Self-Organization of Multi-Agent Systems Using Markov Chain Models

by

Shiba Biswal

A Dissertation Presented in Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy

Approved April 2020 by the
Graduate Supervisory Committee:

Spring Berman, Chair
Georgios Fainekos
Nicolas Lanchier
Marc Mignolet
Matthew Peet

ARIZONA STATE UNIVERSITY

May 2020

ABSTRACT

The problem of modeling and controlling the distribution of a multi-agent system has recently evolved into an interdisciplinary effort. When the agent population is very large, i.e., at least on the order of hundreds of agents, it is important that techniques for analyzing and controlling the system scale well with the number of agents. One scalable approach to characterizing the behavior of a multi-agent system is possible when the agents' states evolve over time according to a Markov process. In this case, the density of agents over space and time is governed by a set of difference or differential equations known as a *mean-field model*, whose parameters determine the stochastic control policies of the individual agents. These models often have the advantage of being easier to analyze than the individual agent dynamics. Mean-field models have been used to describe the behavior of chemical reaction networks, biological collectives such as social insect colonies, and more recently, swarms of robots that, like natural swarms, consist of hundreds or thousands of agents that are individually limited in capability but can coordinate to achieve a particular collective goal.

This dissertation presents a control-theoretic analysis of mean-field models for which the agent dynamics are governed by either a continuous-time Markov chain on an arbitrary state space, or a discrete-time Markov chain on a continuous state space. Three main problems are investigated. First, the problem of stabilization is addressed, that is, the design of transition probabilities/rates of the Markov process (the agent control parameters) that make a target distribution, satisfying certain conditions, invariant. Such a control approach could be used to achieve desired multi-agent distributions for spatial coverage and task allocation. However, the convergence of the multi-agent distribution to the designed equilibrium does not imply the convergence of the individual agents to fixed states. To prevent the agents from continuing

to transition between states once the target distribution is reached, and thus potentially waste energy, the second problem addressed within this dissertation is the construction of feedback control laws that prevent agents from transitioning once the equilibrium distribution is reached. The third problem addressed is the computation of optimized transition probabilities/rates that maximize the speed at which the system converges to the target distribution.

DEDICATION

To my Parents and Karthik

ACKNOWLEDGEMENTS

I would first like to thank my advisor Dr. Spring Berman whose own Ph.D. research formed the basis of this research. She has been a constant source of encouragement throughout my student years. I am especially grateful to her for allowing me the flexibility to take this research in a direction of my choosing.

Secondly, I would like to thank Dr. Matthew Peet for teaching a fundamental course in controls. The second chapter of this dissertation was an outcome of a class project that I completed in this class. I would also like to thank Dr. Georgios Fainekos, for teaching a wonderful class, and being on my committee.

A heartful thanks to Dr. Nicolas Lanchier for our collaboration, our many conversations, and for teaching me probability theory. The last chapter of this dissertation is an outcome of our collaboration. I feel immense gratitude toward Dr. Marc Mignolet, who has truly been an inspiration, as a mentor, researcher and a person. He has been constant source of support since my first semester in ASU.

I would like to acknowledge Dr. Mittelmann's valuable inputs that contributed to the research presented in Chapter six. This thesis would not have been possible without the collaboration with my fiancee Karthik Elamvazhuthi. Many of the projects discussed within this dissertation were inspired by his own PhD.

I am thankful to my friends, Deepthi, Zz, Zahi, Ragesh, Aniket, Rakshith, Hamed, Azadeh, Sean, for making these years enjoyable. I will forever cherish our Friday evening beer sessions.

Finally, I would like to acknowledge my parent's and sister's unconditional love and support.

TABLE OF CONTENTS

CHAPTER Page

CHAPTER Page

LIST OF SYMBOLS

$\bar{\mathbb{R}}_+,\ \mathbb{R}_+$	$[0,\infty),\ (0,\infty)$ respectively
$\bar{\mathbb{Z}}_+,\ \mathbb{Z}_+$	Set of all non-negative, positive integers respectively
$\mathbb{C}$	Set of complex numbers
$B_\delta(x)$	Closed ball in $\mathbb{R}^d$ of radius δ centered at x
$\partial E,\ \text{int}(E)$	Boundary and interior, respectively, of a set E
$\lvert E\rvert$	Cardinality of a set E
$L^p(\mathcal{X},\nu)$, $p\in[1,\infty)$	$\{f:\mathcal{X}\to\mathbb{C} : f \text{ is measurable and } \lVert f\rVert_p = \int \lvert f\rvert^p d\nu^{1/p} < \infty\}$
$L^\infty(\mathcal{X},\nu)$	$\{f:\mathcal{X}\to\mathbb{C} : f \text{ is measurable and } \lVert f\rVert_\infty = \operatorname*{ess\,sup}_{x\in\mathcal{X}} \lvert f(x)\rvert < \infty\}$
$C(\mathcal{X})$	Space of continuous functions f on $\mathcal{X}$
$C_0(\mathcal{X})$	Space of continuous functions f on $\mathcal{X}$ that *vanish at* infinity i.e., for every $\epsilon>0,\ \{x : \lvert f(x)\rvert \geq \epsilon\}$ is compact
$\text{supp}(f)$	Closure of the set of points where f is nonzero
$\chi_A(\cdot)$	Characteristic function on a set $A :=$ $\chi_A(x)=1$ if $x\in A,\ \chi_A(x)=0$ otherwise
$\delta_x(A)$	Dirac measure at $x := \delta_x(A)=1$ if $x\in A,\ \delta_x(A)=0$ otherwise
$\nu \ll \mu$	Measure ν is absolutely continuous with respect to measure μ; that is, $\nu(E)=0$ whenever $\mu(E)=0$
m	Lebesgue measure on $\mathbb{R}$
a.e.	Almost Everywhere; a property holds almost everywhere if it holds for all elements in a set except a subset of measure zero

Chapter 1

INTRODUCTION

Over approximately the past two decades, research on the control of multi-agent systems has gained increasing attention and has evolved into an established, interdisciplinary field that is of interest to engineers, mathematicians, and biologists alike. Graph-theoretic approaches have been widely used in the analysis and control of multi-agent systems, and there are numerous works on such approaches; (Bullo *et al.*, 2009; Lewis *et al.*, 2013; Mesbahi and Egerstedt, 2010) are excellent books dedicated to such methods. However, many of these control approaches do not scale well to very large agent populations. When all agents follow the same dynamics and these dynamics are independent of agents' identities, an alternative approach is to study a fluid approximation of the swarm. This approach is inspired by the *Eulerian* modeling methodology used in fluid dynamics. This entails looking at fluid motion focusing on specific locations in space through which the fluid flows, or a *macroscopic* perspective. Contrary to this approach, the *Lagrangian* method entails observing an individual fluid particle as it moves through space and time, or a *microscopic* perspective. A similar modeling methodology is also employed in statistical mechanics (Van Kampen, 1992) and biology (Brauer *et al.*, 2012) to treat a large collection of entities as a continuum. The term *self-organization* has been used in biology to denote the emergence of ordered behavior in groups of interacting entities. Schools of fish display such self-organized behavior. For example, without 'leader fish,' each individual fish in a school is able to localize with respect to its neighbors, and the school can form a variety of complex patterns (Hemelrijk and Hildenbrandt, 2012).

Taking inspiration from the models discussed in the previous paragraph, in the

models considered within this dissertation, the primary assumption is that every agent follows an identical Markov process, and then the macroscopic behavior of the population is determined by the *Kolmogorov forward equation* corresponding to the Markov process. For a population of N agents, the state space of the forward equation is dependent on N; however, in the limit N tends to infinity, the N-agent forward equation is replaced by a single forward equation defined on the set of probability densities, called the *mean-field model*. Because such models are independent of the agent population size, the analysis and control of these models is scalable with the number of agents, which is an advantage over many existing multi-agent control techniques. A survey on mean-field models in the context of control of multi-agent systems is presented in Elamvazhuthi and Berman (2019). Moreover, a range of tools are available to analyze and control mean-field dynamical models, which have the advantage of linearity in the absence of agent interactions, or are easier to analyze than the corresponding agent-based models. A drawback of using these models is that for low values of N, they are not accurate representations of the N-agent system; as N becomes larger, the approximation tends to be better. We will observe that although stability and convergence results are proven for the mean-field model, simulations of the corresponding N-agent system demonstrate that for relatively small numbers of agents (specific estimates of the order are provided within the chapters), the agent distribution does in fact follow the prediction of the corresponding mean-field model.

We will now state the main contributions of this dissertation.

1.1 Contributions

The goals of this research are threefold. We shall list these next.

1. It is well-established that a Markov chain that is *irreducible* and *positive recurrent* has a unique stationary distribution. A question that is investigated for all

models considered within this dissertation is the following inverse problem: *Given a target probability distribution, can a Markov chain be constructed such that its stationary distribution is this target distribution?* From a control-theoretic perspective, this implies requires proving the *stabilizability* of the Kolmogorov forward equation. A potential application of this work is to control the distribution of large-scale multi-agent systems over discrete or continuous state spaces that could, for instance, represent a set of distinct locations or a domain of interest. Such a control approach could be used, for example, in environmental monitoring, surveillance, disaster response, and autonomous construction. Chapters 2-6 discuss the construction of control methodologies for allocating a swarm of robots among a set of states in a decentralized fashion, using minimal local information or none at all.

The stabilization of macroscopic models to arbitrary measures can alternatively be seen as an extension of the classical control problem of asymptotically stabilizing a system to a point, which corresponds to the Dirac measure in the measure stabilization framework.

In Chapter 2, the time evolution of each robot's state is modeled as a continuous-time Markov chain (CTMC) on a finite graph. We prove that by using the *transition rates* as control parameters, any distribution that is strictly positive everywhere can be stabilized. Specifically, for bidirected graphs, we construct polynomial density feedback laws that stabilize strictly positive stationary distributions. In Chapters 3 and 5, we address the problem of stabilizing a discrete-time Markov chain (DTMC) to a target invariant distribution, using the *transition kernel* as control parameter. The transition kernel is a continuous state space analogue of the transition probability matrix of a DTMC on a discrete state space. We investigate the distribution stabilization problem in two settings: in Chapter 3, the state space is restricted to compact subsets of the Euclidean space, and in Chapter 5, the state space is restricted

to compact manifolds. In the former case, we assume that the target measure is *supported* a.e. on the entire state space of the system and is *absolutely continuous* with respect to the Lebesgue measure, with $L^\infty(\cdot)$ density. We additionally assume that agents evolve according to discrete-time deterministic nonlinear dynamics. In the latter case where the state space is a manifold, we address the stabilization problem for a DTMC that evolves on a compact, smooth, connected manifold without boundary. We consider target measures that are supported a.e. on the entire state space of the system and are absolutely continuous with respect to the *Riemannian volume* with $L^\infty(\cdot)$ density.

2. Having obtained the above stabilizability results, the next problem that we investigate is the construction of transition kernel/rates, the control parameters of the Markov process, that maximize the system's convergence rate to the target distribution. In Chapters 3 and 5, we formulate an infinite-dimensional convex optimization problem to construct feedback control laws that stabilize the system to the target invariant measure at a maximized rate of convergence. This problem has been previously solved in Boyd *et al.* (2004) and Berman *et al.* (2009) for DTMCs and CTMCs, respectively, that are defined on a finite graph with the uniform distribution as their invariant distribution. Similar to these earlier works, we pose the optimization problem in terms of maximizing the spectral gap of the operator that pushes forward measures, also known from here on as the *forward operator*. Since the first eigenvalue of the forward operator is 1 or 0 when time is discrete or continuous, respectively, maximizing the spectral gap reduces to the problem of minimizing the modulus of the second-largest eigenvalue. The optimization problem admits an exact solution in the case where the manifold is a *Lie group* and the target measure is uniform. In Chapter 6, we invoke the *min-max principle* to characterize the modulus of the second-largest eigenvalue of the forward operator for a Markov process that evolves in continuous

time on a continuous state space. Hence, unlike in Chapters 3 and 5, the optimization problem admits an exact solution for arbitrary distributions that are in $L^\infty(\cdot)$.

3. The convergence of a Markov process to an equilibrium distribution does not necessarily imply that the agents evolving according to the process also converge to equilibrium states. In fact, agents may continue to transition between states, which can cause them to waste energy. To prevent agents from continuing to switch between states at the equilibrium distribution, the third goal of this dissertation is to construct a Markov process such that its forward operator is the identity or zero operator at equilibrium, depending on whether time is discrete or continuous. This results in a density-dependent operator that is a function of the distribution and gives rise to a nonlinear Markov process. Moreover, since the operator must be dependent on the distribution, we construct the operator to have a decentralized structure. An operator with this structure corresponds to control laws that require each agent to estimate the population only in its local neighborhood, rather than obtain feedback on the entire agent distribution. In Chapter 2, in which the state space is a finite graph, we construct distribution-dependent transition rates (feedback control laws) that evaluate to zero at the desired equilibrium distribution. This enables the construction of decentralized density feedback controllers, using tools from Sum-of-Squares (SOS)-based polynomial optimization (Prajna *et al.*, 2002), that globally stabilize a swarm of agents modeled as a CTMC to a target state distribution with no state-switching at equilibrium. In Chapter 4, we address this problem for a DTMC that evolves on a compact, connected subset of $\mathbb{R}^d$. The desired distribution is assumed to be in $L^\infty(\cdot)$.

Finally, in Chapter 7, we address the problem of redistributing a swarm of agents over a discrete state space while accounting for the physical constraint that agents must avoid collisions with one another. In this case, we develop an agent-based

model of the swarm using tools from the theory of *interacting particle systems (IPS).* We consider a swarm of K agents that move randomly on the vertex set of a finite connected graph, with at most one agent per vertex, i.e., two agents cannot occupy the same vertex. An agent at vertex x chooses one of the $\deg(x)$ neighboring vertices uniformly at random at a defined transition rate, and jumps to that vertex if and only if it is empty. Using standard probability techniques, the set of invariant measures of this process so defined is identified in order to study the *occupation time* of each vertex, defined as the expected time during which the vertex will be occupied at equilibrium.

1.2 Literature Review

We will now review some of the important works that have contributed to this field of research.

We begin with the discrete state space setting; we will review works wherein the time variable is both discrete and continuous. To stabilize DTMCs or CTMCs to target distributions, several approaches have been developed in which the agents are programmed to switch stochastically between states at tunable transition probabilities or rates; therefore, the transition probabilities or rates are the control parameters of the Markov process. For DTMCs on discrete state spaces, the types of distributions that can be stabilized are well-understood (Açıkmeşe and Bayard, 2015; Elamvazhuthi *et al.*, 2017); this stabilizability result follows from the classical *Perron-Frobenius* theorem (Berman and Plemmons, 1994), which gives a sufficient condition for the uniqueness of the stationary distribution of a Markov chain. In the continuous time setting, optimized transition rates were computed to maximize the rate of convergence to target distributions in Berman *et al.* (2009) for a multi-robot stochastic task allocation scenario. There is a limitation on the kind of probability distributions

that can be stabilized using open-loop control; specifically, distributions that have strongly connected support can be stabilized (Elamvazhuthi *et al.*, 2019a). This assumption, however, was relaxed in Elamvazhuthi *et al.* (2018a) by using discontinuous control laws. By allowing the control laws to be time-dependent, additional control objectives can be achieved. For example, to reduce agent transitions between states at equilibrium Hsieh *et al.* (2008); Mather and Hsieh (2014) construct control laws for a CTMC that are dependent on the population density. A similar problem was addressed for the discrete-time case in Bandyopadhyay *et al.* (2017). In Demir *et al.* (2015), the authors construct time-dependent feedback laws for a DTMC in order to impose safety constraints on the system in terms of an upper bound on the agent density at each time step.

We will next consider the continuous state space setting. The distribution stabilization problem that we address in Chapters 3-5 is closely related to a class of control problems that have been investigated in the context of *mean-field games* (Lasry and Lions, 2007; Bensoussan *et al.*, 2013; Carmona and Delarue, 2018) and optimal transport theory (Villani, 2003; Santambrogio, 2015). In mean-field games, the control problem is to design a feedback control law that is a function of the agent's state, with the goal of optimizing an objective functional that is a function of the agent's state and the probability density of its position over time. The mean-field game problem for agent dynamics evolving in discrete time and continuous space is considered in Saldi *et al.* (2018). Works on mean-field game theory, which has only recently been applied to problems in swarm robotics (Liu *et al.*, 2018), use optimal control techniques to construct policies for strategic decision-making in very large populations of interacting agents. However, control problems in the mean-field games literature usually do not include constraints on the long-time behavior of this probability density, as we do in Chapters 3-5.

Similar to mean-field games, *optimal transport theory* considers a class of measure control problems in which the goal is to construct a map from the state space to itself that pushes forward an initial measure to a target measure while optimizing a given cost function. According to the *Benamou-Brenier* formulation of optimal transport, when the cost function is quadratic, the problem can be framed as a control problem for an advection equation with the velocity field as the control input. For scenarios where the measure represents the distribution of a swarm of agents, this classical version of the optimal transport problem corresponds to agents with single-integrator dynamics. There has been some recent work on extending results on optimal transport to agents that evolve in continuous time with linear dynamics (Hindawi *et al.*, 2011; Chen *et al.*, 2017) and nonlinear dynamics (Rifford, 2014; Agrachev and Lee, 2009). For discrete-time nonlinear systems, a relaxed version of the optimal transport problem was investigated in Elamvazhuthi *et al.* (2019b), where stochastic feedback laws, instead of deterministic feedback laws, were constructed to transport a system from a given initial measure to a target measure. The problem in Elamvazhuthi *et al.* (2019b) can be considered as the fixed-endpoint control version of the problem addressed in Chapters 3-5.

Chapter 2

CONTINUOUS-TIME MARKOV CHAIN MODELS ON DISCRETE STATE SPACES

In this chapter, we address the problem of redistributing a large number of homogeneous agents that evolve in continuous-time among a finite set of states. These states could represent tasks to be performed or spatial locations to occupy. In recent years, approaches to this problem have been developed in which the agents are programmed to switch stochastically between states at tunable transition rates. In some of these approaches (Berman *et al.*, 2009; Mather and Hsieh, 2014), the agents' states evolve according to a continuous-time Markov chain (CTMC), and their state distribution is controlled using the corresponding mean-field model, given by the Kolmogorov forward equation. These methods enable the scalable design of robot controllers due to the independence of the control methodology from the number of agents.

We consider such an approach for controlling a swarm of robots to allocate among a set of states in a decentralized fashion, that is, using only information that the robots can obtain from their local environment. We model the time evolution of each robot's state as a continuous-time Markov chain (CTMC) and frame the control problem in terms of the mean-field model of the system.

In some existing approaches to this problem that use Markov chain models (Berman *et al.*, 2009; Acikmese and Bayard, 2012), the robots continue switching between states at equilibrium, which could unnecessarily expend energy, even though the swarm distribution among states in the mean-field model is stabilized. This continued switching at equilibrium is due to the time- and density-independence of the

control laws considered in these works. To address this problem, density-feedback laws that use *quorum-sensing strategies* were constructed in Hsieh *et al.* (2008) to stabilize a swarm to a desired distribution and reduce the amount of switching between states at equilibrium. In Mather and Hsieh (2014), the authors pose this as a problem of controlling the variance of the agent distribution at equilibrium using density-feedback laws. As in Mather and Hsieh (2014), we propose density-feedback laws that stabilize a swarm of agents to a strictly positive target distribution. However, a key difference from Mather and Hsieh (2014) is that we consider feedback laws that are polynomial functions of the state that do not violate positivity constraints and converge to zero at equilibrium, preventing further switching between states. In Section 2.3, we prove the existence of a stabilizing polynomial control law by constructing a specific example of such a control law.

In Section 2.4, we present algorithmic procedures for computing polynomial feedback controllers that stabilize the mean-field model to a strictly positive target distribution with no state transitions occurring at equilibrium. For this, we use the Sum-of-Squares (SOS)-based MATLAB toolbox SOSTOOLS (Prajna *et al.*, 2002). These procedures could potentially incorporate additional constraints to improve the properties of the closed-loop system response. Moreover, in contrast to the control approaches presented in Bandyopadhyay *et al.* (2017); Demir *et al.* (2015), all computations for control synthesis are done offline in our procedures rather than onboard the robots in real-time, which reduces the robots' computational burden.

2.1 Notation

We denote by $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ a directed graph with M vertices, $\mathcal{V} = \{1, 2, \ldots, M\}$, and a set of $N_{\mathcal{E}}$ edges, $\mathcal{E} \subset \mathcal{V} \times \mathcal{V}$. We say that $e = (i, j) \in \mathcal{E}$ if there is an edge from vertex $i \in \mathcal{V}$ to vertex $j \in \mathcal{V}$. We define a source map $S : \mathcal{E} \to \mathcal{V}$ and a target map

$T : \mathcal{E} \to \mathcal{V}$ for which $S(e) = i$ and $T(e) = j$ whenever $e = (i, j) \in \mathcal{E}$. We assume that $(i, i) \notin \mathcal{E}$ for all $i \in \mathcal{V}$. There is a *directed path* of length s from node $i \in \mathcal{V}$ to node $j \in \mathcal{V}$ if there exists a sequence of edges $\{e_i\}_{i=1}^{s}$ in $\mathcal{E}$ such that $S(e_1) = i$, $T(e_s) = j$, and $S(e_k) = T(e_{k-1})$ for all $1 \leq k < s - 1$. A directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is called *strongly connected* if for every pair of distinct vertices v_0, $v_T \in \mathcal{V}$, there exists a directed path of edges in $\mathcal{E}$ connecting v_0 to v_T. The graph $\mathcal{G}$ is said to be *bidirected* if $e \in \mathcal{E}$ implies that $\tilde{e} = (T(e), S(e))$ also lies in $\mathcal{E}$.

2.2 Problem Statement

Consider a swarm of N autonomous agents whose states evolve in continuous time according to a Markov chain with finite state space $\mathcal{V}$. As an example application, $\mathcal{V}$ can represent a set of spatial locations that are obtained by partitioning the agents' environment. The graph $\mathcal{G}$ determines the pairs of vertices (states) between which the agents can transition. Denoting the set of admissible control inputs by $U \subset \mathbb{R}$, the agents' transition rules are determined by the control parameters $u_e : [0, \infty) \to U$ for each $e \in \mathcal{E}$, and are known as the *transition rates* of the associated CTMC. An agent in state v_1 at time t decides to switch to state v_2 at probability per unit time $u_e(t)$, $e = (v_1, v_2)$. Here, we have $U \subset \mathbb{R}_+$, i.e., the $u_e(t)$ obey positivity constraints, since transition rates must always be positive for a CTMC. The evolution of the N agents' states over time t on the state space $\mathcal{V}$ is described by N stochastic processes, $X_k(t) \in \mathcal{V}$, $k = 1, \ldots, N$. Each stochastic process $X_k(t)$ evolves according to the following conditional probabilities for every $e \in \mathcal{E}$:

$$\mathbb{P}(X_k(t + h) = T(e) | X_k(t) = S(e)) = \ u_e(t)h + o(h). \tag{2.1}$$

Here, $o(h)$ is the little-oh symbol and $\mathbb{P}$ is the underlying probability measure defined on the space of events Ω that is induced by the stochastic processes $\{X_k(t)\}_{k=1}^{N}$.

Let $\mathcal{P}(\mathcal{V})$ be the $(M-1)$-dimensional simplex of probability densities on $\mathcal{V}$, defined as $\mathcal{P}(\mathcal{V}) = \{\mathbf{y} \in \mathbb{R}_+^M : \ \sum_i y_i = 1\}$. Let $\mathbf{x}(t) = [x_1(t) \ \dots \ x_M(t)]^T \in \mathcal{P}(\mathcal{V})$ be the vector of probability distributions of the random variable $X_k(t)$ at time t, that is,

$$x_i(t) = \mathbb{P}(X_k(t) = i), \quad i \in \mathcal{V}. \tag{2.2}$$

The evolution of probability distributions is determined by the Kolmogorov forward equation, which can be cast as s a bilinear control system,

$$\dot{\mathbf{x}}(t) = \sum_{e \in \mathcal{E}} u_e(t) \mathbf{B}_e \mathbf{x}(t), \quad \mathbf{x}(0) = \mathbf{x}^0 \in \mathcal{P}(\mathcal{V}), \tag{2.3}$$

where $\mathbf{B}_e$, $e \in \mathcal{E}$, are control matrices with entries

$$B_e^{ij} = \begin{cases} -1 & \text{if } i = j = S(e), \\ 1 & \text{if } i = T(e),\ j = S(e), \\ 0 & \text{otherwise.} \end{cases} \tag{2.4}$$

Here, B_e^{ij} denotes the element in row i and column j of the matrix $\mathbf{B}_e$.

Note that (2.3) can be rewritten in the following form, which is common in the Markov chain literature:

$$\dot{\mathbf{x}}(t) = \mathbf{Q}\mathbf{x}(t), \quad \mathbf{x}(0) = \mathbf{x}^0 \in \mathcal{P}(\mathcal{V}),$$

where $\mathbf{Q} : \mathbb{R}^M \to \mathbb{R}^{M \times M}$ is the transition rate matrix, defined as $\mathbf{Q} = \sum_{e \in \mathcal{E}} u_e(t) \mathbf{B}_e$.

In Elamvazhuthi *et al.* (2017, 2019a), the Perron-Frobenius theorem was used to characterize the types of stationary distributions that can be stabilized by an irreducible, positively recurrent CTMC. We can now state the main problem addressed in this section.

Problem 2.2.1. *Given a strictly positive desired equilibrium distribution* $\mathbf{x}^d \in \mathcal{P}(\mathcal{V})$*, compute transition rates* $k_e : \mathcal{P}(\mathcal{V}) \to \mathbb{R}_+$*,* $e \in \mathcal{E}$*, such that the closed-loop system*

$$\begin{aligned} \dot{\mathbf{x}}(t) &= \sum_{e \in \mathcal{E}} k_e(\mathbf{x}(t)) \mathbf{B}_e \mathbf{x}(t), \quad t \in [0, \infty), \\ \mathbf{x}(0) &= \mathbf{x}^0 \in \mathcal{P}(\mathcal{V}) \end{aligned} \tag{2.5}$$

satisfies $\lim_{t \to \infty} \|\mathbf{x}(t) - \mathbf{x}^d\| = 0$ *for all* $\mathbf{x}^0 \in \mathcal{P}(\mathcal{V})$*, with the additional constraint that* $k_e(\mathbf{x}^d) = 0$ *for all* $e \in \mathcal{E}$*. Moreover, the density feedback should have a decentralized structure, in that each* k_e *must be a function only of densities* x_i *for which* $i = S(e)$ *or* $i = S(\tilde{e})$*, where* $T(\tilde{e}) = S(e)$*.*

We specify that each agent knows the desired equilibrium distribution $\mathbf{x}^d$. This assumption is used in other approaches to stabilizing solutions of the mean-field model of a swarm to desired probability distributions, e.g. Acikmese and Bayard (2012); Berman *et al.* (2009); Hsieh *et al.* (2008); Mather and Hsieh (2014).

We note that we were able to describe the state evolution of the agents by system (2.3) when the transition rates were density-independent because the agents' states were independent and identically distributed (i.i.d.) random variables in that case. However, when the density feedback control law $\{u_e(\mathbf{x})\}_{e \in \mathcal{E}}$ is used, the independence of the stochastic processes $X_i(t)$ is lost. This implies that the evolution of the probability distribution cannot be described by system (2.3). However, if we invoke the mean-field hypothesis and take the limit $N \to \infty$, then we can model the evolution of the probability distribution according to a nonlinear Markov chain. In this limit, the discrete number of agents $N_v(t, \omega)$ in state $v \in \mathcal{V}$ at time $t \in [0, T]$, where ω is used to emphasize that N_v is a measurable function of the sample path Ω, converges to the continuous agent population $x_v(t)$ in an appropriate sense, provided that solutions of system (2.5) are defined until a given final time $T > 0$. A rigorous process

for taking this limit in a stochastic process setting is described in Ethier and Kurtz (2009); Kolokoltsov (2010).

2.3 Solution to Problem 2.2.1

We consider a target equilibrium distribution $\mathbf{x}^d$ that is defined to be strictly positive everywhere on the domain; i.e., $\mathbf{x}^d \in \text{int}(\mathcal{P}(\mathcal{V}))$, where $\text{int}(\mathcal{P}(\mathcal{V}))$ denotes the *interior* of the simplex $\mathcal{P}(\mathcal{V})$. To prove the asymptotic stability of $\mathbf{x}^d$ for bidirected graphs, consider the continuously differentiable function $V : \mathbb{R}^M \to \mathbb{R}_{\geq 0}$ given by

$$V(\mathbf{y}) = \frac{1}{2}(\mathbf{y} - \mathbf{x}^d)^T \mathbf{D}(\mathbf{y} - \mathbf{x}^d) \tag{2.6}$$

for all $\mathbf{y} \in \mathbb{R}^M$, where $\mathbf{D} \in \mathbb{R}^{M \times M}$ is defined as $\mathbf{D} = [\text{diag}(\mathbf{x}^d)]^{-1}$. This Lyapunov function is commonly used in multi-agent consensus protocols Lewis *et al.* (2013).

Theorem 2.3.1. *Let $\mathcal{G}$ be a bidirected graph. Suppose that $\mathbf{x}^d \in \text{int}(\mathcal{P}(\mathcal{V}))$. Let $k_e : \mathbb{R}^M \to [0, \infty)$ be given by*

$$k_e(\mathbf{y}) = [(y_{S(e)} - x^d_{S(e)})^2 + (y_{T(e)} - x^d_{T(e)})^2]/x^d_{S(e)} \tag{2.7}$$

in system (2.5)*, for each $e \in \mathcal{E}$ and each $\mathbf{y} \in \mathbb{R}^M$. Then $\mathbf{x}^d$ is the globally asymptotically stable equilibrium point of system* (2.5)*.*

Proof. To facilitate our analysis, we rewrite system (2.5) as

$$\dot{\mathbf{x}}(t) = \mathbf{G}(\mathbf{x}(t))\mathbf{x}(t), \quad \mathbf{x}(0) = \mathbf{x}^0 \in \mathcal{P}(\mathcal{V}), \tag{2.8}$$

where $\mathbf{G} : \mathbb{R}^M \to \mathbb{R}^{M \times M}$ is given by $\mathbf{G}(\mathbf{y}) = \sum_{e \in \mathcal{E}} k_e(\mathbf{y})\mathbf{B}_e$ for all $\mathbf{y} \in \mathbb{R}^M$. It is clear that when $x_{S(e)} = x^d_{S(e)}$ and $x_{T(e)} = x^d_{T(e)}$ for all $e \in \mathcal{E}$, $\mathbf{G}(\mathbf{x}^d) = \mathbf{0}$, which satisfies our requirement that the control inputs equal zero at equilibrium.

To prove the stability of system (2.8), we will again invoke LaSalle's invariance principle (Khalil, 2001). Consider the continuously differentiable function $V : \mathbb{R}^M \to$

$\mathbb{R}_{\geq 0}$ defined in equation (2.6). To apply LaSalle's invariance principle, $\dot{V}(\mathbf{x}(t))$ along the solutions $\mathbf{x}(t)$ of system (2.8) is required to be negative. We can compute this derivative as:

$$\begin{aligned}\dot{V}(\mathbf{x}(t)) &= \frac{1}{2}\dot{\mathbf{x}}(t)^T\mathbf{D}(\mathbf{x}(t)-\mathbf{x}^d) + \frac{1}{2}(\mathbf{x}(t)-\mathbf{x}^d)^T\mathbf{D}\dot{\mathbf{x}}(t)\\ &= \frac{1}{2}\Big(\mathbf{x}(t)^T\mathbf{G}(\mathbf{x}(t))\mathbf{D}\mathbf{x}(t) + \mathbf{x}(t)^T\mathbf{D}\mathbf{G}(\mathbf{x}(t))^T\mathbf{x}(t) -\\ &\qquad \mathbf{x}(t)^T\mathbf{G}(\mathbf{x}(t))\mathbf{D}\mathbf{x}^d - (\mathbf{x}^d)^T\mathbf{D}\mathbf{G}(\mathbf{x}(t))^T\mathbf{x}(t)\Big).\end{aligned}$$

A simple computation shows that the last two terms in the expression above are zero. The sum of the first two terms is strictly negative; this can be confirmed by algebraic manipulation of the sum as follows. Setting $\mathbf{r}(t) = [x_1(t)/x_1^d \; ... \; x_M(t)/x_M^d]^T$, we obtain:

$$\begin{aligned}&\frac{1}{2}\mathbf{x}(t)^T\mathbf{G}(\mathbf{x}(t))\mathbf{D}\mathbf{x}(t) + \frac{1}{2}\mathbf{x}(t)^T\mathbf{D}\mathbf{G}(\mathbf{x}(t))^T\mathbf{x}(t)\\ &= \frac{1}{2}\sum_{e\in\mathcal{E}} -(r_{S(e)}(t) - r_{T(e)}(t))^2\big((x_{S(e)}(t) - x^d_{S(e)})^2 + (x_{T(e)}(t) - x^d_{T(e)})^2\big). \qquad (2.9)\end{aligned}$$

The expression (2.9) is a negative sum-of-squares, and thus equals zero only when $\mathbf{x}(t) = \mathbf{x}^d$. Hence, this function is strictly negative for all $\mathbf{x} \in \mathcal{P}(\mathcal{V})\backslash\{\mathbf{x}^d\}$. Moreover, the set $\mathcal{P}(\mathcal{V})$ is invariant for the closed-loop system (2.8) since $\mathbf{G}(\mathbf{y})$ is an essentially non-negative matrix for which each row sums to 0, for all $\mathbf{y} \in \mathcal{P}(\mathcal{V})$. It follows from LaSalle's invariance principle that $\mathbf{x}^d$ is the globally asymptotically stable equilibrium point of the closed-loop system (2.8) with the control inputs k_e defined in (2.7). □

In Elamvazhuthi *et al.* (2018a), we broadened the class of possible target distributions $\mathbf{x}^d$ to those that are not necessarily positive everywhere on the state space; that is, $\mathbf{x}^d$ could have a discontinuous support. This is rendered possible by using feedback control laws that are strictly local, in the sense that these feedback laws require each agent to know the density of agents only in the current state that it occupies, unlike

the control laws constructed in this section, which also require the agent to know the population densities at adjacent states. However, in contrast to the control laws in Elamvazhuthi *et al.* (2018a), the control laws presented in this section are polynomial in nature and can be computed algorithmically, as we will show next.

2.4 Nonlinear Polynomial Controller

In the previous section, we showed that the set of decentralized nonlinear control laws that solve Problem 2.2.1 is nonempty by explicitly constructing one such control law. Here, we present an algorithmic procedure for constructing nonlinear control laws that solve Problem 2.2.1, using the function (2.6) in the construction. We will construct control laws that are polynomial functions of the system state. This allows us to frame Problem 2.2.1 as a polynomial optimization problem that can be solved using SOSTOOLS (Prajna *et al.*, 2002), a MATLAB toolbox for solving sum-of-squares (SOS) programs. SOSTOOLS is widely used to provide algorithmic solutions to problems with polynomial non-negativity constraints that are otherwise difficult to solve. The non-negativity constraints are relaxed to a test for the existence of an SOS decomposition, and this test is performed using semidefinite programming. We note that our procedure is just one possible method for constructing the control laws.

We first pose Problem 2.2.1 as an optimization problem.

Problem 2.4.1. *Let* $\mathbb{R}[\mathbf{x}]$ *denote the set of polynomials (not necessarily positive), and let* Σ_s *denote the set of SOS polynomials. Given system* (2.5) *with the matrix* $\mathbf{B}_e$ *defined as in* (2.4)*, and given the function* $V(\mathbf{x})$ *in Equation* (2.6)*, find* $k_e(\mathbf{x}) \in \mathbb{R}[\mathbf{x}]$

such that

$$k_e(\mathbf{x}) \geq 0, \tag{2.10}$$

$$k(\mathbf{x}^d) = 0, \tag{2.11}$$

$$\nabla V(\mathbf{x})^T F(\mathbf{x}) \leq 0 \tag{2.12}$$

for all $\mathbf{x} \in \mathcal{P}(\mathcal{V})$*, where* $F(\mathbf{x}) = \sum_{e \in \mathcal{E}} k_e(\mathbf{x}) \mathbf{B}_e \mathbf{x}$.

In Problem 2.4.1, the candidate Lyapunov function $V(\mathbf{x})$ is fixed, and an appropriate control law is constructed such that $V(\mathbf{x})$ is indeed a Lyapunov function for the closed-loop system (2.5). We can easily confirm that $V(\mathbf{x}) = 0$ at the equilibrium $\mathbf{x}^d$, and simple algebraic manipulation shows that $V(\mathbf{x}) > 0$ for all $\mathbf{x} \in \mathcal{P}(\mathcal{V}) \backslash \{\mathbf{x}^d\}$. However, establishing the local negative definiteness of the gradient of $V(\mathbf{x})$ on the simplex $\mathcal{P}(\mathcal{V})$, as specified by the inequality (2.12), is a relatively difficult constraint to encode in SOSTOOLS. We enforce this constraint by using the theory of *positivestellansatz* (Chesi, 2011): *Given a function* $f(\mathbf{z})$*, where* $\mathbf{z} \in \mathbb{R}^n$*, and a set* $S \subset \mathbb{R}^n$*, is* $f(\mathbf{z}) \geq 0$ *(or, alternatively, is* $f(\mathbf{z}) < 0$*) for every* $\mathbf{z} \in S$*?* In this formulation, S is a *semialgebraic set*, which is defined as:

$$S := \{\mathbf{z} \in \mathbb{R}^n \mid p_i(\mathbf{z}) \geq 0,\ q_j(\mathbf{z}) = 0,\ i, j \in \mathbb{N}\}, \tag{2.13}$$

where p_i and q_i are polynomial functions of the state variable $\mathbf{z}$ and $\mathbb{N}$ is the set of natural numbers. S may also include constraints of the form $p_i(\mathbf{z}) < 0$ and $q_j(\mathbf{z}) \neq 0$.

Schmudgen's positivestellansatz (Schmüdgen, 1991), stated below, gives sufficient conditions for the positivity of $f(\mathbf{z})$ on a compact semialgebraic set $S \subset \mathbb{R}^n$.

Theorem 2.4.2. *Suppose that the semialgebraic set* (2.13) *is compact. If* $f(\mathbf{z}) \geq 0$ *for all* $\mathbf{z} \in S$*, then there exist* $t_j \in \mathbb{R}[\mathbf{z}]$ *and* $s_0, s_i, s_{ij}, s_{ijk}, ... \in \Sigma_s$ *such that*

$$f = \sum_j t_j q_j + s_0 + \sum_i s_i p_i + \sum_{i,j:i \neq j} s_{ij} p_i p_j + \sum_{i,j,k:i \neq j \neq k} s_{ijk} p_i p_j p_k + ... \tag{2.14}$$

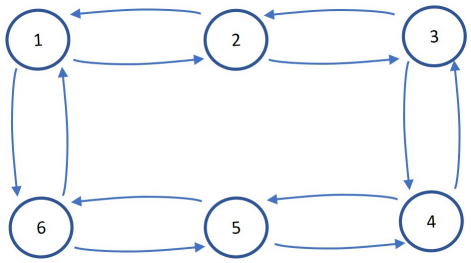

Figure 2.1: Six-vertex bidirected graph.

In our case, the simplex $\mathcal{P}(\mathcal{V})$ is a compact semialgebraic set in $\mathbb{R}^M$ of the form (2.13), in which the inequalities $p_i \geq 0$, $i = 1, ..., M$, are given by $x_1 \geq 0, \ \ldots \ , x_M \geq 0$, and the equality $q_1 = 0$ is given by $1 - x_1 - \ldots - x_M = 0$. Thus, according to Theorem 2.4.2, verifying the inequality (2.12) for all $\mathbf{x} \in \mathcal{P}(\mathcal{V})$ reduces to searching for $t_1 \in \mathbb{R}[\mathbf{x}]$ and $s_0, s_i, s_{ij}, s_{ijk}, ... \in \Sigma_s$ such that

$$-\nabla V(\mathbf{x})^T F(\mathbf{x}) = t_1 q_1 + s_0 + \sum_i s_i p_i + \sum_{i,j:i\neq j} s_{ij} p_i p_j + ... \tag{2.15}$$

We can use SOSTOOLS to find polynomials t_1 and $s_0, s_i, s_{ij}, ...$ that satisfy Equation (2.15).

We note that Problem 2.4.1 could alternatively be formulated to search for both the Lyapunov function and the control law simultaneously. Since this would render the optimization problem bilinear in the variables $V(\mathbf{x})$ and $k_e(\mathbf{x})$, it would be possible to solve the problem by iterating between these variables. However, this approach does not guarantee convergence to a solution.

2.5 Numerical Simulations

In this section, we numerically verify the effectiveness of the decentralized feedback controller (2.7), defined in Theorem 2.3.1, and a controller constructed using SOSTOOLS, as described in Section 2.4. The controllers were designed to redistribute populations of $N = 20$, $N = 80$, and $N = 1200$ agents on the six-vertex bidirected graph shown in Fig. 2.1. The initial distribution is $\mathbf{x}^0 = [0.2 \ \ 0.1 \ \ 0.2 \ \ 0.15 \ \ 0.2 \ \ 0.15]^T$,

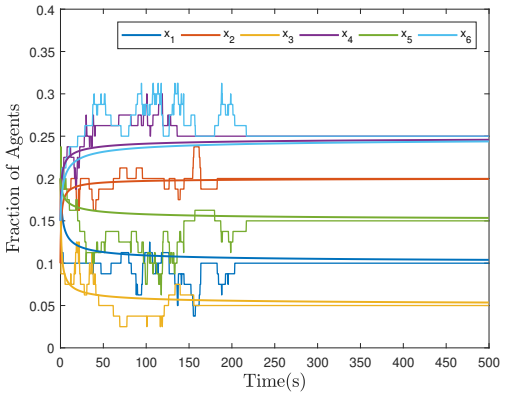

(a) Controller (2.7) with $N = 80$

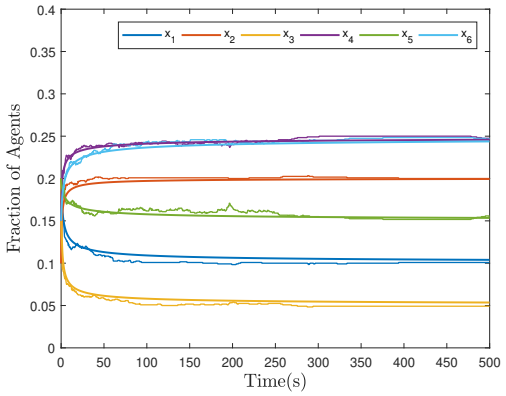

(b) Controller (2.7) with $N = 1200$

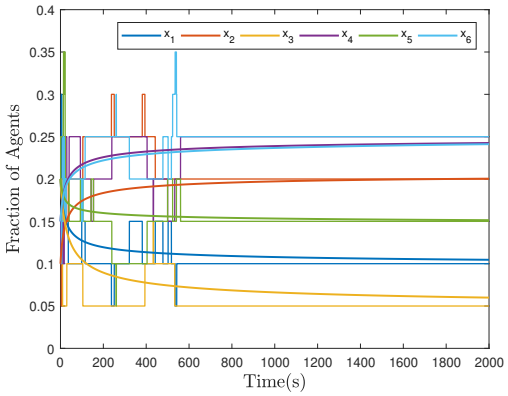

(c) Controller via SOSTOOLS with $N = 20$

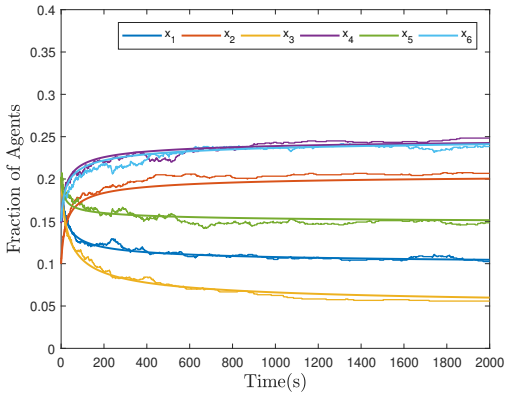

(d) Controller via SOSTOOLS with $N = 1200$

Figure 2.2: Trajectories of the mean-field model *(thick lines)* and the corresponding stochastic simulations *(thin lines)*.

and the desired distribution was $\mathbf{x}^d = [0.1 \;\; 0.2 \;\; 0.05 \;\; 0.25 \;\; 0.15 \;\; 0.25]^T$.

The solution of the mean-field model (2.5) with each of the two controllers and the trajectories of a corresponding stochastic simulation are compared in Fig. 2.2. The probability that an agent i in state (vertex) $S(e)$, $e \in \mathcal{E}$, at time t transitions to state $T(e)$ at time $t + \Delta t$ was set to:

$$\mathbb{P}(X_i(t + \Delta t) = T(e) | X_i(t) = S(e)) = k_e \left(\frac{1}{N} \mathbf{N}^p(t) \right) \Delta t.$$

Here, $\{k_e\}_{e \in \mathcal{E}}$ is the set of feedback laws and $\mathbf{N}^p(t) = [N_1^p(t) \; N_2^p(t) \; \ldots \; N_M^p(t)]^T$, where $N_v^p(t)$ is the number of agents in state $v \in \mathcal{V}$ at time t. We assume that each agent

can measure the agent populations in its current state and in adjacent states. For ease of comparison, the total agent populations were normalized to 1 in both the mean-field model and the stochastic simulation. We observe from the plots that both controllers drive the agent distributions to the desired equilibrium distribution, although it appears that the controller obtained from SOSTOOLS has a much slower convergence rate than the controller (2.7). This is because the inequalities in Problem 2.4.1 only guarantee the asymptotic stability of system (2.5). We note that if faster convergence is desired, then this could be encoded as a constraint in SOSTOOLS.

The underlying assumption of using the mean-field model (2.5) is that the swarm behaves like a continuum. That is, the ODE (2.5) is valid as the number of agents $N \to \infty$ (Kolokoltsov, 2010). Hence, it is imperative to check the performance of the feedback controller for different agent populations. We observe that the stochastic simulation follows the ODE solution closely in all the simulations, and that the stochastic simulation exhibits smaller fluctuations about the ODE solution when the agent population is increased (from $N = 80$ to $N = 1200$ for controller (2.7), and from $N = 20$ to $N = 1200$ for the SOSTOOLS controller). In addition, in all simulations, the numbers of agents in each state remain constant after some time; for the SOSTOOLS controller, the fluctuations stop earlier when $N = 20$ than when $N = 1200$.

The results presented within this chapter are part of (Elamvazhuthi *et al.*, 2017; Deshmukh *et al.*, 2018; Elamvazhuthi *et al.*, 2018a, 2019a).

Chapter 3

DISCRETE-TIME MARKOV CHAIN MODELS ON CONTINUOUS STATE SPACES - PART I

In this chapter, we prove that a particular class of discrete-time nonlinear control systems that evolve on a compact subset of $\mathbb{R}^d$ can be stabilized to target probability measures that are positive almost everywhere on this subset, are absolutely continuous with respect to the Lebesgue measure, with L^∞ density. Moreover, we also pose and numerically solve a relaxed optimization problem to obtain a stochastic feedback law that produces fast exponential convergence of the system to the desired probability measure. The problem of stabilizing a nonlinear control system to a target measure has many potential applications, including the control of large-scale distributed systems, in which these measures could model the distribution of an ensemble of agents such as a swarm of robots (e.g. Açıkmeşe and Bayard (2015); Elamvazhuthi and Berman (2018)) or the distribution of nodes in an electric power grid or a wireless network (Bagagiolo and Bauso, 2014).

Similar measure control problems have also been considered outside the context of mean-field games and optimal transport theory. In Mesquita and Hespanha (2012), piecewise-deterministic Markov processes evolving on $\mathbb{R}^d$ are controlled to make a continuously differentiable probability density invariant and stable. The problem of stabilizing measures that represent swarms of agents with single-integrator dynamics perturbed by Brownian motion was addressed in Elamvazhuthi and Berman (2018); Elamvazhuthi *et al.* (2018c). Other recent works (Vaidya *et al.*, 2010; Raghunathan and Vaidya, 2014; Das *et al.*, 2017) extend classical measure-theoretic studies of deterministic dynamical systems (Ding and Zhou, 2010; Lasota and Mackey, 2013) to

investigate the problem of stabilizing a control system to an attractor set from a measure-theoretic point of view. For the corresponding system evolving on the space of measures/densities, this means that the goal is to make the set of measures that are supported over the attractor set, or a Dirac measure at the desired point, invariant. In contrast, the objective in this section is to asymptotically stabilize a given measure that is subject to particular constraints. A similar measure stabilization problem is addressed in Chen *et al.* (2018), in which the authors consider an optimal control problem that drives a linear system evolving on $\mathbb{R}^d$ to target Gaussian measures. Our approach differs from this work in that we consider both the state space and the set of controls to be compact subsets of $\mathbb{R}^d$, and the set of target measures that can be stabilized is infinite-dimensional rather than finite-dimensional.

We will first identify the types of target measures that can be stabilized by the discrete-time nonlinear control systems that we consider, using stochastic feedback laws. In this case, the closed-loop system defines a discrete-time Markov chain on the continuous state space $\mathbb{R}^d$. For discrete-time Markov chains on discrete state spaces, the types of measures that can be stabilized are well-understood (Açıkmeşe and Bayard, 2015). This stabilizability result follows from the classical Perron-Frobenius theorem (Grobler, 1995), which gives a sufficient condition for the uniqueness of the stationary distribution of a Markov chain. We need an appropriate generalization of the Perron-Frobenius theorem for infinite-dimensional vector spaces. This generalization has been one of the motivating forces in developing the theory of *Banach lattices* and *positive operators* (Schaefer, 1974). At present, this theory has been developed to the point where the classical theorems of Perron-Frobenius are known to hold under very general hypotheses. A review of progress in this field is surveyed in Grobler (1995). We use the *Jentzsch-Perron theorem*, a generalization of the Perron-Frobenius theorem, to prove our results on the stabilizability of measures.

Having obtained these stabilizability results, we next address the problem of constructing feedback control laws that maximize the system's convergence rate to the target measure. For this, we exploit properties of *geometrically ergodic* discrete-time Markov chains, which converge exponentially fast to their target distributions. It is known that a Markov chain is geometrically ergodic if the forward operator that operates on the densities of the process has a spectral gap in L^2; the converse is only true for reversible Markov chains (Roberts and Tweedie, 2001). Therefore, the convergence rate of the Markov chain that describes our system can be characterized using this spectral gap. Thus, to compute feedback controllers that maximize the convergence rate of our system, we first prove the existence of a spectral gap in L^2 and then define an optimization problem that maximizes this spectral gap. Previous work has also addressed the maximization of the convergence rate of discrete-time Markov chains (Boyd *et al.*, 2004; Açıkmeşe and Bayard, 2015) and continuous-time Markov chains (Berman *et al.*, 2009; Deshmukh *et al.*, 2018) to stationary distributions; however, these results are restricted to finite, discrete state spaces.

3.1 Notation

Here, we present notation that will be used in this chapter and the next. We let $\det(\cdot)$ stand for determinant.

We denote the state space by $(\Omega, \mathcal{B}(\Omega))$, a measurable space. Here, $\Omega \subseteq \mathbb{R}^d$ is a compact set and $\mathcal{B}(\Omega)$ represents the Borel sigma algebra on Ω corresponding to the standard topology on $\mathbb{R}^d$. The set of admissible control inputs and its corresponding Borel sigma algebra will be denoted by $(U, \mathcal{B}(U))$. We will assume that U is compact in $\mathbb{R}^d$. The dimension of the set U could be larger than d, but we are restricting it for notational simplicity. We denote the spaces of probability measures on Ω and U by $\mathcal{P}(\Omega)$ and $\mathcal{P}(U)$, respectively. The Lebesgue measure on $\mathbb{R}^d$ will be denoted by

m. For a measure ν on $\mathbb{R}^n$, ν is said to be *absolutely continuous* with respect to m, denoted by $\nu \ll m$, if $\nu(E) = 0$ whenever $m(E) = 0$. In this case, there exists a function $f : \mathbb{R}^n \to \mathbb{R}$ such that $d\nu = f dm$; this function is called the *Radon-Nikodym derivative* of ν with respect to m (Folland, 2013).

$L^+(\mathcal{X}, \nu)$ is the space of all measurable functions from $\mathcal{X}$ to $[0, \infty]$. For topological spaces $\mathcal{X}, \mathcal{Y}$, if $T : \mathcal{X} \to \mathcal{Y}$ is an operator, it will be understood that $\|T\|$ stands for the operator norm, defined as $\sup_x \frac{\|Tx\|_{\mathcal{Y}}}{\|x\|_{\mathcal{X}}}$.

Let $(\mathcal{X}, \mathcal{N}, m)$ be a measure space, where $\mathcal{N}$ is the sigma algebra and m is a measure. A *transition kernel* or *Markov kernel* (or simply *kernel*) is a map $Q : \mathcal{X} \times \mathcal{N} \to [0, 1]$, where $Q(\cdot, E)$ is a measurable function on $\mathcal{X}$ for each fixed $E \in \mathcal{N}$ and $Q(x, \cdot)$ is a measure on $\mathcal{X}$ for each fixed $x \in \mathcal{X}$. Furthermore, for ν on $\mathcal{P}(\mathcal{X})$, the transition kernel Q induces an operator $T : \mathcal{P}(\mathcal{X}) \to \mathcal{P}(\mathcal{X})$ defined as:

$$T\nu(E) = \int_{\mathcal{X}} Q(x, E)\ d\nu(x), \quad E \in \mathcal{N}. \tag{3.1}$$

Similarly, Q can be used to define an operator on $L^2(\mathcal{X}, m)$. Suppose that ν is absolutely continuous with respect to m, denoted as $\nu \ll m$, and the Radon-Nikodym derivative of ν with respect to m, $d\nu/dm$, is given by $d\nu/dm = f_\nu \in L^2(\mathcal{X}, m)$. In simple terms, f_ν is called the *density* of ν. Then Q induces an operator $T^* : L^2(\mathcal{X}, m) \to L^2(\mathcal{X}, m)$, the adjoint of T, defined as

$$T^* f(x) = \int_{\mathcal{X}} f(y) Q(x, dm(y)), \quad x \in \mathcal{X},\ f \in L^2(\mathcal{X}, m). \tag{3.2}$$

We say that Q is *regular* if there exists a function $q \in L^\infty(\mathcal{X} \times \mathcal{X}, m \times m)$ such that for each $x \in \mathcal{X}$, the measure $Q(x, \cdot)$ is absolutely continuous with respect to m and $Q(x, dy) = q(x, y)dm$. The function $q : \mathcal{X} \times \mathcal{X} \to \bar{\mathbb{R}}_+$ will be called the *kernel function* of the transition kernel Q.

We define a continuous map $F : \Omega \times U \to \mathbb{R}^d$. We also define F_x as the map from $U \to \mathbb{R}^d$ when $x \in \Omega$ is held fixed, and F_u as the map from $\Omega \to \mathbb{R}^d$ when $u \in U$ is

held fixed. We specify that F is *non-singular*, which means that for all $E \in \mathcal{B}(\Omega)$, $m(F_u^{-1}(E)) = 0$ and $m(F_x^{-1}(E)) = 0$ whenever $m(E) = 0$.

The spectrum $\sigma(T)$ of a continuous linear operator T on the Banach space $\mathcal{X}$ is the non-void compact set of complex numbers λ for which $T - \lambda I$ does not have a continuous inverse on $\mathcal{X}$. If $\lambda \in \sigma(T)$ is such that $T - \lambda I$ is not injective, then λ is called an eigenvalue of T and the set $\sigma_p(T)$ of all eigenvalues is called the point spectrum of T. The spectral radius of T will be denoted by $r(T) := \sup\{|\lambda| : \lambda \in \sigma(T)\}$. We denote the complement of $\sigma(T)$ by $\rho(T)$ and call it the *resolvent set* of T.

Given a Banach space $\mathcal{X}$, if $\mathcal{X}^*$ is its dual space, then the duality pairing will be denoted by $\langle f, g\rangle_{(\mathcal{X},\mathcal{X}^*)}$, where $f \in \mathcal{X},\ g \in \mathcal{X}^*$.

A linear operator T on a real ordered vector space $\mathcal{X}$ is said to be *positive*, denoted by $T > 0$, if for $x \in \mathcal{X}$, $x \geq 0$ implies that $Tx \geq 0$.

3.2 Problem Formulation

Now we are ready to state the problems addressed in this section. Consider the nonlinear discrete-time control system,

$$\begin{aligned} x_{n+1} &= F(x_n, u_n),\ \ n = \mathbb{Z}_+, \\ x_0 &\in \Omega, \end{aligned} \tag{3.3}$$

where $x_n \in \Omega$ for each $n \in \mathbb{Z}_+$, and $(u_n)_{n=1}^{\infty}$ is a sequence in U such that $F(x_n, u_n) \in \Omega$. Suppose that x_0 is a random variable with distribution μ_0. Then $(x_n)_{n=1}^{\infty}$ is a *Markov chain* with corresponding sequence of distributions $(\mu_n)_{n=1}^{\infty}$. In particular, the nonlinear control system (3.3) induces a controlled flow on the space of measures $\mathcal{P}(\Omega)$, given by

$$\begin{aligned} \mu_{n+1} &= F(\cdot, u_n)_{\#}\mu_n,\ \ n = 0, 1, 2, ... \\ \mu_0 &\in \mathcal{P}(\Omega), \end{aligned} \tag{3.4}$$

where $F(\cdot, u_n)_{\#} : \mathcal{P}(\Omega) \to \mathcal{P}(\Omega)$ is the induced *forward operator* corresponding to the deterministic map $F(\cdot, u_n)$. This operator is defined as

$$F(\cdot, u_n)_{\#}\mu_n(E) = \mu_n(F_{u_n}^{-1}(E)) = \int_{\Omega} \chi_E(F(x, u_n))dx \tag{3.5}$$

for each $E \in \mathcal{B}(\Omega)$, where $\chi_{(\cdot)}$ denotes the characteristic function of a set.

We are interested in the problem of stabilizing system (3.4) to a given target measure. Toward this end, we must determine whether there exists a sequence of feedback laws such that starting from any initial measure, the system (3.4) converges to the target measure. However, in Elamvazhuthi *et al.* (2019b), a counterexample was provided to show that using deterministic feedback laws, the problem of reaching desired measures in finite time is generally unsolvable. A similar argument shows that this problem is unsolvable even without the finite-time convergence requirement. Hence, we will instead address the relaxed version of this problem, which is formulated as Problem 3.2.1 below.

Problem 3.2.1. ***(Stabilizability of target measures with stochastic control)*** *Given a target measure $\mu_d \in \mathcal{P}(\Omega)$ and a non-singular continuous map $F : \Omega \times U \to \mathbb{R}^d$, determine whether there exists a state-to-control transition kernel $K : \Omega \times \mathcal{B}(U) \to [0, 1]$ such that the closed-loop system*

$$\mu_{n+1} = P\mu_n, \;\; n = 0, 1, 2, \ldots; \quad \mu_0 \in \mathcal{P}(\Omega) \tag{3.6}$$

satisfies $\lim_{n\to\infty} P^n\mu_0 \to \mu_d$ for all initial measures $\mu_0 \in \mathcal{P}(\Omega)$, where the forward operator P that keeps $\mathcal{P}(\Omega)$ invariant is defined for each $E \in \mathcal{B}(\Omega)$ as

$$(P\mu)(E) = \int_{\Omega}\int_{U} \chi_E(F(x, u))K(x, du)d\mu(x). \tag{3.7}$$

This problem will be addressed in Section 3.4, wherein an explicit state-to-control transition kernel, also referred to here as a *stochastic feedback law*, will be constructed

for target measures that satisfy certain properties. Additional constraints will be imposed on F and Ω.

Given that there exists such a state-to-control transition kernel, we then address the problem of choosing the kernel that optimizes the convergence rate (mixing rate) of system (3.6) to the desired measure.

Problem 3.2.2. ***(Optimization of convergence rate)*** *Let $\mathcal{K}$ be the set of all Markov kernels defined on $\Omega \times \mathcal{B}(U) \to \overline{\mathbb{R}}_+$, and define $\|\mu\|_{TV} := \sup_{E \in \mathcal{B}(\Omega)} |\mu(E)|$ as the total variation norm. Given a target measure $\mu_d \in \mathcal{P}(\Omega)$, a non-singular continuous map $F : \Omega \times U \to \mathbb{R}^d$, and a constant $\alpha \in (0, 1)$, determine whether the following optimization problem admits a solution:*

$$\min_{\mathcal{K}} \alpha$$

such that $\|\mu_n - \mu_d\|_{TV} \leq \alpha^n$ for all $n \in \mathbb{Z}_+$, subject to the constraint

$$\mu_{n+1} = P\mu_n,$$

where P is the forward operator (3.7).

Markov chains that satisfy the bound α^n above are called *geometrically ergodic* chains. Different definitions of geometric ergodicity can be posed in terms of the particular norm (e.g., the L^1, L^2, or total variation norm) that is used to quantify the distance between the target and initial measures. The relationships among these definitions are discussed in Roberts *et al.* (1997). In addition, the spectral gap in L^2 is often easier to formulate than the total variation norm. Therefore, instead of framing the optimization problem in terms of the total variation norm, we shall pose it as the maximization of the $L^2(\Omega, m)$ spectral gap.

3.3 Analytical Properties of the Forward Operator

For the case where the state space is discrete, Problem 3.2.2 has been solved in Boyd *et al.* (2004) by maximizing the modulus of the second largest eigenvalue of the corresponding transition probability matrix. Since the largest eigenvalue of any stochastic matrix is 1, the problem reduces to maximizing the spectral gap of the transition probability matrix. When the state space is continuous, the forward operator replaces the matrix, and we can optimize the convergence rate by maximizing the spectral gap in $L^2(\Omega, m)$. Before formulating the optimization problem, we establish the existence of the $L^2(\Omega, m)$ spectral gap.

Let $\mu \in \mathcal{P}(\Omega)$. Suppose that μ is absolutely continuous with respect to the Lebesgue measure m (denoted as $\mu \ll m$). Hence, by the Radon-Nikodym theorem Folland (2013), there exists an m-integrable function $f_\mu : \Omega \to \mathbb{R}$ such that $d\mu = f_\mu dm$ and $f_\mu \in L^1(\Omega, m)$. Note that since μ is restricted to be a probability measure, f_μ is naturally non-negative on Ω. We will further restrict f_μ to be square-integrable with respect to m; that is, $f_\mu \in L^2(\Omega, m)$. This restriction gives us the advantage of being able to analyze the forward operator on a Hilbert space. Let $K : \Omega \times \mathcal{B}(U) \to [0, 1]$ be the transition kernel. We specify that K is regular; that is, if its kernel function is denoted as k, then $k \in L^\infty(\Omega \times U, m \times m)$. Furthermore, we impose the following constraints on k:

$$k(x,u) \begin{cases} \geq 0, \text{ for } m\text{-a.e. } x \in \Omega, u \in U \text{ s.t. } F(x,u) \in \Omega \\ = 0, \text{ otherwise} \end{cases} \tag{3.8}$$

$$\int_{u \in U} k(x,u)du = 1, \text{ for } m\text{-a.e. } x \in \Omega. \tag{3.9}$$

These properties ensure that K is indeed stochastic.

Instead of working with P, which acts on the space of probability measures $\mathcal{P}(\Omega)$, we will instead use P to define two linear operators, $\bar{P}$ and $\widetilde{P}$, that act on the spaces

$L^1(\Omega, m)$ and $L^2(\Omega, m)$, respectively. We note that $L^2(\Omega, m)$ has the advantage of being a Hilbert space. The operator $\bar{P}$ on $L^1(\Omega, m)$ is defined by restricting P to those measures in $\mathcal{P}(\Omega)$ that have integrable Radon-Nikodym derivative functions with respect to m, or equivalently, are absolutely continuous measures, denoted here using the subscript ac, i.e., $P : \mathcal{P}(\Omega)|_{ac} \to \mathcal{P}(\Omega)|_{ac}$. Let $\mathcal{L} \subseteq L^1(\Omega, m)$ be defined such that, if $\mu \in \mathcal{P}(\Omega)|_{ac}$ and $d\mu/dm = f$, then $f \in \mathcal{L}$. Note that since $\mathcal{P}(\Omega)$ is not a vector space, $\mathcal{L}$ is a strict subset of $L^1(\Omega, m)$. Define $\bar{P} : \mathcal{L} \to \mathcal{L}$ such that $d(P\mu)/dm = \bar{P}f$. By the linearity of $\bar{P}$, we extend it to the whole of $L^1(\Omega, m)$, so we can now define $\bar{P} : L^1(\Omega, m) \to L^1(\Omega, m)$. Similarly, by restricting $\mathcal{P}(\Omega)$ to measures that have square-integrable densities with respect to m, we define $\widetilde{P} : L^2(\Omega, m) \to L^2(\Omega, m)$. Shortly, we will establish that these operators are well-defined, in the sense that P preserves absolute continuity and square-integrability of the Radon-Nikodym derivatives, and moreover is bounded. The three operators P, $\bar{P}$, and $\widetilde{P}$ are all referred to as *forward operators*, since they describe the evolution of measures/densities forward in time. We will primarily be working with the operator $\widetilde{P}$, and the title of this section refers to this operator. Note that we cannot write an explicit formula for $\widetilde{P}f(\cdot)$ for $f \in L^2(\Omega)$ directly from (3.7). The backward operator is defined to be the Banach adjoint of the forward operator; hence, we define the backward operator $\widetilde{P}^*$ on $L^2(\Omega, m)$, but again, we cannot write an explicit expression for $\widetilde{P}^* f$.

We will now explore properties of the forward operator $\bar{P}$. First, we need to check whether $\bar{P}$ is bounded, linear, and well-defined. To establish these properties, we need the following definition.

Definition 3.3.1. *(Lasota and Mackey, 2013) Let $(\mathcal{X}, \mathcal{M}, \nu)$ be any measure space. Any linear operator $T : L^1(\mathcal{X}, \nu) \to L^1(\mathcal{X}, \nu)$ that satisfies the following two conditions is called a Markov operator:*

1. $Tf \geq 0$ *for* $f \geq 0,\ f \in L^1(\mathcal{X}, \nu)$;

2. $\|Tf\|_1 = \|f\|_1$ *for* $f \geq 0,\ f \in L^1(\mathcal{X}, \nu)$.

Lemma 3.3.2. *If F is non-singular and continuous, then the operator $\bar{P}$ is well-defined, Markov, and bounded.*

Proof. We begin by proving that P preserves those measures that are absolutely continuous with respect to the Lebesgue measure. Given that $\mu \in \mathcal{P}(\Omega)$ is such that $\mu \ll m$, we must show that $P\mu \ll m$. Indeed, if $E \in \mathcal{B}(\Omega)$ is such that $m(E) = 0$, then $\mu(E) = 0$, which further implies that $(\mu \times m)(F^{-1}(E)) = 0$. The last equality holds true due to the non-singularity of F with respect to both variables x, u. Therefore, we have that $\chi_E(F(x,u)) = \chi_{F^{-1}(E)}(x,u) = 0$ m-a.e. $x \in \Omega, u \in U$. From (3.7), we have

$$\begin{aligned}(P\mu)(E) &= \int_\Omega \int_U \chi_E(F(x,u))k(x,u)dm(u)d\mu(x) \\ &\leq \|k\|_\infty \int_\Omega \int_U \chi_E(F(x,u))dm(u)d\mu(x) = 0.\end{aligned}$$

Therefore $(P\mu)(E) = 0$, and we obtain $P\mu \ll m$. Since P preserves absolutely continuous probability measures, $\bar{P}$ preserves $L^1(\Omega, m)$.

Next, we prove that $\bar{P}$ is Markov. Condition *(1)* of Definition 3.3.1 follows from property (3.8), and accordingly the integrand in (3.7) is non-negative. Condition *(2)* follows from the fact that P preserves probability measures $\mathcal{P}(\Omega)$ that are absolutely continuous with respect to m. Thus, $\bar{P}$ is Markov. Also, from condition (ii), it follows that $\|\bar{P}\|_1 = 1$. □

We will be using the following result, which is straightforward to prove, several times in this section.

Lemma 3.3.3. *Suppose that ν is a measure on $\mathbb{R}^d$ such that $\nu \ll m$. Further, suppose that $\nu(E) \leq Cm(E)$ for any set $E \in \mathcal{B}(\mathbb{R}^d)$, where $C \in \mathbb{R}$ is a constant. Then the derivative of ν with respect to m, $d\nu/dm$, is in $L^\infty(\mathbb{R}^d, m)$.*

We can rewrite P in (3.7) in terms of what we will refer to as the closed-loop transition kernel $Q : \Omega \times \mathcal{B}(\Omega) \to \bar{\mathbb{R}}_+$. For $E \in \mathcal{B}(\Omega)$, (3.7) can be rewritten as

$$(P\mu)(E) = \int_\Omega Q(x, E)\ d\mu(x), \tag{3.10}$$

$$Q(x, E) = (P\delta_x)(E). \tag{3.11}$$

Here δ_x is the Dirac measure at x. Example 9.10 in Schilling (2017) justifies evaluation of the integral in (3.10) against a Dirac measure. It is straightforward to confirm that Q is a well-defined transition kernel, and that $Q(x, \Omega) = 1$. This will aid us in proving our next result in Proposition 3.3.4, which is at the heart of the analysis, in that it proves the compactness of $\widetilde{P}$. This in turn guarantees that the spectrum of $\widetilde{P}$ is discrete and therefore that a spectral gap exists (which is not true for operators with a continuous spectrum). This follows from Theorem VII.7.1 in Conway (2013), which states that for a compact operator T on an infinite-dimensional Hilbert space $\mathcal{H}$, the spectrum of T contains 0 and is discrete; furthermore, if the eigenvalues λ_i exist, they can be arranged in a decreasing order that tends to 0: $|\lambda_1| \geq |\lambda_2| \geq \ldots \geq |\lambda_n| \to 0$. We will also require F to satisfy *Lusin's property* (Bogachev, 2007) in both the x and u variables. For $x \in \Omega$ fixed, this condition is stated as follows: for $F_x : (U, m) \to (\mathbb{R}^d, m)$, we say that F_x satisfies Lusin's property if $m(F_x(E)) = 0$ for every $E \in U$ with $m(E) = 0$. Lusin's property for F_u has a similar definition.

Proposition 3.3.4. *If K is regular and F is C^1 differentiable and satisfies Lusin's property, then $\widetilde{P} : L^2(\Omega, m) \to L^2(\Omega, m)$ is well-defined, bounded, and compact.*

Proof. The proof will be divided into the following key steps.

1. Prove that the closed-loop kernel Q defined in (3.11) is regular; that is, its kernel function is in $L^\infty(\Omega \times \Omega, m \times m)$.

2. Prove that the operator $\widetilde{P}$ is an *integral operator*, as defined in Conway (2013), on $L^2(\Omega \times \Omega, m \times m)$.

3. Apply Proposition II.4.7 of Conway (2013) to prove that $\widetilde{P}$ indeed satisfies all the properties stated in the proposition.

Fix $z \in \Omega$ and $E \in \mathcal{B}(\Omega)$. Setting $\mu = \delta_z$ in (3.7), we obtain:

$$\begin{aligned}
(P\delta_z)(E) &= \int_\Omega \int_U \chi_E(F(x,u))\ k(x,u)\ du\ d\delta_z \\
&= \int_U \chi_E(F(z,u))\ k(z,u)\ du \qquad (3.12) \\
&\leq \|k\|_\infty \int_U \chi_E(F(z,u))\ du \qquad (3.13) \\
&= \|k\|_\infty\ m(F_z^{-1}(E)). \qquad (3.14)
\end{aligned}$$

The equality (3.12) follows from *Fubini's theorem* (Folland, 2013). We note that by the non-singularity of F, $m(F_z^{-1}(E)) = 0$ if $\mu(E) = 0$. Therefore, $(P\delta_z)(E)$ is an absolutely continuous measure with respect to m. Recall the generalized version of the *change of variables* theorem (Theorem 5.8.30, Bogachev (2007)) Since the change of variables theorem can only be applied to an open set, we restrict F to $\text{int}(U)$ (i.e., $F|_{\text{int}(U)}$). The boundary ∂U can be excluded, since by Lusin's property, the fact that the measure of ∂U is 0 implies that $m(F_z(\partial U)) = 0$. Since F_x is C^1 differentiable, the derivative of F_x with respect to u, denoted by $D_u(F_x)$, is bounded on U uniformly over all $x \in \Omega$. Hence, the quantity $|\det D_u F_x|$ has both upper and lower bounds, both positive. Let $c_1 = \inf_{x,u} |\det D_u F_x|$. The integral in (3.13) can be bounded from above as follows:

$$\int_U \chi_E(F(z,u))du\ \leq\ c_1 \int_U \chi_E(F(z,u))|\det D_u F_z|du.$$

Since F_z satisfies Lusin's property, we can now apply the change of variables theorem to the right-hand side of the above inequality to obtain,

$$\int_U \chi_E(F(z,u))du \leq c_1 \int_{F_z(U)} \chi_E(y)dy = c_1 \int_{E\cap F_z(U)} dy = c_1\ m(E\cap F_z(U)) \leq c_1\ m(E).$$

Combining this result with (3.13), we obtain $Q(z,E) = (P\delta_z)(E) \leq \|k\|_\infty\ c_1\ m(E)$. The constant c_1 is independent of $z \in \Omega$ and $E \in \mathcal{B}(\Omega)$. Therefore, Q is regular. Let $\widetilde{P}f_\mu$ be the density function of $P\mu$ with respect to the Lebesgue measure. Therefore we have,

$$\begin{aligned}(P\mu)(E) &= \int_E (\widetilde{P}f_\mu)(x)dx = \int_\Omega Q(x,E)f_\mu(x)dx \\ &\leq \int_\Omega \|k\|_\infty c_1 m(E) f_\mu(x)dx = Cm(E),\end{aligned} \tag{3.15}$$

where C is a constant. The second equality follows from (3.10), and the inequality follows from computations above. Hence, we have achieved a uniform bound on $(P\mu)(E)$, which by Lemma 3.3.3 means that $\widetilde{P}$ in fact takes $L^2(\Omega,m)$ functions to $L^\infty(\Omega,m)$. Now we can apply Theorem 1.3 of Arendt and Bukhvalov (1994), which claims that if $\mathcal{X}$ is any σ-finite measure space, any bounded operator from $L^p(\mathcal{X})$ $(1 \leq p < \infty)$ into $L^\infty(\mathcal{X})$ is an integral operator. This proves that our $\widetilde{P}$ is indeed an integral operator.

By Lemma 3.3.3, the kernel function of Q is in $L^\infty(\Omega\times\Omega, m\times m) \subseteq L^2(\Omega\times\Omega, m\times m)$. Denote the kernel function of Q by $q : \Omega\times\Omega \to \overline{\mathbb{R}}_+$. For each x, $q_x : \Omega \to \overline{\mathbb{R}}_+$ is such that $Q(x,dy) = q_x dm$. We can now give an integral representation of $\widetilde{P}$. From (3.15), we have

$$\int_E (\widetilde{P}f_\mu)(x)dx = \int_\Omega \int_E q(x,y)f_\mu(x)dydx.$$

Using Fubini's theorem and comparing the integrands of the two integrals over E yields $(\widetilde{P}f_\mu)(y) = \int_\Omega q(x,y)f_\mu(x)dx$. Since $\widetilde{P}$ is an integral operator on $L^2(\Omega,m)$

with its *integral kernel*, as defined in Conway (2013), given by $q \in L^\infty(\Omega \times \Omega, m \times m)$, we can apply Proposition II.4.7 of Conway (2013) to obtain our result, namely, that $\widetilde{P}$ is well-defined, bounded, and compact. $\square$

From the statement of Proposition 3.3.4, it can be inferred that there exists a $q \in L^\infty(\Omega \times \Omega, m \times m)$ such that we can define the forward operator $\widetilde{P} : L^2(\Omega, m) \to L^2(\Omega, m)$ as:

$$(\widetilde{P} f_\mu)(y) = \int_\Omega q(x, y) f_\mu(x) dx. \tag{3.16}$$

Using the definition of the adjoint operator, the backward operator $\widetilde{P}^* : L^2(\Omega, m) \to L^2(\Omega, m)$ is,

$$(\widetilde{P}^* f)(x) = \int_\Omega q(x, y) f(y) dy. \tag{3.17}$$

In the case of finite-dimensional Markov chains, 1 is the largest eigenvalue of the transition probability matrix, and $\mathbf{1}$ is its corresponding (right) eigenvector. Similarly, obtaining the adjoint of P from (3.10), we evaluate $(P^*\mathbf{1})(x) = \int_\Omega Q(x, dy) = \mathbf{1}$. This is true for every $x \in \Omega$, and therefore $P^*\mathbf{1} = \mathbf{1}$. Thus, 1 is an eigenvalue of P (P^*). Corresponding to 1, we assume that P has a eigenvector or stationary measure at $\pi \in \mathcal{P}(\Omega)$; that is, $P\pi = \pi$. We will further assume that π has a density function $\frac{d\pi}{dm} = f_\pi$ which is strictly positive on Ω, and additionally, $f_\pi, f_\pi^{-1} \in L^\infty(\Omega, m)$. The reason for this choice will become clear shortly. It can be confirmed that f_π is an eigenvector of the operator $\widetilde{P}$ corresponding to eigenvalue 1. Therefore, from (3.16) and (3.17), we have the following properties of q:

$$\int_\Omega q(x, y) dy = 1 \tag{3.18}$$

$$\int_\Omega q(x, y) f_\pi(x) dx = f_\pi(y) \tag{3.19}$$

We now show that 1 is the largest eigenvalue of $\widetilde{P}$. Toward this end, one could prove that $\widetilde{P}$ is a contraction, i.e. $\|\widetilde{P}\| \leq 1$, using the fact that $|\lambda| \leq \|T\|$ for any

bounded linear operator T. However, $\widetilde{P}$ is not necessarily a contraction in the L^2 norm. We introduce a new bounded operator $\widehat{P}$ on a Hilbert space that is isomorphic to $L^2(\Omega, m)$, such that $\widehat{P}$ is a contraction on this new space. We will show that the spectrum of $\widetilde{P}$ is invariant under the transformation $\widetilde{P} \mapsto \widehat{P}$. Recall that π is a stationary measure of P that satisfies both $\pi \ll m$ and $m \ll \pi$. We define $\widehat{P} : L^2(\Omega, \pi) \to L^2(\Omega, \pi)$. Since m and π are mutually absolutely continuous, $L^2(\Omega, m) \cong L^2(\Omega, \pi)$ as Hilbert spaces. To express $\widehat{P}$ as an integral operator, we carry out the following computations. Let $\hat{f}_\mu = \frac{d\mu}{d\pi}$. Then:

$$d(P\mu) = (\widetilde{P} f_\mu) dm = \frac{(\widetilde{P} f_\mu)}{f_\pi} f_\pi \, dm = \left(\frac{\widetilde{P} f_\mu}{f_\pi} \right) d\pi = (\widehat{P} \hat{f}_\mu) d\pi,$$

where the last equality follows from the fact that $\frac{f_\mu}{f_\pi} = \frac{d\mu}{dm} \frac{dm}{d\pi} = \frac{d\mu}{d\pi}$. The operator $\widehat{P}$ is well-defined because π and m are mutually absolutely continuous and because we have assumed that $f_\pi, \frac{1}{f_\pi} \in L^\infty(\Omega, m)$. Now we define a *multiplication operator* $M_{f_\pi} : L^2(\Omega, m) \to L^2(\Omega, \pi)$, $M_{f_\pi} g = f_\pi g$. The operator M_{f_π} is well-defined and bounded according to Theorem II.1.5 of Folland (2013). $\widehat{P}$ can be expressed as,

$$\widehat{P} \hat{f}_\mu = \left(\frac{\widetilde{P} f_\mu}{f_\pi} \right) d\pi = M_{f_\pi}^{-1} (\widetilde{P} f_\mu) = M_{f_\pi}^{-1} \widetilde{P} \left(\frac{f_\mu f_\pi}{f_\pi} \right).$$

From this, we conclude that

$$\widehat{P} = M_{f_\pi}^{-1} \widetilde{P} M_{f_\pi}. \tag{3.20}$$

Finally, from (3.16), (3.20), and the definition of M_{f_π}, we are able to express $\widehat{P}$ as an integral operator:

$$(\widehat{P} \hat{f}_\mu)(y) = \int_\Omega \frac{1}{f_\pi(y)} q(x, y) \hat{f}_\mu(x) f_\pi(x) dx = \int_\Omega \frac{q(x, y)}{f_\pi(y)} \hat{f}_\mu(x) d\pi(x). \tag{3.21}$$

Note that the integral kernel for the above integral operator is $\frac{q(x,y)}{f_\pi(y)}$.

Proposition 3.3.5. *$\widehat{P}$ as defined in* (3.21) *is bounded with $\|\widehat{P}\|_{L^2(\pi)} = 1$, and as a result, $r(\widetilde{P}) \leq 1$.*

Proof. We will use Theorem 6.18 of Folland (2013) to prove the first part of the proposition. In order to check the conditions of this theorem, we need the kernel $\frac{q(x,y)}{f_\pi(y)}$ to be in $L^1(\Omega, \pi)$ with respect to each variable x and y when the other variable is fixed. First, we fix x and evaluate the integral $\int_\Omega \frac{q(x,y)}{f_\pi(y)} d\pi$. By property (3.18), we have

$$\int_\Omega q(x,y)dy = \int_\Omega \frac{q(x,y)}{f_\pi(y)} f_\pi(y)dy = \int_\Omega \frac{q(x,y)}{f_\pi(y)} d\pi(y) = 1. \tag{3.22}$$

Next, we evaluate the integral $\int_\Omega \frac{q(x,y)}{f_\pi(y)} d\pi(x)$. Using property (3.19), we have

$$\int_\Omega q(x,y) f_\pi(x)dx = f_\pi(y) \implies \int_\Omega \frac{q(x,y)}{f_\pi(y)} d\pi(x) = 1.$$

Therefore, the constant C in Theorem 6.18 of Folland (2013) is 1 in this case, and therefore $\|\widehat{P}\|_{L^2(\pi)} \leq 1$. This implies that $r(\widehat{P}) \leq \|\widehat{P}\|_{L^2(\pi)} \leq 1$. Recall that $\widehat{P} = M_{f_\pi}^{-1} \widetilde{P} M_{f_\pi}$. If $\lambda \in \sigma(\widetilde{P})$, then $(\widetilde{P} - \lambda I)$ is not invertible, and further, $M_{f_\pi}^{-1}(\widetilde{P} - \lambda I) M_{f_\pi}$ is not invertible, which implies that $\lambda \in \sigma(\widehat{P})$. From this, we also note that the converse holds true; that is, if $\lambda \in \sigma(\widehat{P})$ then $\lambda \in \sigma(\widetilde{P})$. As a consequence, we conclude that $r(\widehat{P}) \leq 1$. □

In fact, $\widehat{P}$ is *bistochastic*, which implies that $\mathbf{1}$ is both a right and left eigenvector of $\widehat{P}$. This follows from the equations below:

$$\widehat{P}\mathbf{1} = M_{f_\pi}^{-1} \widetilde{P} M_{f_\pi} \mathbf{1} = M_{f_\pi}^{-1} \widetilde{P} f_\pi = M_{f_\pi}^{-1} f_\pi = \mathbf{1}. \tag{3.23}$$

The adjoint equation $\widehat{P}^* \mathbf{1} = \mathbf{1}$ follows from (3.22) in the Appendix.

In conclusion, in this section we showed that the forward operator (3.16) defined on $L^2(\Omega, m)$ is compact and its largest eigenvalue is 1.

3.4 Existence of a Solution to Problem 3.2.1

In this section, we will construct the forward operator $P : \mathcal{P}(\Omega) \to \mathcal{P}(\Omega)$ and the analogous operator on densities, $\widetilde{P} : L^2(\Omega, m) \to L^2(\Omega, m)$, that solve Problem 3.2.1.

This will be achieved in several steps, which are enumerated below. The proofs of the results presented in this section are reserved for the Appendix.

We specify that the measure $\mu_d \in \mathcal{P}(\Omega)$ is such that its density function f_d is strictly positive a.e. on Ω and satisfies $f_d, \frac{1}{f_d} \in L^\infty(\Omega, m)$. Next, we state our assumptions. Suppose that we are given a map $F : \Omega \times U \to \mathbb{R}^d$ that satisfies the conditions stated in Problem 3.2.1. Further, as noted in Proposition 3.3.4, for compactness of the to-be-constructed operator $\widetilde{P}$ to hold, we require F to be C^1 differentiable and to satisfy Lusin's property. Moreover, as we will see, this process of construction will require us to impose additional restrictions on Ω. Specifically, Ω must be *path connected* and satisfy the *cone condition*, to be defined in this section. Lastly, for the system (3.3) to be controllable, we need the following local controllability condition.

Definition 3.4.1. *The system* (3.3) *is said to be locally controllable if there exists* $r > 0$ *such that, for every* $x \in \Omega$, $B_r(x) \cap \Omega \subseteq F(x, U)$.

From here on, we will consider r to be fixed as per this definition.

The steps for constructing P and $\widetilde{P}$ are as follows.

1. Construct a reference transition kernel, or stochastic feedback law, $K : \Omega \times \mathcal{B}(U) \to [0, 1]$. See (3.24). Prove that K is a well-defined Markov kernel; that is, it is a measurable function on Ω in the first variable and a measure on U in the second variable. See Proposition 3.4.2. Prove that K is regular; that is, it has an $L^\infty(\Omega \times U, m \times m)$ kernel function. See Proposition 3.4.3.

2. Using the constructed K, formulate an operator $\widetilde{S} : L^2(\Omega, m) \to L^2(\Omega, m)$. By Proposition 3.3.4, $\widetilde{S}$ should be well-defined, bounded, and compact.

3. Prove that $\widetilde{S}$ is *irreducible* and that $r(\widetilde{S}) = 1$, and moreover, that $\widetilde{S}^* \mathbf{1} = \mathbf{1}$. See Propositions 3.4.5 and 3.4.6. Corresponding to the eigenvalue 1, there must

be an eigenvector of $\widetilde{S}$, say f_π. Prove that f_π is in $L^\infty(\Omega, m)$ and is positive on Ω m-a.e. See Proposition 3.4.7. Prove the uniqueness of the eigenvalue 1, which in turn will guarantee that the eigenvector f_π is the unique equilibrium of system (3.6). See Theorem 3.4.9.

4. Using $\widetilde{S}$, construct a new operator $\widetilde{P} : L^2(\Omega, m) \to L^2(\Omega, m)$ such that the desired function f_d is its eigenvector. See (3.30).

5. Obtain an expression for the closed-loop transition kernel for the operator P, which we shall call $\widehat{Q} : \Omega \times \mathcal{B}(\Omega) \to [0, 1]$. See (3.32).

6. Prove that the discreteness of the spectrum of $\widetilde{S}$ is preserved under the transformation of $\widetilde{S}$ to $\widetilde{P}$. Further, prove that 1 is the spectral radius and a unique eigenvalue of $\widetilde{P}$. See Theorem 3.4.12.

7. Prove that $\widetilde{P}$ is *primitive*. This is to ensure that f_d is the unique asymptotically stable equilibrium of the system (3.6). See Theorem 3.4.13.

8. The final step is to confirm that there exists a state-to-control transition kernel of P, which we will call $\widehat{K} : \Omega \times \mathcal{B}(U) \to [0, 1]$, such that $\widehat{Q}$ is the closed-loop transition kernel of P. See Theorem 3.4.14.

Step 1: We now construct a suitable reference kernel K. Given $x \in \Omega$, define $U_x := F_x^{-1}(\Omega)$. Since F is continuous in both variables, it is clear that the set U_x is Borel measurable for each $x \in \Omega$. Let $W \in \mathcal{B}(U)$. Then K is defined as

$$K(x, W) = \frac{m(W \cap U_x)}{m(U_x)}. \tag{3.24}$$

We note that in general, $U_x \neq U$. We illustrate this with the following example. Let $\Omega = [-1, 1]$ and $U = [-0.5, 0.5]$. Suppose that $F(x, u) = x + u$. Then for $x = 1$ fixed, we do not have that for all $u \in U$, $F(1, u) \in \Omega$; in fact, any $u \in (0, 0.5]$ will result in

$F(1,u) \in (1, 1.5]$, which is outside our defined Ω. Therefore, the appropriate subset of U that ensures that $F(1,u) \in \Omega$ is $U_{x=1} = [-0.5, 0] \subsetneq U$.

To check that K is well-defined, we must confirm that $m(U_x)$ is non-zero. This requires the concept of *cone condition* (Definition 4.6, Adams and Fournier (2003)). A domain $\mathcal{D}$ is said to satisfy the cone condition if there exists a finite cone $\mathcal{C}$ such that each $x \in \Omega$ is the vertex of a finite cone $\mathcal{C}_x$ that is contained in Ω and congruent to $\mathcal{C}$. Note that $\mathcal{C}_x$ need not be obtained from $\mathcal{C}$ by parallel translation, but simply by rigid motion.

Proposition 3.4.2. *If Ω satisfies the cone condition and F is a C^1 function and satisfies Lusin's property, then $m(U_x)$ is non-zero for all $x \in \Omega$, and hence K in* (3.24) *is a well-defined Markov kernel.*

Proof. Let x be an arbitrary point in Ω. In order to show that $m(U_x)$ is non-zero, we will use the fact that $F_x^{-1}(B_r(x) \cap \Omega) \subseteq U_x = F_x^{-1}(\Omega)$ and show that $m(F_x^{-1}(B_r(x) \cap \Omega))$ cannot be arbitrarily small. For clarity in the expressions below, we denote $B_r(x) \cap \Omega$ by B_x. We note that by the non-singularity of F_x, $m(F_x^{-1}(B_x)) > 0$ if $m(B_x) > 0$.

There are two possible conditions under which $m(F_x^{-1}(B_x))$ is arbitrarily small. First, $m(B_x)$ could be arbitrarily small. To show that this is not true, we estimate the lower bound of $m(B_x)$ using the cone condition as follows. According to this condition, there is a cone $\mathcal{C}$ that is completely contained in Ω with x at its vertex. Accordingly, the intersection of this cone and B_x has a positive measure. Denoting this intersection by $\mathcal{C}_B$, we have that $m(\mathcal{C}_B) \leq m(B_x)$. Note that the lower bound $m(\mathcal{C}_B)$ is independent of x.

The second way in which $m(F_x^{-1}(B_x))$ could be arbitrarily small is if the measure of F_x^{-1} of a set of positive measure is arbitrarily small. We show that this is not

true by obtaining a lower bound on $m(F_x^{-1}(B_x))$, given that $m(B_x)$ is bounded from below. By definition,

$$m(F_x^{-1}(B_x)) = \int_U \chi_{B_x}(F(x,u))du$$

Since F is C^1, the determinant of its derivative with respect to each variable is bounded; let

$\sup_{x,u} |\det(D_u F_x)|)^{-1} = c_2 < \infty$. We bound the integral from above and apply the generalized change of variables formula (Bogachev, 2007), as was done in the proof of Proposition 3.3.4, to obtain the following lower bound on $m(F_x^{-1}(B_x))$:

$$\begin{aligned}\int_U \chi_{B_x}(F(x,u))du &\geq c_2 \int_U \chi_{B_x}(F(x,u))|\det(D_u F_x)|du \\ &= c_2 \int_{F_x(U)} \chi_{B_x}(y)dy = c_2\ m(B_x \cap F_x(U)) \\ &= c_2\ m(B_x) \geq m(\mathcal{C}_B) \end{aligned} \tag{3.25}$$

Therefore, $m(F_x^{-1}(B_x))$ is non-zero, and consequently $m(U_x)$ is non-zero, for all $x \in \Omega$.

Next, we confirm that K is a well-defined Markov kernel. Toward this end, we first fix $W \in \mathcal{B}(U)$ and check whether $K(\cdot, W)$ is a measurable function on Ω. Let $G = \{(x,u) \in \Omega \times W : F(x,u) \in \Omega\}$. G is Borel measurable because F is continuous in both variables. Since χ_G is a Borel measurable function, the *Tonelli theorem* Folland (2013) implies that $(\chi_G)_x$ is Borel measurable for each $x \in \Omega$, and therefore that $x \mapsto \int (\chi_G)_x du$ is Borel measurable. Since $(\chi_G)_x(u) = \chi_G(x,u)$, we have that $\int (\chi_G)_x du = m(F_x^{-1}(\Omega)) = m(U_x)$. That is, $x \mapsto m(U_x)$ is Borel measurable, which implies that $x \mapsto m(W \cap U_x)/m(U_x) = K(x,W)$ is Borel measurable. Next, we check that $K(x,\cdot)$ is a measure on $(U, \mathcal{B}(U))$ for each fixed $x \in \Omega$. This is a straightforward consequence of the fact that the Lebesgue measure restricted to U_x, $m|_{U_x}$, is a measure on U. Therefore, K is a well-defined Markov kernel. □

The following result ensures that K is regular; that is, it has a kernel function in $L^\infty(\Omega \times U, m \times m)$.

Proposition 3.4.3. *The transition kernel K defined in* (3.24) *is regular. We denote the kernel function by $k : \Omega \times U \to \mathbb{R}_+$; $k \in L^\infty(\Omega \times U, m \times m)$. For each $x \in \Omega$, $k_x : U \to \mathbb{R}_+$ is such that $K(x, du) = k_x dm$.*

We need a lemma before presenting the proof. To begin, let the transition kernel K induce an operator, say $A : \mathcal{P}(\Omega) \to \mathcal{P}(U)$, as follows. For each measure μ on Ω,

$$(A\mu)(W) = \int_\Omega K(x, W) d\mu(x), \ W \in \mathcal{B}(U) \tag{3.26}$$

defines a measure on $(U, \mathcal{B}(U))$. Similar to our definition of $\bar{P}$, we define $\bar{A} : L^1(\Omega, m) \to L^1(U, m)$.

Lemma 3.4.4. *The operators A and $\bar{A}$ are well-defined; that is, they preserve probability measures on U and $L^1(U, m)$, respectively. Moreover, A and $\bar{A}$ are bounded, and $\bar{A} : L^1(\Omega, m) \to L^\infty(U, m)$.*

Proof. Let $\mu \in \mathcal{P}(\Omega)$ such that $\mu \ll m$. We will first show that $A\mu \in \mathcal{P}(U)$ and $A\mu \ll m$. A straightforward computation shows that $A\mu$ defines a measure and $(A\mu)(U) = 1$, and therefore $A\mu \in \mathcal{P}(U)$. We now check absolute continuity of $A\mu$ with respect to m. Let $W \in \mathcal{B}(U)$ be such that $m(W) = 0$. Then we have that,

$$(A\mu)(W) = \int_\Omega K(x, W) d\mu(x) = \int_\Omega \frac{m(W \cap U_x)}{m(U_x)} d\mu = 0.$$

Hence, $A\mu \ll m$. This shows that A, and therefore $\bar{A}$, is well-defined.

To prove the boundedness of A, we carry out the following computation. Recall that we used the cone condition in the proof of Proposition 3.3.4 to establish that, for any $x \in \Omega$, there exists a cone $\mathcal{C}_x$, congruent to a cone $\mathcal{C}$, that is completely contained in Ω with x at its vertex. The intersection of $\mathcal{C}_x$ and $B_r(x) \cap \Omega$, denoted by $\mathcal{C}_B$, has

a positive measure. In Lemma 3.4.2, we showed that $m(U_x) > m(F_x^{-1}(B_r(x) \cap \Omega)) > m(\mathcal{C}_B)$; that is, $m(U_x)$ is lower-bounded by a constant $m(\mathcal{C}_B)$ that is independent of x. Further, since $W \cap U_x \subseteq W$, $m(W \cap U_x) \leq m(W)$. Combining these results, we obtain the following inequality:

$$(A\mu)(W) = \int_\Omega \frac{m(W \cap U_x)}{m(U_x)} d\mu(x) \leq \int_\Omega \frac{m(W)}{m(\mathcal{C}_B)} d\mu(x) \leq \frac{m(W)}{m(\mathcal{C}_B)}.$$

This shows that $A\mu$ is equivalent to the Lebesgue measure, and therefore A is bounded. Consequently, $\bar{A}$ is also bounded. Finally, since $A\mu$ has a uniform upper bound, by Lemma 3.3.3, the operator $\bar{A}$ takes $L^1(\Omega, m)$ to $L^\infty(U, m)$. □

Proof of Proposition 3.4.3. The proof follows from an application of Theorem 1.3 of Arendt and Bukhvalov (1994) in combination with the approach used in Proposition 3.3.4. □

We note that the kernel function k satisfies the properties (3.8)-(3.9).

Step 2: With the given map F and the constructed kernel K in (3.24), we define a forward operator $S : \mathcal{P}(\Omega) \to \mathcal{P}(\Omega)$ as per (3.7) as follows:

$$(S\mu)(E) = \int_\Omega \int_U \chi_E(F(x,u)) K(x, du) d\mu(x), \quad E \in \mathcal{B}(\Omega).$$

Let $Q : \Omega \times \mathcal{B}(\Omega) \to [0,1]$ be the closed-loop transition kernel of S, defined as

$$Q(x, E) = \int_U \chi_E(F(x,u)) K(x, du).$$

Denote the kernel function of Q by q. Note that q must satisfy the properties (3.18)-(3.19). By Proposition 3.3.4, $q \in L^\infty(\Omega \times \Omega, m \times m)$. By restricting S to those probability measures that have $L^2(\Omega, m)$ derivatives w.r.t m, we define $\widetilde{S} : L^2(\Omega, m) \to L^2(\Omega, m)$ as per (3.16):

$$(\widetilde{S} f)(y) = \int_\Omega q(x, y) f(x) dx, \quad f \in L^2(\Omega).$$

Since K is regular from *Step 1*, we can apply Proposition 3.3.4 to establish that $\widetilde{S}$ is well-defined, bounded, and compact and that it preserves $L^2(\Omega)$.

Step 3: By the Perron-Frobenius theorem, the transition matrix of a finite dimensional Markov chain must be irreducible to have a unique stationary distribution. Similarly, we establish this important property for $\widetilde{S}$. First, we present a few definitions from Eisner *et al.* (2015). A *Banach lattice* is a Banach space with an order defined on it. In our case, $L^2(\Omega, m)$ is a Banach lattice. A linear subspace I of a Banach lattice is a *lattice ideal* if the following condition holds: if $|g| \leq |h|$ pointwise and $h \in I$, then $g \in I$.

A positive operator T on a Banach lattice $\mathcal{X}$ is called *irreducible* if the only T-invariant closed lattice ideals of $\mathcal{X}$ are the trivial ones; that is, if $I \subseteq \mathcal{X}$ is a closed lattice ideal, then $T(I) \subseteq I$ implies that either $I = \{0\}$ or $I = \mathcal{X}$. A topological space $\mathcal{X}$ is *path connected* if any two points $x, y \in \mathcal{X}$ are connected by a *path* in $\mathcal{X}$, which is a continuous map $p : [0, 1] \to \mathcal{X}$ with $p(0) = x$, $p(1) = y$.

Proposition 3.4.5. *If Ω is path connected and system* (3.3) *is locally controllable, then $\widetilde{S}$ is irreducible.*

Before presenting the proof, we need the following characterization of ideals on a finite-dimensional measure space from Eisner *et al.* (2015). On a finite-dimensional measurable space $(\mathcal{X}, \mathcal{M})$, for $1 \leq p < \infty$, each closed lattice ideal $I \subseteq L^p(\mathcal{X})$ has the form I_E shown below for some $E \in \mathcal{M}$:

$$I_E := \big\{ g : E \subseteq \big\{ g = 0 \big\} \big\} \tag{3.27}$$

Proof of Proposition 3.4.5. For the sake of contradiction, let $\widetilde{S}$ be reducible. Then, let I be an $\widetilde{S}$-invariant, non-trivial, closed ideal of $\widetilde{S}$; that is, $\widetilde{S}(I) \subseteq I$. Furthermore, I must have the form (3.27) for some non-trivial $E \in \mathcal{B}(\Omega)$, with $m(E) > 0$. Let $g = \chi_{E^c}$. Then $g \in L^2(\Omega, m)$ and $g = 0$ on E. Therefore, $g \in I$. Now, let $d\mu = g dm$.

Then $\mu = m$ on E^c and $\mu = 0$ on E. By our claim, $\widetilde{S}g \in I$. The idea of the proof is to prove existence of a non μ-null set in E^c from which the measure gets push-forwarded to E, thereby obtaining a contradiction.

According to our claim, $\widetilde{S}g \in I$, and so we have that for all $y \in E$, $(\widetilde{S}g)(y) = \int_{E^c} q(x,y)g(x)dx = 0$. Since $g(x) = 1$ only if $x \in E^c$, this computation implies that $q(x,y) = 0$ for almost all $y \in E$. This further implies that

$$Q(x,E) = (S\delta_x)(E) = \int_U \chi_E(F(x,u))k(x,u)du = 0 \text{ for } \mu\text{-a.e. } x \in E^c \tag{3.28}$$

Next, we examine the two possible ways that the integral in (3.28) can be zero; namely, when $m(U) = 0$ or when the integrand is zero.

First, we show the existence of a subset $A \subseteq E^c$ of positive measure, such that for all $x \in A$, there exist $u \in U$ such that $F(x,u) \in E$. Since Ω is compact, there exists a finite number of points $x_1, \ldots, x_N \in \Omega$, such that Ω can be covered by a finite number of balls, each with positive radius δ and centered at $x_i, i \in \{1, \ldots, N\}$. Therefore, we have that $\Omega = E \cup E^c \subseteq \cup_{i=1}^N B_\delta(x_i)$. We choose δ small enough such that for every i, $B_\delta(x_i) \cap \Omega \subsetneq B_r(z) \cap \Omega$ for all $z \in B_\delta(x_i)$. Since Ω is path connected, these balls cannot be disjoint, and furthermore, there exists at least one ball which intersects both E and E^c in sets of positive measure. That is, there exists $j \in \{1, \ldots, N\}$ such that $m(B_\delta(x_j) \cap E) > 0$ and $m(B_\delta(x_j) \cap E^c) > 0$. By our choice of δ, for any point x within $B_\delta(x_j)$, $B_r(x) \cap \Omega$ strictly contains $B_\delta(x_j) \cap \Omega$. Thus, as illustrated in Figure 3.1, $B_r(\cdot)$ of all points in $B_\delta(x_j) \cap E^c$ must also contain $B_\delta(x_j) \cap E$, which has a strictly positive measure. Define $A := B_\delta(x_j) \cap E^c$, as shown in Figure 3.1.

Returning to the integral (3.28), if the measure of the domain of integration is zero, then the integral evaluates to zero. In this case, although $m(U) > 0$, it is not true that $F(x,u) \in E$ for all $u \in U$ and all $x \in E^c$. Fix $y \in A$. Then $B_r(y) \cap \Omega$ contains $B_\delta(x_j) \cap E$. Because F is non-singular, $m(F_y^{-1}(B_\delta(x_j) \cap E)) > 0$. Letting

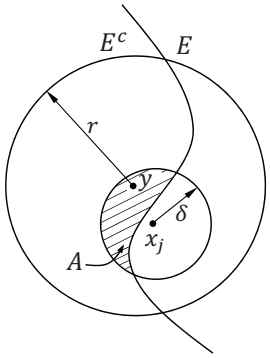

Figure 3.1: Illustration of the subset A (shaded region) used in the proof of Proposition 3.4.5.

$V := F_y^{-1}(B_\delta(x_j) \cap E)$ and restricting the domain of integration in (3.28) to V, observe that $Q(y, E) \geq \int_V \chi_E(F(y,u))k(y,u)du > 0$. Since $y \in A$ and $m(A) > 0$, we arrive at a contradiction with (3.28), and hence $\widetilde{S}$ is indeed irreducible. $\square$

It is easy to see that $\widetilde{S}$ is a positive operator. Next, we present several properties of the spectrum of $\widetilde{S}$ ($\widetilde{S}^*$). For the result below, we note that S is Markov.

Proposition 3.4.6. $\widetilde{S}^*\mathbf{1} = \mathbf{1}$ *and the spectral radius* $r(\widetilde{S})$ *is 1.*

Similar to the procedure in Section 3.3, the proof of this proposition follows from the fact that $\int_\Omega Q(x, dy) = 1$.

Let $f_\pi \in L^\infty(\Omega, m)$ be the eigenvector of $\widetilde{S}$ corresponding to the eigenvalue 1, and let $\pi \in \mathcal{P}(\Omega)$ be the measure such that f_π is its density. It is easy to see that π must be an eigenvector of S corresponding to the (uniform) measure defined by constant function $\mathbf{1}$. We establish properties of π, f_π below.

Proposition 3.4.7. *We have that* $f_\pi, f_\pi^{-1} \in L^\infty(\Omega, m)$ *and* f_π *is positive on* Ω m*-a.e.*

Proof. We have that,

$$(\widetilde{S}f_\pi)(y) = f_\pi(y) = \int_\Omega q(x,y)f_\pi(x)dx \leq \|q\|_\infty \int_\Omega f_\pi(x)dx \leq \|q\|_\infty$$

for m-a.e. $y \in \Omega$. The last inequality follows from the fact that since f_π is the density of a probability measure, its integral over Ω is 1. Therefore, f_π is bounded uniformly by $\|q\|_\infty$. Hence, $f_\pi \in L^\infty(\Omega, m)$.

The irreducibility of $\widetilde{S}$ (proven in Proposition 3.4.5) guarantees that f_π is positive almost everywhere on Ω. However, there could be cases where, for some $x \in \Omega$, $\lim_{\epsilon \to 0} \pi(B_\epsilon(x))/m(B_\epsilon(x)) = 0$, which would lead to $f_\pi^{-1} \notin L^\infty(\Omega, m)$. To show that this is indeed not the case, it is sufficient to prove that for $x \in \Omega$, there exists a measurable set $\mathcal{N}(x)$ of positive measure, containing x, and a constant $c > 0$ such that, for all $z \in \mathcal{N}(x)$,

$$(S\delta_z)(B_\epsilon(x)) \geq cm(B_\epsilon(x)). \tag{3.29}$$

First, we will assume that (3.29) is true. To see why this is a sufficient condition, we compute the following. Fix $x \in \Omega$. We evaluate

$$\begin{aligned}
\pi(B_\epsilon(x)) = (S\pi)(B_\epsilon(x)) &= \int_\Omega Q(z, B_\epsilon(x)) d\pi(z) \\
&= \int_\Omega (S\delta_z)(B_\epsilon(x)) d\pi(z) \geq \int_{\mathcal{N}(x)} cm(B_\epsilon(x)) f_\pi(z) dz \\
&= cm(B_\epsilon(x)) \int_{\mathcal{N}(x)} f_\pi(z) dz = ca(x) m(B_\epsilon(x)),
\end{aligned}$$

where $a(x) \in (0, 1]$ is the integral of f_π over $\mathcal{N}(x)$. Combining the constants $ca(x)$ into one constant c_x, we see that $\pi(B_\epsilon(x)) \geq c_x m(B_\epsilon(x))$. This implies that it will never be true that $\lim_{\epsilon \to 0} \pi(B_\epsilon(x))/m(B_\epsilon(x)) = 0$. Therefore, this shows that $f_\pi^{-1} \in L^\infty(\Omega, m)$.

Now we show that the condition (3.29) indeed holds true for every $x \in \Omega$. Let $x \in \Omega$ and $0 < \epsilon < r/2$. Choose $\mathcal{N}(x)$ to be $B_r(x) \cap \Omega$. Then $\mathcal{N}(x)$ is measurable and has positive measure. Further, for all $z \in \mathcal{N}(x)$, $B_\epsilon(x) \cap \Omega \subseteq B_r(z) \cap \Omega$. This follows from Definition 3.4.1.

From (3.24), we note that $k(\cdot, \cdot)$ is lower bounded by $1/m(U)$ on its support, and we denote this lower bound as c_3. Fix $z \in \mathcal{N}(x)$. For notational simplicity, denote

$B_\epsilon(x) \cap \Omega$ by B_ϵ. The computations below closely follow those preceding (3.25); hence, we have omitted a few steps here.

$$(S\delta_z)(B_\epsilon) = \int_U \chi_{B_\epsilon}(F(z,u))k(z,u)du \quad \text{(from (3.12))}$$
$$\geq c_3 \int_U \chi_{B_\epsilon(x)}(F(z,u))|\det D_u F_z|du = c_3 \int_{F_z(U)} \chi_{B_\epsilon(x)}(y)dy = c_3 m(B_\epsilon(x))$$

This shows that for every $x \in \Omega$, $(S\delta_z)(B_\epsilon(x)) \geq c_3 m(B_\epsilon(x))$. This proves that (3.29) holds true, and thus it is indeed true that $f_\pi^{-1} \in L^\infty(\Omega, m)$. □

We now establish the simplicity of the eigenvalue 1, which in turn will guarantee the uniqueness of the eigenvector π. Toward this end, we now state the generalized Perron-Frobenius theorem for infinite-dimensional compact operators. For a detailed proof, the reader is referred to Grobler (1995). First, we require the following definition. An element x of an L^p space with $p \in [0, \infty)$ (our case) is called a *quasi-interior point* if $x > 0$. A more general definition for quasi-interior points on general Banach lattices is given in Schaefer (1974).

Theorem 3.4.8. ***(Jentzsch-Perron)****(Grobler, 1995) Let T be a linear operator on a Banach lattice $\mathcal{X}$. Suppose that $T > 0$ and compact. If T is irreducible, then $r(T)$ is a positive eigenvalue of algebraic multiplicity one and its eigenspace is spanned by $x \in \mathcal{X}$, a unique normalized quasi-interior point.*

Finally, we summarize the above results in the following theorem.

Theorem 3.4.9. *The operator $\widetilde{S}$ is irreducible, and its spectral radius 1 is simple. Further, the eigenvector corresponding to 1, f_π, is positive on Ω m-a.e. and is in $L^\infty(\Omega, m)$.*

Step 4: Our goal is to construct an operator that has f_d as its fixed point. Toward this end, we define a multiplication operator $D : L^2(\Omega, m) \to L^2(\Omega, m)$ by $D(g) = \frac{g f_\pi}{f_d}$.

Note that $\frac{f_\pi}{f_d} \in L^\infty(\Omega, m)$, and therefore D is well-defined and bounded. Now we construct $\widetilde{P}$ as,

$$\widetilde{P} = (\widetilde{S} - I)\varepsilon D + I, \quad 0 < \varepsilon << 1 \tag{3.30}$$

Remark 3.4.10. *For ε small enough, $\widetilde{P}$ is a positive operator.*

Remark 3.4.11. *The transformation* (3.30) *is the discrete-time analogue of a transformation of the Laplacian Δ, which is the generator of a Brownian motion, into the generator ΔD of a new stochastic process for which the target measure μ_d is invariant. We previously used such a transformation to construct stochastic coverage strategies for robotic swarms in Elamvazhuthi* et al. *(2016); Elamvazhuthi and Berman (2018).*

Step 5: Similar to the pair $S, \widetilde{S}$, corresponding to $\widetilde{P}$ we can define an operator P that acts on $\mathcal{P}(\Omega)$. We note that $\widetilde{P}$ is not compact, since the identity operator I is not compact, so it cannot be represented as an integral operator with an L^2 integral kernel as in (3.16). Instead, we will show that P can be represented as (3.10) with a Markov kernel (that does not have an $L^\infty(\cdot)$ density function). To obtain the Markov kernel, we carry out the following computation.

Let $\mu \in \mathcal{P}(\Omega)$ be such that $\mu \ll m$, and let f_μ be its derivative with respect to m. Let $E \in \mathcal{B}(\Omega)$. We have that $(P\mu)(E) = \int_E (\widetilde{P} f_\mu)(x)dx$. Using (3.30), we evaluate the right-hand side of this equation:

$$\int_E \int_\Omega q(x, y)a(x)f_\mu(x)dydx + \int_E (1 - a(x))f_\mu(x)dx, \tag{3.31}$$

where $a(x) = \frac{\varepsilon f_\pi(x)}{f_d(x)}$.

We will also suppose that $(P\mu)(E) = \int_\Omega \widehat{Q}(x, E)d\mu(x)$ for some $\widehat{Q} : \Omega \times \mathcal{B}(\Omega) \to \mathbb{R}_+$. From (3.30), we will assume that $\widehat{Q}$ is of the following form,

$$\widehat{Q}(x, E) = \int_E q(x, y)a(x)dy + (1 - a(x))\delta_x(E). \tag{3.32}$$

This can be easily confirmed to be a Markov transition kernel. Now we evaluate $\int_\Omega \widehat{Q}(x,E)f_\mu(x)dx$, which equals:

$$\begin{aligned}&\int_\Omega\int_E q(x,y)a(x)f_\mu(x)dydx+\int_\Omega(1-a(x))f_\mu(x)\delta_x(E)dx\\&=\int_\Omega\int_E q(x,y)a(x)f_\mu(x)dydx+\int_E(1-a(x))f_\mu(x)dx.\end{aligned}$$

By applying Fubini's theorem to the first term, we observe that the above expression is exactly equal to (3.31). Hence, $\widehat{Q}$ is indeed the Markov kernel of P, and P can be represented in a form similar to (3.10) as shown below. For all $\mu\in\mathcal{P}(\Omega)$ and all $E\in\mathcal{B}(\Omega)$,

$$(P\mu)(E)=\int_\Omega \widehat{Q}(x,E)d\mu(x) \tag{3.33}$$

$$=\int_\Omega\int_U \chi_E(F(x,u))\widehat{K}(x,du)d\mu(x), \tag{3.34}$$

where $\widehat{K}$ is the state-to-control kernel of P. The existence of $\widehat{K}$ will be ensured by Theorem 3.4.14 below. As we will show in the next two theorems, Theorems 3.4.12 and 3.4.13, this constructed P is our solution to Problem 3.2.1 for μ_d that satisfy the constraints mentioned at the beginning of this section.

Step 6: We now use straightforward computations to demonstrate that the constructed operator $\widetilde{P}$ has 1 in its spectrum with f_d and $\mathbf{1}$ as the corresponding eigenvectors of $\widetilde{P}$ and $\widetilde{P}^*$, respectively. Since $\widetilde{S}^*\mathbf{1}=\mathbf{1}$, we have that $\widetilde{P}^*\mathbf{1}=(\varepsilon D^*(\widetilde{S}^*-I)+I)\mathbf{1}=\mathbf{1}$. In addition, since $\widetilde{S}f_\pi=f_\pi$, $\widetilde{P}f_d=((\widetilde{S}-I)\varepsilon D+I)f_d=f_d$. It is also easy to see that for P, we similarly have that $P^*\mathbf{1}=\mathbf{1}$, where $\mathbf{1}$ is the uniform measure, and $P\mu_d=\mu_d$.

Theorem 3.4.12. *The operator $\widetilde{P}$ defined in* (3.30) *has* 1 *as its largest eigenvalue, and this eigenvalue is algebraically simple and isolated (i.e., is not a limit point).*

We require the following definitions from Hislop and Sigal (2012) for this proof. Let T be a bounded, linear operator on a Hilbert space $\mathcal{H}$ with a nonempty resolvent

set $\rho(T)$. An operator A is called *relatively T-compact* if $AR_T(z) := A(T - zI)^{-1}$, where $R_T(z)$ is the *resolvent* of T, is compact for some $z \in \rho(T)$. The *essential spectrum* σ_{ess} of T is defined as the complement of $\sigma_p(T)$ in $\sigma(T)$. The operator T is said to be *closed* if its graph $\Gamma(T)$, defined as $\Gamma(T) := \{(x, Ax) : x \in \mathcal{H}\}$, is a closed subset of $\mathcal{H} \times \mathcal{H}$.

Proof of Theorem 3.4.12. First, we prove that the eigenvalue 1 is not an accumulation point. The resolvent of D, $R_D(z)$, is bounded for all $z \in \rho(D)$ by definition. Further, $\widetilde{S}$ is compact, and since the product of a compact operator and a bounded operator is always compact, $\widetilde{S}R_D(z)$ is compact; this implies that $\widetilde{S}$ is relatively D-compact. Moreover, by the well-known *closed graph theorem*, D is a closed operator. Now, we can apply *Weyl's theorem* (Theorem 18.8, Hislop and Sigal (2012)), which states that if T is a closed operator on a Hilbert space $\mathcal{H}$ and A is a relatively T-compact operator, then $\sigma_{ess}(T) = \sigma_{ess}(T + A)$. Accordingly, we have that $\sigma_{ess}(D) = \sigma_{ess}(\widetilde{S}D - D)$. Since D is a multiplication operator, its spectrum is the *essential range* of f_π/f_d. Recall our assumption that $f_d, f_d^{-1} \in L^\infty(\Omega, m)$, and Proposition 3.4.7 ensures that $f_\pi, f_\pi^{-1} \in L^\infty(\Omega, m)$. Let $\sigma_{ess}(D) \subseteq [a, b]$, for $a, b > 0$. Therefore we have,

$$\sigma_{ess}(\widetilde{S}D - \varepsilon D + I) \ \subseteq \ [1 - \varepsilon a, 1 - \varepsilon b].$$

Note that $1 \in \sigma_p(\widetilde{P})$. The computation above proves that there is a strict gap between $\sigma_{ess}(\widetilde{P})$ and 1. By Remark 1.5 (2) of Hislop and Sigal (2012), for a linear operator T on a Banach space, $\sigma_{ess}(T)$ and $\sigma_p(T)$ form a complete decomposition of the spectrum. Further, by definition (Hislop and Sigal, 2012), the eigenvalues, constituting the discrete spectrum, are isolated points. Therefore, 1 must be an isolated eigenvalue.

Second, we show that 1 is the spectral radius and an isolated eigenvalue of $\widetilde{P}$. As

per Proposition 4.1 of Schaefer (1974), for a positive operator T on a Banach lattice $\mathcal{X}$, the spectral radius $r(T)$ is an eigenvalue of T. Moreover, Theorem 2.1 of Marek (1970) guarantees that there exists at least one eigenvector x_0 in the *positive cone* (a subset $\mathcal{X}^+ = \{x \in \mathcal{X} : x \geq 0\}$) corresponding to $r(T) : Tx_0 = r(T)x_0,\ x_0 \neq 0$. In addition, there exists at least one eigenfunction x_0' in the positive dual cone corresponding to $r(T) : T^*x_0' = r(T)x_0',\ x_0' \neq 0$. In our case, since $\widetilde{P}$ is a positive operator on $L^2(\Omega, m)$, we therefore have that $r(\widetilde{P}) \in \sigma(\widetilde{P})$, and the eigenvector corresponding to $r(\widetilde{P})$ is positive. Let $\bar{r}$ be the spectral radius of $\widetilde{P}$, and define f_r as the corresponding positive eigenvector. Since the eigenvector f_r is known to be positive, by renormalizing, we can assume that the integral of f_r over Ω is 1. Let μ_r be the measure on Ω defined by f_r. Then it follows that $\mu_r(\Omega) = 1$. Note that 1 is also an eigenvalue for P. We then have that $\bar{r}\mu_r(\Omega) = \int_\Omega \widehat{Q}(x, \Omega)d\mu_r(x) = 1$, which implies that $\bar{r} = 1$. Thus, we conclude that 1 is the largest eigenvalue of $\widetilde{P}$ and that $f_r = f_d$.

Finally, we show that 1 is algebraically simple, which will enable us to conclude that f_d is indeed the unique eigenvector of $\widetilde{P}$ (up to a normalization) corresponding to 1. Theorem 5.2 of Schaefer (1974) states that if a positive, irreducible operator T on a Banach lattice $\mathcal{X}$ with $r(T) = 1$ has a non-void point spectrum and $x_0 = T^*x_0$ for some $x_0 \in \mathcal{X}$, then 1 is the unique eigenvalue of T and is algebraically simple. We note that $\widetilde{P}$ satisfies all these properties, and thus we have the result that 1 is a simple eigenvalue of $\widetilde{P}$, and therefore f_d is its unique positive fixed point. $\square$

Step 7: For (3.6) to be asymptotically stable, we need 1 to be the only eigenvalue of $\widetilde{P}$ that has modulus 1. *Primitivity* of $\widetilde{P}$ is precisely the condition that ensures this. A positive operator T is called primitive if $r(T)$ is the only eigenvalue on the spectral circle (the set $\{\lambda \in \mathbb{C} : |\lambda| = r(T)\}$). Note that primitivity of $\widetilde{P}$ implies *aperiodicity* of the associated Markov chain.

Theorem 3.4.13. *For all ε small enough, if f_π, f_d are bounded from below, then $\widetilde{P}$ in* (3.30) *is primitive.*

To prove that $\widetilde{P}$ is primitive, we require the following theorem from Grobler (1995). Let $\mathcal{X}$ be a Banach lattice and $T > 0$ be an operator on $\mathcal{X}$. Suppose there exists a positive linear functional $\phi \in \mathcal{X}^*$ such that $T^*\phi = \phi$. Then T is primitive if for each $x > 0$, there exists a $d \in \mathbb{N}$ such that $T^d x$ is a quasi-interior point in $\mathcal{X}$ (Schaefer, 1974). Here, the fact that $T^d x$ is a quasi-interior point implies that $T^d x > 0$. In the proof below, for $\nu \in \mathcal{P}(\Omega)$, $\nu > 0$ indicates that the Radon-Nikodym derivative of ν with respect to m, if one exists, is positive m-a.e. on Ω.

Proof of Theorem 3.4.13. We first check whether S is primitive. This is true if for any $x \in \Omega$, there exists an $n_0 \in \mathbb{N}$ such that for all $n \geq n_0$, $S^n \delta_x > 0$ m-a.e. on Ω. Here, we require a uniform n_0 that satisfies this condition for all $x \in \Omega$, so that we can extend the condition to arbitrary probability measures on Ω, which in turn could be constructed from Dirac measures.

We denote the open ball of radius δ centered at z by $V_\delta(z)$. Since Ω is compact, there exists a finite set $\{x_1, \ldots, x_N\} \subseteq \Omega$ such that $\Omega \subseteq \cup_{i=1}^N V_{r/2}(x_i)$. Fix $x \in \Omega$ and let $\mu^1 := S\delta_x$. Then $\mu^1 \ll m$ and $d\mu^1/dm = f_{\mu^1} \in L^\infty(\Omega, m)$ by the proof of Proposition 3.3.4. Furthermore, the support of f_{μ^1} contains $B_r(x)$. Now, we must have $x \in V_{r/2}(x_i)$ for some $i \in \{1, \ldots, N\}$. Without loss of generality, let $x \in V_{r/2}(x_1)$. Then $\mu^1 > 0$ a.e. on $V_{r/2}(x_1)$. Since Ω is path connected, the sets $\{V_{r/2}(x_i)\}_{i=1}^N$ cannot be pairwise disjoint; therefore, there exists another open ball, say $V_{r/2}(x_2)$, that intersects $V_{r/2}(x_1)$. Choose $y \in V_{r/2}(x_1) \cap V_{r/2}(x_2)$. Note that $y \in B_r(x)$. Now let $\mu^2 := S\mu^1 = S^2\delta_x$. Then $\mu^2 \ll m$ and $d\mu^2/dm = f_{\mu^2} \in L^\infty(\Omega, m)$. Furthermore, the support of f_{μ^2} is $E := \cup_{z \in B_r(x)} B_r(z)$. We have that $V_{r/2}(x_2) \subseteq B_r(y) \subseteq E$. Therefore, $\mu^2 > 0$ a.e. on $V_{r/2}(x_2)$. Repeating this procedure of evaluating $\mu^j := S\mu^{j-1}$ at each iteration j, we observe that μ^j is positive a.e. on $V_{r/2}(x_j)$. Since there are only N such

balls that cover Ω, this iterative procedure must stop at N, at which point we have that $\mu^N := S^N \delta_x$ is positive a.e. on Ω. Hence, we have proved that S is primitive, which implies the same for $\widetilde{S}$. From this discussion, we have demonstrated how S acts on Dirac measures. Extending this argument, we can show how S acts on any measure in $\mathcal{P}(\Omega)$ by noting that, for any $x \in \Omega$, $Q(x, \cdot) = (S\delta_x)(\cdot)$. In particular, we have that $(S^n \mu)(\cdot) = \int_\Omega S^n \delta_x(\cdot) d\mu(x)$.

Finally, we establish the primitivity of $\widetilde{P}$. Let $\mu \in \mathcal{P}(\Omega)$. From the definition of $\widetilde{P}$ in (3.30), we have that

$$\widetilde{P}^n = ((\widetilde{S} - I)\varepsilon D + I)^n = (\varepsilon \widetilde{S} D + (I - \varepsilon D))^n.$$

Consider the second expression for $\widetilde{P}^n$ above. Since $\widetilde{S}$ is primitive, the product $\widetilde{S}D$ preserves primitivity. In addition, by choosing a small enough ε, we can ensure strict positivity of the term $I - \varepsilon D$ (also see Remark 3.4.10). This in turn shows that $\widetilde{P}^n$ is a strictly positive operator for all $n \geq N$. Thus, the operator $\widetilde{P}$ is primitive. $\square$

Step 8: Finally, we prove that $\widehat{K}$, the state-to-control kernel of P in (3.34), is well-defined.

Theorem 3.4.14. *Let system* (3.6) *be locally controllable everywhere on* Ω. *Then there exists a Markov kernel* $\widehat{K} : \Omega \times \mathcal{B}(U) \to [0, 1]$ *such that* $\widehat{Q}$ *defined in* (3.32) *is the Markov kernel of the corresponding closed-loop system, and hence the equalities in* (3.33)-(3.34) *hold true.*

Proof. Consider the identity map $G : \Omega \to \Omega$ given by $G(x) = x$ for all $x \in \Omega$. We will also need the *set-valued map* $\hat{F} : \Omega \hookrightarrow U$ defined as $\hat{F}(x) = U$ for all $x \in \Omega$. The map $\hat{F}$ is a measurable set-valued map in the sense of Definition 8.1.1 in Aubin and Frankowska (2009). Since system (3.6) is locally controllable everywhere, we have that $F(x, \hat{F}(x)) \cap \{G(x)\}$ is non-empty for every $x \in \Omega$. Hence, from Theorem 8.2.8 in Aubin and Frankowska (2009), it follows that there exists a measurable function

$v : \Omega \to U$ such that $F(x, v(x)) = G(x) = x$ for every $x \in \Omega$. Then, we define $\widehat{K} : \Omega \times U \to \bar{\mathbb{R}}_+$ as follows. For all $W \in \mathcal{B}(U)$,

$$\widehat{K}(x, W) = a(x)K(x, W) + (1 - a(x))\delta_{v(x)}(W), \tag{3.35}$$

where $a(x) = \frac{\varepsilon f_\pi(x)}{f_d(x)}$. For a fixed $x \in \Omega$, it is easy to see that all terms in (3.35) are Borel measurable functions on Ω, except for the term $\delta_{v(x)}(W)$. The map $x \to \delta_{v(x)}(\cdot)$ can be written as a composition of two Borel measurable functions, $x \to v(x) \to \delta_{v(x)}(W)$, making it measurable in turn. Therefore, $x \to \widehat{K}(x, W)$ is a Borel measurable function on Ω. Furthermore, it is straightforward to show that $W \to \widehat{K}(x, W)$ is a measure on U for each $x \in \Omega$.

Now we evaluate $(P\mu)(E)$ for some $E \in \mathcal{B}(\Omega)$:

$$\begin{aligned}
(P\mu)(E) &= \int_\Omega \int_U \chi_E(F(x,u))\widehat{K}(x, du)d\mu(x) \\
&= \int_\Omega \Big(\int_U \chi_E(F(x,u))a(x)K(x,du) \; + \int_U \chi_E(F(x,u))(1-a(x))d\delta_{v(x)}(u)\Big)d\mu(x) \\
&= \int_\Omega \left(Q(x,E)a(x) + (1-a(x))\delta_x(E)\right) d\mu(x) \quad (\dagger) \\
&= \int_\Omega \left(\int_E q(x,y)a(x)dy + (1-a(x))\delta_x(E) \right) d\mu(x) \\
&= \int_\Omega \widehat{Q}(x,E)d\mu(x)
\end{aligned}$$

We obtained the first term in the integral (†) by using (3.12) as follows:

$$\int_U \chi_E(F(z,u))K(z,du) = (P\delta_z)(E) = Q(z,E).$$

We obtained the second term in this integral by noting that for fixed x, $v(x)$ is the set of u such that $F(x,u) = x$. Hence, we have our required result. □

3.5 Formulation of the Optimization Problem

In this section, we present a solution to a relaxed version of Problem 3.2.2. The reason for this relaxation will be explained shortly. In the previous section, we proved

the existence of an operator P that satisfies the following properties: it has a spectral gap, the desired measure μ_d is its unique eigenvector, and it makes μ_d an asymptotically stable equilibrium point for the system (3.6). In this section, we investigate whether we can pose an optimization problem to search for such an operator P such that the system (3.6) converges exponentially fast to the equilibrium μ_d. The spectral gap of P will determine the rate of convergence of system (3.6); the larger the gap, the faster the convergence. Recall the assumptions on μ_d as stated in Section 3.4: $\mu_d \ll m$, with f_d as its density, and f_d, f_d^{-1} are in $L^\infty(\Omega, m)$ and are strictly positive a.e. on Ω. Instead of P, we will formulate the optimization problem in terms of the operator $\widetilde{P}$ that acts on $L^2(\Omega, m)$. Specifically, we formulate an optimization problem that maximizes the spectral gap of $\widetilde{P}$. Similar to Boyd *et al.* (2004), we can then formulate a convex optimization problem that minimizes the second largest eigenvalue modulus of the operator. We begin with the formulation of the objective function in this problem.

We will pose the optimization problem for $\widehat{P} = M_{f_d}^{-1} \widetilde{P} M_{f_d}$, which has the same spectrum as $\widetilde{P}$. The advantage here is that $\widehat{P}$ is bistochastic, as proved in (3.23), which simplifies the formulation of the optimization problem as explained next. We know that given an operator T on a Hilbert space $\mathcal{H}$, for all $\lambda \in \sigma(T)$, we have that $|\lambda(T)| \leq \|T\|$. Unless the operator is self-adjoint or normal, there is no convex formula, that we know of, to characterize the moduli of the eigenvalues. Since we are not searching for a self-adjoint or normal operator $\widehat{P}$, the second largest eigenvalue modulus of $\widehat{P}$ can only be bounded from above. We obtain this bound by restricting $\widehat{P}$ to the subspace obtained after removing the eigenspace span($\mathbf{1}$) corresponding to its largest eigenvalue 1:

$$\lambda_2(\widehat{P}) = \lambda_1(\widehat{P} \circ \mathrm{Proj}_{\mathbf{1}^\perp}) \leq \|\widehat{P} \circ \mathrm{Proj}_{\mathbf{1}^\perp}\|_2, \tag{3.36}$$

where $\text{Proj}_{(.)}$ is the *projection operator* onto a subspace, and $\|\cdot\|_2$ denotes the $L^2(\Omega, \mu_d)$ norm. The optimization objective is then to minimize the right-hand side of the equation above, knowing that it will be an upper bound for the moduli of all eigenvalues of $\widehat{P}$. This is the relaxation that we mentioned at the beginning of the section.

The projection of an arbitrary vector $v \in L^2(\Omega, \mu_d)$ onto the eigenspace $\mathbf{1}$ is $\text{Proj}_{\mathbf{1}}(v) = \frac{\langle v, \mathbf{1} \rangle}{\|\mathbf{1}\|_2^2}\mathbf{1}$, and the projection of v onto $\mathbf{1}^\perp$ is $\text{Proj}_{\mathbf{1}^\perp} = I - \text{Proj}_{\mathbf{1}}$. Therefore, we have

$$\left(\widehat{P} \circ \text{Proj}_{\mathbf{1}^\perp}\right) v = \widehat{P}\left(v - \frac{\langle v, \mathbf{1} \rangle}{\|\mathbf{1}\|_2^2}\mathbf{1}\right) = \widehat{P}v - \frac{\langle v, \mathbf{1} \rangle}{\|\mathbf{1}\|_2^2}\mathbf{1}.$$

We now formulate the optimization problem. The optimization variable is the state-to-control transition kernel K. Using the variable K, the operator P from (3.7) is defined in constraint (3.39) below. The relationship between $\widehat{P}$ and $\widetilde{P}$ is enforced as constraint (3.38) in the optimization problem, defined as follows:

$$\min_{K} \quad \left\| \widehat{P}(K) \circ \text{Proj}_{1^\perp} \right\| \tag{3.37}$$

subject to

$$\widehat{P} = M_{f_d}^{-1} \widetilde{P} M_{f_d}, \tag{3.38}$$

$$P\mu(A) = \int_\Omega \int_U \chi_A(F(x,u)) K(x, du) d\mu, \;\; \forall A \in \mathcal{B}(\Omega), \;\; \mu \in \mathcal{P}(\Omega), \tag{3.39}$$

$$K(x, E) \geq 0 \quad \forall x \in \Omega, \;\; \forall E \in \mathcal{B}(U), \tag{3.40}$$

$$\int_\Omega K(x, U) dx = 1 \quad \forall x \in \Omega, \tag{3.41}$$

$$Q_K(x, A) = \int_U \chi_A(F(x,u)) K(x, du) \;\; \forall A \in \mathcal{B}(\Omega), \tag{3.42}$$

$$\int_\Omega f_d(y) Q_K(x, dy) = f_d(x) \quad \forall x \in \Omega. \tag{3.43}$$

The constraints (3.40)-(3.41) ensure that K is indeed a Markov kernel. Constraint (3.42) defines the closed-loop transition kernel Q_K in terms of K, and constraint (3.43) ensures that f_d is the stationary distribution of $\widetilde{P}$.

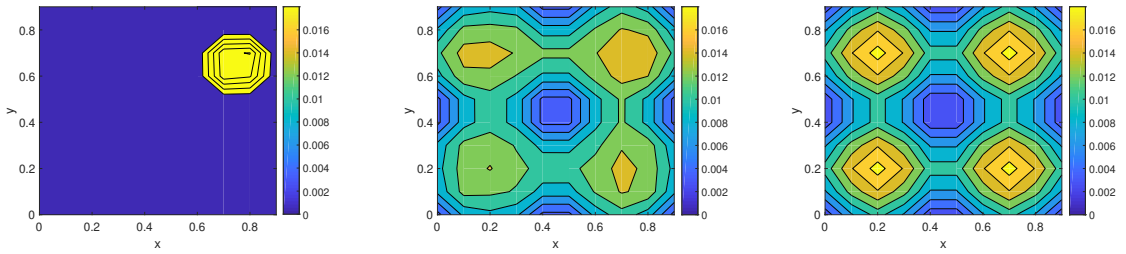

(a) $n = 2$ (b) $n = 7$ (c) $n = 15$

Figure 3.2: Simulation of the additive system (3.52) in Example 1 at three times n.

We end this section by showing that the optimization problem posed above is convex. Let $\mathcal{K}$ be the set of closed-loop transition kernels, defined as follows:

$$\mathcal{K} = \Big\{ K : \Omega \times \mathcal{B}(U) \to \bar{\mathbb{R}}_+ : \int_\Omega K(x, U)dx = 1, \int_\Omega f_d(y) Q_K(x, dy) = f_d(x) \ \forall x \in \Omega \Big\}.$$

Then $\mathcal{K}$ is the set of decision variables. We note that each constraint in this set is convex, therefore making $\mathcal{K}$ a convex set. Furthermore, the objective function is a norm of an operator, which makes it convex.

Finally, we note that it is not immediately clear whether an optimal solution to this problem exists. We reserve this investigation for future work.

3.6 Numerical Optimization

In this section, we present a numerical approach to solving the optimization problem (3.37)-(3.43). Our approach can be applied to control systems of the form (3.3), in which the state space Ω and the control set U are compact subsets of $\mathbb{R}^2$. The subset Ω is partitioned into $n_x \in \bar{\mathbb{Z}}_+$ sets, $\widetilde{\Omega} = \{\Omega_1, \ldots, \Omega_{n_x}\}$, where $\Omega = \cup_{i=1}^{n_x} \Omega_i$ and the sets Ω_i have intersections of zero Lebesgue measure. The set of control inputs U is approximated as a set of $n_u \in \bar{\mathbb{Z}}_+$ discrete elements, $\widetilde{U} = \{v_1, \ldots, v_{n_u}\}$, where $v_i \in U$ for each i. Define index sets $\mathcal{I} = \{1, \ldots, n_x\}$ and $\mathcal{J} = \{1, \ldots, n_u\}$. We define an equivalent of the state-to-control transition kernel K, with kernel function k, in

the discrete-time case. Let $\tilde{k}_{il}$ be the probability of choosing the control variable v_l, given that the system state is in Ω_i. This probability is given by,

$$\tilde{k}_{il} = \int_{\Omega_i} k(x, v_l) dx.$$

Let $\mathbf{K}$ be the matrix $[\tilde{k}_{il}]_{i\in\mathcal{I}, l\in\mathcal{J}}$. Using this definition, we construct an approximating controlled Markov chain on the finite state space $\mathcal{I}$. For $i \in \mathcal{I}$, when the system state is in the set Ω_i, we will consider the Markov chain state to be i. We use *Ulam's method* (Ding and Zhou, 2010) to construct this approximation. In the uncontrolled setting, Ulam's method is a classical technique to construct approximations of the *pushforward map* (Perron-Frobenius operators) induced by dynamical systems. Let p_{ij}^l denote the probability of the system state being in the set Ω_j in the next time step, given that the system state is uniformly randomly distributed over the set Ω_i and the selected control input is v_l. We define the transition probabilities of the controlled Markov chain as follows:

$$p_{ij}^l = \frac{m(\Omega_i \cap F_l^{-1}(\Omega_j))}{m(\Omega_i)}, \tag{3.44}$$

where $F_l(\cdot) = F(\cdot, v_l)$.

Let $\mu \in \mathcal{P}(\widetilde{\Omega})$ and $j \in \mathcal{I}$. Let $\mathbf{P}$ be the equivalent expression for the operator P, defined in (3.7), in matrix form. Then $\mathbf{P}$ is given by:

$$\mathbf{P}\mu(j) = \sum_{i\in\mathcal{I}} \sum_{l\in\mathcal{J}} \tilde{k}_{il}\ p_{ij}^l \mu(i). \tag{3.45}$$

Let $\mu_d \in \mathcal{P}(\widetilde{\Omega})$ be a desired distribution that is positive on $\widetilde{\Omega}$, and define a diagonal matrix $\mathbf{M}_d = \mathrm{diag}(\mu_d)$.

We can now formulate the finite-dimensional quadratic program that is equivalent to optimization problem (3.37)-(3.43). We define a bistochastic matrix $\widehat{\mathbf{P}}$ according to (3.20). This equation is enforced as constraint (3.47) in the quadratic program, defined as follows:

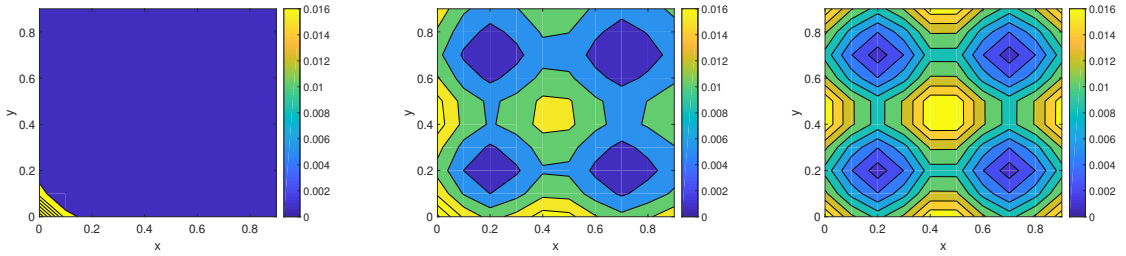

(a) $n = 2$ (b) $n = 7$ (c) $n = 15$

Figure 3.3: Simulation of the unicycle system (3.53) in Example 2 at three times n.

$$\min_{\mathbf{K}} \left\| \widehat{\mathbf{P}} - \frac{\mathbf{1}\mathbf{1}^T}{n_x} \right\| \tag{3.46}$$

subject to

$$\widehat{\mathbf{P}} = \mathbf{M}_d^{-1}\mathbf{P}\mathbf{M}_d, \tag{3.47}$$

$$\mathbf{P}\mu(j) = \sum_{i\in\mathcal{I}} \sum_{l\in\mathcal{J}} \tilde{k}_{il}\ p_{ij}^l \mu(i) \quad \forall j \in \mathcal{I},\ \forall \mu \in \mathcal{P}(\widetilde{\Omega}), \tag{3.48}$$

$$\tilde{k}_{il} \geq 0 \quad \forall i \in \mathcal{I},\ \forall l \in \mathcal{J}, \tag{3.49}$$

$$\mathbf{K}\mathbf{1} = \mathbf{1}, \tag{3.50}$$

$$\mathbf{P}\mu_d = \mu_d. \tag{3.51}$$

The constraint (3.48) above is written from (3.45), where $p_{ij}^l,\ i, j \in \mathcal{I},\ l \in \mathcal{J}$, is obtained via Ulam's method as per (3.44). Note that $\mathbf{1}$ in (3.46) is a vector in $\mathbb{R}^{n_x}$. We observe that this problem is convex and similar to the optimization problem solved in Boyd *et al.* (2004).

3.7 Simulation Results

In this section, we apply the numerical optimization procedure to two control systems of the form (3.3) evolving in $\mathbb{R}^2$. To solve the optimization problem (3.46)-(3.51), we used CVX, a MATLAB package for solving convex programs (Grant and

Boyd, 2014). Since the optimization problem is a quadratic program, it becomes computationally intractable for very fine discretizations of the domain Ω. Therefore, in Examples 1 and 2 below, we use a relatively coarse discretization. In Example 3, we solve a feasibility problem for which a finer discretization is possible. In all three cases, the quadratic program (3.46)-(3.51) was solved to obtain a state-to-control transition probability matrix $\mathbf{K}$. Defining $\mathbf{P}$ from the resulting $\mathbf{K}$ according to (3.48), we simulated system (3.6) with the initial measure μ_0 set to be a Dirac measure concentrated at the lower left corner of the domain. To quantify the degree of convergence of the simulated measure μ_n to the target measure μ_d, we computed the 2-norm $\gamma_n = ||\mu_n - \mu_d||_2$ at selected times n.

Example 1: Additive Model

We first consider a linear additive vector field F in system (3.3):

$$x_{n+1} = x_n + u_n, \tag{3.52}$$

where $x_n = [x_n^1 \; x_n^2]^T \in \Omega$ and $u_n = [u_n^1 \; u_n^2]^T \in U$. The state space is $\Omega = [0,1]^2$, and the set of control inputs is $U = [-1,1]^2$. The target measure is set to $\mu_d = \sin^2(2\pi x^1) + \sin^2(2\pi x^2) + \epsilon$, where $[x^1 \; x^2]^T \in \Omega$ and $\epsilon > 0$ is chosen to ensure a strictly positive measure over Ω. We use a 10×10 grid for Ω ($n_x = 100$) and a 20×20 grid for U ($n_u = 400$). Figures 3.2a-3.2c show snapshots of the simulation of system (3.6) at three times. Figure 3.5 plots the natural logarithm of the error metric γ_n during the simulation. It is evident from the time evolution of the snapshots, along with the accompanying decrease in γ_n, that the measure converges asymptotically to the target measure.

Example 2: Unicycle Model

We next consider a nonlinear vector field F in system (3.3) that represents a

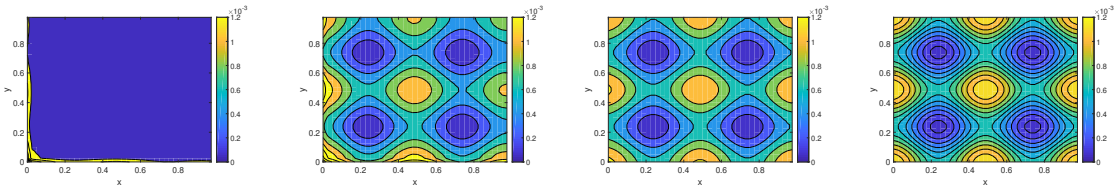

(a) $n = 5$ (b) $n = 10$ (c) $n = 15$ (d) $n = 20$

Figure 3.4: Simulation of the unicycle system (3.53) in Example 3 at four times n.

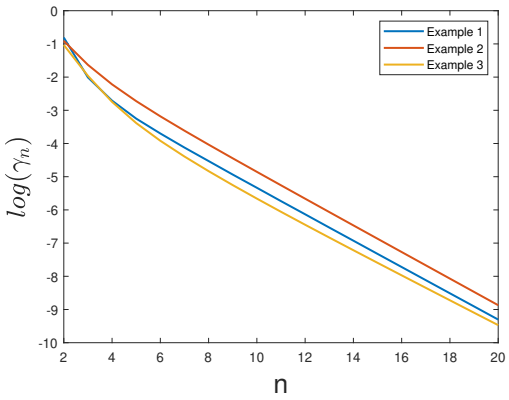

Figure 3.5: The time evolution of the natural logarithm of the error between the simulated and target measures in each example.

unicycle model:

$$x^1_{n+1} = x^1_n + u^1_n \cos(u^2_n),$$
$$x^2_{n+1} = x^2_n + u^1_n \sin(u^2_n). \tag{3.53}$$

In this case as well, $x_n = [x^1_n \; x^2_n]^T \in \Omega$ and $u_n = [u^1_n \; u^2_n]^T \in U$. The state space is $\Omega = [0, 1]^2$, and the set of control inputs is $U = [-1, 1] \times [0, 2\pi]$. The target measure in this case is set to $\mu_d = \cos^2(2\pi x^1) + \cos^2(2\pi x^2) + \epsilon$, where $[x^1 \; x^2]^T \in \Omega$ and $\epsilon > 0$ is chosen to ensure a strictly positive measure over Ω. We use a 10×10 grid for Ω ($n_x = 100$) and a 20×20 grid for U ($n_u = 400$). Figures 3.3a-3.3c show snapshots of the simulation of system (3.6) at three times, and Figure 3.5 plots the natural logarithm of the error metric γ_n over time. Again, the measure converges asymptotically to the target measure.

Example 3: Feasibility Problem

In the two previous cases, the optimization problem (3.46)-(3.51) was found to be computationally intractable for grid sizes $n_x > 100$, due to the fact that the problem is quadratic in nature. Here, instead of optimizing the convergence rate of system (3.6), we solve the feasibility problem. This entails searching for any matrix $\mathbf{K}$ such that the steady-state distribution of system (3.6) is μ_d; i.e., any $\mathbf{K}$ that satisfies the constraints (3.47)-(3.51). This serves to demonstrate that the feasibility problem can be solved for larger grid sizes than $n_x = 100$. The simulated system (3.3) is defined as the unicycle model (3.53) in Example 2. We use a 40×40 grid for Ω ($n_x = 1600$) and a 45×45 grid for U ($n_u = 2025$). Figures 3.4a-3.4d and 3.5 show snapshots of the simulation of system (3.6) at four times and the natural logarithm of the error metric γ_n over time. We see that the measure again converges asymptotically to the target measure.

Figure 3.5 shows that the system in this case exponentially converges to the target distribution at a rate close to the rates observed in Examples 1 and 2. Note that the convergence rate of Example 3 cannot be directly compared to the those of Examples 1 and 2, since the grid sizes n_x and n_u in Example 3 are much larger than in those two examples. We also make note of the fact that in Examples 1 and 2, we are only solving a relaxation of the optimization problem, as mentioned at the beginning of Section 3.5. The relaxation is due to the fact that we are only optimizing the norm of the operator, which is an upper bound on the moduli of all the eigenvalues of the operator.

The results presented within this chapter are part of Biswal *et al.* (2019b).

Chapter 4

DISCRETE-TIME MARKOV CHAIN MODELS ON CONTINUOUS STATE SPACES - PART II

In this chapter, we address the problem of stabilizing a multi-agent system evolving on a compact, connected subset of $\mathbb{R}^d$ to a target distribution. We consider groups of agents that all follow the same dynamics and control policies, which are independent of the agents' identities. We assume that each agent can obtain local measurements of the agent population but do not require inter-agent communication. The goals of this chapter are threefold:

1. We design the transition kernel to stabilize the mean-field model to target measures that have $L^\infty(\cdot)$ densities, a larger class of measures than we previously considered in Chapter 3. In the previous chapter, we considered measures that have $L^\infty(\cdot)$ densities that are strictly positive a.e. (almost everywhere) on the domain. In general, discrete-time Markov chains cannot be stabilized to distributions that do not have connected supports; we showed this for continuous-time Markov chains in Elamvazhuthi *et al.* (2019a), and similar arguments can be applied to discrete-time Markov chains (Açıkmeşe and Bayard, 2015). However, in this paper, we are able to stabilize the mean-field model to distributions that are not supported everywhere, due to the fact that, unlike in the previous chapter, the control law considered here is density-dependent; the reason for this will be explained next.

2. The convergence of a Markov process to an equilibrium distribution does not necessarily imply that the agents evolving according to the process also converge to equilibrium states. To prevent agents from continuing to switch between states at the equilibrium distribution, we construct the Markov process such that its forward

operator is the identity operator at the desired equilibrium. This results in a time-dependent transition kernel that is a function of the distribution and gives rise to a *nonlinear Markov process*. Such stochastic processes, called *density-dependent population processes*, are used to model the dynamics of logistic growth, epidemics, and chemical reaction networks (Ethier and Kurtz, 2009).

3. Since we establish that the transition kernel must depend on the distribution, our third goal is to construct the kernel to have a decentralized structure. A kernel with this structure corresponds to agent control policies that require each agent to estimate the population only in its local neighborhood. Toward this end, we construct a kernel for the mean-field model that is defined pointwise; that is, it is a function of the value of the distribution at the current state. We proved the existence of such feedback control laws in the case of continuous-time Markov chains evolving on finite graphs in Chapter 2. A similar problem is addressed in Mather and Hsieh (2014), which develops a decentralized control approach by *a priori* restricting the controller to have a decentralized structure. Another related work Demir *et al.* (2015) designs a centralized controller and uses estimation algorithms to determine the entire agent distribution in a decentralized manner.

4.1 Notation

The reader is referred to the notation section of Chapter 3. Some basic notations used in this chapter are mentioned here. As in Chapter 3, we denote the state space by $\Omega \subset \mathbb{R}^d$, a compact set. The set of admissible control inputs is denoted by $U \subset \mathbb{R}^d$. U is assumed to be compact. The Borel sigma algebras, corresponding to the standard topology on $\mathbb{R}^d$, on Ω and U are $\mathcal{B}(\Omega)$ and $\mathcal{B}(U)$ respectively. We denote the space of probability measures on Ω and U by $\mathcal{P}(\Omega)$. We define a continuous map $F : \Omega \times U \to \mathbb{R}^d$. We also define F_x as the map from $U \to \mathbb{R}^d$ when $x \in \Omega$ is

held fixed, and F_u as the map from $\Omega \to \mathbb{R}^d$ when $u \in U$ is held fixed. We specify that F is *non-singular*, which means that for all $E \in \mathcal{B}(\Omega)$, $m(F_u^{-1}(E)) = 0$ and $m(F_x^{-1}(E)) = 0$ whenever $m(E) = 0$. We also assume that $F(x, 0) = x$.

4.2 Problem Formulation

We now state the problem addressed in this section. Consider a system of N agents evolving in discrete time on the set $\Omega \subset \mathbb{R}^d$. We suppose that the dynamics of each agent $k \in \{1, \ldots, N\}$ is governed by the following nonlinear discrete-time control system:

$$\begin{aligned} \xi_{n+1}^k &= F(\xi_n^k, u_n^k), \;\; n = 0, 1, 2, ... \\ \xi_0^k &\in \Omega, \end{aligned} \tag{4.1}$$

where $\xi_n^k \in \Omega$, and $(u_n^k)_{n=1}^{\infty}$ is a sequence in U such that $F(\xi_n^k, u_n^k) \in \Omega$ for each $n \in \mathbb{Z}_+$. Let ξ_0^k be a random variable with distribution $\mu_0 \in \mathcal{P}(\Omega)$.

The *empirical distribution* of the N-agent system over Ω at time n is given by $\frac{1}{N}\sum_{k=1}^N \delta_{\xi_n^k}$. Our goal is to design a feedback control law u_n^k that redistributes the agents from their initial empirical distribution $\frac{1}{N}\sum_{k=1}^N \delta_{\xi_0^k}$ to a desired empirical distribution $\frac{1}{N}\sum_{i=1}^N \delta_{\xi^{k,d}}$ that "closely approximates" a target density $f^d \in L^\infty(\Omega)$ as $n \to \infty$, where $\frac{1}{N}\sum_{i=1}^N \delta_{\xi^{k,d}}$ is a sample of f^d. Since we assume that the agents are identity-free, we will define the control law as a function of the current empirical distribution $\frac{1}{N}\sum_{k=1}^N \delta_{\xi_n^k}$ rather than the individual agent states ξ_n^k. However, $\frac{1}{N}\sum_{k=1}^N \delta_{\xi_n^k}$ is not a state variable of the system (4.1). In order to treat $\frac{1}{N}\sum_{k=1}^N \delta_{\xi_n^k}$ as the state, we consider the *mean-field limit* of this quantity as $N \to \infty$.

Suppose that every agent $k \in \{1, \ldots, N\}$ uses the same control law $u_n^k = u_n$ at each time n; that is, the control law is independent of the agent identity k. In this case, when $N \to \infty$, the empirical distribution $\frac{1}{N}\sum_{k=1}^N \delta_{\xi_n^k}$ converges to a deterministic

quantity $\mu_n \in \mathcal{P}(\Omega)$, which evolves according to the following forward equation,

$$\mu_{n+1} = F^{\#}(\cdot, u_n)\mu_n, \quad \mu_0 \in \mathcal{P}(\Omega), \tag{4.2}$$

where $F^{\#}(\cdot, u_n) : \mathcal{P}(\Omega) \to \mathcal{P}(\Omega)$ is the induced *forward operator* corresponding to the deterministic map $F(\cdot, u_n)$. This operator is defined as

$$(F^{\#}(\cdot, u_n)\mu_n)(E) = \mu_n(F_{u_n}^{-1}(E)) = \int_{\Omega} \chi_E(F(x, u_n))dx$$

for each $E \in \mathcal{B}(\Omega)$. Since we are interested in the problem of stabilizing system (4.2) to a given target measure μ^d with density f^d, we must determine whether there exists a sequence of feedback laws u_n such that starting from any initial measure, the system (4.2) converges to μ^d. In general, this problem cannot be solved using deterministic feedback laws, as was shown in Elamvazhuthi *et al.* (2019b). Therefore, we will construct a *stochastic* feedback law using a state-to-control transition kernel $K : \Omega \times \mathcal{B}(U) \to [0, 1]$. On a continuous state space, the transition kernel plays the role of the transition probability matrix on a discrete state space. That is, given that an agent is at state $x \in \Omega$, it chooses a subset of control inputs $W \subset U$ with probability $K(x, W)$. We note that deterministic control laws $v : \Omega \to U$ are a special type of stochastic control law in that $K(x, du) = \delta_{v(x)}$; that is, given the state x, the probability of choosing the control $v(x)$ is 1. The transition kernel K induces a forward Kolmogorov operator $P : \mathcal{P}(\Omega) \to \mathcal{P}(\Omega)$, defined as

$$(P\mu)(E) = \int_{\Omega} \int_{U} \chi_E(F(x, u))K(x, du)d\mu(x)$$

for each $E \in \mathcal{B}(\Omega)$. The mean-field model that governs the time evolution of μ_n can then be written as

$$\mu_{n+1} = P\mu_n, \;\; \mu_0 \in \mathcal{P}(\Omega). \tag{4.3}$$

Hence, taking the mean-field limit of the empirical distribution enables us to treat the N-agent system as a continuum, as described in the Introduction. Moreover, to achieve our goal of redistributing agents over Ω, we will construct K to be a function of the current distribution μ_n.

Using K, we can define a closed-loop transition kernel $Q : \Omega \times \mathcal{B}(U) \to [0,1]$. That is, if the Markov chain $(\xi_n^k)_n$ induces a probability measure $\mathbb{P}$ on Ω^∞, then an agent k evolves on Ω according to the following conditional probability,

$$\mathbb{P}(\xi_{n+1}^k \in E | \xi_n^k = x) = Q(x, E), \tag{4.4}$$

for each $x \in \Omega$ and $E \in \mathcal{B}(\Omega)$. For $\mu \in \mathcal{P}(\Omega)$ and $E \in \mathcal{B}(\Omega)$, P can be redefined as

$$(P\mu)(E) = \int_\Omega Q(x, E) d\mu(x). \tag{4.5}$$

In this paper, instead of arbitrary measures in $\mathcal{P}(\Omega)$, we will consider those measures that have L^1 *densities* (derivatives with respect to m). By restricting P to this subset of $\mathcal{P}(\Omega)$, we can define an operator $\widetilde{P}$ on $L^1(\Omega)$; the exact construction will be carried out in the next section. Then (4.3) can be rewritten as

$$f_{n+1} = \widetilde{P} f_n, \quad f_0 \in L^1(\Omega). \tag{4.6}$$

We are now ready to state the problem that we address in this paper rigorously.

Problem 4.2.1. *Let $\widetilde{P}$ be the forward operator induced by the operator P defined in* (4.5)*. Given a target distribution $\mu^d \in \mathcal{P}(\Omega)$ with density $f^d \in L^\infty(\Omega)$ and a non-singular continuous map $F : \Omega \times U \to \mathbb{R}^d$, determine whether there exists a transition kernel $Q : \Omega \times \mathcal{B}(\Omega) \to [0,1]$ such that (a) equation* (4.6) *satisfies $\lim_{n\to\infty} \widetilde{P}^n f_0 = f^d$ for all initial measures $f_0 \in L^1(\Omega)$, and (b) $\widetilde{P}(f^d) = I$, where I is the identity operator.*

The operator $\widetilde{P}$ governs the stochastic transitions of individual agents between states. Thus, the condition $\widetilde{P}(f^d) = I$ ensures that all agents stop transitioning between states once the density f^d of the target equilibrium distribution is reached. This condition leads to a nonlinear operator $\widetilde{P}$ that depends on f. We will address Problem 1 in Section 4.3, where we show that the construction of $\widetilde{P}$ requires additional conditions on Ω and F.

Having proven the existence of such an operator $\widetilde{P}$, in Section 4.4 we will introduce the system of N agents that evolve according to the N-agent Markov process that is an approximation of the mean-field model (4.3). Since P (via Q) can be constructed such that μ^d is an equilibrium of the system (4.3), we observe that in simulations of the corresponding N-agent system, presented in Section 4.5, the empirical distribution $\frac{1}{N}\sum_{k=1}^{N}\delta_{\xi_n^k}$ converges to an empirical distribution that approximates f^d as $n \to \infty$.

4.3 Stability Result

In this section, an operator $\widetilde{P}$ that solves Problem 1 will be constructed. As stated in Problem 1, $f^d \in L^\infty(\Omega)$ is the density of the target measure. In the previous chapter, we assumed that f^d is supported m almost everywhere on Ω; in this chapter, we relax this assumption. The cost of this generality comes at the price of working with a nonlinear operator $\widetilde{P}$, which is necessary to ensure that agent transitions between states stop once the equilibrium distribution is reached.

Within this chapter we impose the same assumptions that we had in Chapter 6. That is, we assume that Ω is a path connected, compact subset of $\mathbb{R}^d$. We also require Ω to satisfy the cone condition which ensures that the boundary of Ω is regular enough. Lastly, for the system (4.1) to be controllable, we need the local controllability condition Definition 3.4.1. From here on, we will consider r to be fixed as per Definition 3.4.1.

Let $\mu \in \mathcal{P}(\Omega)$ be such that $\mu \ll m$. Further, if f_μ is the derivative of μ with respect to m, we assume that $f_\mu \in L^1(\Omega)$. For an arbitrary $f \in L^1(\Omega)$, define a function a_f on Ω as

$$a_f(x) = \begin{cases} \frac{f(x)-f^d(x)}{f(x)} \text{ for } m\text{-a.e. } x \text{ if } f(x) - f^d(x) > 0; \\ 0 \quad \text{otherwise.} \end{cases} \tag{4.7}$$

We note that $a_f \in L^\infty(\Omega)$ with norm 1.

Define $k : \Omega \times U \to [0, 1]$ to be a bounded function that satisfies the following properties:

$$k(x,u) \begin{cases} \geq 0 \text{ for } m\text{-a.e. } x \in \Omega, u \in U \text{ st. } F(x,u) \in \Omega; \\ = 0 \text{ otherwise;} \end{cases} \tag{4.8}$$

$$\int_U k(x,u)du = 1 \text{ for } m\text{-a.e. } x \in \Omega. \tag{4.9}$$

Before we proceed, we must determine whether we can construct a measurable $k \in L^\infty(\Omega \times U, m \times m)$ that satisfies these properties. We note that due to the first condition (4.8), the integral in the second condition (4.9) is computed over the set $U_x := F_x^{-1}(\Omega)$. This integral can therefore be expressed as $\int_U k(x,u)\chi_{U_x}(u)du = 1$. Since F_x is continuous, the set U_x is measurable for each x. The following lemma proves that k is measurable.

Lemma 4.3.1. *For $\forall x \in \Omega$, we have the following results:*

1. *There exists an $\epsilon > 0$ such that $m(U_x) > \epsilon$.*
2. *The map $x \mapsto m(U_x)$ is measurable.*
3. *The characteristic function $\chi_{U_x}(u)$ is jointly measurable in x and u.*

Proof. Result (1) is proved in Theorem 3.4.2. The idea of the proof is as follows. To prove results (1) and (3), let $G = \{(x,u) \in \Omega \times U : F(x,u) \in \Omega\}$. G is Borel

measurable because F is continuous in both variables. Since χ_G is a Borel measurable function, the *Tonelli theorem* Folland (2013) implies that $(\chi_G)_x$ is Borel measurable for each $x \in \Omega$. Since, $(\chi_G)_x(u) = \chi_{U_x}(u)$, we have that $\chi_{U_x}(u)$ is a measurable function in both variables, proving result (2). Then, by the Tonelli theorem, we have that $x \mapsto \int_U (\chi_G)_x du$ is Borel measurable. Since $(\chi_G)_x(u) = \chi_G(x,u)$, we have that $\int_U (\chi_G)_x du = m(F_x^{-1}(\Omega)) = m(U_x)$. That is, $x \mapsto m(U_x)$ is Borel measurable. □

The existence of a measurable function k then trivially follows from the fact that one can set k to be the uniform kernel, $k(x,u) = \frac{1}{m(U_x)}$.

Next, we define a transition kernel $K : \Omega \times \mathcal{B}(U) \to [0,1]$. For $W \in \mathcal{B}(U)$,

$$
\begin{aligned}
K(x,W) = K(x, W \cap U_x) &= K_1 + K_2, \quad \text{where} \\
K_1 &= a_{f_\mu}(x) \int_W k(x,u) du, \\
K_2 &= (1 - a_{f_\mu}(x)) \delta_0(W).
\end{aligned}
\tag{4.10}
$$

Recall that we have assumed that $F(x,0) = x$. Since this kernel is a function of a_{f_μ}, it depends on the density f_μ. The kernel is defined such that the corresponding Markov chain stays at control 0 with probability $1 - a_{f_\mu}(x)$ and moves to a control in the set U_x with probability $a_{f_\mu}(x)$, and when it moves, the distribution is given by the density $k(x, du)$. The integral term K_1 is regular because its kernel function $k(x,u)$ is in $L^\infty(\Omega \times U)$.

Remark 4.3.2. *We note that the local controllability assumption, Definition 3.4.1, implies that there exists a measurable control $V(x) \in U$ such that $F(x, V(x)) = x$ (see Proposition 3.4.14 in Chapter 3. Therefore, the condition $F(x,0) = x$ is not restrictive. However, we impose this condition here for the sake of simplicity and note that we can extend our results even when this condition is not satisfied following the steps in Chapter 3.*

Lemma 4.3.3. *The kernel K is well-defined. That is, $K(\cdot, W)$ is a measurable function on Ω for each fixed $W \in \mathcal{B}(U)$ and $K(x, \cdot)$ is a probability measure on U for each fixed $x \in \Omega$.*

Using K, we define a closed-loop kernel $Q : \Omega \times \mathcal{B}(\Omega) \to [0, 1]$. For $E \in \mathcal{B}(\Omega)$,

$$\begin{aligned} Q(x,E) &= \int_U \chi_E(F(x,u))K(x,du) \qquad (4.11)\\ &= a_{f_\mu}(x) \int_U \chi_E(F(x,u))k(x,u)du \;+\; (1 - a_{f_\mu}(x)) \int_U \chi_E(F(x,u))d\delta_0 \\ &= Q_1 + Q_2, \text{ where} \qquad (4.12)\\ Q_1 &= a_{f_\mu}(x) \int_U \chi_E(F(x,u))k(x,u)du, \\ Q_2 &= (1 - a_{f_\mu}(x))\delta_x(E). \end{aligned}$$

Lemma 4.3.4. *The kernel Q is well-defined; that is, $Q(\cdot, E)$ is a measurable function on Ω for each $E \in \mathcal{B}(\Omega)$ and $Q(x, \cdot)$ is a probability measure on Ω for each $x \in \Omega$. Further, if F satisfies Lusin's property, then Q_1 is regular.*

Proof. The proof that Q is well-defined is similar to the proof that K is well-defined (Lemma 4.3.3). To prove that Q_1 is regular, we first require that $Q_1(x, \cdot) \ll m$ for every x. Indeed, if $E \in \mathcal{B}(\Omega)$ is such that $m(E) = 0$, then due to the non-singularity of F with respect to both variables x and u, we have that $(m \times m)(F^{-1}(E)) = 0$. Therefore, for $x \in \Omega, u \in U$, we have that $\chi_E(F(x,u)) = \chi_{F^{-1}(E)}(x,u) = 0$ in the integral that defines Q_1 . Hence, $Q_1(x, E) = 0$.

The full proof that Q_1 has a kernel function $q \in L^\infty(\Omega \times \Omega)$ is given in 3.3.4. Here, we provide a brief idea of the proof. If q exists, then for $E \in \mathcal{B}(\Omega)$ and $x \in \Omega$,

$$\begin{aligned} Q_1(x,E) &= \int_E q(x,y)dy = \int_U \chi_E(F(x,u))K_1(x,du) \\ &= \int_U \chi_E(F(x,u))k(x,u)du \le c\; m(E), \end{aligned}$$

where $c > 0$ is a constant that is independent of x and E. The proof of Proposition 3.3.4 shows the existence of the uniformly bound c. That is, the measure $Q_1(x, \cdot)$ has a uniform upper bound, and hence, we must have that its kernel function q is in $L^\infty(\Omega \times \Omega)$ (Lemma 3.3.3). □

Next, we define an operator $P : \mathcal{P}(\Omega) \to \mathcal{P}(\Omega)$ in terms of Q as follows:

$$
\begin{aligned}
(P\mu)(E) &= \int_\Omega Q(x, E) d\mu(x) && (4.13)\\
&= \int_\Omega \int_U \chi_E(F(x,u)) K(x, du) d\mu(x) \\
&= \int_\Omega \int_U a_{f_\mu}(x) \chi_E(F(x,u)) k(x,u) du \; d\mu(x) \; + \; \int_\Omega (1 - a_{f_\mu}(x)) \delta_x(E) d\mu(x) \\
&= \int_\Omega \int_U a_{f_\mu}(x) \chi_E(F(x,u)) k(x,u) du \; d\mu(x) \; + \; \int_E (1 - a_{f_\mu}(x)) d\mu(x). && (4.14)
\end{aligned}
$$

Using expression for Q in (4.12), it is straightforward to show that $Q(x, E)$ can also be expressed as

$$Q(x, E) = (P\delta_x)(E). \tag{4.15}$$

Due to the properties of Q (Lemma 4.3.4), we immediately have the following lemma.

Lemma 4.3.5. *Operator P preserves $\mathcal{P}(\Omega)$, and furthermore, it preserves absolutely continuous measures.*

By restricting P to those measures that are absolutely continuous w.r.t m, that is, measures that have L^1 densities, we can define $\widetilde{P} : L^1(\Omega) \to L^1(\Omega)$. The next few steps will be toward this effort. Lemma 4.3.5 implies that $P\mu \ll m$; let $\widetilde{P} f_\mu$ be the density, that is, for $E \in \mathcal{B}(\Omega)$, $(P\mu)(E) = \int_E (\widetilde{P} f_\mu)(y)$. We note that since Q_1 is regular, there must exist a function $q \in L^\infty(\Omega \times \Omega)$. Therefore, from (4.12), we have that

$$Q_1(x, E) = \int_E q(x, y) dy = \int_U \chi_E(F(x,u)) K_1(x, du).$$

Using this expression, (4.14) can be rewritten as follows. For $E \in \mathcal{B}(\Omega)$,

$$(P\mu)(E) = \int_{\Omega}\int_{E} a_{f_\mu}(x)q(x,y)dy f_\mu(x)dx \; + \int_{E}(1 - a_{f_\mu}(x))f_\mu(x)dx = \int_{E}(\widetilde{P}f_\mu)(y)dy.$$

Applying Fubini's theorem Folland (2013) to the equation above, we obtain an expression for an operator $\widetilde{P}$ defined on $L^1(\Omega)$ as follows. For $f \in L^1(\Omega)$,

$$\begin{aligned} \widetilde{P} &= \widetilde{P}_1 + \widetilde{P}_2, \text{ where} \\ (\widetilde{P}_1 f)(y) &= \int_{\Omega} a_f(x)q(x,y)f(x)dx, \\ (\widetilde{P}_2 f)(y) &= (1 - a_f(y))f(y). \end{aligned} \tag{4.16}$$

Since Q_1 was proven to be regular in Lemma 4.3.4, we have that $\widetilde{P}$ preserves $L^1(\Omega)$, as we state in the proposition below.

Proposition 4.3.6. *We have the following two results.*

1. $\widetilde{P} : L^1(\Omega) \to L^1(\Omega)$ *is well-defined. Moreover,* $\widetilde{P}$ *preserves probability densities; in other words, it is a Markov operator Lasota and Mackey (2013).*

2. *In fact,* $\widetilde{P} : L^2(\Omega) \to L^2(\Omega)$ *is well-defined.*

The second result above is a consequence of Proposition II.4.7 of Conway (2013), detailed in 3.3.4. We will require this result in Section 4.4.

When we need to emphasize the fact that P and $\widetilde{P}$ are nonlinear operators which depend on μ and f, respectively, we will write these operators as $P(\mu)$ and $\widetilde{P}(f)$. Note that for each fixed f, $\widetilde{P}(f)$ is a linear operator. In the following result, $\mathbb{B}(L^2(\Omega))$ stands for the set of bounded linear operators on $L^2(\Omega)$.

Lemma 4.3.7. *The map from* $L^2(\Omega) \to \mathbb{B}(L^2(\Omega))$*, defined as* $f \mapsto \widetilde{P}(f)$*, is uniformly bounded; that is, for every* $f \in L^2(\Omega)$*,* $\|\widetilde{P}(f)\| \leq C$ *for some* $C > 0$*. Moreover, this result also holds true for* $\widetilde{P}$ *as an operator on* $L^1(\Omega)$*.*

Proof. This follows from the fact that $\widetilde{P}$ depends on f through the a_f function, which is in L^∞ for any $f \in L^1(\Omega)$ or $L^2(\Omega)$. An application of Theorem 6.18 in Folland (2013) then proves the result for $\widetilde{P}_1$. The result holds true trivially for $\widetilde{P}_2$, since it is a *multiplication operator* □

We will use this result in Section 4.4.

Clearly, the operator $\widetilde{P}$ satisfies $\widetilde{P}f^d = f^d$. Further, note that $\widetilde{P}$ is constructed to satisfy $\widetilde{P}(f^d) = I$, in order to ensure that all agents stop transitioning between states when the target density f^d is reached.

Next, we will show that f^d is a globally asymptotically stable equilibrium of system (4.6).

Theorem 4.3.8. *For the system* (4.6)*, f^d is globally asymptotically stable in the $L^1(\Omega, m)$ norm; that is,*

$$\|f_n - f^d\|_1 \to 0 \quad as \quad n \to \infty.$$

Before presenting the proof of this theorem, we make the following observation. Consider the case when for some $y \in \Omega$, $f_n(y) > f^d(y)$. Then, it follows that $a_{f_n}(y) > 0$. Expression (4.16) then becomes:

$$\widetilde{P}f_n(y) = \int_\Omega a_{f_n}(x)k(x,y)f_n(x)dx + f^d(y).$$

The first term in the equation above is non-negative. Therefore, one of the following conditions must be true:

$$\begin{aligned} f_{n+1}(y) &\geq f_n(y) > f^d(y); \\ f_n(y) &> f_{n+1}(y) \geq f^d(y). \end{aligned} \tag{4.17}$$

Consequently, it is not possible that $f_{n+1}(y) < f^d(y)$ for any value of n. Next, consider the case when $y \in \Omega$ is such that $f_n(y) \leq f^d(y)$. In this case, $a_{f_n}(y) = 0$. Expression

(4.16) then reduces to:

$$\widetilde{P}f_n(y) = \int_\Omega a_{f_n}(x)k(x,y)f_n(x)dx + f_n(y). \tag{4.18}$$

Similar to the previous case, given that the first term in the equation above is non-negative, one of the following conditions must be true:

$$\begin{aligned} f_{n+1}(y) \geq f^d(y) > f_n(y); \\ f^d(y) > f_{n+1}(y) \geq f_n(y). \end{aligned} \tag{4.19}$$

Therefore, in this case, we observe that $f_{n+1}(y)$ monotonically increases with n.

Define the sets

$$\begin{aligned} E_n^1 &= \{y \in \Omega : f_n(y) < f^d(y)\}, \\ E_n^2 &= \{y \in \Omega : f_n(y) = f^d(y)\}, \\ E_n^3 &= \{y \in \Omega : f_n(y) > f^d(y)\}. \end{aligned}$$

We note that $\Omega = E_n^1 \sqcup E_n^2 \sqcup E_n^3$, where $\sqcup$ denotes a disjoint union.

We can now state the proof of Theorem 4.3.8. To summarize, the proof employs an argument by contradiction that if the density f_n converges to a function other than f^d, then the measure μ_n is pushed from sets where its density f_n is greater than f^d to sets where $f_n < f^d$. This is straightforward to conclude from the definitions of the transition kernels K and Q; however, to prove the convergence of f_n to f^d, it is necessary to precisely quantify the measure that is pushed during each time step, which is computed in the proof.

Proof of Theorem 4.3.8. To prove this result, it is sufficient to show that on the set E_n^1, $\|f_n - f^d\|_1 \to 0$ as $n \to \infty$. This follows from the fact that each f_n is a probability density on Ω. On E_n^1, by (4.19), we have that $f_{n+1} \geq f_n$, and hence that $f^d - f_n \geq f^d - f_{n+1}$. Set $F_n = (f^d - f_n)^+$, where for an arbitrary function $h : \mathbb{R}^d \to \mathbb{R}$, h^+ denotes the positive part of h. Then F_n is monotonically decreasing on Ω. The sequence $(F_n)_n$ is bounded, and monotonically decreasing, which implies

that F_n converges pointwise to a function, say g. By the *monotone convergence theorem* (Folland, 2013), we then have that $\int_\Omega F_n \to \int_\Omega g$. If $g = 0$, then we have our result. If $g \neq 0$, then since f_n is a probability density on Ω, $\int_\Omega F_n \to \int_\Omega g$ implies that$\int_\Omega (f_n - f^d)^+ \not\to 0$. We will next prove by contradiction that g is in fact 0.

We suppose that $g \neq 0$. Let $\int_\Omega g \geq \gamma$, where $\gamma > 0$. Define $S = \{x \in \Omega : g(x) > 0\}$. We note that the definition of S is independent of time. Given the conditions in (4.17) and (4.19), it follows that $E_n^1 \supset E_{n+1}^1$ for all n. Due to the convergence of F_n to g, we must have that for all n, $S \subset E_n^1$. Moreover, $\lim_{n\to\infty} m(E_n^1) \to m(S)$. Note that,

$$\int_S f^d(x) - f_n(x)dx \geq \int_S g(x)dx > \gamma. \tag{4.20}$$

Since Ω is compact, Ω can be covered by a finite number M of balls of radius ε, where $4\varepsilon < r$. That is, $\Omega \subset \cup_{i=1}^M B_\varepsilon(x_i)$ for some $x_i \in \Omega$. We will denote $B_\varepsilon(x_i) \cap \Omega$ by $B(x_i)$. Choose a ball $B(x_j)$ from this cover that intersects both E_n^1 and $(E_n^1)^c$. Then,

$$m(B(x_j)) = \; m(B(x_j) \cap S) + m(B(x_j) \cap (E_n^1 \setminus S)) + m(B(x_j) \cap (E_n^1)^c). \tag{4.21}$$

Let $m(B(x_j) \cap S) \geq \epsilon_0$, for some $\epsilon_0 > 0$. If $m(B(x_j) \cap (E_n^1)^c) = 0$ at the current time n, then we look for a large enough time $T \in \mathbb{Z}_+$ such that $m(E_T^1 \setminus S)) \leq \epsilon_1 << \epsilon_0$. At times $n \geq T$, (4.21) shows that $m(B(x_j) \cap (E_n^1)^c) > 0$, ensuring the existence of at least one ball from the cover that has intersections of positive measure with both S and $(E_n^1)^c$.

Next, let $J = \{1, \ldots, M\}$ and define the following sets:

$$N_1 = \bigcup_{\substack{i\in J \\ m(B(x_i)\cap S)>0}} B(x_i),$$

$$N_k = \bigcup_{\substack{i\in J \\ m(B(x_i)\cap N_{k-1})>0}} B(x_i) \setminus N_{k-1}, \quad k > 1.$$

Let $n > T$. If $\int_{N_1 \cap (E_n^1)^c} f_n - f^d$ is not tending to 0 with increasing n, then we must have that $\int_{N_1 \cap (E_n^1)^c} f_n - f^d \geq \delta$ infinitely often $(i.o)$, for some $\delta > 0$. Moreover, each time the integral exceeds δ, the measure that is pushed from $N_1 \cap (E_n^1)^c$ to S can be quantified as

$$\begin{aligned}
&\int_{N_1 \cap (E_n^1)^c} Q_1(x, S) d\mu_n(x) \\
&= \int_{N_1 \cap (E_n^1)^c} \int_S a_{f_n}(x) q(x, y) dy f_n(x) dx \\
&= \int_{N_1 \cap (E_n^1)^c} (f_n(x) - f^d(x)) \int_S q(x, y) dy dx \\
&= C_1 \int_{N_1 \cap (E_n^1)^c} f_n(x) - f^d(x) dx,
\end{aligned}$$

where the constant C_1 in the last expression is $\int_S q(x, y) dy$. Therefore, the measure that gets pushed onto S from $N_1 \cap (E_n^1)^c$ is $C_1 \delta$ at every time n when $\int_{N_1 \cap (E_n^1)^c} f_n - f^d \geq \delta$. Let $\{t_n\}_n$ be a sequence in $\mathbb{Z}_+$ of all such times n, with $t_0 > T$. When the integral exceeds δ, we have that

$$\int_S f_{n+1}(x) dx = \int_S f_n(x) dx + C_1 \delta.$$

Consequently, for each t_n we have,

$$\int_S f_{t_n}(x) dx = \int_S f_n(x) dx + C_1 n \delta,$$

which implies that

$$\int_S f^d(x) - f_{t_n}(x) dx = \int_S f^d(x) - f_n(x) dx - C_1 n \delta.$$

As $n \to \infty$, the integral on the right-hand side of the equation above tends to $-\infty$, contradicting the fact that this integral is an upper bound on the integral of g over S, as per (4.20). Thus, we must have that $\int_{N_1 \cap (E_n^1)^c} f_n - f^d \to 0$ as $n \to \infty$.

We will now use an induction argument to show that $\int_{(E_n^1)^c} f_n - f^d \to 0$. We have just shown that this was true for the neighborhood of S given by $N_1 \cap (E_n^1)^c$. We

assume that $\int_{N_k \cap (E_n^1)^c} f_n - f^d \to 0$ for some $k > 1$. We will prove that this also holds true for $N_{k+1} \cap (E_n^1)^c$. Suppose that it is not true; then, $\int_{N_{k+1} \cap (E_n^1)^c} f_n - f^d \geq \delta_1$ *i.o* for some $\delta_1 > 0$. Again, denote the sequence of times when this happens by $\{t_n\}_n$. By construction, N_{k+1} does not intersect S; however, N_{k+1} may intersect E_n^1 (possibly a subset of N_k), to which it can push measure. We now demonstrate that N_{k+1} pushes most of its measure to $N_k \cap (E_n^1)^c$. We have established that for any $n \geq T$, $m(N_k \cap E_n^1) \leq m(E_n^1 \setminus S) \leq \epsilon_1$, which is arbitrarily small. Hence, $m(N_k \cap E_n^1)$ must be arbitrarily small, and therefore $m(N_k \cap (E_n^1)^c)$ must have positive measure. Consequently, we have that

$$\begin{aligned}
&\int_{N_{k+1} \cap (E_n^1)^c} Q_1(x, N_k \cap (E_n^1)^c) d\mu_n(x) \\
&= \int_{N_{k+1} \cap (E_n^1)^c} (f_n(x) - f^d(x)) \int_{N_k \cap (E_n^1)^c} q(x,y) dy dx \\
&= C_k \int_{N_{k+1} \cap (E_n^1)^c} (f_n(x) - f^d(x)) dx,
\end{aligned}$$

where $C_k = \int_{N_k \cap (E_n^1)^c} q(x,y) dy$. That is, the measure pushed from $N_{k+1} \cap (E_n^1)^c$ to $N_k \cap (E_n^1)^c$ is $C_k \delta_1$ for every t_n. Using similar arguments, we can conclude that $\int_{N_{k+1} \cap (E_n^1)^c} f_n - f^d \to 0$ as $n \to \infty$. Since Ω is compact, this process of induction must stop at a finite k. Therefore, we have that $\int_{(E_n^1)^c} f_n - f^d \to 0$, and consequently, $g = 0$, proving that f^d is globally attractive. Since, $f^d - f_n$ is strictly decreasing on the set E_n^1 and $\int f_n = 1$ for all n, we can conclude that, in fact, the equilibrium distribution f^d is stable in the sense of Lyapunov. This concludes the proof. $\square$

4.4 The N-Agent System

In this section, we will define the microscopic description of the system, i.e., the model of individual agents' state transitions, and study how it relates to the macroscopic or mean-field model (4.3). The following mathematical definitions are adapted from Del Moral *et al.* (1998). Consider a population of N agents evolving on the state

space Ω. Let the state of each agent k at time n be given by the random variable $\xi_n^k \in \Omega,\ k = \{1, \ldots, N\}$. Each agent transitions between states on Ω according to the transition kernel Q defined in (3.11). The N-agent system can therefore be described as a Markov chain $\xi_n = (\xi_n^1, \ldots, \xi_n^N)$ with state space Ω^N. To a measure $\nu \in \mathcal{P}(\Omega)$, we associate a measure $\nu^{\otimes N} = \nu \times \ldots \times \nu \in \mathcal{P}(\Omega^N)$. The empirical measure $m^N(x)$ associated with the point $x = (x^1, \ldots, x^N) \in \Omega^N$, where each entry x^k is the state of agent k, is given by a normalized sum of Dirac measures associated with each agent,

$$m^N(x) = \frac{1}{N} \sum_{k=1}^{N} \delta_{x^k}. \tag{4.22}$$

The corresponding Markov process $(\xi_n)_n$ on $(\Omega^N, \mathcal{F}_n, \mathbf{P})$ is defined by

$$\mathbf{P}(\xi_0 \in dx) = \mu_0^{\otimes N}(dx), \tag{4.23}$$

$$\mathbf{P}(\xi_n \in dx | \xi_{n-1} = z) = (Pm^N(z))^{\otimes N}(dx), \tag{4.24}$$

where $dx = dx^1 \times \ldots \times dx^N$, and P is as defined in (4.14). At time $n = 0$, the N-agent system can be modeled as N independent random variables $\xi_0^1, \ldots, \xi_0^N$ with common distribution μ_0. At time $n \geq 1$, define $\mu_n^N := m^N(\xi_n)$. Then, μ_{n+1}^N is evaluated as

$$\mu_{n+1}^N = Pm^N(\xi_n). \tag{4.25}$$

Thus, from the equation above, at time n the N-agent system is modeled as N random variables $\xi_n^1, \ldots, \xi_n^N$ that are conditional on ξ_{n-1} and distributed according to $Pm^N(\xi_{n-1})$. The agents' states are therefore not independent of one another; their distribution is dependent on the system configuration at time $n-1$. Although the evolution of each agent's state is not Markovian, the distribution of the N-agent system evolves according to an *interacting* Markov chain. At time $n = 0$, $\mu_0^N \to \mu_0$ as $N \to \infty$. At times $n \geq 1$, due to the aforementioned interaction between agents, the *law of large numbers* does not apply. Thus, another method must be used to establish

the limit $\mu_n^N \overset{N\to\infty}{\longrightarrow} \mu_n$, where μ_n evolves according to (4.3). This limit is called the *mean-field limit*. The work Del Moral *et al.* (1998) proved this limit for systems of the form (4.3) in which the right-hand side is continuous. In Kolokoltsov (2010), this limit is referred to as the *dynamic law of large numbers*; it is proven for Markov processes whose evolution is governed by a partial differential equation (PDE).

Since the empirical measure m^N is a sum of Dirac measures, it is not absolutely continuous with respect to the Lebesgue measure. We will "mollify" the Dirac measures in order to be able to use results from the previous section and to apply the operators $\widetilde{P}$ and P defined in (4.14) and (4.16), respectively, to absolutely continuous measures. Mathematically, this means that the measure m^N is convolved with a smooth function $\phi : \mathbb{R} \to \mathbb{R}$, a *mollifier*, to obtain a smooth function (density). The convolution of m^N and ϕ is carried out as

$$\phi * m^N = \int_\Omega \phi(x) dm^N = \frac{1}{N} \sum_{i=1}^{N} \phi(x - x_i). \tag{4.26}$$

The result of this convolution is a sum of smooth functions, which is smooth. Loosely speaking, this convolution replaces each Dirac measure by a measure with smooth density ϕ. We can now apply $\widetilde{P}$ and P to the right-hand side of this equation. In our simulations, we have defined ϕ as the standard *bump function* with a compact support:

$$\phi(x) = \begin{cases} e^{-\left(\frac{1}{1-\|x\|^2}\right)}, & x \in (-1, 1), \\ 0, & \text{otherwise.} \end{cases} \tag{4.27}$$

To change the support of ϕ, we define a function ϕ_h on $\mathbb{R}^d$ for some $h > 0$ as follows Folland (2013):

$$\phi_h(x) = h^{-d} \phi\left(\frac{x}{h}\right). \tag{4.28}$$

Note that $\int \phi_h = 1$, which is independent of h. Moreover, the "mass" of ϕ_h becomes concentrated at the origin as $h \to 0$; that is, ϕ_h tends to a Dirac measure as $h \to 0$.

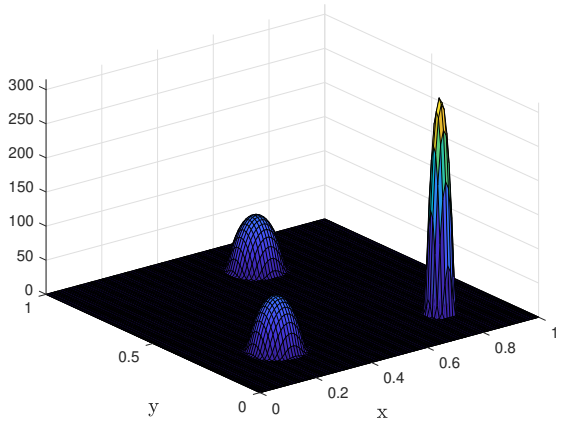

Figure 4.1: Visualization of two bump functions ϕ_h with $h = 0.1$ and one bump function with $h = 0.05$.

Figure 4.1 shows a visualization of two bump functions with $h = 0.1$ and one with $h = 0.05$. Since the integral of all bump functions is 1, to compensate for the decrease in h, the peak of the bump function with $h = 0.05$ is significantly higher than the peaks of the functions with $h = 0.1$.

The introduction of the mollifier also has implications for the implementation of the N-agent system in practice. For an agent with state x, given a distribution f, the transition kernel K in (4.10) is defined such that it requires pointwise evaluation of the function $f(x)$ in the term $a_f = (f - f^d)/f$ from (4.7). However, to evaluate the density ϕ at its state x using (4.26), the agent must know the states x_i of all other agents whose states are within a distance h of its own. For example, if the agents' states are their positions in space, mollification of the empirical measure implies that each agent must estimate the density ϕ in (4.26) based on its relative distance to all agents that are located within a neighborhood whose size is determined by the parameter h. As $h \to 0$, this neighborhood shrinks, and the density tends to the Dirac measure, which is singular.

In order to derive the macroscopic (mean-field) model from the microscopic de-

scription of the system, i.e. the dynamics of N individual agents, one typically needs to take the mean-field limit, as described earlier in this section. Since we have introduced the mollifier, a second limit needs to be proven as well. Both limits are defined below.

1. $N \to \infty$: We now introduce a measure μ_n^h that evolves according to the deterministic difference equation

$$\mu_{n+1}^h = P(\phi_h * \mu_n^h)\mu_n^h, \quad \mu_0^h \in \mathcal{P}(\Omega). \tag{4.29}$$

   Due to the introduction of the mollifier, we expect the N-agent system (4.25) to converge to the system above, which is different from (4.3). That is,

$$\mu_n^N \to \mu_n^h \ \text{ as } \ N \to \infty.$$

   This limit is usually proven in the *weak topology* and can be established for discrete-time systems using results from Del Moral *et al.* (1998). Applying these results requires proving that the right-hand side of (4.3) is continuous in the weak topology, which is significantly challenging for our system. Thus, we will reserve this investigation for future work.

2. $h \to 0$: The second limit proves that the solution of (4.29) converges to the solution of (4.3); that is, for all $n \in \mathbb{Z}_+$,

$$\mu_n^h \to \mu_n \ \text{ as } \ h \to 0. \tag{4.30}$$

   We shall prove this convergence in the $L^1(\cdot)$ norm in the next subsection.

4.4.1 The limit as $h \to 0$

We prove the limit (4.30) for a dense subset of $L^1(\Omega)$; specifically, we consider distributions $\mu \in \mathcal{P}(\Omega)$ that have $L^2(\Omega)$ densities. Moreover, we require f^d to be

bounded from below a.e. on Ω.

Let $\mu_0 \ll m$ with density $f_0 \in L^2(\Omega)$. In Proposition 4.3.6, we proved that $\widetilde{P}$ preserves $L^2(\Omega)$; that is, $f_n = P^n f_0 \in L^2(\Omega)$ for $n \in \mathbb{Z}_+$ Therefore, system (4.29) can be rewritten on $L^2(\Omega)$ as

$$f_{n+1}^h = \widetilde{P}(\phi_h * f_n^h) f_n^h, \quad f_0^h \in L^2(\Omega). \tag{4.31}$$

Since $m(\Omega) < \infty$, $L^2(\Omega) \subset L^1(\Omega)$, and therefore we will consider system (4.6) to be a system on $L^2(\Omega)$ instead of $L^1(\Omega)$. We will show that solutions of the above system converge to those of (4.6) in the $L^1(\Omega)$ norm.

Theorem 4.4.1. *Suppose the initial conditions f_0^h and f_0 are in $L^2(\Omega)$. Let f_n^h and f_n be solutions of (4.31) and (4.6), respectively. If f^d is bounded from below a.e. on Ω, then*

$$\|f_n^h - f_n\|_1 \to 0$$

for any $n \in \mathbb{Z}_+$.

To prove this result, we need the following proposition,

Proposition 4.4.2. *Let $g \in L^2(\Omega)$. If f^d is bounded from below a.e. on Ω, then we have the following convergence results:*

1. *For $f \in L^2(\Omega)$,*
$$\|\widetilde{P}(\phi_h * f)g - \widetilde{P}(f)g\|_1 \overset{h\to 0}{\longrightarrow} 0$$
2. *If $f_i \overset{i\to\infty}{\longrightarrow} f$ in the $L^1(\Omega)$ norm, then*
$$\|\widetilde{P}(f_i)g - \widetilde{P}(f)g\|_1 \overset{i\to\infty}{\longrightarrow} 0$$

Proof. (1) Let $f \in L^2(\Omega)$. Then, $\phi_h * f \in L^2(\Omega)$. By Theorem 8.14 of Folland (2013), $\phi_h * f \overset{h\to 0}{\longrightarrow} f$ in the L^2 norm. To prove convergence of $\widetilde{P}(\phi_h * f)$ to $\widetilde{P}(f)$ as operators

on $L^1(\Omega)$, choose $g \in L^2(\Omega)$ (since Ω has finite measure, $g \in L^1(\Omega)$), and compute the following:

$$\begin{aligned} &\|\widetilde{P}(\phi_h * f)g - \widetilde{P}(f)g\|_1 \\ &= \int_\Omega \left|\widetilde{P}(\phi_h * f)g(y) - \widetilde{P}(f)g(y)\right| dy. \end{aligned} \tag{4.32}$$

Recall that according to (4.16), $\widetilde{P} = \widetilde{P}_1 + \widetilde{P}_2$. We will now evaluate the integral (4.32) in terms of the two operators $\widetilde{P}_1$ and $\widetilde{P}_2$.

In (4.32), the component of the integrand that depends on $\widetilde{P}_1$ is given by:

$$\widetilde{P}_1(\phi_h * f)g(y) - \widetilde{P}_1(f)g(y) = \int_\Omega A(x,y)dx, \text{ where}$$
$$A(x,y) = a_{\phi_h * f}(x)q(x,y)g(x) - a_f(x)q(x,y)g(x).$$

We now define the following sets:

$$E_1 = \{x \in \Omega : \phi_h * f(x) > f^d(x)\},\ E_2 = \{x \in \Omega : \phi_h * f(x) \le f^d(x)\},$$
$$E_3 = \{x \in \Omega : f(x) > f^d(x)\},\ E_4 = \{x \in \Omega : f(x) \le f^d(x)\}.$$

We will split the integral $\int_\Omega A$ over four sets constructed from these sets, namely, $S_1 = \{E_1 \cap E_3\}, S_2 = \{E_2 \cap E_3\}, S_3 = \{E_1 \cap E_4\}$, and $S_4 = \{E_2 \cap E_4\}$. Note that $S_1 \sqcup S_2 \sqcup S_3 \sqcup S_4 = \Omega$. Consider the integral of A over S_1:

$$\begin{aligned} \int_{E_1 \cap E_3} A &\le \|q\|_\infty \int_{E_1 \cap E_3} a_{\phi_h * f}g - a_f g \\ &= \|q\|_\infty \int_{E_1 \cap E_3} \frac{\phi_h * f - f^d}{\phi_h * f} g - \frac{f - f^d}{f} g \\ &= \|q\|_\infty \int_{E_1 \cap E_3} \frac{\phi_h * f - f}{\phi_h * f} \frac{f^d}{f} g. \end{aligned} \tag{4.33}$$

Note that on $E_1 \cap E_3$, $\phi_h * f \xrightarrow{\|\cdot\|_2} f > f^d > 0$ and $\left\|\frac{f^d}{f}\right\|_\infty < 1$. Since f^d is bounded from below a.e. on Ω, we must have that $\left\|\frac{1}{\phi_h * f}\right\|_\infty < \infty$. Continuing the computation

from above,

$$\begin{aligned}
&\|q\|_\infty \int_{E_1\cap E_3} \frac{\phi_h * f - f}{\phi_h * f}\frac{f^d}{f} g \\
&\leq \|q\|_\infty \left\| \frac{1}{\phi_h * f} \right\|_\infty \left\| \frac{f^d}{f} \right\|_\infty \int\limits_{E_1\cap E_3} \phi_h * f - fg \\
&\leq \|q\|_\infty \left\| \frac{1}{\phi_h * f} \right\|_\infty \left\| \frac{f^d}{f} \right\|_\infty \|\phi_h * f - f\|_2 \|g\|_2.
\end{aligned}$$

The second inequality above follows from *Hölder's inequality*. Since we have established that $\|\phi_h * f - f\|_2 \to 0$ as $h \to 0$, the integral of A over S_1 must converge to 0. Next, we consider the integral of A over S_2:

$$\begin{aligned}
\int\limits_{E_2\cap E_3} A &\leq \|q\|_\infty \int\limits_{E_2\cap E_3} -\frac{f - f^d}{f} g \\
&\leq \|q\|_\infty \left\| \frac{f - f^d}{f} \right\|_\infty \|g\|_2 \; m(E_2 \cap E_3).
\end{aligned} \tag{4.34}$$

The second inequality follows from Hölder's inequality. In this case, we will establish that $m(E_2 \cap E_3) \to 0$ as $h \to 0$, which would imply that the integral of A over S_2 converges to 0. We can compute $m(E_2 \cap E_3)$ as:

$$\begin{aligned}
&m(E_2 \cap E_3) = m(\{\phi_h * f - f^d \leq 0\} \cap \{f - f^d > 0\}) \\
&= m(\{(\phi_h * f - f) + (f - f^d) \leq 0\} \cap \{f - f^d > 0\}).
\end{aligned}$$

Note that,

$$\{(\phi_h * f - f) + (f - f^d) \leq 0\} \subset \{(\phi_h * f - f) \leq 0\}.$$

Continuing the computation from above,

$$\begin{aligned}
m(E_2 \cap E_3) &\leq m(\{\phi_h * f - f \leq 0\} \cap \{f - f^d > 0\}) \\
&= m(\{\phi_h * f - f < 0\} \cap \{f - f^d > 0\})
\end{aligned}$$

By Proposition 2.29 of Folland (2013), since $\phi_h * f - f \to f$ in the L^2 norm as $h \to 0$, then $\phi_h * f - f \to f$ in measure; that is, $m(\{\phi_h * f - f \leq \delta\}) \to 0$ as $h \to 0$ for every

$\delta > 0$. Therefore, we must have that $m(E_2 \cap E_3) \to 0$ as $h \to 0$, and consequently, the integral of A over $E_2 \cap E_3$ must converge to 0. Now, consider the integral of A over S_3:

$$\begin{aligned} \int\limits_{E_1 \cap E_4} A \;&\leq\; \|q\|_\infty \int_{E_1 \cap E_4} \frac{\phi_h * f - f}{\phi_h * f} g \\ &\leq\; \|q\|_\infty \left\| \frac{1}{\phi_h * f} \right\|_\infty \|g\|_2 \; \|\phi_h * f - f\|_2 \; m(E_1 \cap E_4). \end{aligned} \tag{4.35}$$

The second inequality follows from Hölder's inequality. Since we have that $\|\phi_h * f - f\|_2 \to 0$ as $h \to 0$, the integral of A over $E_1 \cap E_4$ converges to 0. Finally, the integral of A over S_4 is trivially zero:

$$\int_{E_2 \cap E_4} A = \int_{E_2 \cap E_4} a_{\phi_h * f} g - a_f g = 0. \tag{4.36}$$

Thus, we have shown that $\int_\Omega A \to 0$ as $h \to 0$.

Returning to the integral (4.32), the component of the integrand that depends on $\widetilde{P}_2$ is given by:

$$\begin{aligned} &\widetilde{P}_2(\phi_h * f) g(y) - \widetilde{P}_2(f) g(y) = \left(1 - a_{\phi_h * f}(y)\right) g(y) - \left(1 - a_f(y)\right) g(y) \\ &\quad = a_f(y) g(y) - a_{\phi_h * f}(y) g(y) := B(y). \end{aligned} \tag{4.37}$$

This term is equal to the integrand of each of the four integrals considered in (4.33)-(4.36). Since we showed that each of these integrands tends to 0 as $h \to 0$, we must have that $B(y) \to 0$ as well.

We can now evaluate (4.32) as

$$\begin{aligned} \|\widetilde{P}(\phi_h * f) g - \widetilde{P}(f) g\|_1 \;&= \int_\Omega \left| \int_\Omega A(x,y) dx + B(y) \right| dy \\ &= \int_\Omega \left| \int_{S_1 \sqcup S_2 \sqcup S_3 \sqcup S_4} A(x,y) dx + B(y) \right| dy. \end{aligned}$$

Since we have shown that both $\int_\Omega A \to 0$ and $B(y) \to 0$ as $h \to 0$, the outer integral converges to 0 as well, and we have our result.

(2) The proof of this result is similar to the proof of result (1). $\square$

Remark 4.4.3. *Proving convergence of the operators $\widetilde{P}(\phi_h * f)$ and $\widetilde{P}(f_i)$ (as operators on L^1) in the standard topologies, such as the operator norm or the strong operator, would have required us to prove both results in the above proposition for any $g \in L^1(\Omega)$, which is quite challenging for our system. Since we have established that $f_n = \widetilde{P}^n f_0$ is restricted to $L^2(\Omega)$ for any initial density $f_0 \in L^2(\Omega)$, choosing $g \in L^2(\Omega)$ is sufficient to prove Theorem 4.4.1.*

We can now prove Theorem 4.4.1.

Proof of Theorem 4.4.1. To prove this result, we will use an induction argument. For $n = 1$, we have that

$$f_1^h = \widetilde{P}(\phi_h * f_0) f_0,$$
$$f_1 = \widetilde{P}(f_0) f_0.$$

Then, by statement (1) of Proposition 4.4.2, $\|f_1^h - f_1\|_1 = \|\widetilde{P}(\phi_h * f_0) f_0 - \widetilde{P}(f_0) f_0\|_1 \to 0$ as $h \to 0$. Assume that this is true for some $n > 1$; i.e., $\|f_n^h - f_n\|_1 \to 0$ as $h \to 0$. We will show that this limit holds true for $n+1$ using the following computation:

$$\begin{aligned}
\|f_{n+1}^h - f_{n+1}\| &= \left\| \widetilde{P}(\phi_h * f_n^h) f_n^h - \widetilde{P}(f_n) f_n \right\|_1 \\
&= \left\| (\widetilde{P}(\phi_h * f_n^h) f_n^h - \widetilde{P}(f_n) f_n^h) \ + \ (\widetilde{P}(f_n) f_n^h - \widetilde{P}(f_n) f_n) \right\|_1 \\
&= \left\| (\widetilde{P}(\phi_h * f_n^h) - \widetilde{P}(f_n)) f_n^h + \widetilde{P}(f_n)(f_n^h - f_n) \right\|_1 \\
&\leq \left\| (\widetilde{P}(\phi_h * f_n^h) - \widetilde{P}(f_n)) f_n^h \right\|_1 + \left\| \widetilde{P}(f_n)(f_n^h - f_n) \right\|_1
\end{aligned}$$

The bracket $(f_n^h - f_n)$ in the second term converges to 0 as $h \to 0$ due to our assumption. Considering the first term, we observe that:

$$\lim_{f_n^h \to f_n} \lim_{h \to 0} \widetilde{P}(\phi_h * f_n^h) = \widetilde{P}(f_n)$$

This follows from the fact that the inner limit tends to $\widetilde{P}(f_n^h)$ by statement (1) of Proposition 4.4.2, and the outer limit $\lim_{f_n^h \to f_n} \widetilde{P}(f_n^h)$ tends to $\widetilde{P}(f_n)$ by statement

(2) of Proposition 4.4.2. Therefore, the bracket $\widetilde{P}(\phi_h * f_n^h) - \widetilde{P}(f_n)$ in the first term tends to 0 as $h \to 0$, and hence we have our result. $\square$

4.5 Simulations

In this section, we present numerical solutions of the mean-field model (4.3) and simulations of the corresponding N-agent system. We provide verification via these simulations that as $N \to \infty$, the simulations of the N-agent system (stochastic simulations) approach the solution of the deterministic system (4.3).

In the example below, we define the agent state space $\Omega \subset \mathbb{R}^2$ as the unit square $[0,1] \times [0,1]$, representing a physical domain in which the agents move. The target distribution, shown in Fig. 4.2, is set to $f^d = \sin^2(2\pi x^1) + \sin^2(2\pi x^2)$, where $[x^1\ x^2]^T \in \Omega$. The initial distribution is set to the Dirac measure at $(0,0)$. We consider a nonlinear vector field F in system (4.1) that represents a unicycle model:

$$\begin{aligned} x_{n+1}^1 &= x_n^1 + u_n^1 \cos(u_n^2), \\ x_{n+1}^2 &= x_n^2 + u_n^1 \sin(u_n^2). \end{aligned} \tag{4.38}$$

Here, $x_n = [x_n^1\ x_n^2]^T \in \Omega$ and $u_n = [u_n^1\ u_n^2]^T \in U$. The set of control inputs is defined as $U = [0, 0.1] \times [0, 2\pi]$. This map F satisfies all the required conditions stated in Section 4.3.

To simulate the mean-field model (4.6), we use Ulam's method that was introduced in Chapter 3. We need to discretize both Ω and U. The set Ω is partitioned into $n_x \in \mathbb{Z}_+$ sets, $\widetilde{\Omega} = \{\Omega_1, \ldots, \Omega_{n_x}\}$, where $\Omega = \cup_{i=1}^{n_x} \Omega_i$ and the sets Ω_i have intersections of zero Lebesgue measure. The set of control inputs U is approximated as a set of $n_u \in \mathbb{Z}_+$ discrete elements, $\widetilde{U} = \{v_1, \ldots, v_{n_u}\}$, where $v_i \in U$ for each i.

We now define the discretization of the mean-field model (4.3). Let $\mu \in \mathcal{P}(\widetilde{\Omega})$ and $j \in \mathcal{I}$, and let μ^d be the discretization of f^d on $\tilde{\Omega}$. Let $\mathbf{P} \in \mathbb{R}^{n_x \times n_x}$ be the

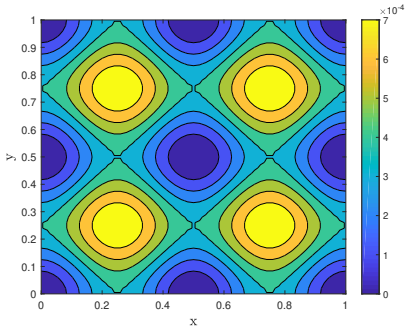

Figure 4.2: Target distribution f^d

discretization of the operator P defined in (4.14). Then the discretization of system (4.3) is given by:

$$\mu_{n+1} = \mathbf{P}\mu_n, \tag{4.39}$$
$$\mathbf{P}\mu(j) = \sum_{i\in\mathcal{I}} a_\mu(i) \sum_{l\in\mathcal{J}} \tilde{k}_{il}\ p_{ij}^l \mu(i) + (1 - a_\mu(j))\mu(j),$$

where $a_\mu(i) = (\mu(i) - \mu^d(i))/\mu(i)$ if $\mu(i) - \mu^d(i) > 0$, and $a_\mu(i) = 0$ otherwise. Figure 4.3 shows snapshots of the simulation of this system at several times n.

Algorithm 1 presents the program that simulates the evolution of agents over a domain Ω with a control set U, until a specified final time T_f. An agent considers another agent to be its neighbor if their relative distance is less than h, the parameter of the bump function ϕ_h described in the previous section. We denote the set of neighbors of agent k at any given time by $\mathcal{N}(k)$. At every time step, each agent computes the value of the bump function based on the relative distances of its neighbors. Computation of the bump function is provided in Algorithm 2. Note that C in Line 4 is a normalizing constant which is chosen to ensure that Φ is a probability density. Figures 4.4-4.7 show snapshots of the N-agent simulation for agent population sizes of N =100, 500, and 1000, with $h = 0.1$ in the first three figures and 0.05 in the last figure.

We first investigate the effect of increasing N while keeping h fixed on the time

Algorithm 1 Simulation of N agents

Input: $\Omega, U, k, F, N, f^d, h, T_f$

Initialize $n = 0, a^k = 0, x_0^k \in \Omega$ for all $k = 1, \ldots, N$

while $n \leq T_f$ **do**

  for $k = 1 : N$ **do**

    $s = 0$, $y = x_n^k$ ▷ y denotes current location of agent k

    for all $j \in \mathcal{N}(k)$ **do** ▷ $\mathcal{N}(k) :=$ {agents within distance h of k}

      $z = x_n^j$

      $s = s + \mathrm{PHI}(y, z, h)$

    end for

    $f_n(y) = \frac{1}{|\mathcal{N}(k)|} s$

    if $f_n(y) > f^d(y)$ **then**

      $a^k = \frac{f_n(y) - f^d(y)}{f_n(y)}$

    end if

    if $a^k > 0$ **then**

      Draw v uniformly from $(0, 1)$

      if $v \leq a^k$ **then**

        Draw $u \sim k(y, \cdot)$ from U

        $y = F(y, u)$

      end if

    end if

    $x_{n+1}^k = y$

  end for

  $n = n + 1$

end while

Algorithm 2 Evaluation of the local bump function

function PHI(y, z, h)
    $d = \|y - z\|_2$
    if $\frac{d}{h} < 1$ **then**
        $\Phi = \frac{1}{C}\frac{1}{h^2}\exp\left(\frac{-1}{1-(d/h)^2}\right)$ ▷ $C :=$ Normalizing constant
    end if
    return Φ
end function

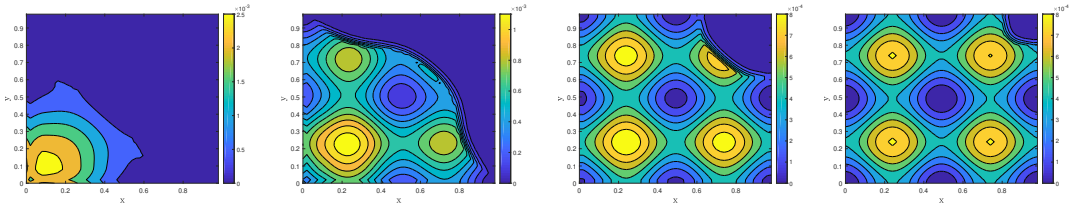

(a) $n = 50$ (b) $n = 200$ (c) $n = 500$ (d) $n = 2000$

Figure 4.3: Snapshots of the simulation of system (4.39) at several times n.

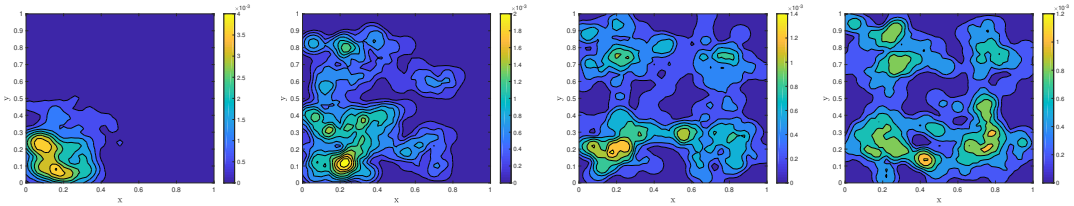

(a) $n = 50$ (b) $n = 200$ (c) $n = 500$ (d) $n = 2000$

Figure 4.4: Snapshots of a stochastic simulation of $N = 100$ agents, with $h = 0.1$, at several times n.

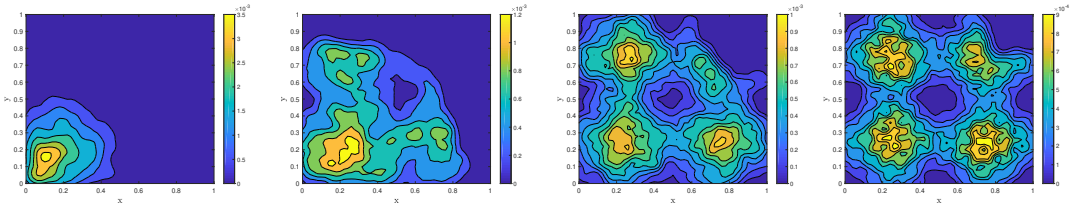

(a) $n = 50$ (b) $n = 200$ (c) $n = 500$ (d) $n = 2000$

Figure 4.5: Snapshots of a stochastic simulation of $N = 500$ agents, with $h = 0.1$, at several times n.

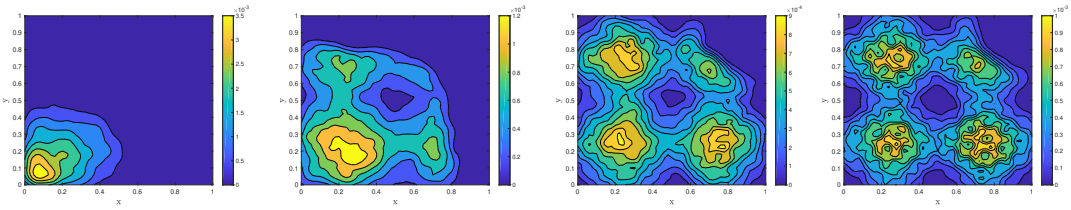

(a) $n = 50$ (b) $n = 200$ (c) Density at time $n = 500$ (d) $n = 2000$

Figure 4.6: Snapshots of a stochastic simulation of $N = 1000$ agents, with $h = 0.1$, at several times n.

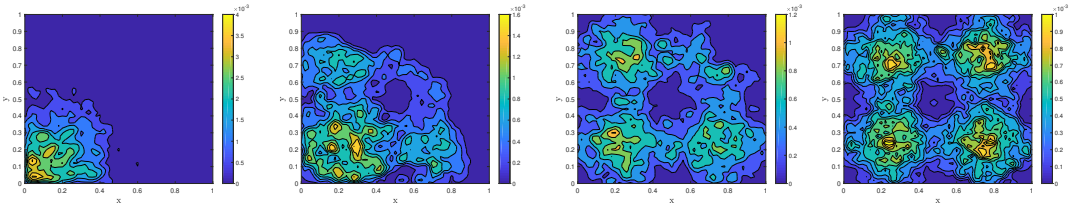

(a) $n = 50$ (b) $n = 200$ (c) $n = 500$ (d) $n = 2000$

Figure 4.7: Snapshots of a stochastic simulation of $N = 1000$ agents, with $h = 0.05$, at several times n.

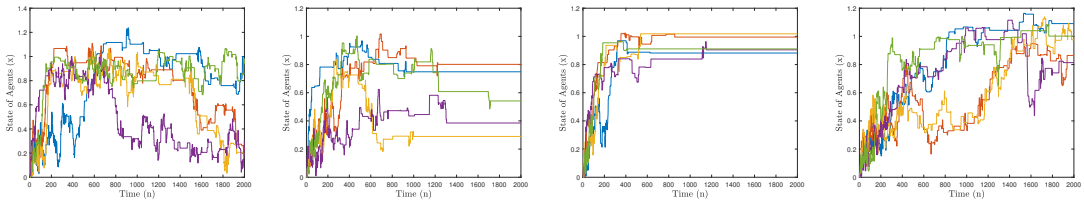

(a) $N = 100$, $h = 0.1$ (b) $N = 500$, $h = 0.1$ (c) $N = 1000$, $h = 0.1$ (d) $N = 1000$, $h = 0.05$

Figure 4.8: Time evolution of the 2-norm of five randomly selected agents' states in each of the N-agent simulations.

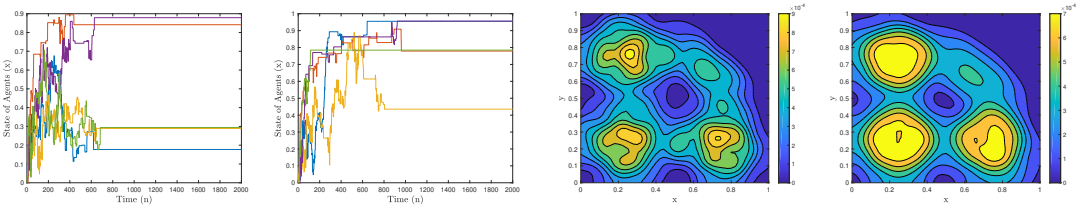

(a) $N = 500$, $h = 0.2$ (b) $N = 100$, $h = 0.25$ (c) $N = 1000$, $h = 0.2$ (d) $N = 1000$, $h = 0.25$

Figure 4.9: First and second sub-figures show the time evolution of the 2-norm of five randomly selected agents' states in two N-agent simulations with different values of N and h (snapshots of corresponding stochastic simulations not shown); Third and fourth sub-figures show snapshots at time $n = 2000$ of stochastic simulations of $N = 1000$ agents with different values of h.

evolution of the simulated N-agent system. Figure 4.3 shows that as time n increases, the mean-field model indeed converges asymptotically to the target distribution in Fig. 4.2. We observe that the convergence slows down significantly after time $n = 500$. Following our discussion in Section 4.4, we expect the stochastic simulations of the N-agent system to converge to the discretization of system (4.29) in the limit $N \to \infty$. Although system (4.29) is different from the system (4.3), note that the solutions of the two systems (4.3) and (4.29) converge in the limit $h \to 0$. The snapshots in Figs. 4.4-4.6 show that as the population size N is increased with a fixed value of h, the agent distribution in the N-agent simulation approaches the solution of system (4.39), plotted in Fig. 4.3. In all three figures, the agent distribution converges to a discrete approximation of the continuous target distribution.

Next, we study the effect of N on the frequency of agent state transitions. For each of the N-agent simulations shown in Figs. 4.4-4.7, Figs. 4.8a-4.8d plot the time evolution of the 2-norm of five randomly selected agents' states. The plots in Figs. 4.8a-4.8c show that the agents' frequency of state transitions significantly decreases with increasing N; the agents eventually stop transitioning between states (i.e., stop moving) for both $N = 500$ and $N = 1000$. This trend can be attributed to our approximation of a continuous distribution by a discrete function representing the state of the N-agent system. For low values of N, the resulting coarse discretization of f^d might yield an operator $\widetilde{P}(\phi_h * \mu_n^N)$ that is not a sufficiently accurate approximation of $\widetilde{P}(f^d) = I$, the condition that stops agent state transitions. Higher values of N produce a finer discretization of f^d, which improves the accuracy of the approximation of $\widetilde{P}(f^d) = I$. This validates our claim that control policies designed for the mean-field model can be implemented on a population of individual agents to achieve a target distribution, as long as this population is sufficiently large.

Finally, we investigate the effect of reducing h while keeping N fixed. Similar to the

agent distribution in Fig. 4.6 ($h = 0.1$), the agent distribution in Fig. 4.7 ($h = 0.05$) approaches the solution of (4.39) shown in Fig. 4.3 as N increases. However, the relative closeness of the distributions in Figs. 4.6 and 4.7 to the distribution in Fig. 4.3 is not apparent from the figures. This can be explained by noting that in this case, we are holding N constant and decreasing h, thereby reversing the order of limits that we considered in Section 4.4. There is no mathematical guarantee that the limits commute, and hence, we do not necessarily expect that with reduced h, the N-agent simulations will more closely approach the solution of (4.39). Moreover, a lower value of h for a fixed N yields a smaller neighborhood in which each agent evaluates the local density, which can produce a less accurate approximation of $\widetilde{P}(f^d) = I$. As explained previously, this can result in persistent agent state transitions, which are evident in the simulation of $N = 1000$ agents when h is reduced from 0.1 (Fig. 4.8c) to 0.05 (Fig. 4.8d). Increasing h, on the other hand, can result in the eventual cessation of agent transitions in smaller agent populations N. This is demonstrated in Figs. 4.9a-4.9b, which show that when h is increased from 0.2 to 0.25, the population of N agents stops transitioning for a lower value of N. The snapshots of stochastic simulations for $N = 1000$ at time $n = 2000$ in Figs. 4.6d, 4.7d, 4.9c, and 4.9d demonstrate that the agent distribution becomes smoother as h is increased, due to the smoothening effect of the mollification.

The results presented within this chapter are part of Biswal *et al.* (2020c,a).

Chapter 5

DISCRETE-TIME MARKOV CHAIN MODELS ON COMPACT MANIFOLDS

In this chapter, we consider discrete-time Markov chains (DTMC) that evolve on compact, smooth, connected manifolds without boundary. We focus on stabilizing and optimizing the convergence rate of a DTMC to target probability measures that are positive almost everywhere on the manifold and that are absolutely continuous with respect to the *Riemannian volume*, with L^∞ Radon-Nikodym derivatives. Our primary motivation stems from applications in multi-agent control systems; specifically, the problem of distributing an ensemble of identity-independent agents on a state space of choice.

A standard assumption in the literature on multi-agent control is that the state space of the agents is Euclidean. However, the state spaces of many mechanical systems are naturally represented as manifolds (Marsden and Ratiu, 2013). Some works have extended multi-agent control approaches on Euclidean spaces to manifolds; for example, consensus strategies on manifolds are presented in (Sarlette and Sepulchre, 2009), respectively.

From here on, therefore, we consider the problems addressed in the Chapter 3, wherein the state space was a compact subset of a Euclidean space, for state spaces that are manifolds. As we will see, most of our results in Chapter 3 carry straightforwardly to the results here.

5.1 Notation

Here, we introduce notation that is specific to the case when the state space is a manifold. We denote the state space by $\mathcal{M}$, a d-dimensional smooth, compact,

connected manifold without boundary. Let $T_x\mathcal{M}$ denote the tangent space of the manifold at $x \in \mathcal{M}$. We assume that $\mathcal{M}$ is equipped with a *bi-invariant Riemannian metric* of $g : T\mathcal{M} \times T\mathcal{M} \to \bar{\mathbb{R}}_+$, where $T\mathcal{M} = \cup_{x\in\mathcal{M}} T_x\mathcal{M}$ denotes the tangent bundle of the manifold $\mathcal{M}$. In particular, the natural measure associated with the Riemannian manifold, known as the Riemannian volume, will be denoted by m_g. Let $d_g : \mathcal{M} \times \mathcal{M} \to \bar{\mathbb{R}}_+$ denote the Riemannian distance on $\mathcal{M} \times \mathcal{M}$. For $x \in \mathcal{M}$ and $h > 0$, let $B_g(x, h) = \{y \in \mathcal{M};\ d_g(x, y) \leq h\}$ denote the ball of radius h centered at x. We denote the space of probability measures on $\mathcal{M}$ by $\mathcal{P}(\mathcal{M})$.

Consider the Hilbert space of real-valued square integrable functions $L^2(\mathcal{X}, m)$. The dual space of this space is itself. For $x, y \in L^2(\mathcal{X}, m)$, $\langle x, y\rangle = \int xydm$ defines an inner product on $L^2(\mathcal{X}, m)$. The *weak topology* on $L^2(\mathcal{X}, m)$, denoted as w, is the topology defined by the family of seminorms $\{p_{x^*} : x^* \in L^2(\mathcal{X}, m)\}$, where $p_{x^*}(x) = |\langle x, x^*\rangle|$. The *weak* topology*, denoted as w^*, is defined on the dual space $L^2(\mathcal{X}, m)$ by the family of seminorms $\{p_x : x \in L^2(\mathcal{X}, m)\}$, where $p_x(x^*) = |\langle x, x^*\rangle|$. The *weak operator topology* (WOT) on $\mathbb{B}(L^2(\mathcal{X}, m))$, the space of linear bounded operators that map $L^2(\mathcal{X}, m)$ to $L^2(\mathcal{X}, m)$, is defined by seminorms $\{p_{x,y} : x, y \in L^2(\mathcal{X}, m)\}$, where $p_{x,y}(T) = |\langle Tx, y\rangle|$ for $T \in \mathbb{B}(L^2(\mathcal{X}, m))$. Convergence in this topology is as follows:

$$(T_i)_i \xrightarrow{WOT} T \iff \langle T_i x, y\rangle \to \langle Tx, y\rangle,\ \forall x, y \in L^2(\mathcal{X}, m).$$

5.2 Problem Formulation

We begin by stating our assumptions. We consider measures in $\mathcal{P}(\mathcal{M})$ that have square integrable Radon-Nikodym derivatives with respect to m_g; this assumption gives us the advantage of working on a Hilbert space, $L^2(\mathcal{M}, m_g)$, which significantly simplifies the analysis. Consider the following discrete-time flow on the space of

probability densities $L^2(\mathcal{M}, m_g)$:

$$f_{n+1} = Pf_n, \;\; n = 0, 1, 2, ...$$
$$f_0 \in L^2(\mathcal{M}, m_g), \tag{5.1}$$

where $P : L^2(\mathcal{M}, m_g) \to L^2(\mathcal{M}, m_g)$ is the induced forward operator. To define P, let $K : \mathcal{M} \times \mathcal{B}(\mathcal{M}) \to [0, 1]$ be a transition kernel. To ensure that P preserves probability densities, we impose the following property on K:

$$\int_{\mathcal{M}} K(x, \mathcal{M}) dm_g(x) = 1, \quad \text{for } m_g\text{-a.e. } x \in \mathcal{M}. \tag{5.2}$$

Using system (5.1), we define a discrete-time Markov chain (DTMC) $\Phi = \{\Phi_0, \Phi_1, \ldots\}$ on $\mathcal{M}$ that describes an agent's dynamics on the state space. The Markov chain induces a probability measure $\mathbb{P}$ on $\mathcal{M}^\infty$, defined as follows: $\mathbb{P}(E)$ is the probability of the event $\{\Phi \in E\}$, where $E \in \bigvee_{i=0}^{\infty} \mathcal{B}(\mathcal{M}_i)$ (the product sigma algebra) with $\mathcal{M}_i = \mathcal{M}$ for each $i \in \bar{\mathbb{Z}}_+$. For every $n \in \mathbb{Z}_+$, we say that the random variable Φ_n is distributed according to μ_n, the measure corresponding to f_n, if $\mathbb{P}(\Phi_n \in E) = \mu_n(E)$. Suppose that Φ_n is the current agent state and is distributed according to μ_n. Then the Markov property implies that Φ_{n+1} is distributed according to μ_{n+1}, where the density f_{n+1} corresponding to μ_{n+1} is given by (5.1).

The action of P on a function $f \in L^2(\mathcal{M}, m_g)$ can be represented as follows: for $E \in \mathcal{B}(\mathcal{M})$,

$$\int_E (Pf)(x) dm_g(x) = \int_{\mathcal{M}} K(x, E) f(x) dm_g(x). \tag{5.3}$$

If K is regular, then we can obtain an explicit expression for P, rather than defining P through (5.3). Defining $k : \mathcal{M} \times \mathcal{M} \to \mathbb{R}_+$ as the kernel function of K, we have that $k \in L^\infty(\mathcal{M} \times \mathcal{M}, m_g \times m_g)$. From (5.3), we obtain the following: for $y \in \mathcal{M}$ and $f \in L^2(\mathcal{M}, m_g)$,

$$Pf(y) = \int_{\mathcal{M}} k(x, y) f(x) dm_g(x). \tag{5.4}$$

Operators of this form are called *integral operators* (Conway, 2013). The function k is called the *kernel of the integral operator*.

We will first consider the problem of stabilizing system (5.1) to a target density.

Problem 5.2.1. *Given a target density $f_d \in L^\infty(\mathcal{M}, m_g)$, determine whether there exists a transition kernel $K : \mathcal{M} \times \mathcal{B}(\mathcal{M}) \to [0, 1]$ such that* (5.1) *satisfies $\lim_{n\to\infty} P^n f_0 = f_d$ for all initial densities $f_0 \in L^2(\mathcal{M}, m_g)$, where the forward operator P is defined in* (5.3).

This problem will be addressed in Section 5.3. When the state space $\mathcal{M}$ is a Lie group, we can in fact show the existence of a regular transition kernel K, in which case P is defined in (5.4).

Given that there exists such a transition kernel, we then address the problem of choosing the transition kernel that optimizes the convergence rate (mixing rate) of system (5.1) to the target density. As in the Euclidean state space case, the convergence rate is characterized by the L^2 spectral gap. Toward this goal, we will prove the existence of a spectral gap for P. Further, we will prove in the next section that 1 is the unique largest eigenvalue of P, which implies that P is stochastic, as in the case of DTMCs that evolve on a discrete state space. Let $\lambda_2(P)$ be the eigenvalue of P with the second-largest modulus.

Problem 5.2.2. ***(Optimization of convergence rate)*** *Let $\mathcal{K}$ be the set of all Markov kernels defined on $\mathcal{M} \times \mathcal{B}(\mathcal{M}) \to [0, 1]$ that each correspond to a well-defined bounded operator on $L^2(\mathcal{M}, m_g)$. Given a target density $f_d \in L^\infty(\mathcal{M}, m_g)$, determine whether the following optimization problem admits a solution:*

$$\min_{\mathcal{K}} |\lambda_2(P)|$$

subject to the constraint $Pf_d = f_d$, where P is the forward operator (5.3).

5.3 Existence of a Solution to Problem III.1

We closely follow the steps that are laid out in Section 3.4 for the case of a Euclidean state space. Let a target density f_d be given that is strictly positive almost everywhere on $\mathcal{M}$ and satisfies $f_d, \frac{1}{f_d} \in L^\infty(\mathcal{M}, m_g)$. In this section, we will prove the existence of an operator P that has f_d as its fixed point, i.e., $Pf_d = f_d$.

We will first address the problem for the case where the state space is $\mathcal{M}$, a compact manifold. We define a regular kernel K associated with a kernel function k as

$$K(x, dm_g(y)) = k(x, y)dm_g(y) = \frac{\chi_{B_g(x,h)}}{m_g(B_g(x, h))}dm_g(y) \tag{5.5}$$

for all $x, y \in \mathcal{M}$, where $\chi_{(.)}$ is the characteristic function. This transition kernel induces the Markov chain known as the "ball walk" (Lebeau *et al.*, 2010). Let $S : L^2(\mathcal{M}, m_g) \to L^2(\mathcal{M}, m_g)$ be the operator defined by this transition kernel as per (5.4),

$$Sf(y) = \int_{\mathcal{M}} k(x, y)f(x)dm_g(x), \ f \in L^2(\mathcal{M}, m_g).$$

The resulting Markov chain has the invariant measure $m_g(B_g(x, h))dm_g(y)$, i.e., $Sf_\pi = f_\pi$ with $f_\pi(x) = m_g(B_g(x, h))$ for all $x \in \mathcal{M}$. Note that in the case when $\mathcal{M}$ is a Lie group, S is a self-adjoint operator.

Recall the definitions of irreducibility and primitivity of an operator from Section 3.4. We will now establish some fundamental spectral properties of the operator S. For small $h > 0$, these results were demonstrated in Lebeau *et al.* (2010) using the theory of pseudodifferential operators, along with precise quantitative estimates of $\lambda_2(S)$. Here we sketch an alternative proof, using basic functional analytic principles, that S has a spectral gap for arbitrary $h > 0$. We again invoke the Jentzsch-Perron theorem, Theorem 3.4.8, to establish the simplicity of the eigenvalue 1.

We can now state the following theorem.

Theorem 5.3.1. *The operator S is compact and $Sf_\pi = f_\pi$, where $f_\pi \in L^\infty(\mathcal{M}, m_g)$ is given by*

$$f_\pi(x) = Cm_g(B_g(x,h)) \tag{5.6}$$

for all $x \in \mathcal{M}$. Here, C is a normalizing constant such that $\int_{\mathcal{M}} f_\pi(x) dm_g(x) = 1$. Moreover, $r(S) = 1$ is a simple eigenvalue of the operator S, and $|\lambda_2(S)| < 1$.

Proof. The operator S is compact since it is an integral operator with an essentially bounded integral kernel k (Proposition II.4.7, Conway (2013)). Let π denote a measure that is absolutely continuous with respect to m_g with density f_π. In addition, let $M_{f_\pi} : L^2(\mathcal{M}, m_g) \to L^2(\mathcal{M}, \pi)$ be a *multiplication operator*, defined as $M_{f_\pi} g = f_\pi g$. Since $f_\pi \in L^\infty(\mathcal{M}, m_g)$, M_{f_π} is bounded and well-defined (Theorem II.1.5, Conway (2013)). Consider the operator $\hat{S} : L^2(\mathcal{M}, m_g) \to L^2(\mathcal{M}, \pi)$ that is given by

$$\hat{S} = M_{f_\pi}^{-1} S M_{f_\pi}. \tag{5.7}$$

The operator $\hat{S}$ is an integral operator with integral kernel $\hat{q}$, defined as

$$\hat{q} = \frac{k(x,y) f_\pi(x)}{f_\pi(y)}. \tag{5.8}$$

Then it follows from the proof of Proposition 3.3.5 that $\hat{S}$ is a contraction, and hence $r(S) = r(\hat{S}) = 1$.

To establish that $|\lambda_2(S)| < 1$, we consider the set $U_0^x = \{x\}$ for each $x \in \mathcal{M}$ and inductively define the sets $U_m^x = \cup_{y \in U_{m-1}^x} B_g(y,h)$ for each $m \in \mathbb{Z}_+$. Since the manifold $\mathcal{M}$ is compact, there exists $n \in \mathbb{Z}_+$, independent of $x \in \mathcal{M}$, such that $\mathcal{M} = U_n^x$. From this, it follows that if $f \in L^2(\mathcal{M}, m_g)$ is a non-zero non-negative function, then $S^n f$ is positive almost everywhere on $\mathcal{M}$. Hence, from Theorem 6.1 of Grobler (1995), it follows that the operator S is primitive, and the only eigenvalue of S with modulus 1 is $r(S)$. This last observation can also be concluded from the fact that $\hat{S}$ is self-adjoint. Using the fact that $S^n f$ is positive almost everywhere on

$\mathcal{M}$ if $f \in L^2(\mathcal{M}, m_g)$ is a non-zero non-negative function, we can conclude that S is irreducible. We can now invoke 3.4.8 to conclude that the eigenvalue 1 is simple, and hence that $|\lambda_2(S)| < 1$. $\square$

In the case that $\mathcal{M}$ is a *Lie Group* (Lee, 2001; Kirillov, 2008), the invariant measure characterized by f_π can be described explicitly under a particular condition on the metric g. Let $x \cdot y$ denote the *right-translation* of $x \in \mathcal{M}$ by y. Similarly, $y \cdot x$ denotes the *left-translation* of x by y. If the metric g is *bi-invariant*, then the distance d_g is invariant under translations, i.e., $d_g(x \cdot y, y \cdot z) = d_g(x, y) = d_g(z \cdot x, z \cdot y)$ for all $x, y, z \in \mathcal{M}$. In this case, the invariant measure coincides with the Riemannian volume m_g, or more specifically, the *Haar volume*. Due to the bi-invariance of the metric, it follows that for each $x, y \in \mathcal{M}$, $B_g(x \cdot y, h) = B_g(x, h) \cdot y := \{z \cdot y,\ z \in B_g(x, h)\}$. Similarly, for each $x, y \in \mathcal{M}$, $B_g(y \cdot x, h) = y \cdot B_g(x, h) := \{y \cdot z,\ z \in B_g(x, h)\}$. Hence, we have that $m_g(B_g(x, h)) = m_g(B_g(e, h)) = x^{-1} \cdot m_g(B_g(x, h)) = m_g(B_g(x, h)) \cdot x^{-1}$ for all $x \in \mathcal{M}$, where $e \in \mathcal{M}$ is the unique identity element of $\mathcal{M}$. Therefore, Theorem 5.3.1 can be rewritten for Lie groups as follows.

Theorem 5.3.2. *Let $\mathcal{M}$ be a Lie Group such that the metric g is bi-invariant. Then $S : L^2(\mathcal{M}, m_g) \to L^2(\mathcal{M}, m_g)$ is a compact operator, and $S\mathbf{1} = \mathbf{1}$. Moreover, S has a spectral gap, and hence $|\lambda_2(S)| < 1$.*

Our goal is to construct an operator P that has f_d as its fixed point. Toward this end, we define a multiplication operator $D : L^2(\mathcal{M}, m_g) \to L^2(\mathcal{M}, m_g)$ as $D(g) = \frac{g f_\pi}{f_d}$. Since $\frac{f_\pi}{f_d} \in L^\infty(\mathcal{M}, m_g)$ (note that f_d is bounded from below), D is well-defined and bounded. We define P as

$$P = (S - I)\varepsilon D + I, \quad 0 < \varepsilon << 1, \tag{5.9}$$

where I is the identity operator on $L^2(\mathcal{M}, m_g)$.

Remark 5.3.3. *For ε small enough, P is a positive operator.*

We note that since the identity operator I is not compact, P is not compact, and therefore it cannot be represented as an integral operator (5.4) with an L^2 kernel. Instead, we will show that P can be represented as (5.3) with a Markov kernel $Q : \mathcal{M} \times \mathcal{B}(\mathcal{M}) \to [0, 1]$. Unlike the kernel K in (5.5), Q is not regular. From (5.9), we can write this kernel as follows:

$$Q(x, E) = \int_E k(x, y)a(x)dy + (1 - a(x))\delta_x(E), x \in \mathcal{M}, \tag{5.10}$$

where $E \in \mathcal{B}(\mathcal{M})$, $a(x) = \frac{\varepsilon f_\pi(x)}{f_d(x)}$, and $\delta_{(\cdot)}$ is the Dirac measure. This can be easily confirmed to be a Markov transition kernel. Next, we establish properties of the spectrum of the new operator P. The proof of the result closely follows the proofs of Theorems 3.4.12 and 3.4.13, and hence we will omit the proof here.

Proposition 5.3.4. *The operator P defined in* (5.9) *satisfies $P^*\mathbf{1} = \mathbf{1}$, $Pf_d = f_d$. The eigenvalue* 1 *is algebraically simple, isolated (i.e., is not a limit point), and coincides with the spectral radius of P. Furthermore, for ε small enough, and with f_π, f_d bounded from below, P is primitive.*

The construction of P concludes our discussion on the existence of an operator that has a unique fixed point at f_d. We note that such an operator P is not necessarily unique. Next, we move on to optimizing over all such operators in order to maximize the convergence rate of system (5.1) to f_d.

5.4 Formulation of the Optimization Problem

As we did for the case of a Euclidean state space, in this section, we present a solution to a relaxed version of Problem 5.2.2. Recall our assumptions that f_d is in $L^\infty(\mathcal{M}, m_g)$ and is a.e. strictly positive on $\mathcal{M}$. Let μ_d be a measure that is absolutely continuous with respect to m_g with density f_d.

Instead of constructing P, we will pose this optimization problem for $\hat{P} = M_{f_d}^{-1} P M_{f_d}$, defined as in (5.7), which has the same spectrum as P. The advantage here is that $\hat{P}$ is doubly stochastic, which simplifies the formulation of the optimization problem. Recall that we are bounding from above the modulus of the second-largest eigenvalue, as is expressed in (3.36).

We now formulate the optimization problem. The optimization variable is the transition kernel function K in the definition (5.3) of P. The relationship between $\hat{P}$ and P is enforced as constraint (5.12) in the optimization problem, defined below.

$$\min_{K} \quad \left\| \hat{P}(K) \circ \mathrm{Proj}_{\mathbf{1}^\perp} \right\| \tag{5.11}$$

subject to

$$\hat{P} = M_{f_d}^{-1} P M_{f_d}, \tag{5.12}$$

$$K(x, E) \geq 0 \quad \forall x \in \mathcal{M},\ E \in \mathcal{B}(\mathcal{M}), \tag{5.13}$$

$$\int_{\mathcal{M}} K(x, \mathcal{M}) dx = 1 \quad \forall x \in \mathcal{M}, \tag{5.14}$$

$$\int_{\mathcal{M}} f_d(y) K(x, dy) = f_d(x) \quad \forall x \in \mathcal{M}, \tag{5.15}$$

$$K(x, \mathcal{M} \backslash B_g(x, r)) = 0, \quad \forall x \in \mathcal{M}. \tag{5.16}$$

The constraints (5.13), (5.14) ensure that K is indeed a transition kernel. Constraint (5.15) ensures that f_d is the stationary distribution of P. Equation (5.16) imposes a localization constraint on the corresponding Markov chain; that is, starting from any point $x \in \mathcal{M}$, the probability of choosing a point lying outside a ball of specified radius r is zero. This constraint captures physical limitations on an agent's motion as it traverses the state space, which restrict the agent to moving a distance bounded by r in a single time step.

We will discuss the convexity of the optimization problem in the next section.

5.5 Optimal Solution

In this section, we show that an optimal solution to the optimization problem (5.11)-(5.16) exists. In order to show this, we must prove that the set of decision variables, which will be defined shortly, is compact in some topology and that the objective function (5.11) is continuous on this set with respect to the chosen topology.

We begin with a definition. Operators that are described by expression (3.2), where the kernel is not necessarily regular, are called *pseudo-integral operators* (Sourour, 1979); integral operators form a subset. Suppose that $(\mathcal{X}, \mathcal{N}, \mu)$ is a finite Borel measure space.

Definition 5.5.1. *A bounded linear operator $T : L^2(\mathcal{X}, \mu) \to L^2(\mathcal{X}, \mu)$, where μ is a kernel, is called a pseudo-integral operator if T is given by the expression*

$$(Tf)(x) = \int f(y)\mu(x, dy), \ m\text{-a.e.} \tag{5.17}$$

for every $f \in L^2(\mathcal{X}, \mu)$.

Remark 5.5.2. *In fact, the kernel is uniquely determined by the operator in the sense that if $\nu(x, dy)$ satisfies* (5.17)*, then $\mu(x, \cdot) = \nu(x, \cdot)$ for m-almost every x.*

The following result (Theorem 3.1, Sourour (1979)) will be used in our upcoming discussion.

Theorem 5.5.3. *Let T be defined as in* (5.17)*. Then T is a pseudo-integral operator with a positive kernel if and only if T is a positive operator.*

The decision variable in the optimization problem is the transition kernel K. However, in view of the remark and theorem above, we shall instead define a set of operators that satisfy the optimization constraints and identify the transition kernels

with these operators. This formulation will significantly simplify our analysis. To begin, we define the following set.

$$\mathscr{P} = \Big\{ P : L^2(\mathcal{M}, m_g) \to L^2(\mathcal{M}, m_g) \text{ is a pseudo-integral operator with a kernel}$$
$$K : \mathcal{M} \times \mathcal{B}(\mathcal{M}) \to [0,1],\ P\mathbf{1} = \mathbf{1},\ P^* f_d = f_d \ \text{ for } f_d \in L^2(\mathcal{M}, m_g),$$
$$Pf(x) = \int_{B_g(x,r)} f(y) K(x, dm_g(y)) \ \text{ for } x \in \mathcal{M} \Big\}. \tag{5.18}$$

Here, the constraint $Pf(\cdot) = \int_{B_g(\cdot,r)} fK(\cdot, dm_g)$ is equivalent to the condition $K(x, \mathcal{M} \backslash B_g(x,r)) = 0$ in (5.16). For sufficiency, we obtain this condition by choosing f to be the constant function $\mathbf{1}$. The necessary direction is straightforward to prove. Note that the set $\mathscr{P}$ is defined in terms of operators that are of the form (3.2) due to the statement of Theorem 5.5.3, which involves operators of the form (5.17).

Proposition 5.5.4. *The set $\mathscr{P}$ is closed in the WOT topology.*

Proof. Let $(P_i)_i$ be a sequence in $\mathscr{P}$, and suppose that $(P_i)_i$ converges to P in WOT. We will show that $P \in \mathscr{P}$. WOT convergence implies that $\langle P_i f, g \rangle \overset{i \to \infty}{\longrightarrow} \langle Pf, g \rangle$ for all $f, g \in L^2(\mathcal{M}, m_g)$. In particular, take $f = \mathbf{1}$. Since $P_i \mathbf{1} = \mathbf{1}$ for all i, we have that $\langle P\mathbf{1}, g \rangle = \langle \mathbf{1}, g \rangle$ for all $g \neq 0$, which implies that $P\mathbf{1} = \mathbf{1}$. Similarly, $P^* f_d = f_d$. We now show that P is a positive operator. Suppose that $f, g \in L^2(\mathcal{M}, m_g)$ are positive functions. Then, $P_i f$ is non-negative for every i. Since $\bar{\mathbb{R}}_+$ is closed, the limit $\langle Pf, g \rangle$ must be non-negative, which implies that P must be a positive operator. From Theorem 5.5.3 and the condition $P\mathbf{1} = \mathbf{1}$, we conclude that P must be a pseudo-integral operator with a kernel K taking values in $[0,1]$. We now consider the last constraint in the set (5.18). Let $f, g \in L^2(\mathcal{M}, m_g)$ be positive functions. Again, from the definition of the WOT topology, we have that $\lim_i \langle (P_i - P) f, g \rangle = 0$; that is, $\lim_i \int_{\mathcal{M}} (P_i - P) f g dm_g = 0$. Since P_i, P are positive operators, this implies that $(P_i - P) fg \to 0$ in L^1. Therefore, there exists a subsequence such that $\lim_j ((P_i)_j - P) fg = 0$

a.e. (Folland, 2013). Since g is positive, this implies that $\lim_j((P_i)_j - P)f = 0$ a.e. Finally, for all i and fixed $x \in \mathcal{M}$, $(P_i f)x$ is non-zero over the set $B_g(x, r)$, and since f is positive, we conclude that the limit $Pf(x)$ must be zero everywhere outside the ball $B_g(x, r)$. Hence, $P \in \mathscr{P}$. □

Remark 5.5.5. *We note that $\mathscr{P}$ is a set of operators of the form* (3.2)*, whereas we are interested in the adjoints of these operators, which are of the form* (3.1)*. Therefore, we require that $\mathscr{P}^* := \{T^*, T \in \mathscr{P}\}$ be closed in the WOT topology, which follows from the fact that the map $P \to P^*$ is WOT continuous.*

Since the optimization problem is a minimization problem, it is sufficient for us to prove that the objective function is only lower-semicontinuous, rather than continuous. We prove this in the following proposition.

Proposition 5.5.6. *The map $P \mapsto \|M_{f_d}^{-1} P M_{f_d} \circ \mathrm{Proj}_{\mathbf{1}^\perp}\|$ is lower-semicontinuous on $\mathscr{P}$ and convex.*

Proof. It is clear that the map $P \to M_{f_d}^{-1} P M_{f_d} \circ \mathrm{Proj}_{\mathbf{1}^\perp}$ is continuous. Further, by Problem 109 in Halmos (2012), the operator norm is weak* lower-semicontinuous on the dual space $L^2(\mathcal{M}, m_g)$. We observe that the objective function is a composition of a lower-semicontinuous function and a continuous function, and is therefore a lower-semicontinuous function on $\mathscr{P}$. Convexity follows from the fact that the objective function is defined as a norm function. □

We can now state the following result, which proves the existence of an optimal solution.

Theorem 5.5.7. *The optimization problem* (5.11)-(5.16) *has an optimal solution.*

Proof. We know that the unit ball in $\mathcal{B}(\mathcal{H})$ is compact in WOT (Theorem 5.1.3, Kadison and Ringrose (1997)). In addition, Theorem 5.3.1 guarantees the existence of an

operator P that satisfies the constraints of the optimization problem, and is therefore an element of $\mathscr{P}$ and is bounded. The optimization algorithm will hence generate bounded operators with norms that do not exceed the norm of P. Accordingly, these operators will form a bounded subset of $\mathscr{P}$, which we will refer to as $\mathscr{P}'$. Since $\mathscr{P}$ is closed in WOT by Proposition 5.5.4, we conclude that $\mathscr{P}'$ is closed and bounded in WOT, and is therefore compact in WOT. By Theorem 5.5.3, we can identify a positive kernel K with each pseudo-integral operator P. Therefore, the set of kernels K that satisfy the constraints of the optimization problem is compact in the topology induced by this bijective identification. We denote this set by $\mathcal{K}'$.

Finally, in view of Proposition 5.5.6, the infimum of the map $K \mapsto \|M_{f_d}^{-1} P(K) M_{f_d} \circ \mathrm{Proj}_{\mathbf{1}^1} \|$ over the set $\mathcal{K}'$ can indeed be achieved; that is, there exists an optimal (minimal) solution. □

5.6 Special Case

In the case where $\mathcal{M}$ is a Lie group and f_d corresponds to $\mathbf{1}$, the uniform distribution, we know that there exists a regular transition kernel K, defined in (5.5). Then, the optimization problem (5.11)-(5.16) can be posed in terms of the kernel function k. In this case, i.e. when $\mathcal{M}$ is a Lie group and $f_d = \mathbf{1}$, we can expect an exact minimization of the second largest eigenvalue modulus of $\hat{P}$. This is because the inequality in (3.36) is in fact an equality when the operator $\hat{P}$ is self-adjoint, and we have shown that there exists at least one solution to the optimization problem, the operator S constructed in Section 5.3, that is self-adjoint. In this subsection, we outline a proof of the existence of the optimal solution in this case.

Let c be a positive constant and r be the radius in constraint (5.16). In contrast to our approach for the general case, in which we identified the kernels K (the decision variables) with a set of operators P, here we can directly define a set of L^∞ kernel

functions, denoted by $\mathcal{K}$, that satisfy the constraints of the optimization problem:

$$\mathcal{K} = \Big\{ k \in L^\infty(\mathcal{M} \times \mathcal{M}, m_g \times m_g) : 0 \leq \|k\|_\infty \leq c, \ \int_{\mathcal{M}} k(x,z) dm_g(z) = \mathbf{1}, \\ \int_{\mathcal{M}} k(z,y) dm_g(z) = \mathbf{1}, k(x,y) = 0 \text{ if } d_g(x,y) > r \ \ \forall x, y \in \mathcal{M} \Big\}. \quad (5.19)$$

Now, $\mathcal{K}$ is the set of decision variables.

Proposition 5.6.1. *$\mathcal{K}$ is compact in the weak* topology and is convex.*

Proof. Since $\mathcal{M}$ has finite measure, we have that $\mathcal{K} \subseteq L^\infty(\mathcal{M} \times \mathcal{M}, m \times m) \subseteq L^2(\mathcal{M} \times \mathcal{M}, m \times m)$. First, we will show that $\mathcal{K}$ is closed and bounded in the topology induced by the $\|\cdot\|_2$ norm. Let $(k_i)_i \in \mathcal{K}$ be such that $k_i \xrightarrow{\|\cdot\|_2} \bar{k}$. That is, $\int |k_i - \bar{k}|^2 \to 0$. We will show that $\bar{k} \in \mathcal{K}$. It is straightforward to show that the limit $\bar{k}$ must satisfy $0 \leq \|\bar{k}\|_\infty \leq c$. Next, we observe that $\left|\int (k_i - \bar{k})\right| \leq \int |k_i - \bar{k}| \leq \|k_i - k\|_2$ (by Holder's inequality). Therefore, we have that $\left|\int (k_i - \bar{k})\right| \to \mathbf{1}$, which implies that $\int \bar{k} = \mathbf{1}$. We now consider the last constraint in (5.19). Note that $k_i(x,y) = 0$ for all $x, y \in \mathcal{M}$ such that $d_g(x,y) > r$. Since $k_i \to \bar{k}$ in $\|\cdot\|_2$, there exists a subsequence $(k_{ij})_j$ that converges to $\bar{k}$ m_g-a.e. It then follows that $\bar{k}(x,y) = 0$. Hence, $\mathcal{K}$ is closed in the $\|\cdot\|_2$ norm. The boundedness of $\mathcal{K}$ follows trivially from the condition $\|k(\cdot,\cdot)\|_\infty \leq c$.

The convexity of $\mathcal{K}$ follows from brief algebraic computations which show that the constraints in (5.19) are convex. An application of *Mazur's theorem* (Proposition 12.2, DiBenedetto (2002)) proves that $\mathcal{K}$ is closed in the weak topology. On the real-valued function space L^2, this implies that $\mathcal{K}$ is in fact closed in the weak* topology. Finally, we obtain our result by applying *Alaoglu's theorem* (Corollary 3.15, Clarke (2013)), which states that a set that is weak* closed and bounded is also weak* compact. □

In the proposition below, we will prove lower-semicontinuity of the map considered in Proposition 5.5.6.

Proposition 5.6.2. *The map* $k \mapsto \|M_{f_d}^{-1} P(k) M_{f_d} \circ \mathrm{Proj}_{\mathbf{1}^1}\|$ *is weakly lower-semicontinuous on* $\mathcal{K}$ *and convex.*

Proof. Let $(k_i)_i \in \mathcal{K}$ be such that $(k_i)_i \to k \in \mathcal{K}$ in the weak* topology. Let $P_i, P \in \mathbb{B}(L^2(\mathcal{M}, m_g))$ be the corresponding operators defined by k_i and k, respectively. Consider the WOT topology on $\mathbb{B}(L^2(\mathcal{M}, m_g))$. We will show that $P_i \to P$ in WOT. Convergence in WOT entails showing that $\langle P_i f, g\rangle \overset{i\to\infty}{\longrightarrow} \langle Pf, g\rangle$ for all $f, g \in L^2(\mathcal{M}, m_g)$, which implies that:

$$\int_{\mathcal{M}}\int_{\mathcal{M}} k_i(x,y) f(x) g(y) dm_g(x) dm_g(y) \overset{i\to\infty}{\longrightarrow} \int_{\mathcal{M}}\int_{\mathcal{M}} k(x,y) f(x) g(y) dm_g(x) dm_g(y).$$

The *tensor product* for functions $f, h \in L^2(\mathcal{M}, m_g)$ is denoted by $f \otimes h : \mathcal{M} \times \mathcal{M} \to \mathbb{R}$, which is defined as $f \otimes h(x,y) := f(x)h(y)$. By exercise 1.4.25 of Tao (2010), $f \otimes h \in L^2(\mathcal{M} \times \mathcal{M}, m_g \times m_g)$. Therefore, the equation above can also be written as

$$\int_{\mathcal{M}\times\mathcal{M}} k_i(z)(f \otimes g)(z) dz \overset{i\to\infty}{\longrightarrow} \int_{\mathcal{M}\times\mathcal{M}} k(z)(f \otimes g)(z) dz.$$

This is exactly the definition of convergence of $(k_i)_i$ to k in weak*. Therefore, we have that $(P_i)_i \to P$ in WOT. The rest of the proof is similar to the proof of Proposition 5.5.6. □

In conclusion, the objective function is weak* lower-semicontinuous on the compact set $\mathcal{K}$, and therefore the infimum of the objective function can indeed be achieved.

5.7 Numerical Optimization

In this section, we present a numerical approach to solving the optimization problem (5.11)-(5.16). As stated previously, we assume that the state space $\mathcal{M}$ is a compact smooth connected manifold, without boundary, of dimension d. The subset $\mathcal{M}$ is partitioned into $N \in \bar{\mathbb{Z}}_+$ sets, $\widetilde{\mathcal{M}} = \{\mathscr{M}_1, \ldots, \mathscr{M}_N\}$, where $\mathcal{M} = \cup_{i=1}^N \mathscr{M}_i$ and the sets $\mathscr{M}_i$ have intersections of zero Riemannian volume. We define an equivalent

of the transition kernel K for this discretized state space. Let $\tilde{k}_{ij}$ be the probability of jumping to $\mathscr{M}_j$, given that the system state is in $\mathscr{M}_i$. This probability is given by,

$$\tilde{k}_{ij} = \int_{\mathscr{M}_i} K(x, \mathscr{M}_j) dx.$$

We define $\mathbf{K}$ as the matrix $[\tilde{k}_{ij}]_{i,j\in\mathcal{I}}$, where $\mathcal{I} = \{1, \ldots, N\}$. We use $\mathbf{K}$ to construct an approximating Markov chain on the finite state space $\mathcal{I}$. Let $\mathcal{G} = (\mathcal{I}, \mathcal{E})$ be a graph defined on $\mathcal{I}$ with edge set $\mathcal{E} = \{(i,j) : i,j \in \mathcal{I},\ \tilde{k}_{ij} > 0\}$, which specifies the transitions of the Markov chain. An edge (i,j) is in the edge set $\mathcal{E}$ if the distance between the centers of $\mathscr{M}_i$ and $\mathscr{M}_j$ does not exceed r, as per the constraint (5.16).

Let $\mu \in \mathcal{P}(\widetilde{\mathcal{M}})$ and $\mathbf{P} \in \mathbb{M}(\mathbb{R}^N)$, the space of real-valued matrices. Then $\mathbf{P}$ defined below is equivalent to the operator defined in (5.3):

$$(\mathbf{P}\mu)(j) = \sum_{i\in\mathcal{I}} \tilde{k}_{ij}\mu(i), \quad j \in \mathcal{I}. \tag{5.20}$$

Let $\mu_d \in \mathcal{P}(\widetilde{\mathcal{M}})$ be a desired distribution that is positive on $\widetilde{\mathcal{M}}$, and define a diagonal matrix $\mathbf{M}_d = \mathrm{diag}(\mu_d)$.

We can now formulate a finite-dimensional quadratic program that is equivalent to optimization problem (5.11)-(5.16) as follows:

$$\min_{\mathbf{K}} \left\| \widehat{\mathbf{P}} - \frac{\mathbf{1}\mathbf{1}^T}{N} \right\| \tag{5.21}$$

subject to

$$\widehat{\mathbf{P}} = \mathbf{M}_d^{-1}\mathbf{P}\mathbf{M}_d, \tag{5.22}$$

$$(\mathbf{P}\mu)(j) = \sum_{i\in\mathcal{I}} \tilde{k}_{ij}\mu(i) \quad \forall j \in \mathcal{I},\ \forall \mu \in \mathcal{P}(\widetilde{\mathcal{M}}), \tag{5.23}$$

$$\tilde{k}_{ij} \geq 0 \quad \forall i,j \in \mathcal{I}, \tag{5.24}$$

$$\mathbf{K}\mathbf{1} = \mathbf{1}, \tag{5.25}$$

$$\mu_d\mathbf{K} = \mu_d, \tag{5.26}$$

$$\tilde{k}_{ij} = 0 \quad \forall (i,j) \notin \mathcal{E}. \tag{5.27}$$

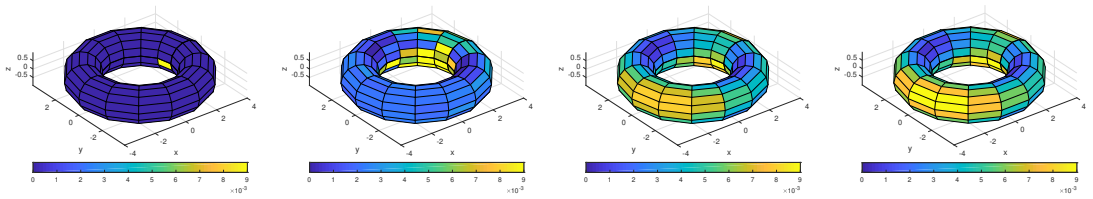

(a) $n = 0$ (b) $n = 10$ (c) $n = 35$ (d) Target measure μ_d

Figure 5.1: Simulation of the system (5.28) at different times n.

Note that $\mathbf{1}$ in (5.21) is a vector in $\mathbb{R}^N$. The constraint in (5.22) ensures that the matrix $\widehat{\mathbf{P}}$ is bistochastic. The constraint (5.27) is equivalent to (5.16). Observe that this optimization problem is convex and is similar to the optimization problem solved in Boyd *et al.* (2004).

5.8 Simulation Results

In this section, we apply our numerical optimization procedure to an example system evolving on a torus in $\mathbb{R}^3$. The numerical optimization problem is solved with the MATLAB package CVX (Grant and Boyd, 2014). Recall that, as described in Section 3.5, we are solving a relaxation of the optimization problem in Problem 5.2.2.

The state space $\mathcal{M}$ in our example is the torus in $\mathbb{R}^2$ embedded in $\mathbb{R}^3$. This state space is discretized into a 15×15 grid ($N = 225$). We define the initial and target measures on the discretized space as shown in Figs. 5.1a and 5.1d, respectively. The value of r in constraint (5.16) is chosen to be 0.2, a number greater than the partition size 1/15. We solve the optimization problem (5.21)-(5.27) to obtain a transition probability matrix $\mathbf{K}$. Defining $\mathbf{P}$ from the resulting $\mathbf{K}$, we simulate the following version of system (5.1):

$$\mu_{n+1} = \mathbf{P}\mu_n. \tag{5.28}$$

Figures 5.1a-5.1c show snapshots of the simulation of system (5.28) at three different times. It is evident from the time evolution of the snapshots that the simulated

measure μ_n converges asymptotically to the target measure μ_d. To quantify this degree of convergence, we computed the 2-norm error metric $\gamma_n = ||\mu_n - \mu_d||_2$ at the times of the snapshots. The corresponding values of γ_n for $n = 0$, 10, and 35 are 0.7611, 0.0824, and 6.7382×10^{-4}, which clearly shows that μ_n is tending toward μ_d.

The results presented within this chapter are part of Biswal *et al.* (2019a).

Chapter 6

CONTINUOUS-TIME MARKOV CHAIN MODELS ON CONTINUOUS STATE SPACES

In Chapters 3 and 5, we constructed operators that stabilize a given target distribution for a discrete-time Markov process and also optimize the convergence rate to this distribution. However, except for a special case of the state space (see Section 5.6), the optimization problem posed in both chapters was not exact, meaning that we were able to minimize only an upper bound on the eigenvalues of the operator rather than the modulus of the second largest eigenvalue.

As a next step, in this chapter we consider swarms of agents whose dynamics evolve over a continuous state space in continuous time; in particular, the dynamics of each agent can be modeled as a stochastic differential equation (SDE). The Kolmogorov forward equation in this case is a partial differential equation (PDE) that is commonly known as the *Fokker-Planck equation*. We construct a partial differential operator that stabilizes target swarm distributions that are bounded and positive almost everywhere on the domain. In particular, the operator that we construct has a structure similar to the divergence form operator, which is known to be self-adjoint. The advantage in this case is that we can invoke the min-max principle to characterize the modulus of the second largest eigenvalue of the operator, which characterizes the asymptotic convergence rate of the swarm to the target distribution. Hence, unlike in our previous works, the optimization problem is exact. However, not all divergence form operators are Fokker-Planck equations; that is, they do not all give rise to SDEs. Instead of working with a restricted class of divergence form operators that do give rise to SDEs, we will work with general divergence form operators, and consider the SDE

description of agent dynamics only formally. See Stroock (1988) for a discussion on divergence type operators that correspond to diffusion semigroups.

We begin by briefly reviewing literature on the topic of using PDE models to control the distribution of a swarm of agents. In the field of multi-robot systems, a number of works utilize PDEs to model and control the collective behaviors of robotic swarms. A PDE model with a constant velocity field to simulate a swarm of small robots performing an inspection task is investigated in in Prorok *et al.* (2011). In Kingston and Egerstedt (2011), the authors design swarm control strategies that mimic fluid flow behavior by constructing state-feedback laws that are piecewise constant with respect to space. PDEs with feedback laws that are functions of population densities are used in Hamann and Wörn (2008) to model collective migration and collective perception in swarms. To simulate the phenomenon of emergent taxis, the authors construct *mean-field feedback laws* in the sense that the velocity fields and diffusion coefficients are functions of the population densities, as in biological models of chemotaxis. The work Elamvazhuthi *et al.* (2018b) applies optimal control of PDEs and PDE-constrained optimization to design time-dependent robot controllers for problems of stochastic spatial coverage and feature mapping by robotic swarms.

With regard to characterizing the spectral gap of Fokker-Planck equations or in general, diffusion equations, the *Bakry-Emery method* allows one to establish convex Sobolev inequalities and to compute exponential decay rates toward equilibrium for solutions of diffusion equations (Bakry *et al.*, 2013). In Arnold *et al.* (2001), the authors quantify convergence rates of Fokker-Planck equations using convex Sobolev inequalities. In Jafarizadeh (2018), the author poses the spectral optimization problem for a Fokker-Planck equation in $\mathbb{R}^1$ using the min-max principle. In contrast to our work, the domain is restricted to $\mathbb{R}^1$ and the constraint is posed in terms of minimizing the variance of the corresponding Markov process.

6.1 Notation

We denote the state space by $\Omega \subset \mathbb{R}^d$, an open, bounded connected set. The boundary of Ω is denoted by $\partial\Omega$, which is assumed to be *Lipschitz continuous* (Grisvard, 2011).

The space $L^2(\Omega)$ is a Hilbert space equipped with the standard inner product $\langle \cdot, \cdot \rangle_2 : L^2(\Omega) \times L^2(\Omega) \to \mathbb{R}$ given by $\langle f, g \rangle_2 = \int_\Omega f(x)g(x)dx$, for all $f, g \in L^2(\Omega)$. The symbol $\| \cdot \|_2$ will be reserved for the $L^2(\cdot)$ norm. For a given real-valued function $h \in L^\infty(\Omega)$, (weighted)

$$L^2_h(\Omega) = \left\{ f : \int_\Omega |f(x)|^2 |h(x)| dx < \infty \right\}.$$

In this case, $L^2_h(\Omega)$ is a Hilbert space with respect to the weighted inner product $\langle \cdot, \cdot \rangle_h : L^2_h(\Omega) \times L^2_h(\Omega) \to \mathbb{R}$ given by

$$\langle f, g \rangle_h = \int_\Omega f(x)g(x)h(x)dx.$$

We let $\| \cdot \|_F$ stand for the weighted L^2_F norm. Let f_{x_i} denote the first-order weak partial derivative of the function f with respect to coordinate x_i. We define the Sobolev space $H^1(\Omega) \subset L^2(\Omega)$ functions whose partial derivatives, in the weak sense, are also in $L^2(\Omega)$. This is a Hilbert space with the norm:

$$\|f\|_{H^1} = \|f\|_2^2 + \left(\sum_{i=1}^d \|f_{x_i}\|_2^2 \right)^{1/2} \quad \text{for } f \in H^1(\Omega).$$

Correspondingly, for $h \in L^\infty(\Omega)$ we define the space

$$H^1_h(\Omega) = \left\{ f \in L^2_h(\Omega) : (fh)_{x_i} \in L^2(\Omega) \text{ for } 1 \le i \le d \right\},$$

equipped with the norm

$$\|f\|_{H^1_a} = \left(\|f\|_a^2 + \sum_{i=1}^d \|(af)_{x_i}\|_2^2 \right)^{1/2}.$$

Let $\mathcal{X}$ be a Hilbert space. Let A be a closed linear operator that is densely defined on a subset $\mathcal{D}(A) \subset \mathcal{X}$, the domain of the operator. If A is a bounded operator, then $\|A\|_{op}$ will denote the operator norm induced by the norm defined on $\mathcal{X}$. The spectrum $\sigma(A)$ of A is the non-void compact set of complex numbers λ for which $A - \lambda I$ does not have a continuous inverse on $\mathcal{X}$. The operator A is said to be *positive*, denoted by $A > 0$, if for $x \in \mathcal{X}$, $x \geq 0$ implies that $Ax \geq 0$.

6.2 Problem Formulation

We begin by setting up the problem that we address in this section. Let $F \in L^\infty(\Omega)$ such that $F(x) > 0$ *a.e.* be the target steady-state probability density function for a swarm of robots. Then F must satisfy the condition $\int_\Omega F(x)dx = 1$. Define $\Omega_{T_f} = \Omega \times (0, T_f)$ for some fixed final time T_f. Let $p : \Omega_{T_f} \to \mathbb{R}^n$ denote a probability density function. The forward Kolmogorov equation, also called the Fokker-Planck equation, gives the evolution of probability densities on the state space Ω. In continuous time and continuous space, this equation is a partial differential equation (PDE) of the form:

$$\frac{\partial}{\partial t}p(x,t) = \frac{1}{2}\sum_{i,j=1}^{d}\frac{\partial^2}{\partial x_i \partial x_j}[D_{ij}(x,t)p(x,t)] - \sum_{i=1}^{d}\frac{\partial}{\partial x_i}[a_i(x,t)p(x,t)]. \tag{6.1}$$

Here, the coefficients D_{ij} and a_{ij} represent diffusion and advection parameters, respectively. In this paper, however, we will not be working with such a general formulation. For reasons that will be made clear later, we require the partial differential operator associated with the PDE to be self-adjoint. This is not true for the PDE (6.1). We will therefore introduce an operator, formally, which is self-adjoint.

Let $a_{ij} : \Omega \to \mathbb{R}^d$ for $i, j = 1, \ldots, d$, with $a_{ij} = a_{ji}$, be in $L^\infty(\Omega)$. Further, we assume that the coefficients satisfy the uniform *ellipticity* condition; that is, there exists a constant α such that for every vector $\xi \in \mathbb{R}^d$ and every $x \in \Omega$,

$\sum_{i,j=1}^{d} a_{ij}(x)\xi_i\xi_j \geq \alpha|\xi|^2$. Consider the following *unbounded* operator,

$$\mathcal{L}_F u = \sum_{i,j=1}^{d} \frac{\partial}{\partial x_i}\left(a_{ij}(x)\frac{\partial(u/F)}{\partial x_j}\right). \tag{6.2}$$

We note that this operator has the advantage of being self-adjoint. Moreover, the operator is almost in the standard divergence form (Evans, 2010); however, the inclusion of F makes it non-standard. It is clear that the inclusion of F ensures that $\mathcal{L}_F F = 0$; that is, the PDE generated by this operator has F as an equilibrium point. Without the inclusion of F, there are only a few special cases in which Eq. (6.2) can be rewritten as Eq. (6.1), and vice versa.

Remark 6.2.1. *We note that the operator* (6.2) *is not defined rigorously. This is because the L^∞ condition on the coefficients a_{ij} makes it impossible to define the operator on $H^1(\Omega)$ or $H^2(\Omega)$ (defined similarly to $H^1(\Omega)$, but comprised of functions that are twice weakly differentiable and are in L^2); either space is not necessarily preserved under the multiplication of an L^∞ function and an H^1 or H^2 function. Therefore, in order to proceed, we will instead define a weak formulation of the operator* (6.2) *via forms.*

Note that according to our notation, as per standard definitions, H^1_F and L^2_F norms entail a multiplication by F; that is, $\|f\|_F = \int_\Omega |f|^2|F|$. However, in this paper, the norm entails a division by F; that is, $\|f\|_F = \int_\Omega |f|^2|1/F|$.

We define a *bilinear form* $B_F[u,v] : H^1_F(\Omega) \times H^1_F(\Omega) \to \mathbb{R}$ as follows:

$$B_F[u,v] = \int_\Omega \sum_{i,j=1}^{d} a_{ij}(x)\frac{\partial(u/F)}{\partial x_j}\frac{\partial(v/F)}{\partial x_i}dx \tag{6.3}$$

The space $H^1_F(\Omega)$ is called the domain of B_F, $\mathcal{D}(B_F)$. We associate with the form B an operator $\hat{\mathcal{L}}_F : \mathcal{D}(\hat{\mathcal{L}}_F) \subset L^2_F(\Omega) \to L^2_F(\Omega)$, defined as $\hat{\mathcal{L}}_F u = f$ if $B_F[u,v] = \langle f, v\rangle_F$ for all $v \in H^1_F(\Omega)$ and $u \in \mathcal{D}(\hat{\mathcal{L}}_F) = \{g \in H^1_F(\Omega) : \exists h \in L^2_F(\Omega)\ s.t.\ B_F[g,\varphi] =$

$\langle h, \varphi \rangle \ \forall \ \varphi \in H_F^1(\Omega)\}$. The operator $\hat{\mathcal{L}}_F$ so defined is a weak formulation of the operator (6.2). Defining $\hat{\mathcal{L}}_F$ via the bilinear form B_F is similar in spirit to the formulation of weak solutions to elliptic equations. A detailed treatment of the interplay between forms and operators is provided in Schmüdgen (2012). In the specific case where the coefficients a_{ij} and the function F are uniformly Lipschitz functions, then $\hat{\mathcal{L}}_F$ coincides with the operator (6.2), with $H^2(\Omega)$ as its domain (Grisvard, 2011).

Although the bilinear form (6.3) simplifies the analysis, we lose the guarantee that the generated PDE corresponds to a stochastic differential equation. Only in the case where the coefficients a_{ij} are uniformly Lipschitz continuous does the operator (6.2) give rise to a forward equation (Friedman, 2006). However, one can make sense of the stochastic differential equations that divergence form operators give rise to in a non-classical way; see Lejay (2006) for this description.

We consider the following PDE generated by the operator $\mathcal{L}_F$ in (6.2). Note that this is only a formal statement because of the explanation in the previous paragraphs.

$$\frac{\partial p}{\partial t} = -\mathcal{L}_F p \text{ on } \Omega_{T_f} \tag{6.4}$$

$$\sum_{i,j=1}^{d} a_{ij} \frac{\partial (p/F)}{\partial x_j} n_i = 0 \text{ on } \partial\Omega \times (0, T_f) \tag{6.5}$$

$$p(x, 0) = p_0(x) \text{ on } \Omega. \tag{6.6}$$

Equation (6.5) represents the zero flux boundary condition, also called the *Neumann* boundary condition; n_i is the i^{th} unit normal vector to Ω, pointing outward.

We now state the problem that we solve in this paper. To address this problem, in the next section we will prove that 0 is the unique largest eigenvalue of the operator $-\hat{\mathcal{L}}_F$, with all other eigenvalues located in the left half-plane. Therefore, the convergence rate of the PDE (6.4) to its equilibrium is characterized by the L_F^2 spectral gap. First, however, we will need to prove the existence of this spectral gap for $\hat{\mathcal{L}}_F$.

Problem 6.2.2. *Given F, determine whether there exist time-independent, spatially-dependent parameters $a_{ij} : \Omega \to \mathbb{R}^n$, for $i, j \in 1 \dots d$, such that F is an exponentially stable equilibrium point for the PDE* (6.4)*. Toward this end, determine whether the following optimization problem admits a solution.*

$$\min_{a_{ij}} |\lambda_2(\hat{\mathcal{L}}_F)|$$

Due to the definition of the operator $\hat{\mathcal{L}}$, we need not impose the condition $\hat{\mathcal{L}}_F F = 0$ as a constraint. In Section 6.4, we will characterize the eigenvalues of $\hat{\mathcal{L}}_F$ via the min-max principle, which is only true for a self-adjoint operator. We chose to work with divergence form operators in order to be able to characterize their eigenvalues via this principle.

6.3 Analytical Properties of $\hat{\mathcal{L}}_F$

We begin by proving a few properties of the operator $\hat{\mathcal{L}}_F$; proofs for general functions F and general domains Ω are given in Elamvazhuthi *et al.* (2018c). Therefore, only those parts of the proofs that are specific to our case are detailed below.

Proposition 6.3.1. *The operator $\hat{\mathcal{L}}_F$ is closed, densely defined, self-adjoint, and positive. Moreover, the operator $\hat{\mathcal{L}}_F$ has a purely discrete spectrum.*

Proof. First we prove that the bilinear form (6.3) is closed; that is, the space $\mathcal{D}(B_F) = H^1_F(\Omega)$ equipped with the norm $\|u\|_B = (\|u\|^2_F + B_F[u,u])^{1/2}$ for each $u \in \mathcal{D}(B_F)$ must be complete (Schmüdgen, 2012). To see this, we note that by the uniform ellipticity condition on the coefficients a_{ij}, we have that

$$B_F[u,u] = \int_\Omega \left[\tfrac{\partial(u/F)}{\partial x}\right] A \left[\tfrac{\partial(u/F)}{\partial x}\right]^T \geq \int_\Omega \alpha \left|\frac{\partial(u/F)}{\partial x}\right|^2,$$

where $A = [a_{ij}]$. We also have that

$$\|A\|_\infty \int_\Omega \left|\frac{\partial(u/F)}{\partial x}\right|^2 \geq B_F[u,u].$$

Therefore, the norm $\|\cdot\|_B$ is equivalent to $\|u\|_{H^1_F}$. It has been shown in Elamvazhuthi *et al.* (2018c) that $H^1_F(\Omega)$ is complete. Therefore, B_F is closed.

Next, from Elamvazhuthi *et al.* (2018c), we can show that B_F is densely defined; that is, $\mathcal{D}(B_F)$ must be dense in $L^2_F(\Omega)$, which is true in this case. Furthermore, B_F is *symmetric*, that is, $B_F[u,v] = \overline{B_F[v,u]}$ for each $u, v \in \mathcal{D}(B_F)$, and B_F is *semibounded*, that is, $B_F[u,u] \geq m\|u\|^2_F$ for some $m \in \mathbb{R}$, for each $u \in \mathcal{D}(B_F)$. The latter property is true for $m = 0$. By Theorem 10.7 of (Schmüdgen, 2012), these properties imply that $\hat{\mathcal{L}}_F$ is self-adjoint, which further implies that $\hat{\mathcal{L}}_F$ is also closed and densely defined.

Finally, we have that $H^1_F(\Omega) = \mathcal{D}(B_F)$ equipped with the norm $\|\cdot\|_B$ is *compactly embedded* in $L^2(\Omega)$. By Proposition 10.6 of Schmüdgen (2012), this condition is sufficient for the operator $\hat{\mathcal{L}}_F$ to have a discrete spectrum. □

Proposition 6.3.2. *The spectrum of the operator $\hat{\mathcal{L}}_F$ satisfies $\sigma(\hat{\mathcal{L}}_F) \in (\infty, 0]$. Furthermore, 0 is a unique eigenvalue of $\hat{\mathcal{L}}_F$.*

Proof. From the definition of the bilinear form, we observe that the operator $-\hat{\mathcal{L}}_F$ must be negative semidefinite. Hence, $\sigma(-\hat{\mathcal{L}}_F) \in (\infty, 0]$. Consider the bilinear form (6.3) with $F = \mathbf{1}$. In this case, it is clear that $\hat{\mathcal{L}}_{\mathbf{1}}\mathbf{1} = 0$; that is, $\mathbf{1}$ is an eigenvector corresponding to the eigenvalue 0. To prove the uniqueness of 0, we use the *Poincaré inequality* (Evans, 2010): there exists a constant C such that $\int_\Omega |u(x) - u_\Omega| dx \leq C \int_\Omega |\nabla u(x)|^2$, where $u_\Omega = \frac{1}{m(\Omega)} \int_\Omega u(x) dx$, and $m(\Omega)$ stands for the Lebesgue measure of the set Ω. Using the uniform ellipticity condition and assuming that $\alpha \geq C$, we have that

$$\begin{aligned} \int_\Omega |u(x) - u_\Omega| dx &\leq \alpha \int_\Omega |\nabla u(x)|^2 \\ &\leq \alpha B_F[u,u] = \alpha \int_\Omega \begin{bmatrix} \frac{\partial u}{\partial x} \end{bmatrix} A(x) \begin{bmatrix} \frac{\partial u}{\partial x} \end{bmatrix}^T. \end{aligned} \tag{6.7}$$

If u is an eigenvector other than $\mathbf{1}$, then the right-hand side of the inequality above

evaluates to 0 while the left-hand side is positive, leading to a contradiction. Therefore, the eigenvalue 0 must be unique. For general F we define the multiplication map $M_F : L^2(\Omega) \to L^2_F(\Omega)$ that takes a function $u \in L^2_F(\Omega)$ to $u/F \in L^2(\Omega)$. Note that $\hat{\mathcal{L}}_F = \hat{\mathcal{L}}_{\mathbf{1}} M_F$. From this observation we can infer that $\mathbf{1}$ is an eigenvector of $\hat{\mathcal{L}}_{\mathbf{1}}$ for the eigenvalue 0 if and only if F is an eigenvector of $\hat{\mathcal{L}}_F$.

In the case where $\alpha \leq C$, we can replace α by C/α in equation (6.7), and the analysis remains the same. □

Due to the lack of smoothness of the functions a_{ij} and F, the PDE (6.4) might not have solutions that are continuously differentiable in the classical sense, or even solutions that are weakly twice differentiable. Using the above properties, one can show that the PDE (6.4) has a *mild solution* (Engel and Nagel, 2000), which can be represented as a *semigroup* of linear operators. This follows from the *Lumer-Phillips theorem* by noting that the operator $\hat{\mathcal{L}}_F$ is self-adjoint and dissipative. See Elamvazhuthi *et al.* (2018c) for details. Since $\mathcal{D}(\hat{\mathcal{L}}_F)$ is a subset of $H^1_F(\Omega)$, it follows that if the initial condition is in $\mathcal{D}(\hat{\mathcal{L}}_F)$, then the mild solution lies in $H^1_F(\Omega)$ for all time $t \geq 0$. One can also show that the semigroup is analytic, and hence has regularizing properties. This implies that even if the initial condition is known to be only in $L^2(\Omega)$, the solution of the PDE (6.4) lies in $H^1_F(\Omega)$ for all $t > 0$.

6.4 Formulation of the Optimization Problem

Recall the conditions on F, the desired density function: it is in $L^\infty(\Omega)$ and is strictly positive almost everywhere. We have established that F is a unique eigenvector of the operator $-\hat{\mathcal{L}}_F$ corresponding to the largest eigenvalue 0. Furthermore, we have showed the existence of a spectral gap of $\hat{\mathcal{L}}_F$. In this section, we solve Problem 6.2.2.

The *Courant-Fisher min-max principle* provides a way to formulate the objec-

tive function of the optimization problem in Problem 6.2.2. Let $(T, \mathcal{D}(T))$ be a lower-semibounded, self-adjoint operator on a Hilbert space H with a purely discrete spectrum. Let $(\lambda_n(T))_n$ be the increasing sequence of eigenvalues of T, counted with multiplicities. The min-max principle gives a variational characterization for the eigenvalues that are below the bottom of the essential spectrum (Schmüdgen, 2012). Let E_k be a linear subspace of H of dimension k. Then the eigenvalues λ_k can be defined as:

$$\lambda_k(T) = \max_{E_k} \min_{\substack{v \in \mathcal{D}(T), \|v\|=1, \\ v \in E_k^{\perp}}} \langle Tv, v \rangle.$$

The inner product in this definition is called the Rayleigh quotient.

In our case, the operator $\hat{\mathcal{L}}_F$ satisfies the properties listed above, and therefore we can characterize the second largest eigenvalue of $\hat{\mathcal{L}}_F$ by restricting $\hat{\mathcal{L}}_F$ to the subspace obtained after removing the eigenspace F corresponding to the eigenvalue 0. The objective function is hence formulated as,

$$\lambda_2(-\hat{\mathcal{L}}_F) = \lambda_1(-\hat{\mathcal{L}}_F \circ \mathrm{Proj}_{F^{\perp}}) \; = \min_{v \in \mathcal{D}(\hat{\mathcal{L}}_F),\ \|v\|_F=1 \int_\Omega v=0} \langle -\hat{\mathcal{L}}_F v, v \rangle_F. \tag{6.8}$$

Here $\mathrm{Proj}(\cdot)$ is the projection operator onto a subspace. We note that removing the negative sign changes the minimization problem to a maximization problem. Further, the outer optimization, that is, the maximization can be omitted, since $E_0 \subset \mathcal{D}(\mathcal{L})$ is just $\{0\}$. The integral constraint in the equation above represents the projection onto $F^{\perp}$. To see this, let $v \in F^{\perp}$; then $\langle v, F \rangle_F = 0$, and this is exactly the integral $\int_\Omega v = 0$.

The constraints of the optimization problem are listed below.

$$a_{ij} \leq c, \text{ for some } c > 0 \tag{6.9}$$

$$a_{ij} = a_{ji} \tag{6.10}$$

$$\sum_{i,j=1}^{d} a_{ij}(x)\xi_i\xi_j \geq \alpha|\xi|^2, \ \forall \xi \in \mathbb{R}^d. \tag{6.11}$$

Constraint (6.9) ensures that the coefficients are bounded in the L^∞ norm. Constraint (6.11) ensures that the coefficients satisfy the uniform ellipticity condition. Equations (6.8)-(6.11) formulate the optimization problem that we will solve in this paper.

The set of decision variables is given by

$$\mathscr{A} = \{(a_{ij}) \in (L^\infty(\Omega))^{\frac{d(d+1)}{2}} : a_{ij} \leq c, \sum_{i,j=1}^{d} a_{ij}(x)\xi_i\xi_j \geq \alpha|\xi|^2, \forall \xi \in \mathbb{R}^d,\ i,j \in 1,\ldots,n\},$$

where $d(d+1)/2$ is the number of upper triangular elements in the coefficient matrix. In the next result we prove the continuity of eigenvalues of the operator $\hat{\mathcal{L}}_F$ with respect to the coefficients following the approach outlined in Henrot (2006), where the authors consider the special case when $F = \mathbf{1}$.

Theorem 6.4.1. *Let $\hat{\mathcal{L}}_F^n$ be the sequence of operators corresponding to a sequence of functions a_{ij}^n that is bounded in $L^\infty(\Omega)$ for each i and j, such that the functions converge almost everywhere to a function a_{ij} for each i and j. Let $\hat{\mathcal{L}}_F$ be the elliptic operator as defined in* (6.2) *by the functions a_{ij}. Then each eigenvalue of $\hat{\mathcal{L}}_F^n$ converges to the corresponding eigenvalue of $\hat{\mathcal{L}}_F$.*

Proof. From Henrot (2006)[Theorem 2.3.3] it is known that under the convergence conditions on the function a_{ij}^n, for each fixed $f \in L^2(\Omega)$, $(\hat{\mathcal{L}}_{\mathbf{1}}^n)^{-1}f$ converges to $(\hat{\mathcal{L}}_{\mathbf{1}})^{-1}f$ in norm. To prove the result in our modified case we let $M_F : L^2(\Omega) \to L_F^2(\Omega)$ be the multiplication map that takes $u \in L_F^2\Omega$ to $u/F \in L^2(\Omega)$. Since $\hat{\mathcal{L}}_F^n = \hat{\mathcal{L}}_{\mathbf{1}}^n M_F$, we can infer that for each fixed $f \in L^2(\Omega)$, $(\hat{\mathcal{L}}_F^n)^{-1}f$ converges to $(\hat{\mathcal{L}}_F)^{-1}f$ in norm. From this, we can conclude that the resolvents of the operators $\hat{\mathcal{L}}_F^n$ strongly converge to the resolvent of the operator $\hat{\mathcal{L}}_F$ (Theorem 2.3.2, Henrot (2006)). The operators $\hat{\mathcal{L}}_F^n$ and $\hat{\mathcal{L}}_F$ have a compact resolvent since $H_F^1(\Omega)$ is compactly embedded in $L_F^2(\Omega)$. Therefore, it follows from (Theorem 2.3.1, Henrot (2006)) that the eigenvalues of the operators $\hat{\mathcal{L}}_F^n$ converge to the respective eigenvalues of the operator $\hat{\mathcal{L}}_F$. □

6.5 Numerical Optimization

In this section, we numerically solve the optimization problem. Instead of discretrizing the operator $\hat{\mathcal{L}}_F$, we discretize the inner product in the objective function (6.8). Discretizing the inner product, rather than discretizing the operator $\hat{\mathcal{L}}_F$ and substituting it into the objective function, significantly reduces the computational complexity of solving the optimization problem. From the bilinear form (6.3), we have that for $u \in \mathcal{D}(\hat{\mathcal{L}}_F)$, $B_F[u,u] = \langle \hat{\mathcal{L}}_F u, u \rangle_F$. Therefore, the objective function can be recast as the following expression:

$$\langle -\hat{\mathcal{L}}_F u, u \rangle_F = -\int_\Omega \sum_{i,j=1}^{d} \left(a_{ij}(x) \frac{\partial (u/F)}{\partial x_j} \frac{\partial (u/F)}{\partial x_i} \right) dx \tag{6.12}$$

We demonstrate our numerical optimization procedure for a domain $\Omega \subset \mathbb{R}^2$. In this case, the above equation can be simplified to:

$$-\int_\Omega \begin{bmatrix} \frac{\partial v(x,y,t)}{\partial x} & \frac{\partial v(x,y,t)}{\partial y} \end{bmatrix} A(x,y) \begin{bmatrix} \frac{\partial v(x,y,t)}{\partial x} \\ \frac{\partial v(x,y,t)}{\partial y} \end{bmatrix} dxdy, \tag{6.13}$$

where $v = u/F$ and $A = [a_{ij}]$ is the coefficient matrix in $\mathbb{R}^{2\times 2}$.

In our example, we define $\Omega = [0,1] \times [0,1]$. We partition Ω into an $N \times N$ grid and define $h = 1/N$. Let I be the index set $\{1, \ldots, N\}$. Then $\Omega = \cup_{i,j \in I} \widetilde{\Omega}_{ij}$, where $\widetilde{\Omega}_{ij} = [x_{i-\frac{1}{2}}, x_{i+\frac{1}{2}}] \times [y_{j-\frac{1}{2}}, y_{j+\frac{1}{2}}]$ for $i,j \in I$. Let $w_{ij}(t) = v(x_i, y_j, t)$ be evaluated at the midpoint of each grid cell $\widetilde{\Omega}_{ij}$. Let $\widetilde{F}(i,j) = F(x_i, y_j)$. Note that we can remove the negative sign in the objective function (6.8) and pose the optimization problem as a maximization problem. The finite-dimensional optimization problem that is equivalent to (6.8)-(6.11) can be stated as:

$$\max_w \frac{1}{N^2} \sum_{i=1}^{N-1} \sum_{j=1}^{N-1} \begin{bmatrix} \frac{w_{i+1,j} - w_{i,j}}{h} \\ \frac{w_{i,j+1} - w_{i,j}}{h} \end{bmatrix}^T A(i,j) \begin{bmatrix} \frac{w_{i+1,j} - w_{i,j}}{h} \\ \frac{w_{i,j+1} - w_{i,j}}{h} \end{bmatrix} \tag{6.14}$$

subject to

$$\sum_{i,j}^{N} w_{ij}\widetilde{F}(i,j) = 0 \tag{6.15}$$

$$\|w\widetilde{F}\|_F^2 = \|w\sqrt{\widetilde{F}}\|_2^2 = 1 \tag{6.16}$$

$$\text{Equations } (6.9) - (6.11)$$

Constraint (6.15) ensures that the vector u (before discretization) is perpendicular to F. Constraint (6.16) ensures that the weighted 2-norm of u is 1. The objective function (6.14) is nonlinear. Further, it is difficult to prove that it is convex. Therefore, the nonlinear optimization solver KNITRO (Nocedal, 2006) was used to solve this problem. This solver implements both interior-point and active-set methods for solving nonlinear optimization problems. The problem was solved in AMPL (A Mathematical Programming Language) (Fourer *et al.*, 2003). We ran two test cases, described below.

In the first case, F was defined as the uniform distribution $\mathbf{1}$. Four different grid sizes $N \times N$ and two different values of c, $c = 1$ and $c = 10$, were tested. The eigenvalue $-\lambda_2$ was computed for each combination of grid size and c value, and the results are tabulated in Table 6.1. This table shows that as the discretization becomes finer, the eigenvalue converges. Note that for $c = 1$, the operator (6.2) corresponds to the Neumann Laplacian, accordingly, the computed second-largest eigenvalue, which is close to -12, is closer to $-\pi^2 \approx -9.87$ (the second largetst eigenvalue of the Neumann Laplacian).

In the second case, F was defined as the non-uniform distribution $F = (\sin(2\pi i/N))^2 + (\sin(2\pi j/N))^2 + \epsilon$, where ϵ was chosen to be 0.1 to ensure strict positivity of F over Ω. In this case, we also investigate how the eigenvalue changes in magnitude with respect to the $L^\infty(\Omega)$ bound c on the parameters a_{ij}. Table 6.2 shows the eigenvalue $-\lambda_2$ that was computed for each combination of nine different

Table 6.1: Eigenvalue $-\lambda_2$ for the case $F = \mathbf{1}$

$N \times N$	$c = 1$	$c = 10$
20×20	11.9	154.604
40×40	11.97	155.65
80×80	11.99	155.9
100×100	11.995	155.94

Table 6.2: Eigenvalue $-\lambda_2$ for the non-uniform F case

$N \times N$	$c = 1$	$c = 2$	$c = 5$
20×20	71.82	163.15	437.386
40×40	91.74	211.5	570.883
60×60	102.64	237.58	642.505
80×80	108.96	252.72	684.042
100×100	113.06	262.52	710.958
140×140	118.03	274.44	743.7
200×200	121.979	283.91	769.75
300×300	125.189	291.62	790.95

grid sizes $N \times N$ and three different values of c, and Fig. 6.1 graphs the data in this table. We observe that the magnitude of the eigenvalue depends on the magnitude of the parameter c, and that the convergence rate of the eigenvalue in this case is much slower than in the case where $F = \mathbf{1}$.

The results presented within this chapter are part of Biswal *et al.* (2020b).

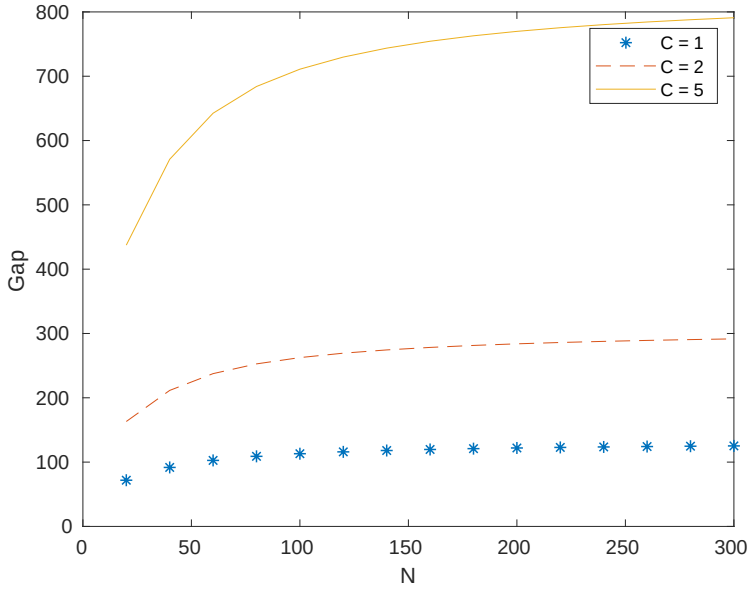

Figure 6.1: Plots of the eigenvalue $-\lambda_2$ as a function of c and N for the non-uniform F case.

Chapter 7

THE SIMPLE EXCLUSION PROCESS ON FINITE CONNECTED GRAPHS

The main objective of this chapter is to initiate the study of *simple exclusion processes*, defined below, on general finite connected graphs: an agent at vertex x jumps to a vertex chosen uniformly at random among the $\deg(x)$ neighbor(s) of x. If the chosen vertex is occupied by another agent, then the agent at x does not move. We assume that the rate at which agents jump depends on their location, with the agent at x jumping at rate ρ_x. This modeling approach is motivated by applications in multi-agent systems, such as the spatial redistribution of robotic swarms, in which agents represent robots moving on a finite graph and avoiding each other. As a first step, we study the invariant measures of the system and prove a monotonicity property for the occupation times at different vertices that holds for all finite connected graphs and all choices of the rates ρ_x.

The simple exclusion process, introduced in Spitzer (1970), is one of the most popular *interacting particle systems*, along with the voter model (Clifford and Sudbury, 1973; Holley and Liggett, 1975) and the contact process (Harris, 1974). These three models can be viewed as spatial stochastic models of diffusion, competition, and invasion, respectively. In particular, the simple exclusion process consists of a system of symmetric random walks by agents that move independently on a connected graph except that jumps onto already occupied vertices are suppressed (exclusion rule) so that each vertex is occupied by at most one agent.

All three models have been studied extensively on infinite lattices, and we refer to Liggett (1985, 1999) for a review of their main properties. The voter model and the contact process have also been studied on the torus in $\mathbb{Z}^d$ (the time to consensus

of the voter model on the torus is studied in Cox (1989) and the time to extinction of the contact process on the torus is studied in Durrett and Liu (1988); Durrett and Schonmann (1988); Durrett *et al.* (1989)), as well as on various finite deterministic and random graphs. In contrast, to the best of our knowledge, there is no work about the simple exclusion process on *finite graphs* in the probability literature, with the notable exception of the asymmetric nearest neighbor exclusion process on the finite path (each agent jumps to its immediate left or right with different probabilities) introduced in Liggett (1975) and reviewed in (Liggett, 1999, Chapter III.3). The primary motivation in that work, however, was to study the properties of the stationary distribution on a long path and relate them to those of an infinite path.

7.1 Main results

Here we present the theorems that constitute the main results of this section. The proofs of these theorems are given in Section 7.2.

Letting $\mathscr{G} = (\mathscr{V}, \mathscr{E})$ be a finite connected graph on N vertices, with vertex set $\mathscr{V}$ and edge set $\mathscr{E}$, the process is a continuous-time Markov chain whose state at time t is a configuration

$$\eta_t : \mathscr{V} \to \{0,1\} \quad \text{where} \quad \eta_t(x) = \begin{cases} 0 & \text{if vertex } x \text{ is empty} \\ 1 & \text{if vertex } x \text{ is occupied by a agent.} \end{cases}$$

Motivated by potential applications in robotic swarms, we assume that agents may jump at a rate that depends on their location, and denote by ρ_x the rate attached to vertex x. To describe the dynamics, for all $x, y \in \mathscr{V}$, we let $\tau_{x,y}\,\eta$ be the configuration

$$(\tau_{x,y}\,\eta)(z) = \eta(x)\,\mathbf{1}\{z = y\} + \eta(y)\,\mathbf{1}\{z = x\} + \eta(z)\,\mathbf{1}\{z \neq x, y\}$$

obtained from η by exchanging the states at x and y. Then, for all $\eta, \xi \in \{0,1\}^{\mathscr{V}}$, the

process jumps from configuration η to configuration ξ at rate

$$q(\eta, \xi) = \frac{\rho_x}{\deg(x)} \, \mathbf{1}\{\eta_t(x) = 1 \text{ and } \xi = \tau_{x,y}\, \eta \text{ for some } (x, y) \in \mathscr{E}\}.$$

In words, the agent at x, if there is one, chooses one of the neighbors of x uniformly at random at rate ρ_x, and jumps to this vertex if and only if it is empty. It will be convenient later to identify each configuration η with the subset of vertices occupied by an agent:

$$\eta \equiv \{x \in \mathscr{V} : \eta(x) = 1\} \subset \mathscr{V}.$$

This defines a natural bijection between the set of configurations and the subsets of the vertex set, and it will be obvious from the context whether η refers to a configuration or a subset.

The main objective of this section is to study the fraction of time each vertex is occupied in the long run. We can prove that these limits exist and only depend on the initial configuration through its number of agents, so we will write from now on

$$p_K(x) = \lim_{t \to \infty} P(x \in \eta_t \,|\, \operatorname{card}(\eta_0) = K) \quad \text{for all} \quad 0 < K \leq N \text{ and } x \in \mathscr{V}. \tag{7.1}$$

To state our results, for all $0 < K \leq N$ and $B \subset \mathscr{V}$, we define

$$\Lambda_K^+(B) = \{\eta \in \Lambda_K : B \cap \eta = B\} \quad \text{and} \quad \Lambda_K^-(B) = \{\eta \in \Lambda_K : B \cap \eta = \varnothing\}$$

where Λ_K is the set of configurations with K agents. In addition, for each vertex $z \in \mathscr{V}$, each subset $\eta \subset \mathscr{V}$, and each collection $\mathscr{C}$ of subsets of $\mathscr{V}$, we let

$$D(z) = \frac{\deg(z)}{\rho_z}, \quad D(\eta) = \prod_{z \in \eta} D(z) \quad \text{and} \quad \Sigma(\mathscr{C}) = \sum_{\eta \in \mathscr{C}} D(\eta).$$

Using that the sets Λ_K are closed communication classes as well as reversibility to identify the stationary distribution on each Λ_K, we prove that the limits in (7.1) are characterized as follows.

Theorem 7.1.1 (Occupation time)**.** – *For all* $0 < K \leq N$ *and* $x \in \mathscr{V}$*,*

$$p_K(x) = \frac{\Sigma(\Lambda_K^+(x))}{\Sigma(\Lambda_K)}. \tag{7.2}$$

In particular, the limits in (7.1) *are characterized by*

$$\frac{p_K(x)}{p_K(y)} = \frac{\Sigma(\Lambda_K^+(x))}{\Sigma(\Lambda_K^+(y))} \quad \textit{and} \quad \sum_{z \in \mathscr{V}} p_K(z) = K. \tag{7.3}$$

Even though the right-hand side of (7.2) cannot be simplified in general, some interesting properties can be deduced from this expression for arbitrary finite connected graphs. It is intuitively clear that, the graph being fixed, the probability $p_K(x)$ increases with K. It can also be proved that this probability increases with $D(x)$ while a more precise and challenging analysis shows that, though the occupation time at x is smaller than the occupation time at y when $D(x) < D(y)$, the ratio of the two occupation times converges monotonically to one as the number of agents increases.

Theorem 7.1.2 (Monotonicity)**.** – *For all* $1 < K < N - 1$*,*

$$\frac{D(x)}{D(y)} = \frac{p_1(x)}{p_1(y)} < \frac{p_K(x)}{p_K(y)} < \frac{p_{K+1}(x)}{p_{K+1}(y)} < \frac{p_N(x)}{p_N(y)} = 1 \quad \textit{when} \quad D(x) < D(y).$$

It follows from (the proof of) the theorem that, when all the $D(x)$ are equal, all the vertices are equally likely to be occupied at equilibrium. In particular, assuming for simplicity that the agents always jump at rate $\rho_x \equiv 1$, all the vertices are occupied with the same probability K/N for the process on finite regular graphs. Along these lines, the probabilities in (7.2) can be computed explicitly when $\rho_x \equiv 1$ and most of the vertices have the same degree.

7.2 Proof of Theorem 7.1.1

Note that the simple exclusion process is not irreducible because configurations with different numbers of agents do not communicate. However, the set of configurations with K agents forms a closed communication class, so the process restricted

to Λ_K is irreducible and converges to a unique stationary distribution. To find this stationary distribution and prove the theorem, we will also use reversibility. To show that Λ_K is a communication class, we first prove that any two configurations in Λ_K that only differ in two vertices communicate.

Lemma 7.2.1. – *For all* $\eta \in \{0,1\}^{\mathscr{V}}$ *and* $x, y \in \mathscr{V}$,

$$P(\eta_t = \tau_{x,y}\,\eta \,|\, \eta_0 = \eta) > 0 \quad \textit{for all} \quad t > 0.$$

Proof. The result is obvious when $\eta(x) = \eta(y)$ because in this case $\tau_{x,y}\,\eta = \eta$. The result is also clear when x and y are not in the same state but connected by an edge because

$$q(\eta, \tau_{x,y}\,\eta) = \lim_{\epsilon \downarrow 0} \frac{P(\eta_{t+\epsilon} = \tau_{x,y}\,\eta \,|\, \eta_t = \eta)}{\epsilon} = \frac{\rho_x}{\deg(x)} \quad \text{for all} \quad (x,y) \in \mathscr{E}. \tag{7.4}$$

To deal with the nontrivial case when the vertices are neither in the same state nor connected by an edge, we may assume without loss of generality that, in configuration η, vertex x is occupied and vertex y empty. Because the graph $\mathscr{G}$ is connected, there exists a self-avoiding path

$$(z_1, z_2, \ldots, z_l) \subset \mathscr{V} \quad \text{with} \quad (z_i, z_{i+1}) \in \mathscr{E},\ z_1 = x \text{ and } z_l = y$$

connecting x and y. Due to the absence of cycles, we have also $l \leq N$. To remove the agent at x and put a agent at y without changing the state of the other vertices, we let

$$\{i : z_i \in \eta\} = \{z_{i(1)}, z_{i(2)}, \ldots, z_{i(k)}\} \quad \text{with} \quad 1 = i(1) < i(2) < \cdots < i(k) < l$$

be the set of vertices along the self-avoiding path that are occupied in configuration η. It is also convenient to set $i(k+1) = l$. To obtain configuration $\tau_{x,y}\,\eta$ from η, the basic idea, is to move the agent at $z_{i(k)}$ to $z_{i(k+1)}$, then the agent at $z_{i(k-1)}$ to $z_{i(k)}$, and

so on. To prove that this sequence of events indeed occurs with positive probability, note that, because the vertices $z_{i(k)+1}, z_{i(k)+2}, \dots, z_{i(k+1)}$ are empty,

$$\tau_{z_{i(k)},z_i(k+1)}\,\eta = (\tau_{z_{i(k+1)-1},z_{i(k+1)}} \circ \cdots \circ \tau_{z_{i(k)+1},z_{i(k)+2}} \circ \tau_{z_{i(k)},z_{i(k)+1}})(\eta).$$

This, together with (7.4), implies that

$$P(\eta_t = \tau_{z_{i(k)},z_{i(k+1)}}\,\eta \,|\, \eta_0 = \eta) > 0 \quad \text{for all} \quad t > 0.$$

Similarly, we prove by induction that, for $j = 1, 2, \dots, k$,

$$\begin{aligned} &P(\eta_t = \tau_{z_{i(j)},z_{i(k+1)}}\,\eta \,|\, \eta_0 = \tau_{z_{i(j+1)},z_{i(k+1)}}\,\eta) \\ &\qquad = P(\eta_t = \tau_{z_{i(j)},z_{i(j+1)}}(\tau_{z_{i(j+1)},z_{i(k+1)}}\,\eta) \,|\, \eta_0 = \tau_{z_{i(j+1)},z_{i(k+1)}}\,\eta) > 0 \end{aligned} \tag{7.5}$$

for all $t > 0$. In addition, for $j = 1, 2, \dots, k$,

$$\tau_{z_{i(j)},z_{i(k+1)}}\,\eta = (\tau_{z_{i(j)},z_{i(j+1)}} \circ \cdots \circ \tau_{z_{i(k-1)},z_{i(k)}} \circ \tau_{z_{i(k)},z_{i(k+1)}})(\eta). \tag{7.6}$$

Using (7.5) and (7.6), and that $z_{i(1)} = x$ and $z_{i(k+1)} = y$, we deduce that

$$\begin{aligned} P(\eta_{kt} = \tau_{x,y}\,\eta \,|\, \eta_0 = \eta) &\geq \prod_{j=1}^{k} P(\eta_{(k-j+1)t} = \tau_{z_{i(j)},z_{i(k+1)}}\,\eta \,|\, \eta_{(k-j)t} = \tau_{z_{i(j+1)},z_{i(k+1)}}\,\eta) \\ &= \prod_{j=1}^{k} P(\eta_t = \tau_{z_{i(j)},z_{i(k+1)}}\,\eta \,|\, \eta_0 = \tau_{z_{i(j+1)},z_{i(k+1)}}\,\eta) > 0 \end{aligned}$$

for all $t > 0$. This completes the proof.

□

Lemma 7.2.2. – *For all $K = 0, 1, \dots, N$, the set Λ_K is a closed communication class.*

Proof. The fact that Λ_K is closed is an immediate consequence of the fact that the number K of agents is preserved by the dynamics. To prove that this set is also a communication class, fix two configurations η and ξ with K agents, and define the sets

$$S = \{z \in \mathscr{V} : \eta(z) = 1,\, \xi(z) = 0\} \quad \text{and} \quad T = \{z \in \mathscr{V} : \eta(z) = 0,\, \xi(z) = 1\}$$

that we call respectively the source set and the target set. Because η and ξ have the same number of agents, these sets have the same number of vertices, and we write

$$S = \{x_1, x_2, \ldots, x_k\} \subset \mathscr{V} \quad \text{and} \quad T = \{y_1, y_2, \ldots, y_k\} \subset \mathscr{V}.$$

By definition of the source and target sets,

$$\xi = (\tau_{x_1,y_1} \circ \tau_{x_2,y_2} \circ \cdots \circ \tau_{x_k,y_k})(\eta).$$

In particular, letting $\sigma_j = \tau_{x_1,y_1} \circ \cdots \circ \tau_{x_j,y_j}$ and using Lemma 7.2.1, we get

$$\begin{aligned} P(\eta_{kt} = \xi \,|\, \eta_0 = \eta) &\geq \prod_{j=1}^{k} P(\eta_{jt} = \sigma_j\,\eta \,|\, \eta_{(j-1)t} = \sigma_{j-1}\,\eta) \\ &= \prod_{j=1}^{k} P(\eta_t = \tau_{x_j,y_j}(\sigma_{j-1}\,\eta) \,|\, \eta_0 = \sigma_{j-1}\,\eta) > 0 \end{aligned}$$

for all $t > 0$. Since this holds for any two configurations in Λ_K and since configurations outside Λ_K cannot be reached from Λ_K, the result follows. □

We now use reversibility to identify the stationary distributions.

Lemma 7.2.3. – *For all $K = 0, 1, \ldots, N$, the distribution*

$$\pi_K(\eta) = \frac{D(\eta)}{\Sigma(\Lambda_K)} \quad \textit{for all} \quad \eta \in \Lambda_K$$

is a reversible distribution concentrated on Λ_K.

Proof. Let $\eta, \xi \in \Lambda_K$, $\eta \neq \xi$. Then,

$$q(\eta, \xi) = q(\xi, \eta) = 0 \quad \text{when} \quad \xi \neq \tau_{x,y}\,\eta \quad \text{for all} \quad (x, y) \in \mathscr{E}$$

in which case it is clear that

$$\pi_K(\eta)\, q(\eta, \xi) = \pi_K(\xi)\, q(\xi, \eta) = 0. \tag{7.7}$$

When $\xi = \tau_{x,y}\,\eta$ for some $(x, y) \in \mathscr{E}$, with say $\eta(x) = \xi(y) = 1$,

$$q(\eta, \xi) = \frac{\rho_x}{\deg(x)} \quad \text{and} \quad q(\xi, \eta) = \frac{\rho_y}{\deg(y)}.$$

In this case, because $\eta \setminus \{x\} = \xi \setminus \{y\}$,

$$\begin{aligned} \pi_K(\eta)\, q(\eta, \xi) &= \frac{D(\eta)}{\Sigma(\Lambda_K)}\, q(\eta, \xi) = \frac{D(\eta \setminus \{x\})}{\Sigma(\Lambda_K)} \left(\frac{\deg(x)}{\rho_x} \right) q(\eta, \xi) \\ &= \frac{D(\eta \setminus \{x\})}{\Sigma(\Lambda_K)} = \frac{D(\xi \setminus \{y\})}{\Sigma(\Lambda_K)} = \pi_K(\xi)\, q(\xi, \eta). \end{aligned} \tag{7.8}$$

Combining (7.7) and (7.8) gives the result. □

Proof of Theorem 7.1.1. According to Lemma 7.2.3, π_K is a reversible distribution so this is also a stationary distribution. See e.g. (Lanchier, 2017, Sec. 10.3) for a proof. Now, according to Lemma 7.2.2, the set Λ_K is a finite closed communication class for the simple exclusion process therefore there is a unique stationary distribution that concentrates on Λ_K, and this distribution is the limit of the process starting from $\eta_0 \in \Lambda_K$. See e.g. (Lanchier, 2017, Sec. 10.4) for a proof. In particular,

$$\lim_{t \to \infty} P(\eta_t = \eta \,|\, \eta_0 = \xi) = \pi_K(\eta) = \frac{D(\eta)}{\Sigma(\Lambda_K)} \quad \text{for all} \quad \eta, \xi \in \Lambda_K$$

from which it follows that

$$p_K(x) = \sum_{\eta \in \Lambda_K : x \in \eta} \pi_K(\eta) = \sum_{\eta \in \Lambda_K : x \in \eta} \frac{D(\eta)}{\Sigma(\Lambda_K)} = \sum_{\eta \in \Lambda_K^+(x)} \frac{D(\eta)}{\Sigma(\Lambda_K)} = \frac{\Sigma(\Lambda_K^+(x))}{\Sigma(\Lambda_K)}.$$

This shows the second part of the theorem, and

$$\frac{p_K(x)}{p_K(y)} = \frac{\Sigma(\Lambda_K^+(x))}{\Sigma(\Lambda_K)}\, \frac{\Sigma(\Lambda_K)}{\Sigma(\Lambda_K^+(x))} = \frac{\Sigma(\Lambda_K^+(x))}{\Sigma(\Lambda_K^+(y))}.$$

Finally, using that the expected value is linear, we get

$$\sum_{z \in \mathscr{V}} p_K(z) = \lim_{t \to \infty} \sum_{z \in \mathscr{V}} E(\mathbf{1}\{z \in \eta_t\}) = \lim_{t \to \infty} E\left(\sum_{z \in \mathscr{V}} \mathbf{1}\{z \in \eta_t\} \right) = E(K) = K.$$

This completes the proof. □

In the presence of one or N agents, we have

$$\frac{\Sigma(\Lambda_1^+(x))}{\Sigma(\Lambda_1^+(y))} = \frac{D(x)}{D(y)} = \frac{\rho_y \deg(x)}{\rho_x \deg(y)} \quad \text{and} \quad \frac{\Sigma(\Lambda_N^+(x))}{\Sigma(\Lambda_N^+(y))} = \frac{D(\mathscr{V})}{D(\mathscr{V})} = 1. \tag{7.9}$$

In all the other cases, however, the ratios above become much more complicated. Also, the theorem cannot be proved using direct calculations. The main ingredient is given by the next lemma whose proof relies on a somewhat sophisticated construction.

Lemma 7.3.1. – *For all* $0 < K < N$, *we have* $\Sigma(\Lambda_{K+1})\,\Sigma(\Lambda_{K-1}) < (\Sigma(\Lambda_K))^2$.

Proof. The key is to find partitions $\mathcal{P}$ of $\Lambda_{K+1} \times \Lambda_{K-1}$ and $\mathcal{Q}$ of $\Lambda_K \times \Lambda_K$ such that

1. partition $\mathcal{P}$ has less elements than partition $\mathcal{Q}$,

2. each $A_i \in \mathcal{P}$ can be paired with a $B_i \in \mathcal{Q}$ such that $\text{card}(A_i) \leq \text{card}(B_i)$,

3. for all $(\eta, \eta') \in A_i$ and $(\xi, \xi') \in B_i$, we have $D(\eta)\,D(\eta') = D(\xi)\,D(\xi')$.

Let $x_1, x_2, \ldots, x_N$ denote the N vertices. To construct the partitions, let

$$S_{2K} = \{(u_1, u_2, \ldots, u_N) \in \{0,1,2\}^N : u_1 + \cdots + u_N = 2K\}$$

and $\phi : \Lambda_{K+1} \times \Lambda_{K-1} \to S_{2K}$ and $\psi : \Lambda_K \times \Lambda_K \to S_{2K}$ defined as

$$\begin{aligned} \phi(\eta, \eta') &= u = (u_1, u_2, \ldots, u_N) \quad \text{where} \quad u_i = \mathbf{1}\{x_i \in \eta\} + \mathbf{1}\{x_i \in \eta'\} \\ \psi(\xi, \xi') &= u = (u_1, u_2, \ldots, u_N) \quad \text{where} \quad u_i = \mathbf{1}\{x_i \in \xi\} + \mathbf{1}\{x_i \in \xi'\}. \end{aligned} \tag{7.10}$$

The two functions have the same expression but differ in that they are not defined on the same sets of configurations. The two partitions are then given by

$$\begin{aligned} \mathcal{P} &= \{\phi^{-1}(u) : u \in S_{2K} \text{ and } \phi^{-1}(u) \neq \varnothing\} \\ \mathcal{Q} &= \{\psi^{-1}(u) : v \in S_{2K} \text{ and } \psi^{-1}(u) \neq \varnothing\}. \end{aligned}$$

The function ψ is surjective. In contrast, $\phi(\eta, \eta')$ has $\mathrm{card}(\eta \setminus \eta') \geq 2$ coordinates equal to one therefore it is not surjective:

$$S^*_{2K} = \{u \in S_{2K} : \phi^{-1}(u) \neq \varnothing\} \neq S_{2K}$$

from which it follows that

$$\mathrm{card}(\mathcal{P}) = \mathrm{card}(S^*_{2K}) < \mathrm{card}(S_{2K}) = \mathrm{card}(\mathcal{Q}). \tag{7.11}$$

This proves the first item above. Now, let

$$K_1 = \mathrm{card}\{i : u_i = 1\} \quad \text{and} \quad K_2 = \mathrm{card}\{i : u_i = 2\},$$

with $K_1 \geq 2$. To count the number of preimages (η, η') and (ξ, ξ') in (7.10), note that the vertices that are either empty or occupied in all four configurations are uniquely determined by u. This leaves K_1 vertices that are occupied in two of the four configurations, and we have that:

- the number of choices for η is the number of choices of $K + 1 - K_2$ vertices among K_1 vertices to be occupied in configuration η but not η',

- the number of choices for ξ is the number of choices of $K - K_2$ vertices among K_1 vertices to be occupied in configuration ξ but not ξ'.

Using also that $K_1 + 2K_2 = 2K < 2K + 1$, we get

$$\begin{aligned} \mathrm{card}(\phi^{-1}(u)) &= \binom{K_1}{K + 1 - K_2} = \left(\frac{K_1 - K + K_2}{K + 1 - K_2}\right)\binom{K_1}{K - K_2} \\ &< \binom{K_1}{K - K_2} = \mathrm{card}(\psi^{-1}(u)), \end{aligned} \tag{7.12}$$

which shows the second item above. Finally, for all $(\eta, \eta') \in \phi^{-1}(u)$,

$$\begin{aligned} D(\eta)\, D(\eta') &= \prod_{z \in \eta} \frac{\deg(z)}{\rho_z} \prod_{z \in \eta'} \frac{\deg(z)}{\rho_z} = \prod_{z \in \eta \Delta \eta'} \frac{\deg(z)}{\rho_z} \prod_{z \in \eta \cap \eta'} \left(\frac{\deg(z)}{\rho_z}\right)^2 \\ &= \prod_{i=1}^{N} \left(\frac{\deg(x_i)}{\rho_{x_i}}\right)^{u_i} = \widehat{D}(u) \end{aligned}$$

is a function $\widehat{D}(u)$ of the vector u only. The same holds for $(\xi, \xi') \in \psi^{-1}(u)$. This implies that the third item above is also satisfied in the sense that

$$D(\eta)\, D(\eta') = D(\xi)\, D(\xi') = \widehat{D}(u) \quad \text{for all} \quad (\eta, \eta') \in \phi^{-1}(u), (\xi, \xi') \in \psi^{-1}(u). \quad (7.13)$$

Combining (7.11)–(7.13), we conclude that

$$\begin{aligned}
\Sigma(\Lambda_{K+1})\, \Sigma(\Lambda_{K-1}) &= \sum_{u \in S^*_{2K}} \sum_{(\eta,\eta') \in \phi^{-1}(u)} D(\eta)\, D(\eta') \overset{(7.13)}{=} \sum_{u \in S^*_{2K}} \sum_{(\eta,\eta') \in \phi^{-1}(u)} \widehat{D}(u) \\
&\overset{(7.12)}{<} \sum_{u \in S^*_{2K}} \sum_{(\xi,\xi') \in \psi^{-1}(u)} \widehat{D}(u) \overset{(7.13)}{=} \sum_{u \in S^*_{2K}} \sum_{(\xi,\xi') \in \psi^{-1}(u)} D(\xi)\, D(\xi') \\
&\overset{(7.11)}{<} \sum_{u \in S_{2K}} \sum_{(\xi,\xi') \in \psi^{-1}(u)} D(\xi)\, D(\xi') = (\Sigma(\Lambda_K))^2.
\end{aligned}$$

This completes the proof. □

With the previous technical lemma, we can now prove the theorem.

Proof of Theorem 7.1.2. To simplify the notations, we write

$$A_K = \Sigma(\Lambda_K^+(\{x, y\})) \quad \text{and} \quad B_K = \Sigma(\Lambda_{K-1}^-(\{x, y\})).$$

Then, we can rewrite

$$\begin{aligned}
\Sigma(\Lambda_K^+(x)) &= \Sigma(\Lambda_K^+(x) \cap \Lambda_K^+(y)) + \Sigma(\Lambda_K^+(x) \setminus \Lambda_K^+(y)) \\
&= \Sigma(\Lambda_K^+(\{x, y\})) + D(x)\, \Sigma(\Lambda_{K-1}^-(\{x, y\})) = A_K + D(x)\, B_K.
\end{aligned}$$

Using some obvious symmetry, we deduce that

$$\frac{p_K(x)}{p_K(y)} = \frac{\Sigma(\Lambda_K^+(x))}{\Sigma(\Lambda_K^+(y))} = \frac{A_K + D(x)\, B_K}{A_K + D(y)\, B_K}. \quad (7.14)$$

Now, applying Lemma 7.3.1 to the configurations on $\mathscr{V} - \{x, y\}$, we get

$$\begin{aligned}
A_K B_{K+1} &= \Sigma(\Lambda_K^+(\{x, y\}))\, \Sigma(\Lambda_K^-(\{x, y\})) \\
&= D(\{x, y\})\, \Sigma(\Lambda_{K-2}^-(\{x, y\}))\, \Sigma(\Lambda_K^-(\{x, y\})) \\
&< D(\{x, y\})\, \Sigma(\Lambda_{K-1}^-(\{x, y\}))\, \Sigma(\Lambda_{K-1}^-(\{x, y\})) \\
&= \Sigma(\Lambda_{K+1}^+(\{x, y\}))\, \Sigma(\Lambda_{K-1}^-(\{x, y\})) = A_{K+1} B_K.
\end{aligned}$$

This, together with $D(x) < D(y)$, implies that

$$D(x)\,A_{K+1}B_K + D(y)\,A_K B_{K+1} < D(x)\,A_K B_{K+1} + D(y)\,A_{K+1}B_K$$

which is equivalent to

$$\begin{aligned}(A_K + D(x)\,B_K)(A_{K+1} + D(y)\,B_{K+1}) \\ < \ (A_{K+1} + D(x)\,B_{K+1})(A_K + D(y)\,B_K).\end{aligned} \tag{7.15}$$

Combining (7.14) and (7.15) gives

$$\frac{P_K(x \in \eta_t)}{P_K(y \in \eta_t)} = \frac{A_K + D(x)\,B_K}{A_K + D(y)\,B_K} = \frac{A_{K+1} + D(x)\,B_{K+1}}{A_{K+1} + D(y)\,B_{K+1}} = \frac{p_{K+1}(x)}{p_{K+1}(y)}. \tag{7.16}$$

The theorem is then a combination of (7.9) and (7.16). □

The results presented within this chapter are part of Biswal and Lanchier (2020).

Chapter 8

CONCLUSION AND FUTURE WORK

In this dissertation, we have used both discrete- and continuous-time mean-field models to describe the evolution of a multi-agent system on Euclidean spaces and manifolds. Specifically, decentralized agent control laws were designed to drive the agents asymptotically to a target state distribution. In the case where the agents are required to stop transitioning at the equilibrium distribution, control laws that only require knowledge of the local agent density, for example the density of agents within an agent's sensing range, were designed. The mean-field model considered within this dissertation is the forward Kolmogorov equation of a Markov process according to which the agents' states evolve. In the case of a continuous state space, the mean-field models were stabilized to arbitrary distributions that have $L^\infty(\cdot)$ densities with respect to the Lebesgue measure. Moreover, in some cases, the Markov process can be constructed such that its forward operator is the identity operator at the desired distribution. This prevents agents from switching between states once the equilibrium distribution is reached. In Chapter 2 and 4, although stability and convergence results were proven for the mean-field model, simulations of the corresponding N-agent system demonstrate that even when the number of agents is relatively small, the agents indeed redistribute themselves to the target distribution and thereafter cease switching between states.

8.1 Future Research

In most of the topics that are discussed within this dissertation, we have almost exclusively concentrated on mean-field, or macroscopic, models. There are several

reasons to work with mean-field models: usually, the mean-field model is linear or much more amenable to analysis and control than the agent dynamics, which could be quite complex. Second, when the number of agents is large, one is usually interested in predicting or controlling some macroscopic quantity; for example, the agent distribution. The underlying assumption one makes when designing feedback control laws for macroscopic models is that the these control laws would produce the same population dynamics in the corresponding N-agent system. This assumption is justified by taking the mean-field limit (i.e., the limit as N tend to infinity). However, there may not actually exist a correspondence between the mean-field model and the microscopic, or agent-level, model. One such example is given in the Chapter 6, wherein we assumed that the macroscopic dynamics are governed by a divergence-type operator, but only under stricter conditions on the coefficients a_{ij} does the operator correspond to a stochastic differential equation (SDE). One possible research direction, therefore, is to develop more physically realistic microscopic models and characterizing the resulting macroscopic behavior. Here, we provide two examples of scenarios that could be explored.

1. First, the agent-level models used in this dissertation ignore the constraints that robots occupy a finite volume and that the paths of any two robots should be separated by a certain distance in order to avoid inter-robot collisions. Introducing these constraints into the microscopic model would render the i.i.d. assumption invalid. Searching for a closed-form expression for the corresponding mean-field model would be extremely challenging. In this case, at best we can hope to establish bounds on the difference between the trajectories of this microscopic model and those of an idealized model that ignores such constraints.

2. Achieving the goal wherein the agents must stop transitioning between states when the stationary distribution is reached leads to the construction of a nonlinear

Markov process. In the models considered here, the transition kernels were designed to have a decentralized structure. However, there may be scenarios in which a particular parameter should be optimized, wherein a centralized feedback law may be preferred. At the agent level, this would imply that each agent needs to know the distribution of the swarm over the entire domain. Evaluation of this quantity requires a central observer. However, one of the challenges in the control of multi-agent systems is to develop decentralized strategies for the agents. In fact, some researchers, e.g. Musco *et al.* (2016), have drawn inspiration from ant models to develop a decentralized algorithm for density estimation. Therefore, the second question that that could be addressed is the following: using measurements of local population densities by each agent, can the entire swarm distribution be constructed? This question is known in the control theory literature as the *observability* problem. One could also add noise to the agents' estimates to model more realistic scenarios.

Additionally, one could work with the SDE description of the microscopic models. These result in forward equations that are partial differential equations (PDEs), which offer much richer dynamics than ODEs or difference equations. The models arising from PDEs also more accurately describe many physical multi-agent scenarios of interest. However, the main challenge in this approach is that the resulting control systems can have infinite degrees of freedom. Much of the control theory literature addresses the simple case in which the control inputs of a PDE are linearly dependent on the state. On the other hand, in many practical cases, the control parameter could be a reaction rate, velocity field, or a diffusion parameter. In such scenarios, the control-to-state dependence is nonlinear. Thus, there are a number of extensions one could potentially investigate, this approach to multi-agent systems is still in its nascency.

REFERENCES

Acikmese, B. and D. S. Bayard, "A Markov chain approach to probabilistic swarm guidance", in "American Control Conference (ACC)", pp. 6300–6307 (IEEE, 2012).

Açıkmeşe, B. and D. S. Bayard, "Markov chain approach to probabilistic guidance for swarms of autonomous agents", Asian Journal of Control **17**, 4, 1105–1124 (2015).

Adams, R. A. and J. J. Fournier, *Sobolev spaces*, vol. 140 (Elsevier, 2003).

Agrachev, A. and P. Lee, "Optimal transportation under nonholonomic constraints", Transactions of the American Mathematical Society **361**, 11, 6019–6047 (2009).

Arendt, W. and A. V. Bukhvalov, "Integral representations of resolvents and semigroups", in "Forum Mathematicum", vol. 6, pp. 111–136 (Walter de Gruyter, Berlin/New York, 1994).

Arnold, A., P. Markowich, G. Toscani and A. Unterreiter, "On convex Sobolev inequalities and the rate of convergence to equilibrium for Fokker-Planck type equations", Communications in Partial Differential Equations **26**, 1-2, 43–100 (2001).

Aubin, J.-P. and H. Frankowska, *Set-valued analysis* (Springer Science & Business Media, 2009).

Bagagiolo, F. and D. Bauso, "Mean-field games and dynamic demand management in power grids", Dynamic Games and Applications **4**, 2, 155–176 (2014).

Bakry, D., I. Gentil and M. Ledoux, *Analysis and geometry of Markov diffusion operators*, vol. 348 (Springer Science & Business Media, 2013).

Bandyopadhyay, S., S. J. Chung and F. Y. Hadaegh, "Probabilistic and distributed control of a large-scale swarm of autonomous agents", IEEE Transactions on Robotics (2017).

Bensoussan, A., J. Frehse, P. Yam *et al.*, *Mean field games and mean field type control theory*, vol. 101 (Springer, 2013).

Berman, A. and R. J. Plemmons, *Nonnegative matrices in the mathematical sciences* (SIAM, 1994).

Berman, S., Á. Halász, M. A. Hsieh and V. Kumar, "Optimized stochastic policies for task allocation in swarms of robots", IEEE Transactions on Robotics **25**, 4, 927–937 (2009).

Biswal, S., K. Elamvazhuthi and S. Berman, "Fastest mixing Markov chain on a compact manifold.", 2019 IEEE Conference on Decision and Control (CDC), Nice, France. (2019a).

Biswal, S., K. Elamvazhuthi and S. Berman, "Stabilization of nonlinear discrete-time systems to target measures using stochastic feedback laws", Conditionally accepted to *IEEE Transactions on Automatic Control*. Preprint available at: https://www.researchgate.net/publication/330103622 (2019b).

Biswal, S., K. Elamvazhuthi and S. Berman, "Decentralized control of multi-agent systems using local density feedback", Submitted to *IEEE Transactions on Automatic Control* (2020a).

Biswal, S., K. Elamvazhuthi and S. Berman, "Spectral gap optimization of divergence type diffusion operators.", Accepted to the 2020 European Control Conference (ECC), St. Petersburg, Russia. (2020b).

Biswal, S., K. Elamvazhuthi and S. Berman, "Stabilization of multi-agent systems to target distributions using local interactions", Submitted to the 2020 International Symposium on Mathematical Theory of Networks and Systems (MTNS). Preprint available at: https://www.researchgate.net/publication/339227488 (2020c).

Biswal, S. and N. Lanchier, "The simple exclusion process on finite connected graphs", Preprint available at: https://www.researchgate.net/publication/333640331 (2020).

Bogachev, V. I., *Measure theory*, vol. 1 (Springer Science & Business Media, 2007).

Boyd, S., P. Diaconis and L. Xiao, "Fastest mixing Markov chain on a graph", SIAM Review **46**, 4, 667–689 (2004).

Brauer, F., C. Castillo-Chavez and C. Castillo-Chavez, *Mathematical models in population biology and epidemiology*, vol. 2 (Springer, 2012).

Bullo, F., J. Cortes and S. Martinez, *Distributed control of robotic networks: a mathematical approach to motion coordination algorithms*, vol. 27 (Princeton University Press, 2009).

Carmona, R. and F. Delarue, *Probabilistic Theory of Mean Field Games with Applications I-II* (Springer, 2018).

Chen, Y., T. T. Georgiou and M. Pavon, "Optimal transport over a linear dynamical system", IEEE Transactions on Automatic Control **62**, 5, 2137–2152 (2017).

Chen, Y., T. T. Georgiou and M. Pavon, "Steering the distribution of agents in mean-field games system", Journal of Optimization Theory and Applications **179**, 1, 332–357 (2018).

Chesi, G., *Domain of attraction: analysis and control via SOS programming*, vol. 415 (Springer Science & Business Media, 2011).

Clarke, F., *Functional analysis, calculus of variations and optimal control*, vol. 264 (Springer Science & Business Media, 2013).

Clifford, P. and A. Sudbury, "A model for spatial conflict", Biometrika **60**, 581–588, URL https://doi-org.ezproxy1.lib.asu.edu/10.1093/biomet/60.3.581 (1973).

Conway, J. B., *A course in functional analysis*, vol. 96 (Springer Science & Business Media, 2013).

Cox, J. T., "Coalescing random walks and voter model consensus times on the torus in $\mathbf{Z}^{d}$", Ann. Probab. **17**, 4, 1333–1366 (1989).

Das, A. K., A. U. Raghunathan and U. Vaidya, "Transfer operator-based approach for optimal stabilization of stochastic systems", in "2017 American Control Conference (ACC)", pp. 1759–1764 (IEEE, 2017).

Del Moral, P. *et al.*, "Measure-valued processes and interacting particle systems. application to nonlinear filtering problems", The Annals of Applied Probability **8**, 2, 438–495 (1998).

Demir, N., U. Eren and B. Açıkmeşe, "Decentralized probabilistic density control of autonomous swarms with safety constraints", Autonomous Robots **39**, 4, 537–554 (2015).

Deshmukh, V., K. Elamvazhuthi, S. Biswal, Z. Kakish and S. Berman, "Mean-field stabilization of Markov chain models for robotic swarms: Computational approaches and experimental results", IEEE Robotics and Automation Letters **3**, 3, 1985–1992 (2018).

DiBenedetto, E., *Real analysis* (Springer, 2002).

Ding, J. and A. Zhou, *Statistical properties of deterministic systems* (Springer Science & Business Media, 2010).

Durrett, R. and X. F. Liu, "The contact process on a finite set", Ann. Probab. **16**, 3, 1158–1173 (1988).

Durrett, R. and R. H. Schonmann, "The contact process on a finite set. II", Ann. Probab. **16**, 4, 1570–1583 (1988).

Durrett, R., R. H. Schonmann and N. I. Tanaka, "The contact process on a finite set. III. The critical case", Ann. Probab. **17**, 4, 1303–1321 (1989).

Eisner, T., B. Farkas, M. Haase and R. Nagel, *Operator theoretic aspects of ergodic theory*, vol. 272 (Springer, 2015).

Elamvazhuthi, K., C. Adams and S. Berman, "Coverage and field estimation on bounded domains by diffusive swarms", in "2016 IEEE Conference on Decision and Control (CDC)", pp. 2867–2874 (2016).

Elamvazhuthi, K. and S. Berman, "Nonlinear generalizations of diffusion-based coverage by robotic swarms", in "2018 IEEE Conference on Decision and Control (CDC)", (2018).

Elamvazhuthi, K. and S. Berman, "Mean-field models in swarm robotics: A survey", Bioinspiration & Biomimetics **15**, 1, 015001 (2019).

Elamvazhuthi, K., S. Biswal and S. Berman, "Mean-field stabilization of robotic swarms to probability distributions with disconnected supports", in "2018 Annual American Control Conference (ACC)", pp. 885–892 (IEEE, 2018a).

Elamvazhuthi, K., S. Biswal and S. Berman, "Controllability and decentralized stabilization of the Kolmogorov forward equation for Markov chains", Provisionally accepted to *Automatica*. (2019a).

Elamvazhuthi, K., P. Grover and S. Berman, "Optimal transport over deterministic discrete-time nonlinear systems using stochastic feedback laws", IEEE Control Systems Letters **3**, 1, 168–173 (2019b).

Elamvazhuthi, K., M. Kawski, S. Biswal, V. Deshmukh and S. Berman, "Mean-field controllability and decentralized stabilization of Markov chains", in "2017 IEEE 56th Annual Conference on Decision and Control (CDC)", pp. 3131–3137 (IEEE, 2017).

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

Elamvazhuthi, K., H. Kuiper, M. Kawski and S. Berman, "Bilinear controllability of a class of advection-diffusion-reaction systems", IEEE Transactions on Automatic Control DOI: 10.1109/TAC.2018.2885231 (2018c).

Engel, K.-J. and R. Nagel, *One-parameter semigroups for linear evolution equations*, vol. 194 (Springer Science & Business Media, 2000).

Ethier, S. N. and T. G. Kurtz, *Markov Processes: Characterization and Convergence*, vol. 282 (John Wiley & Sons, 2009).

Evans, L. C., *Partial differential equations* (American Mathematical Society, 2010).

Folland, G. B., *Real analysis: modern techniques and their applications* (John Wiley & Sons, 2013).

Fourer, R., D. M. Gay and B. W. Kernighan, *AMPL: A modeling language for mathematical programming* (Thomson, 2003).

Friedman, A., *Stochastic differential equations and applications* (Dover Publications, 2006).

Grant, M. and S. Boyd, "CVX: Matlab software for disciplined convex programming, version 2.1", http://cvxr.com/cvx (2014).

Grisvard, P., *Elliptic problems in nonsmooth domains* (SIAM, 2011).

Grobler, J., "Spectral theory in Banach lattices", in "Operator theory in function spaces and Banach lattices", pp. 133–172 (Springer, 1995).

Halmos, P. R., *A Hilbert space problem book*, vol. 19 (Springer Science & Business Media, 2012).

Hamann, H. and H. Wörn, "A framework of space–time continuous models for algorithm design in swarm robotics", Swarm Intelligence **2**, 2-4, 209–239 (2008).

Harris, T. E., "Contact interactions on a lattice", Ann. Probability **2**, 969–988, URL `https://doi-org.ezproxy1.lib.asu.edu/10.1214/aop/1176996493` (1974).

Hemelrijk, C. K. and H. Hildenbrandt, "Schools of fish and flocks of birds: their shape and internal structure by self-organization", Interface focus **2**, 6, 726–737 (2012).

Henrot, A., *Extremum problems for eigenvalues of elliptic operators* (Springer Science & Business Media, 2006).

Hindawi, A., J.-B. Pomet and L. Rifford, "Mass transportation with LQ cost functions", Acta Applicandae Mathematicae **113**, 2, 215–229 (2011).

Hislop, P. D. and I. M. Sigal, *Introduction to spectral theory: With applications to Schrödinger operators*, vol. 113 (Springer Science & Business Media, 2012).

Holley, R. A. and T. M. Liggett, "Ergodic theorems for weakly interacting infinite systems and the voter model", Ann. Probability **3**, 4, 643–663, URL `https://doi-org.ezproxy1.lib.asu.edu/10.1214/aop/1176996306` (1975).

Hsieh, M. A., Á. Halász, S. Berman and V. Kumar, "Biologically inspired redistribution of a swarm of robots among multiple sites", Swarm Intelligence **2**, 2, 121–141 (2008).

Jafarizadeh, S., "Optimal diffusion processes", IEEE Control Systems Letters **2**, 3, 465–470 (2018).

Kadison, R. V. and J. R. Ringrose, *Fundamentals of the theory of operator algebras*, vol. 1 (American Mathematical Society, 1997).

Khalil, H. K., *Nonlinear systems* (Pearson, 2001), 3rd edn.

Kingston, P. and M. Egerstedt, "Distributed-infrastructure multi-robot routing using a helmholtz-hodge decomposition", in "2011 50th IEEE Conference on Decision and Control and European Control Conference", pp. 5281–5286 (2011).

Kirillov, A. A., *An introduction to Lie groups and Lie algebras*, vol. 113 (Cambridge University Press, 2008).

Kolokoltsov, V. N., *Nonlinear Markov Processes and Kinetic Equations*, vol. 182 (Cambridge University Press, 2010).

Lanchier, N., *Stochastic modeling*, Universitext (Springer, Cham, 2017).

Lasota, A. and M. C. Mackey, *Chaos, fractals, and noise: stochastic aspects of dynamics*, vol. 97 (Springer Science & Business Media, 2013).

Lasry, J.-M. and P.-L. Lions, "Mean field games", Japanese Journal of Mathematics **2**, 1, 229–260 (2007).

Lebeau, G., L. Michel *et al.*, "Semi-classical analysis of a random walk on a manifold", The Annals of Probability **38**, 1, 277–315 (2010).

Lee, J. M., *Introduction to smooth manifolds* (Springer, 2001).

Lejay, A., "Stochastic differential equations driven by processes generated by divergence form operators I: a Wong-Zakai theorem", ESAIM: Probability and Statistics **10**, 356–379 (2006).

Lewis, F. L., H. Zhang, K. Hengster-Movric and A. Das, *Cooperative control of multi-agent systems: optimal and adaptive design approaches* (Springer Science & Business Media, 2013).

Liggett, T. M., "Ergodic theorems for the asymmetric simple exclusion process", Trans. Amer. Math. Soc. **213**, 237–261, URL `https://doi-org.ezproxy1.lib.asu.edu/10.2307/1998046` (1975).

Liggett, T. M., *Interacting particle systems*, vol. 276 of *Grundlehren der Mathematischen Wissenschaften* (Springer-Verlag, New York, 1985), URL `http://dx.doi.org/10.1007/978-1-4613-8542-4`.

Liggett, T. M., *Stochastic interacting systems: contact, voter and exclusion processes*, vol. 324 of *Grundlehren der Mathematischen Wissenschaften* (Springer-Verlag, Berlin, 1999), URL `http://dx.doi.org/10.1007/978-3-662-03990-8`.

Liu, Z., B. Wu and H. Lin, "A mean field game approach to swarming robots control", in "2018 Annual American Control Conference (ACC)", pp. 4293–4298 (IEEE, 2018).

Marek, I., "Frobenius theory of positive operators: Comparison theorems and applications", SIAM Journal on Applied Mathematics **19**, 3, 607–628 (1970).

Marsden, J. E. and T. S. Ratiu, *Introduction to mechanics and symmetry: a basic exposition of classical mechanical systems*, vol. 17 (Springer Science & Business Media, 2013).

Mather, T. W. and M. A. Hsieh, "Synthesis and analysis of distributed ensemble control strategies for allocation to multiple tasks", Robotica **32**, 02, 177–192 (2014).

Mesbahi, M. and M. Egerstedt, *Graph theoretic methods in multiagent networks*, vol. 33 (Princeton University Press, 2010).

Mesquita, A. R. and J. P. Hespanha, "Jump control of probability densities with applications to autonomous vehicle motion", IEEE Transactions on Automatic Control **57**, 10, 2588–2598 (2012).

Musco, C., H.-H. Su and N. Lynch, "Ant-inspired density estimation via random walks", in "Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing", pp. 469–478 (ACM, 2016).

Nocedal, J., "KNITRO: an integrated package for nonlinear optimization", in "Large-Scale Nonlinear Optimization", pp. 35–60 (Springer, 2006).

Prajna, S., A. Papachristodoulou and P. A. Parrilo, "Introducing sostools: A general purpose sum of squares programming solver", in "Decision and Control, 2002, Proceedings of the 41st IEEE Conference on", vol. 1, pp. 741–746 (IEEE, 2002).

Prorok, A., N. Correll and A. Martinoli, "Multi-level spatial modeling for stochastic distributed robotic systems", International Journal of Robotics Research **30**, 5, 574–589 (2011).

Raghunathan, A. and U. Vaidya, "Optimal stabilization using Lyapunov measures", IEEE Transactions on Automatic Control **59**, 5, 1316–1321 (2014).

Rifford, L., *Sub-Riemannian geometry and optimal transport* (Springer, 2014).

Roberts, G. O., J. S. Rosenthal *et al.*, "Geometric ergodicity and hybrid Markov chains", Electronic Communications in Probability **2**, 2, 13–25 (1997).

Roberts, G. O. and R. L. Tweedie, "Geometric L^2 and L^1 convergence are equivalent for reversible Markov chains", Journal of Applied Probability **38**, A, 37–41 (2001).

Saldi, N., T. Basar and M. Raginsky, "Markov–Nash equilibria in mean-field games with discounted cost", SIAM Journal on Control and Optimization **56**, 6, 4256–4287 (2018).

Santambrogio, F., "Optimal transport for applied mathematicians", Birkäuser, NY pp. 99–102 (2015).

Sarlette, A. and R. Sepulchre, "Consensus optimization on manifolds", SIAM Journal on Control and Optimization **48**, 1, 56–76 (2009).

Schaefer, H. H., "Banach lattices", in "Banach Lattices and Positive Operators", pp. 46–153 (Springer, 1974).

Schilling, R. L., *Measures, integrals and martingales* (Cambridge University Press, 2017).

Schmüdgen, K., "The K-moment problem for compact semi-algebraic sets", Mathematische Annalen **289**, 1, 203–206 (1991).

Schmüdgen, K., *Unbounded self-adjoint operators on Hilbert space*, vol. 265 (Springer Science & Business Media, 2012).

Sourour, A., "Pseudo-integral operators", Transactions of the American Mathematical Society **253**, 339–363 (1979).

Spitzer, F., "Interaction of Markov processes", Adv. Math. **5**, 246–290 (1970) (1970).

Stroock, D. W., "Diffusion semigroups corresponding to uniformly elliptic divergence form operators", in "Séminaire de Probabilités XXII", pp. 316–347 (Springer, 1988).

Tao, T., *An Epsilon of Room, I: Real Analysis: pages from year three of a mathematical blog* (American Mathematical Society, Providence, RI, 2010).

Vaidya, U., P. G. Mehta and U. V. Shanbhag, "Nonlinear stabilization via control Lyapunov measure", IEEE Transactions on Automatic Control **55**, 6, 1314–1328 (2010).

Van Kampen, N. G., *Stochastic processes in physics and chemistry*, vol. 1 (Elsevier, 1992).

Villani, C., *Topics in optimal transportation*, vol. 58 (American Mathematical Society, 2003).