically stable case is when $\psi^*=\psi_2^*$ (the second eigenfunction corresponding to $\Omega\setminus B_r(x^*)$), which implies that $J(\psi^*)=J(\psi_2^*)=\lambda_2^*$. And $\mathbf{v}=0$ implies that $\int_{\partial B_r(x^*)} |\nabla \psi_2^*|^2 \mathbf{n} = \lambda_2^* \int_{\partial B_r(x^*)} |\psi_2^*|^2 \mathbf{n}$, the critical point of the functional λ_2 from (6.8).

Consider perturbations δx and $\delta \psi$, about an equilibrium (x^*, ψ_2^*) such that $x^* + \delta x \in \tilde{\Omega}$ and $\tilde{\psi}_2 = \psi_2^* + \delta \psi \in \Psi(x^* + \delta x)$ is the second eigenfunction of the domain $\Omega \setminus B_r(x^* + \delta x)$. In other words, we consider perturbations purely in x to investigate the local asymptotic stability of the critical points of $\lambda_2(x)$. The dynamics in x in this case, referring to (6.13), are given by:

$$\frac{d}{dt}(x^* + \delta x) = -\int_{\partial B_r(x^* + \delta x)} \left(|\nabla \tilde{\psi}_2|^2 - \tilde{\lambda}_2 |\tilde{\psi}_2|^2 \right) \mathbf{n}.$$

This can be reduced to:

$$\frac{d}{dt}(\delta x) = -\frac{\partial}{\partial x}\bigg|_{x=x^*} \left(\int_{\partial B_r(x)} \left(|\nabla \tilde{\psi}_2|^2 - \tilde{\lambda}_2 |\tilde{\psi}|^2 \right) \mathbf{n} \right) \delta x. \tag{6.14}$$

From Equation (6.10), we recognize that $\int_{\partial B_r(x)} \left(|\nabla \tilde{\psi}_2|^2 - \tilde{\lambda}_2 |\tilde{\psi}|^2 \right) \mathbf{n} = \frac{\partial \lambda_2}{\partial x}$. Therefore, the linearized dynamics reduces to:

$$\frac{d}{dt}(\delta x) = -\frac{\partial^2 \lambda_2}{\partial x^2}\bigg|_{x=x^*} \delta x,$$

where $\frac{\partial^2 \lambda_2}{\partial x^2}\Big|_{x=x^*}$ is the Hessian of λ_2 at $x = x^*$. Therefore, we have that the linearized dynamics is asymptotically stable if and only if the Hessian of λ_2 is positive definite, in other words, if and only if x^* is a strict local minimum of λ_2. Therefore, the necessary condition for the local asymptotic stability of the primal-dual dynamics at a critical point of λ_2 is that it is a strict local minimum. $\square$

Remark 14 (Implication of Theorem 21)**.** *Theorem 21 states that we have convergence to the equilibria of the hole-placement dynamics which are also critical points of* $\lambda_2(\Omega \setminus B_r(x))$*. In addition, we have that among the critical points of* $\lambda_2(\Omega \setminus B_r(x))$*, only the strict local minima are locally asymptotically stable. For almost global convergence to these strict local minima, we additionally require non-degeneracy of the saddle points of* $\lambda_2(\Omega \setminus B_r(x))$ *(i.e., that the Hessian* *is non-singular at the critical point), but this additional characterization is not contained in our result.*

We now consider the following question: if an initial failure happens with the removal of a node, what is the most critical node? This is appropriately posed in the continuum setting as the hole placement problem where the size of the hole is very small, i.e., as the radius $r \to 0$. For this, we investigate the minimum of the function $f(x) = \lim_{r\to 0} \frac{1}{|\partial B_r(x)|} \frac{\partial}{\partial r} \lambda_2(\Omega \setminus B_r(x))$, which quantifies as a function of the hole position, the rate of deterioration of the metric as failure begins to occur.

Theorem 22. ***(Connection to the nodal set of eigenfunction).*** *In the limit $r \to 0$ for the radius of the hole, the hole-placement problem reduces to finding the minima $x^* \in \Omega$ of the function:*

$$f(x) = \mu_2^{\Omega} |\psi_2^{\Omega}(x)|^2 - |\nabla \psi_2^{\Omega}(x)|^2,$$

where $(\mu_2^{\Omega}, \psi_2^{\Omega}(x))$ is the second eigenpair of the domain Ω. Moreover, if the family of level sets of ψ_2^{Ω} is locally flat at a point $x^* \in \Omega$, then x^* is a local minimizer of f if and only if $\psi_2^{\Omega}(x^*) = 0$. In other words, under local flatness, the nodal points of ψ_2^{Ω} are the local minimizers of f.

Proof. Let $r : \mathbb{R} \to \mathbb{R}_{\geq 0}$ with $r(0) = 0$ be a smooth non-negative function. Let $\Omega(t) = \Omega \setminus B_{r(t)}(x)$ for some $x \in \Omega \subset \mathbb{R}^N$, be a one parameter family of spatial domains such that $\Omega(0) = \Omega$. Let $\mu_2(t)$ be the second eigenvalue of the domain $\Omega(t)$ and $\psi_2(t, \cdot)$ the corresponding normalized eigenfunction (we assume that the family of spatial domains $\Omega(t)$ have simple eigenvalues). Thus, we have $\mu_2(t) = \int_{x \in \Omega(t)} |\nabla \psi_2(t,x)|^2$. From [40], we have that $\mu_2(t)$ and ψ_2 are real-analytic locally at $t = 0$. Thus, for small $\tau > 0$, we have:

$$\begin{aligned} \mu_2(\tau) &= \mu_2(0) + \frac{d}{dt}\mu_2 \Big|_{t=0} \tau + \dots \\ \psi_2(\tau, x) &= \psi_2(0,x) + \partial_t \psi_2(t,x) \Big|_{t=0} \tau + \dots \end{aligned} \tag{6.15}$$

We note that $\mu_2(0)$ and $\psi_2(0,\cdot)$ are the second eigenpair corresponding to Ω. At a given $t > 0$, let the deformation of the domain be characterized by $\mathbf{v} = -\varepsilon\mathbf{n}$, the velocity of points on the boundary of the hole, $B_{r(t)}(x)$, where $\mathbf{n}$ is the normal to the domain $\Omega(t)$ on the boundary of $B_{r(t)}(x)$, and $\varepsilon > 0$ is a small constant. We have:

$$\begin{aligned}
\frac{d}{dt}\mu_2 &= \frac{d}{dt}\int_{x\in\Omega(t)} |\nabla\psi_2(t,x)|^2 \\
&= 2\int_{\Omega(t)} \nabla\psi_2\,\nabla\partial_t\psi_2 + \int_{\partial B_{r(t)}(x)} |\nabla\psi_2|^2 \mathbf{v}\cdot\mathbf{n} \\
&= -2\int_{\Omega(t)} \Delta\psi_2\,\partial_t\psi_2 + \int_{\partial B_{r(t)}(x)} |\nabla\psi_2|^2 \mathbf{v}\cdot\mathbf{n} \\
&= 2\mu_2(t)\int_{\Omega(t)} \psi_2\,\partial_t\psi_2 + \int_{\partial B_{r(t)}(x)} |\nabla\psi_2|^2 \mathbf{v}\cdot\mathbf{n} \\
&= \mu_2(t)\left(\frac{d}{dt}\int_{\Omega(t)} |\psi_2|^2 - \int_{\partial B_{r(t)}(x)} |\psi_2|^2 \mathbf{v}\cdot\mathbf{n}\right) \\
&\quad + \int_{\partial B_{r(t)}(x)} |\nabla\psi_2|^2 \mathbf{v}\cdot\mathbf{n} \\
&= \mu_2(t)\varepsilon\int_{\partial B_{r(t)}(x)} |\psi_2|^2 - \varepsilon\int_{\partial B_{r(t)}(x)} |\nabla\psi_2|^2,
\end{aligned}$$

since $\int_{\Omega(t)} |\psi_2|^2 = 1$, for all t. For small $\tau > 0$, we then substitute from (6.15) in the above equation, to obtain:

$$\begin{aligned}
\frac{d}{dt}\mu_2 = \varepsilon &\left(\mu_2(0) + \frac{d}{dt}\mu_2\Big|_{t=0}\tau + \dots\right) \times \\
&\int_{y\in\partial B_{r(\tau)}(x)} |\psi_2(0,y) + \partial_t \psi_2(t,y)\Big|_{t=0}\tau + \dots|^2 \\
&- \varepsilon \int_{y\in\partial B_{r(\tau)}(x)} |\nabla(\psi_2(0,y) + \partial_t \psi_2(t,y)\Big|_{t=0}\tau + \dots)|^2 \\
= &\mu_2(0)\varepsilon \int_{y\in\partial B_{r(\tau)}(x)} |\psi_2(0,y)|^2 - \varepsilon \int_{y\in\partial B_{r(\tau)}(x)} |\nabla\psi_2(0,y)|^2 \\
&+ \mathscr{O}(\tau) \\
= &\mu_2(0) S_{N-1} r(\tau)^{N-1} \varepsilon |\psi_2(0,x)|^2 \\
&- S_{N-1} r(\tau)^{N-1} \varepsilon |\nabla\psi_2(0,x)|^2 + \mathscr{O}(r(\tau)^{N-1}\tau),
\end{aligned}$$

where S_N is the surface area of th unit N-sphere. Now, given that $\mathbf{v} = -\varepsilon\mathbf{n}$, we have $r(\tau) = \varepsilon\tau$, and therefore:

$$\begin{aligned}
\frac{d}{dt}\mu_2 = &\mu_2(0) S_{N-1} \varepsilon^N \tau^{N-1} |\psi_2(0,x)|^2 \\
&- S_{N-1} \varepsilon^N \tau^{N-1} |\nabla\psi_2(0,x)|^2 + \mathscr{O}(\tau^N).
\end{aligned}$$

Substituting for $\frac{d}{dt}\mu_2$ from the above equation into $\mu_2(\tau) = \mu_2(0) + \frac{d}{dt}\mu_2\big|_{\bar{\tau}}\tau$ (where $\bar{\tau} \in [0,\tau]$), we get:

$$\begin{aligned}
\mu_2(\tau) = &\mu_2(0) + S_{N-1}\varepsilon^N \bar{\tau}^{N-1}\tau \left(\mu_2(0)|\psi_2(0,x)|^2 - |\nabla\psi_2(0,x)|^2\right) \\
&+ \mathscr{O}(\bar{\tau}^N \tau) \\
\leq &\mu_2(0) + S_{N-1}\varepsilon^N \tau^N \left(\mu_2(0)|\psi_2(0,x)|^2 - |\nabla\psi_2(0,x)|^2\right) \\
&+ \mathscr{O}(\tau^{N+1}) \\
\approx &\mu_2(0) + c(\tau)\left(\mu_2(0)|\psi_2(0,x)|^2 - |\nabla\psi_2(0,x)|^2\right),
\end{aligned}$$

where we have ignored the $\mathscr{O}(\tau^{N+1})$ term in the final expression. We also have $r(\tau) = \varepsilon\tau$, and therefore the above can also be written as $\mu_2(r) \approx \mu_2(0) + c(r)\left(\mu_2(0)|\psi_2(0,x)|^2 - |\nabla\psi_2(0,x)|^2\right)$ as a function of the radius of the hole. We also note that the function $\left(\mu_2(0)|\psi_2(0,x)|^2 - |\nabla\psi_2(0,x)|^2\right) = \lim_{r\to 0}\frac{1}{|\partial B_r(x)|}\frac{\partial}{\partial r}\mu_2(\Omega\setminus B_r(x))$.

We now show that the local minima of $f(x) = \mu_2^\Omega|\psi_2^\Omega|^2 - |\nabla\psi_2^\Omega|^2$ occur along the nodal set of ψ_2^Ω, that is, in the set $\{x\in\Omega|\psi_2^\Omega(x) = 0\}$, in the region where the family of level sets of ψ_2^Ω is locally flat. Let $\{\mathbf{r},\mathbf{t}_1,\ldots,\mathbf{t}_{N-1}\}$ be an orthonormal basis at $x\in\Omega$, where $\mathbf{r}$ is the unit normal to the level set of ψ_2^Ω at x and $\{\mathbf{t}_1,\ldots,\mathbf{t}_{N-1}\}$ the unit tangents. We can express the gradient operator in this coordinate system as $\nabla = \mathbf{r}\frac{\partial}{\partial r} + \sum_{i=1}^{N-1}\mathbf{t}_i\frac{\partial}{\partial t_i}$. We now have $\nabla\psi_2^\Omega = \frac{\partial\psi_2^\Omega}{\partial r}\mathbf{r}$ (since the derivative of ψ_2^Ω vanishes along the tangent space of its level set). Moreover, the eigenvalue equation $\Delta\psi_2^\Omega + \mu_2^\Omega\psi_2^\Omega = 0$ expressed in this coordinate system is given by $\frac{\partial^2\psi_2^\Omega}{\partial r^2} + (N-1)H\frac{\partial\psi}{\partial r} + \mu_2^\Omega\psi_2^\Omega = 0$, where $H(x)$ is the mean curvature at $x\in\Omega$ of the level set of ψ. Following some computation, we get that the gradient of f is given by $\nabla f = 4\mu_2^\Omega\psi_2^\Omega\frac{\partial\psi_2^\Omega}{\partial r}\mathbf{r} + 2(N-1)H\left|\frac{\partial\psi_2^\Omega}{\partial r}\right|^2\mathbf{r}$. Moreover, in computing the entries of the Hessian of f in this coordinate frame, we first have:

$$\frac{\partial^2 f}{\partial r^2} = 4\mu_2^\Omega\left(\left|\frac{\partial\psi_2^\Omega}{\partial r}\right|^2 - \mu_2^\Omega|\psi_2^\Omega|^2\right) - 8(N-1)H\mu_2^\Omega\psi_2^\Omega\frac{\partial\psi_2^\Omega}{\partial r} + \left(2(N-1)\frac{\partial H}{\partial r} - 4(N-1)^2H^2\right)\left|\frac{\partial\psi_2^\Omega}{\partial r}\right|^2.$$

Clearly, for any point x^* where the family of level sets of ψ_2^Ω is locally flat (which in particular implies $H(x^*) = 0$), we have that x^* is a critical point if and only if $\psi_2^\Omega(x^*) = 0$ or $\nabla\psi_2^\Omega(x^*) = 0$. Furthermore, we have $\frac{\partial^2 f}{\partial r^2}(x^*) = 4\mu_2^\Omega\left(\left|\frac{\partial\psi_2^\Omega}{\partial r}(x^*)\right|^2 - \mu_2^\Omega|\psi_2^\Omega(x^*)|^2\right)$. Also, under local flatness of the family of level sets, the off-diagonal entries $\frac{\partial^2 f}{\partial r\partial t_i}$ and $\frac{\partial^2 f}{\partial t_i\partial t_j}$ vanish for all $i\in\{1,\ldots,N-1\}$, and so do the rest of the diagonal entries of the Hessian, i.e. $\frac{\partial^2 f}{\partial t_i^2}(x^*) = 0$ for $i\in\{1,\ldots,N-1\}$. It thereby follows that the Hessian is positive semidefinite when $\psi_2^\Omega(x^*) = 0$ and negative semidefinite when $\nabla\psi_2^\Omega(x^*) = 0$. Therefore, under local flatness of the family of level sets of

ψ_2^Ω, the nodal points of ψ_2^Ω correspond to the local minima of f. □

Remark 15 (Geometry of nodal sets)**.** *The nodal sets of Neumann eigenfunctions have been extensively investigated [11]. It is known that if the domain is symmetric about a subset, then it contains the nodal set of ψ_2. The nodal set for the second eigenfunction ψ_2^Ω divides the domain Ω into no more than two regions Ω_a and Ω_b. Now, λ_2^Ω is the first eigenvalue λ_1 of the Laplacian for Ω_a and Ω_b, with Neumann boundary condition on $\partial\Omega \cap \partial\Omega_a$ and Dirichlet boundary condition on $\partial\Omega_a \cap \partial\Omega_b$.*

Remark 16 (Implication for networks)**.** *Theorem 22 can be used to provide new insight on where the most critical nodes in a network with a finite number of nodes are located, via a continuum approximation. This is based on the fact that the entries v_i^F of the Fiedler eigenvector v^F of the finite graph embedded in Ω can be approximated by the value of the eigenfunction ψ_2^Ω at the location x_i of the node i. That is, $v_i^F \approx \psi_2^\Omega(x_i)$. Then the most critical nodes in the network correspond to the zero entries of the Fiedler eigenvector. The Fiedler eigenvector, however, does not necessarily contain zero entries for general finite graphs (this situation improves with the size of the graph), in which case we may expect the critical nodes to be concentrated at the entries of lowest magnitude. This is a heuristic obtained from the fact that ψ_2^Ω is smooth and that ψ_2^Ω more closely approximates v^F as $n \to \infty$.*

6.4 Simulation results

In this section, we present some numerical simulation results that can illustrate the concepts and algorithms of the previous sections.

First, we consider a disk-shaped domain Ω of unit radius, and the placement of a hole B of radius of 0.1 units. Figure 6.1 shows a plot of λ_2 for the residual domain $\Omega \setminus B$ as a function of h (distance between the center of the disk and the center of the hole). Since the hole is of radius 0.1 units and is contained in Ω, we note that $h \in [0, 0.9)$. We observe from Figure 6.1 that the second (also the smallest positive) eigenvalue of the Laplace operator for a disk-shaped domain with

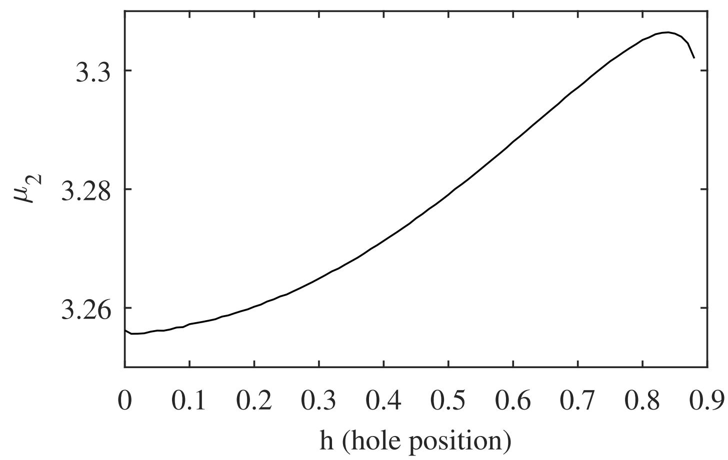

Figure 6.1. λ_2 as a function of h for a disk-shaped domain.

a hole increases with the distance between the centers of the domain and the hole, but also appears to decrease as the hole approaches close to the domain boundary (around $h = 0.85$ units). Moreover, λ_2 as a function of h appears to be a convex in the interval $h \in [0, 0.85]$ and concave for $h \in (0.85, 0.9)$.

We now present simulation results for the projected gradient flow (6.13). For the simulation, we have separated the dynamics into two time scales, with x (the center of the hole) as the slow-scale variable and ψ the fast-scale variable. We first consider the case of the disk-shaped domain, that is, the dynamics (6.13) corresponds to hole placement for the disk-shaped domain to minimize λ_2 of the residual domain.

Figure 6.2 is a plot of $x(t)$, the path of the center of the hole, on the spatial domain, for two different initial conditions $x(0) = (0.4, 0.5)$ and $x(0) = (-0.5, -0.5)$. We observe that the hole center approaches the center of the disk with time, approximately along a straight line.

Figure 6.3 is a plot of $x(t)$, the path of the center of the hole (from the dynamics (6.13)) for a convex polygonal spatial domain. The final location of the hole is also indicated in the figure.

Figure 6.4 contains the results for a non-convex polygonal domain. The outer polygon

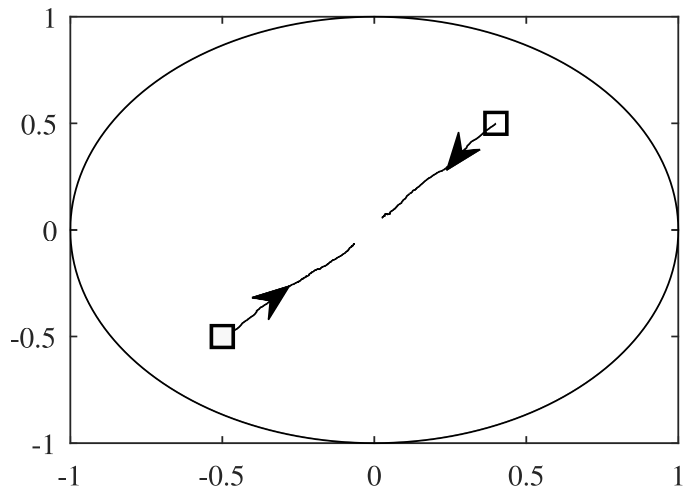

Figure 6.2. Path of the center of the hole, $x(t)$, from two different initial conditions $x(0) = (0.4, 0.5)$ and $x(0) = (-0.5, -0.5)$.

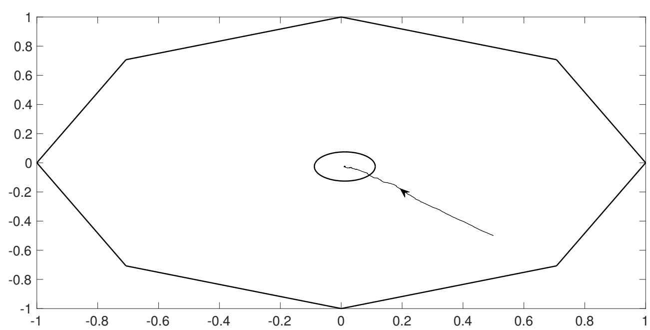

Figure 6.3. Path of the center of the hole, $x(t)$, from initial condition $x(0) = (0.5, -0.5)$ for a convex polygonal domain.

is the spatial domain Ω, while the inner polygon is the domain $\tilde{\Omega}$ (the set of allowed positions for the center of the hole). The heatmap shows the value of λ_2 of the residual domain (which was obtained by first sampling the domain uniformly at random at the points indicated by the tiny circles, placing the hole at those points, computing λ_2 of the residual domain, and then interpolating to obtain the plot). The paths of the center of the hole $x(t)$ (from the dynamics (6.13)) from different initial conditions are also plotted. The paths do not all converge to the same point in this case, but to a broader region (the darker region in the heatmap), which possibly contains more than one local minimum x^*.

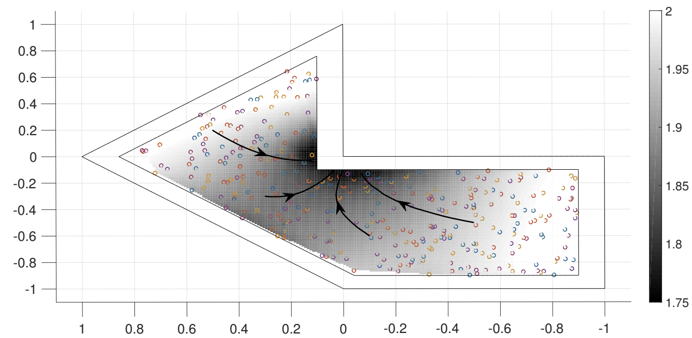

Figure 6.4. Paths of the center of the hole, $x(t)$ from different initial conditions.

In Figure 6.5, we present a numerical validation of the discussion in Remark 16. We first generated a random connected graph G with 50 nodes. We then computed the algebraic connectivities of the residual graphs obtained by the removal of one node from the graph $\lambda_2(L(G \setminus \{i\}))$, for each node, plotting it against the corresponding entry of the Fiedler eigenvector v_i^F (the eigenvector corresponding to the second eigenvalue of the Laplacian, or algebraic connectivity) of the original graph G. From the discussion in Remark 16, we expect that the local minima of $\lambda_2(L(G \setminus \{i\}))$ are concentrated around nodes corresponding to the entries of the Fiedler eigenvector of lowest magnitude, which is illustrated in the figure. We note that in the corresponding hole-placement problem, the nodal sets of the second eigenfunction ψ_2^Ω are only the local minimizers of $f(x) = \lambda_2^\Omega|\psi_2^\Omega(x)|^2 - |\nabla\psi_2^\Omega(x)|^2$. We thereby do not expect all the zero entries of the Fiedler eigenvector to correspond necessarily to global minimizers. However, the figure shows that the global minimum is indeed concentrated around nodes corresponding to the entries of the Fiedler eigenvector of lowest magnitude.

6.5 Summary

This chapter was devoted to the study of robustness of multi-agent networks, particularly to the problem of identifying the critical nodes for consensus in large-scale spatial networks,

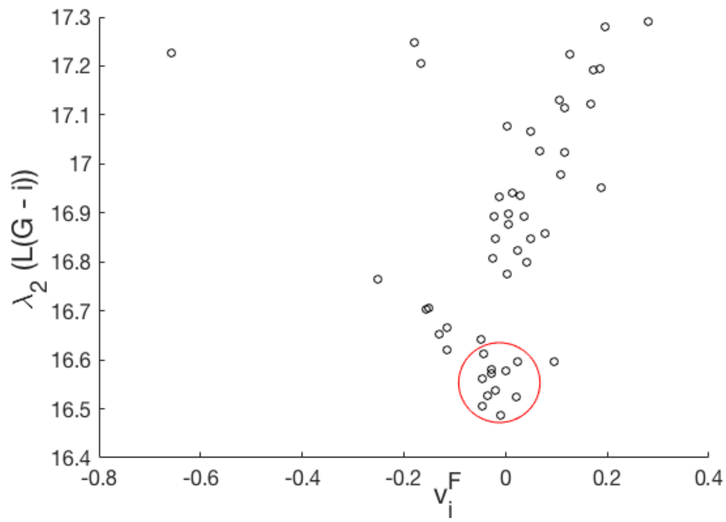

Figure 6.5. Plot of algebraic connectivity of residual network with the removal of one node vs. its corresponding entry in the Fiedler eigenvector, for a network with 50 nodes.

aided by an approximation of the Laplacian matrix of the graph by the Laplace operator on the domain. In addition to being a natural step in the large-N limit, the real advantage of the approximation is that it does not conceal the geometry of the problem, which is important for spatial networks such as swarms and sensor networks.

The material in this chapter, in full, is a reprint of the material in *Identification of critical nodes in large-scale spatial networks*, V. Krishnan and S. Martínez, IEEE Transactions on Control of Network Systems, 6(2), pp. 842–851, 2019. A preliminary version of the work appeared in the proceedings of the IFAC World Congress, Toulouse, France, July 2017 as *Identification of critical node clusters for consensus in large-scale spatial networks*, V. Krishnan and S. Martínez. The dissertation author was the primary investigator and author of these papers.

Chapter 7

Conclusion

In this dissertation, we began with a presentation of a multiscale theory of large-scale transport of multi-agent systems, and a framework for the design of (proximal) gradient descent-based algorithms. The proposed framework in its current form requires that the agents possess complete information about the states of the other agents. While specific instances of incomplete information were dealt with in the design of algorithms in the subsequent chapters, future work involves incorporating the notion of partial information into the formal framework . In subsequent chapters, we proposed and analyzed algorithms for multi-agent optimal transport, swarm self-organization and developed a unifying distributional framework for optimization-based state estimation, with an eye towards applications of multi-agent deployment, navigation and tracking. In the final chapter, we investigated the question of robustness of multi-agent systems, through the critical node identification problem for large-scale spatial networks.

Bibliography

[1] W. Abbas, A. Laszka, Y. Vorobeychik, and X. Koutsoukos. Improving network connectivity using trusted nodes and edges. In *American Control Conference*, pages 328–333, 2017.

[2] F. Albertini and D. D'Alessandro. Observability and forward–backward observability of discrete-time nonlinear systems. *Mathematics of Control, Signals and Systems*, 15(4):275–290, 2002.

[3] A. Alessandri, M. Baglietto, and G. Battistelli. Moving-horizon state estimation for nonlinear discrete-time systems: New stability results and approximation schemes. *Automatica*, 44(7):1753–1765, 2008.

[4] A. Alessandri and M. Gaggero. Moving-horizon estimation for discrete-time linear and nonlinear systems using the gradient and newton methods. In *IEEE Int. Conf. on Decision and Control*, page 2906–2911, 2016.

[5] A. Alessandri and M. Gaggero. Fast moving horizon state estimation for discrete-time systems using single and multi iteration descent methods. *IEEE Transactions on Automatic Control*, 62(9):4499–4511, 2017.

[6] L. Ambrosio. Lecture notes on optimal transport problems. In *Mathematical aspects of evolving interfaces*, pages 1–52. Springer, 2003.

[7] L. Ambrosio. Transport equation and Cauchy problem for non-smooth vector fields. In *Calculus of variations and nonlinear partial differential equations*, page 1–41. Springer, 2008.

[8] L. Ambrosio, N. Gigli, and G. Savaré. *Gradient flows: in metric spaces and in the space of probability measures*. Springer, 2008.

[9] C. Andrieu, A. Doucet, and R. Holenstein. Particle Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society: Series B*, 72(3):269–342, 2010.

[10] A. Arulselvan, C. Commander, L. Elefteriadou, and P. Pardalos. Detecting critical nodes in sparse graphs. *Computers and Operations Research*, 36(7):2193–2200, 2009.

[11] R. Atar and K. Burdzy. On nodal lines of Neumann eigenfunctions. *Electronic Communications in Probability*, 7:129–139, 2002.

[12] B. Açıkmeşe and D. S. Bayard. Markov chain approach to probabilistic guidance for swarms of autonomous agents. *Asian Journal of Control*, 17(4):1105–1124, 2015.

[13] J. Bachrach, J. Beal, and J. McLurkin. Composable continuous-space programs for robotic swarms. *Neural Computing and Applications*, 19(6):825–847, 2010.

[14] S. Bandyopadhyay, S. J. Chung, and F. Y. Hadaegh. Inhomogeneous Markov chain approach to probabilistic swarm guidance algorithms. In *Int. Conf. on Spacecraft Formation Flying Missions and Technologies*, page 1–13, 2013.

[15] S. Bandyopadhyay, S. J. Chung, and F. Y. Hadaegh. Probabilistic swarm guidance using optimal transport. In *IEEE Conf. on Control Applications*, page 498–505, 2014.

[16] 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(5):1103–1123, 2017.

[17] M. Belkin and P. Niyogi. Towards a theoretical foundation for Laplacian-based manifold methods. In *Int. Conf. on Computational Learning Theory*, page 486–500, 2005.

[18] M. Belkin and P. Niyogi. Convergence of Laplacian eigenmaps. *Advances in Neural Information Processing Systems*, 19:129–136, 2007.

[19] M. Belkin and P. Niyogi. Towards a theoretical foundation for Laplacian-based manifold methods. *Journal of Computer and System Sciences*, 74(8):1289 – 1308, 2008.

[20] J. D. Benamou and Y. Brenier. A computational fluid mechanics solution to the Monge-Kantorovich mass transfer problem. *Numerische Mathematik*, 84(3):375–393, 2000.

[21] J. D. Benamou, G. Carlier, M. Cuturi, L. Nenna, and G. Peyré. Iterative Bregman projections for regularized transportation problems. *SIAM Journal on Scientific Computing*, 37(2):A1111–A1138, 2015.

[22] S. Berman, A. Halász, M. A. Hsieh, and V. Kumar. Optimized stochastic policies for task allocation in swarms of robots. *IEEE Transactions on Robotics*, 25, 2009.

[23] P. Billingsley. *Probability and measure*. John Wiley, 2008.

[24] P. Billingsley. *Convergence of probability measures*. John Wiley, 2013.

[25] V. Bogachev. *Measure theory*, volume 2. Springer, 2007.

[26] D. Bourne, B. Schmitzer, and B. Wirth. Semi-discrete unbalanced optimal transport and quantization. *arXiv preprint arXiv:1808.01962*, 2018.

[27] F. Bullo, J. Cortés, and S. Martínez. *Distributed Control of Robotic Networks*. Applied Mathematics Series. Princeton University Press, 2009.

[28] S. Camazine. *Self-organization in biological systems*. Princeton University Press, 2003.

[29] I. Chattopadhyay and A. Ray. Supervised self-organization of homogeneous swarms using ergodic projections of markov chains. *IEEE Transactions on Systems, Man, & Cybernetics. Part B: Cybernetics*, 39, 2009.

[30] Y. Chen, T. Georgiou, and M. Pavon. On the relation between optimal transport and Schrödinger bridges: A stochastic control viewpoint. *Journal of Optimization Theory & Applications*, 169(2):671–691, 2016.

[31] Y. Cheng. Mean shift, mode seeking, and clustering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17, 1995.

[32] A. J. Chorin and J. E. Marsden. *A mathematical introduction to fluid mechanics*, volume 3. Springer, 1990.

[33] J. Cortés. Coverage optimization and spatial load balancing by robotic sensor networks. *IEEE Transactions on Automatic Control*, 55(3):749–754, 2010.

[34] J. Cortés, G. E. Dullerud, S. Han, J. Le Ny, S. Mitra, and G. J. Pappas. Differential privacy in control and network systems. In *IEEE Int. Conf. on Decision and Control*, pages 4252–4272, Las Vegas, NV, 2016.

[35] J. Cortés, S. Martínez, T. Karatas, and F. Bullo. Coverage control for mobile sensing networks. *IEEE Transactions on Robotics and Automation*, 20(2):243–255, 2004.

[36] J. Cortés, S. Martínez, T. Karatas, and F. Bullo. Coverage control for mobile sensing networks. *IEEE Transactions on Robotics and Automation*, 20(2):243–255, 2004.

[37] M. Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in Neural Information Processing Systems*, page 2292–2300, 2013.

[38] M. Cuturi and A. Doucet. Fast computation of Wasserstein barycenters. In *Int. Conf. on Machine Learning*, page 685–693, Beijing, China, 2014.

[39] M. Cuturi and G. Peyré. Semidual regularized optimal transport. *SIAM Review*, 60(4):941–965, 2018.

[40] M. L. de Cristoforis. Simple Neumann eigenvalues for the Laplace operator in a domain with a small hole. *Revista Matematica Complutense*, 25(2):369–412, 2012.

[41] N. Demir, U. Eren, and B. Açıkmeşe. Decentralized probabilistic density control of autonomous swarms with safety constraints. *Autonomous Robots*, 39(4):537–554, 2015.

[42] V. Deshmukh, K. Elamvazhuthi, S. Biswal, Z. Kakish, and S. Berman. Mean-field stabilization of markov chain models for robotic swarms: Computational approaches and experimental results. 3(3):1985–1992, 2018.

[43] C. Dimitrakakis, B. Nelson, A. Mitrokotsa, and B. Rubinstein. Robust and private Bayesian inference. In *Int. Conf. on Algorithmic Learning Theory*, page 291–305, 2014.

[44] C. Dimitrakakis, B. Nelson, Z. Zhang, A. Mitrokotsa, and B. Rubinstein. Bayesian differential privacy through posterior sampling. 1050:23, 2016.

[45] A. Doucet, N. D. Freitas, and N. Gordon. Sequential Monte Carlo methods in practice. Springer, 2001.

[46] C. Dwork and A. Roth. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3–4):211–407, 2014.

[47] M. Dziubiński and S. Goyal. Network design and defence. *Games and Economic Behavior*, 79:30–43, 2013.

[48] J. Eells and L. Lemaire. Deformations of metrics and associated harmonic maps. *Proceedings Mathematical Sciences*, 90, 1981.

[49] K. Elamvazhuthi and S. Berman. Optimal control of stochastic coverage strategies for robotic swarms. In *IEEE Int. Conf. on Robotics and Automation*, pages 1822–1829, 2015.

[50] K. Elamvazhuthi, S. Biswal, and S. Berman. Mean-field stabilization of robotic swarms to probability distributions with disconnected supports. In *American Control Conference*, pages 885–892, 2018.

[51] K. Elamvazhuthi, M. Kawski, S. Biswal, V. Deshmukh, and S. Berman. Mean-field controllability and decentralized stabilization of markov chains. In *IEEE Int. Conf. on Decision and Control*, pages 3131–3137, 2017.

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

[53] U. Eren and B. Açıkmeşe. Velocity field generation for density control of swarms using heat equation and smoothing kernels. In *IFAC Papers Online*, volume 50, page 9405–9411, 2017.

[54] K. Eswaran and R. Tarjan. Augmentation problems. *SIAM Journal on Computing*, 5(4):653–665, 1976.

[55] L. C. Evans. *Partial differential equations*. Graduate studies in mathematics. American Mathematical Society, Providence (R.I.), 1998.

[56] F. Farokhi, J. Milosevic, and H. Sandberg. Optimal state estimation with measurements corrupted by laplace noise. In *IEEE Int. Conf. on Decision and Control*, page 302–307, 2016.

[57] S. Ferrari, G. Foderaro, P. Zhu, and T. A. Wettergren. Distributed optimal control of multiscale dynamical systems: a tutorial. *IEEE Control Systems*, 36(2):102–116, 2016.

[58] M. Fiedler. Algebraic connectivity of graphs. *Czechoslovak Mathematical Journal*, 23(98):298–305, 1973.

[59] G. Foderaro, S. Ferrari, and T. A. Wettergren. Distributed optimal control for multi-agent trajectory optimization. *Automatica*, 50, 2014.

[60] P. Frihauf and M. Krstic. Leader-enabled deployment onto planar curves: A PDE-based approach. *IEEE Transactions on Automatic Control*, 56(8):1791–1806, 2011.

[61] A. Genevay, M. Cuturi, G. Peyré, and F. Bach. Stochastic optimization for large-scale optimal transport. In *Advances in Neural Information Processing Systems*, page 3440–3448, 2016.

[62] A. L. Gibbs and F. E. Su. On choosing and bounding probability metrics. *International Statistical Review*, 70(3):419–435, 2002.

[63] W. Gilks, S. Richardson, and D. Spiegelhalter. *Markov chain Monte Carlo in practice*. Chapman and Hall/CRC, 1995.

[64] P. Goyal and E. Ferrara. Graph embedding techniques, applications, and performance: A survey. *arXiv preprint arXiv:1705.02801*, 2017.

[65] R. Hall, A. Rinaldo, and L. Wasserman. Differential privacy for functions and functional data. *Journal of Machine Learning Research*, 14:703–727, 2013.

[66] V. Hartmann and D. Schuhmacher. Semi-discrete optimal transport-the case p= 1. *arXiv preprint arXiv:1706.07650*, 2017.

[67] W. K. Hastings. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57(1):97–109, 1970.

[68] A. Henrot. Continuity with respect to the domain for the laplacian: a survey. *Control and Cybernetics Journal*, 23(3):427–443, 1994.

[69] A. Henrot. *Extremum problems for eigenvalues of elliptic operators*. Springer Science & Business Media, 2006.

[70] D. Henry. *Geometric theory of semilinear parabolic equations*. Springer, 1981.

[71] J. Hespanha, D. Liberzon, D. Angeli, and E. Sontag. Nonlinear norm-observability notions and stability of switched systems. *IEEE Transactions on Automatic Control*, 50(2):154–168, 2005.

[72] A. Howard, M. J. Matarić, and G. S. Sukhatme. Mobile sensor network deployment using potential fields: A distributed, scalable solution to the area coverage problem. In *Int. Symposium on Distributed Autonomous Robotic Systems*, page 299–308. Springer, 2002.

[73] W. Hu. Robust stability of optimization-based state estimation. *arXiv preprint arXiv:1702.01903*, 2017.

[74] F. Hélein. *Harmonic Maps, Conservation Laws and Moving Frames*. Cambridge University Press, second edition, 2002.

[75] A. Jamakovic and P. V. Mieghem. On the robustness of complex networks by using the algebraic connectivity. page 183–194, 2008.

[76] A. Jazwinski. Limited memory optimal filtering. *IEEE Transactions on Automatic Control*, 13(5):558–563, 1968.

[77] L. Ji, J. Rawlings, W. Hu, A. Wynn, and M. Diehl. Robust stability of moving horizon estimation under bounded disturbances. *IEEE Transactions on Automatic Control*, 61(11):3509–3514, 2016.

[78] A. Kirsch. *An introduction to the mathematical theory of inverse problems*, volume 120. Springer, 2011.

[79] T. Kolokolnikov, M. S. Titcombe, and M. J. Ward. Optimizing the fundamental Neumann eigenvalue for the Laplacian in a domain with small traps. *European Journal of Applied Mathematics*, 16(2):161– 200, 2005.

[80] V. Krishnan and S. Martínez. Distributed control for spatial self-organization of multi-agent swarms. *SIAM Journal on Control and Optimization*, 56(5):3642–3667, 2018.

[81] M. Krstic and A. Smyshlyaev. *Boundary control of PDEs: A course on backstepping designs*, volume 16. SIAM, 2008.

[82] M. Kuang and E. Tabak. Sample-based optimal transport and barycenter problems. *Communications on Pure and Applied Mathematics*, 2017. Submitted.

[83] T. S. Lee, K. P. Dunn, and C. B. Chang. On observability and unbiased estimation of nonlinear systems. In *System Modeling and Optimization*, page 258–266. 1982.

[84] G. Leoni. *A first course in Sobolev spaces*, volume 105. American Mathematical Society, 2009.

[85] C. Liu, S. Chakraborty, and P. Mittal. Dependence makes you vulnberable: Differential privacy under dependent tuples. In *Network and Distributed System Security Symposium (NDSS)*, volume 16, page 21–24, 2016.

[86] L. Lovász. *Large networks and graph limits*, volume 60 of *Coloquium Publications*. American Mathematical Society, 2012.

[87] R. McCann. Existence and uniqueness of monotone measure-preserving maps. *Duke Mathematical Journal*, 80(2):309–324, 1995.

[88] M. Mesbahi and M. Egerstedt. *Graph Theoretic Methods in Multiagent Networks*. Princeton Series in Applied Mathematics. Princeton University Press, 2010.

[89] A. Mesquita, J. Hespanha, and K. Åström. Optimotaxis: A stochastic multi-agent optimization procedure with point measurements. In *Hybrid systems: Computation and Control*, page 358–371, 2008.

[90] N. Metropolis, A. Rosenbluth, M. Rosenbluth, A. Teller, and E. Teller. Equation of state calculations by fast computing machines. *Journal of Chemical Physics*, 21(6):1087–1092, 1953.

[91] T. Mikami and M. Thieullen. Optimal transportation problem by stochastic optimal control. *SIAM Journal on Control and Optimization*, 47(3):1127–1139, 2008.

[92] R. Murray, B. Swenson, and S. Kar. Revisiting normalized gradient descent: Evasion of saddle points. *IEEE Transactions on Automatic Control*, 2019. Available Online.

[93] Q. Mérigot. A multiscale approach to optimal transport. In *Computer Graphics Forum*, volume 30, page 1583–1592, 2011.

[94] M. Müller. Nonlinear moving horizon estimation in the presence of bounded disturbances. *Automatica*, 79:306–314, 2017.

[95] H. Nijmeijer. Observability of autonomous discrete time non-linear systems: a geometric approach. *International Journal of Control*, 36(5):867–874, 1982.

[96] E. Nozari, P. Tallapragada, and J. Cortés. Differentially private distributed convex optimization via functional perturbation. *IEEE Transactions on Control of Network Systems*, 5(1):395–408, 2018.

[97] J. Le Ny and G. Pappas. Differentially private filtering. *IEEE Transactions on Automatic Control*, 59(2):341–354, 2014.

[98] N. Papadakis, G. Peyré, and E. Oudet. Optimal transport with proximal splitting. *SIAM Journal on Imaging Sciences*, 7(1):212–238, 2014.

[99] G. Peyré. Entropic approximation of Wasserstein gradient flows. *SIAM Journal on Imaging Sciences*, 8(4):2323–2351, 2015.

[100] G. Peyré and M. Cuturi. Computational optimal transport. Technical report, 2017.

[101] J. Qi, R. Vazquez, and M. Krstic. Multi-agent deployment in 3-D via PDE control. *IEEE Transactions on Automatic Control*, 60, 2015.

[102] C. Rao, J. Rawlings, and D. Mayne. Constrained state estimation for nonlinear discrete-time systems: Stability and moving horizon approximations. *IEEE Transactions on Automatic Control*, 48(2):246–258, 2003.

[103] C. Robert and G. Casella. *Monte Carlo statistical methods*. Springer, 2013.

[104] M. Rubenstein, C. Ahler, N. Hoff, A. Cabrera, and R. Nagpal. Kilobot: A low cost robot with scalable operation designed for collective behaviors. *Robotics and Autonomous Systems*, 62, 2014.

[105] M. Rubenstein, A. Cornejo, and R. Nagpal. Programmable self-assembly in a thousand-robot swarm. *Science*, 345, 2014.

[106] F. Santambrogio. *Optimal transport for applied mathematicians*. Springer, 2015.

[107] F. Santambrogio. {Euclidean, metric, and Wasserstein} gradient flows: an overview. *Bulletin of Mathematical Sciences*, 7(1):87–154, 2017.

[108] M. Schwager, D. Rus, and J. J. Slotine. Decentralized, adaptive coverage control for networked robots. *International Journal of Robotics Research*, 28(3):357–375, 2009.

[109] V. Seguy, B. Damodaran, R. Flamary, N. Courty, A. Rolet, and M. Blondel. Large-scale optimal transport and mapping estimation. *arXiv preprint arXiv:1711.02283*, 2017.

[110] B. Shaw and T. Jebara. Structure preserving embedding. In *Int. Conf. on Machine Learning*, page 937–944, 2009.

[111] M. Sheng, J. Li, and Y. Shi. Critical nodes detection in mobile ad hoc network. In *Int. Conf. on Advanced Information Networking and Applications*, volume 2, 2006.

[112] G. D. Smith. *Numerical solution of partial differential equations: finite difference methods*. Oxford University Press, 1985.

[113] S. Song, Y. Wang, and K. Chaudhuri. Pufferfish privacy mechanisms for correlated data. page 1291–1306. ACM, 2017.

[114] E. Tabak and G. Trigila. Data-driven optimal transport. *Communications on Pure and Applied Mathematics*, 69(4):613–648, 2016.

[115] J. Tsinias and C. Kitsos. Observability and state estimation for a class of nonlinear systems. *arXiv preprint arXiv:1803.08386*, 2018.

[116] M. Ventresca and D. Aleman. Efficiently identifying critical nodes in large complex networks. *Computational Social Networks*, 2(1), 2015.

[117] C. Villani. *Optimal transport: old and new*, volume 338. Springer, 2008.

[118] J. A. Walker. Some results on Liapunov functions and generated dynamical systems. *Journal of Differential Equations*, 30(3):424–440, 1978.

[119] J.A. Walker. *Dynamical Systems and Evolution Equations: Theory and Applications*, volume 20. Springer, 2013.

[120] G. M. Whitesides and B. Grzybowski. Self-assembly at all scales. *Science*, 295, 2002.

[121] A. Wynn, M. Vukov, and M. Diehl. Convergence guarantees for moving horizon estimation based on the real-time iteration scheme. *IEEE Transactions on Automatic Control*, 59(8):2215–2221, 2014.

[122] I. Zang and M. Avriel. On functions whose local minima are global. *Journal of Optimization Theory & Applications*, 16(3-4):183–190, 1975.

[123] F. Zhang, A. Bertozzi, K. Elamvazhuthi, and S. Berman. Performance bounds on spatial coverage tasks by stochastic robotic swarms. *IEEE Transactions on Automatic Control*, 63(6):1563–1578, 2018.