Stabilization of Nonlinear Discrete-Time Systems to Target Measures Using Stochastic Feedback Laws

Shiba Biswal, Karthik Elamvazhuthi, and Spring Berman

***Abstract*—In this paper, we address the problem of stabilizing a discrete-time deterministic nonlinear control system to a target invariant measure using time-invariant stochastic feedback laws. This problem can be viewed as an extension of the problem of designing the transition probabilities of a Markov chain so that the process is exponentially stabilized to a target stationary distribution. Alternatively, it can be seen as an extension of the classical control problem of asymptotically stabilizing a discrete-time system to a single point, which corresponds to the Dirac measure in the measure stabilization framework. We assume that the target measure is supported on the entire state space of the system and is absolutely continuous with respect to the Lebesgue measure. Under the condition that the system is locally controllable at every point in the state space within one time step, we show that the associated measure stabilization problem is well-posed. Given this well-posedness result, we then frame an infinite-dimensional convex optimization problem to construct feedback control laws that stabilize the system to a target invariant measure at a maximized rate of convergence. We validate our optimization approach with numerical simulations of two-dimensional linear and nonlinear discrete-time control systems.**

I. Introduction

In this paper, 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, and are absolutely continuous with respect to the Lebesgue measure with L^∞ density. This result can be viewed as a generalization of results on the stabilization of ordinary difference equations to points which can be identified with Dirac measures. Since the systems that we consider follow discrete-time Markov chains (DTMC) on continuous state spaces, the problem of stabilizing these systems can also be viewed as an extension of the problem of designing the transition probabilities of a Markov chain on a discrete state space to stabilize a target stationary distribution.

The problem of stabilizing a nonlinear control system to a target measure has many potential applications, including the control of large-scale distributed systems. For example, these measures could model the distribution of nodes in an electric power grid [6] or a wireless network [43], or the distribution of an ensemble of agents such as a swarm of robots (e.g. [1], [19]). Our main motivation is to address problems in the control of very large multi-agent systems. When all agents follow the same control laws and these control laws are independent of the agents' identities, it is possible to apply control techniques to a fluid approximation of the multi-agent system in the form of a *mean-field model* [20]. This approximation is justified by modeling each agent's dynamics as a Markov process, and then the mean-field behavior of the population is governed by the Kolmogorov forward equation corresponding to the Markov process. The mean-field model is independent of the number of agents, and consequently the control approach scales well to very large agent populations.

The problem that we address is closely related to a class of control problems that have been investigated in the context of mean-field games [31], [7], [11] and optimal transport theory [45], [40]. 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 [39]. A few works on mean-field games, including [34], [13], impose final time constraints on the probability density representing the 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 this paper.

Similar to mean-field games, optimal transport theory considers a class of measure control problems where 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 [27], [12] and nonlinear dynamics [23], [36], [3]. For discrete-time nonlinear systems, a relaxed version of the optimal transport problem was investigated in [21], 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 [21] can be considered as the fixed-endpoint control version of the problem addressed in this paper.

This work was supported by ONR Young Investigator Award N00014-16-1-2605.

Shiba Biswal and Karthik Elamvazhuthi are with the Department of Mathematics, University of California, Los Angeles, CA 90095, USA {sbiswal, karthikevaz}@math.ucla.edu.

Spring Berman is with the School for Engineering of Matter, Transport and Energy, Arizona State University, Tempe, AZ, 85281, USA spring.berman@asu.edu.

Similar measure control problems have also been considered outside the context of mean-field games and optimal transport theory. In [33], piecewise-deterministic Markov processes evolving on $\mathbb{R}^d$ are controlled to make a continuously differentiable probability density invariant and stable. For the case of DTMCs, probability distributions are stabilized on a finite graph in [1], and the existence of a stationary distribution is proven for a unicycle model in [29]. We have addressed the problem of stabilizing measures that represent swarms of agents with single-integrator dynamics perturbed by Brownian motion in [19], [22]. Other recent works [44], [35], [15] extend classical measure-theoretic studies of deterministic dynamical systems [17], [30] 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 our paper is to asymptotically stabilize a given measure that is subject to particular constraints. A similar measure stabilization problem is addressed in [13], 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.

In this paper, we 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 DTMC on the continuous state space $\mathbb{R}^d$. For DTMCs on discrete state spaces, the types of measures that can be stabilized are well-understood [1]. This stabilizability result follows from the classical *Perron-Frobenius theorem* [26], which gives a sufficient condition for the uniqueness of the stationary distribution of a Markov chain. In order to generalize the results of [1] to our type of system, 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* [41]. At present, this theory has been developed to the point where the classical theorems of Perron-Frobenius are known to hold under general hypotheses. A review of progress in this field is surveyed in [26]. 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* DTMCs, 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 [38]. 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 works have also addressed the maximization of the convergence rate of DTMCs [10], [1] and continuous-time Markov chains [8], [16] to stationary distributions; however, these results are restricted to finite, discrete state spaces.

II. Preliminaries

In this section, we present notation that will be used throughout the paper. We define $\bar{\mathbb{R}}_+ := [0,\infty)$, $\mathbb{R}_+ := (0,\infty)$, and $\mathbb{C}$ as the set of complex numbers. Similarly, we define $\bar{\mathbb{Z}}_+$ as the set of all non-negative integers and $\mathbb{Z}_+$ as the set of all positive integers. The closed ball in $\mathbb{R}^d$ of radius δ centered at x will be denoted by $B_\delta(x)$. We define ∂E and $\mathrm{int}(E)$ as the boundary and interior, respectively, of a set E. 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 [24]. The Dirac measure concentrated at a point x is denoted as δ_x, where $\delta_x(E) = 1$ if $x \in E$ and $\delta_x(E) = 0$ otherwise.

For a measure space $(\mathcal{X}, \nu)$, we define $L^p(\mathcal{X}, \nu)$, where $p \in [1,\infty)$, as the space $\{f : \mathcal{X} \to \mathbb{R} : f$ is measurable and $\|f\|_p < \infty\}$, where $\|f\|_p = (\int |f|^p d\nu)^{1/p}$. In addition, we define $L^\infty(\mathcal{X}, \nu) = \{f : \mathcal{X} \to \mathbb{R} : f$ is measurable and $\|f\|_\infty < \infty\}$, where $\|f\|_\infty = \operatorname{ess\,sup}_{x\in\mathcal{X}} |f(x)|$. $C_0(\mathcal{X})$ is the space of continuous functions f on $\mathcal{X}$ that vanish at infinity, which implies that for every $\epsilon > 0$, the set $\{x : |f(x)| \geq \epsilon\}$ is compact. For a function $f : \mathcal{X} \to \mathbb{R}$, the *support* of f is the closure of the set of points where f is nonzero. 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 $\sup_x \frac{\|Tx\|_{\mathcal{Y}}}{\|x\|_{\mathcal{X}}}$. The characteristic function over a set A will be denoted as $\chi_A(\cdot)$.

For measurable spaces $(\mathcal{X}, \mathcal{M})$ and $(\mathcal{Y}, \mathcal{N})$, where $\mathcal{M}$ and $\mathcal{N}$ are the respective sigma algebras, a *transition kernel* or *Markov kernel* is a map $\mathcal{T} : \mathcal{X} \times \mathcal{N} \to [0,1]$, where $\mathcal{T}(\cdot, E)$ is a Borel measurable function on $\mathcal{X}$ for each fixed $E \in \mathcal{N}$ and $\mathcal{T}(x, \cdot)$ is a measure on $\mathcal{Y}$ for each fixed $x \in \mathcal{X}$. The transition kernel $\mathcal{T}$ induces an operator $T : \mathcal{P}(\mathcal{X}) \to \mathcal{P}(\mathcal{Y})$ as follows. For each probability measure ν on $\mathcal{X}$,

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

defines a probability measure on $(\mathcal{Y}, \mathcal{N})$. We will say that $\mathcal{T}$ is *regular* if there exists a function $h \in L^\infty(\mathcal{X} \times \mathcal{Y}, m \times m)$ such that for each $x \in \mathcal{X}$, the measure $\mathcal{T}(x, \cdot)$ is absolutely continuous with respect to m and $\mathcal{T}(x, du) = h(x, u)dx$. The function $h : \mathcal{X} \times \mathcal{Y} \to \bar{\mathbb{R}}_+$ will be called the *kernel function* of the transition kernel $\mathcal{T}$.

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 $f \in \mathcal{X},\ g \in \mathcal{X}^*$.

III. Problem Formulation

Now we are ready to state the problems addressed in this paper. We suppose that agent dynamics is governed by the following nonlinear discrete-time control system:

$$\begin{aligned} x_{n+1} &= F(x_n, u_n), \quad n = 0, 1, 2, ... \\ x_0 &\in \Omega, \end{aligned} \tag{1}$$

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 (1) induces a controlled flow on the space of measures $\mathcal{P}(\Omega)$, given by

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

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

$$F(\cdot, u_n)_\# \mu_n(E) = \mu_n(F_{u_n}^{-1}(E)) = \int_\Omega \chi_E(F(x, u_n))dx \tag{3}$$

for each $E \in \mathcal{B}(\Omega)$.

We are interested in the problem of stabilizing system (2) 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 (2) converges to the target measure. However, in [21], 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 III.1 below.

Problem III.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*

$$\mu_{n+1} = P\mu_n, \quad n = 0, 1, 2, \ldots; \quad \mu_0 \in \mathcal{P}(\Omega) \tag{4}$$

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

$$(P\mu)(E) = \int_\Omega \int_U \chi_E(F(x, u))K(x, du)d\mu(x) \tag{5}$$

for each $E \in \mathcal{B}(\Omega)$.

This problem will be addressed in Section V, 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. Note that we will prove convergence of measures in the L^2 norm, which is a much stronger form of convergence than convergence in the weak topology, the Wasserstein metric, and the total variation norm.

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 (4) to the desired measure.

Problem III.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 $\|\mu_n - \mu_d\|_{TV} \le \alpha^n$ for all $n \in \mathbb{Z}_+$, subject to the constraint

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

where P is the forward operator (5).

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 [37]. 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. In Section IV, we establish the existence of this gap. We solve a relaxed version of the optimization problem in Section VI, as explained at the beginning of that section. In Section

VII, we solve this problem numerically. In Section VIII, we apply our optimization approach to example control systems and confirm through simulations that the systems converge to specified target measures.

IV. Analytical Properties of the Forward Operator

In this section, we will establish a few properties of the forward operator, the continuous state space analogue of the transition probability matrix for a discrete state space DTMC.

Let $\mu \in \mathcal{P}(\Omega)$. Suppose that μ is absolutely continuous with respect to the Lebesgue measure m (i.e., $\mu \ll m$). Then, μ has an $L^1(\Omega, m)$ density $f_\mu : \Omega \to \mathbb{R}$, that is, $d\mu = f_\mu dm$. 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 the Hilbert space $L^2(\Omega, m)$. 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:

$$k(x,u) \begin{cases} \geq 0, & \text{for } m\text{-a.e. } x \in \Omega, u \in U \text{ s.t. } F(x,u) \in \Omega \\ = 0, & \text{otherwise} \end{cases} \tag{6}$$

$$\int_{u \in U} k(x,u) du = 1, \text{ for } m\text{-a.e. } x \in \Omega. \tag{7}$$

These properties ensure that K is indeed stochastic.

Instead of working with P defined in (5), 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. The operator $\bar{P}$ on $L^1(\Omega, m)$ is defined by restricting P to those measures in $\mathcal{P}(\Omega)$ that have $L^1(\Omega, m)$ densities with respect to m, or equivalently, are absolutely continuous measures (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 densities, 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. At this point, we cannot write an explicit formula for $\widetilde{P}f(\cdot)$ for $f \in L^2(\Omega)$ directly from (5). Finally, the backward operator is defined to be the *Banach adjoint* of the forward operator $\widetilde{P}$; hence, we define the backward operator $\widetilde{P}^*$ on $L^2(\Omega, m)$.

We will now explore properties of the 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 IV.1. *[30] 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:*

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

The proofs of Lemma IV.2, Proposition IV.4, and Proposition IV.5 below are given in the Appendix.

Lemma IV.2. *If F is non-singular and continuous, then the operator $\bar{P}$ is well-defined (it preserves $L^1(\Omega, m)$), Markov, and bounded.*

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

Lemma IV.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 $L^\infty(\mathbb{R}^d, m)$.*

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

$$Q(x, E) = \int_U \chi_E(F(x,u)) K(x, du). \tag{8}$$

We can rewrite P in (5) in terms of Q as follows:

$$(P\mu)(E) = \int_\Omega Q(x, E)\ d\mu(x). \tag{9}$$

Note that Q can also be expressed as

$$Q(x, E) = (P\delta_x)(E). \tag{10}$$

Example 9.10 in [42] justifies evaluation of the integral in (9) 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 IV.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 [14], 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* [9] 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 $E \in \mathcal{B}(U)$ with $m(E) = 0$. Lusin's property for F_u has a similar definition.

Proposition IV.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.*

From the statement of Proposition IV.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:

$$(\widetilde{P} f_\mu)(y) = \int_\Omega q(x,y) f_\mu(x) dx. \tag{11}$$

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

$$(\widetilde{P}^* f)(x) = \int_\Omega q(x,y) f(y) dy. \tag{12}$$

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 (9), 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, P must have an eigenvector or stationary measure $\pi \in \mathcal{P}(\Omega)$; that is, $P\pi = \pi$. We assume that P is constructed such that π admits a density function $\frac{d\pi}{dm} = f_\pi$ which is strictly positive a.e. on Ω, and additionally, $f_\pi, f_\pi^{-1} \in L^\infty(\Omega, m)$. The reason for this choice will become clear shortly. Then, we must have that f_π is an eigenvector of the operator $\widetilde{P}$ corresponding to eigenvalue 1. Therefore, from (11) and (12), we have the following properties of q:

$$\int_\Omega q(x,y) dy = 1 \tag{13}$$

$$\int_\Omega q(x,y) f_\pi(x) dx = f_\pi(y) \tag{14}$$

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:

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

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 [24]. $\widehat{P}$ 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

$$\widehat{P} = M_{f_\pi}^{-1} \widetilde{P} M_{f_\pi}. \tag{15}$$

Finally, from (11), (15), and the definition of M_{f_π}, we are able to express $\widehat{P}$ as an integral operator:

$$\begin{aligned}(\widehat{P} \hat{f}_\mu)(y) &= \int_\Omega \frac{1}{f_\pi(y)} q(x,y) \hat{f}_\mu(x) f_\pi(x) dx \\ &= \int_\Omega \frac{q(x,y)}{f_\pi(y)} \hat{f}_\mu(x) d\pi(x). \end{aligned} \tag{16}$$

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

Proposition IV.5. *$\widehat{P}$ as defined in* (16) *is bounded with $\|\widehat{P}\|_{L^2(\pi)} = 1$, and as a result, $r(\widetilde{P}) \leq 1$.*

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:

$$\widehat{P}\mathbf{1} = M_{f_\pi}^{-1} \widetilde{P} M_{f_\pi} \mathbf{1} = M_{f_\pi}^{-1} \widetilde{P} f_\pi = M_{f_\pi}^{-1} f_\pi = \mathbf{1}. \tag{17}$$

The adjoint equation $\widehat{P}^*\mathbf{1} = \mathbf{1}$ follows from (45) in the Appendix.

In conclusion, we showed that for a particular choice of K, the forward operator (11) defined on $L^2(\Omega, m)$ is compact and its largest eigenvalue is 1.

V. Existence of a Solution to Problem III.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 III.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 begin by stating our assumptions. Let the measure $\mu_d \in \mathcal{P}(\Omega)$ in Problem III.1 be such that it has a density function $f_d > 0$ a.e. on Ω and satisfies $f_d, \frac{1}{f_d} \in L^\infty(\Omega, m)$. Suppose that we are given a map $F : \Omega \times U \to \mathbb{R}^d$ that satisfies the conditions stated in Problem III.1. Further, as noted in Proposition IV.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 (1) to be controllable, we need the following local controllability condition.

Definition V.1. *The system* (1) *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 (18). 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 V.2. Prove that K is regular; that is, it has an $L^\infty(\Omega \times U, m \times m)$ kernel function. See Proposition V.3.

2) Using the constructed K, formulate an operator $\widetilde{S} : L^2(\Omega, m) \to L^2(\Omega, m)$. By Proposition IV.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 V.4 and V.5. 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 V.6. Prove the uniqueness of the eigenvalue 1, which in turn will guarantee that the eigenvector f_π is the unique equilibrium of system (4). See Theorem V.8.
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 (19).
5) Obtain an expression for $P : \mathcal{P}(\Omega) \to \mathcal{P}(\Omega)$ via a closed-loop transition kernel that we will call $\widehat{Q} : \Omega \times \mathcal{B}(\Omega) \to [0, 1]$. See (21)-(23).
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 V.11.
7) Prove that $\widetilde{P}$ is *primitive*. This is to ensure that f_d is the unique asymptotically (exponentially) stable equilibrium of the system (4). See Theorem V.12.
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 V.13.

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

$$K(x, W) = \frac{m(W \cap U_x)}{m(U_x)}. \tag{18}$$

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, [2]). 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 V.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* (18) *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 V.3. *The transition kernel K defined in* (18) *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 note that the kernel function k satisfies the properties (6)-(7).

Step 2: With the given map F and the constructed kernel K in (18), we define a forward operator $S : \mathcal{P}(\Omega) \to \mathcal{P}(\Omega)$ as per (5) 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 per (8). Since K is regular from *Step 1*, Q is regular; denote the kernel function of Q by q. 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 (11):

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

We can then apply Proposition IV.4 to establish that $\widetilde{S}$ is well-defined, bounded, and compact and that it preserves $L^2(\Omega)$. Moreover, q is in $L^\infty(\Omega \times \Omega, m \times m)$ and satisfies the properties (13)-(14).

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 [18]. 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 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 $Tx \geq 0$ [26]. 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 V.4. *If Ω is path connected and system* (1) *is locally controllable, then $\widetilde{S}$ is irreducible.*

It is easy to see that $\widetilde{S}$ is a positive operator. Next, we present several properties of the spectrum of $\widetilde{S}$ ($\widetilde{S}^*$). The following lemma is straightforward to prove.

Lemma V.5. $\widetilde{S}^*\mathbf{1} = \mathbf{1}$.

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, and the eigenvalue 1 below.

Proposition V.6. *We have that:*

1) $f_\pi, f_\pi^{-1} \in L^\infty(\Omega, m)$ *and* f_π *is positive on* Ω m*-a.e.*
2) *The spectral radius* $r(\widetilde{S})$ *is 1.*

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. 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 [41].

Theorem V.7. ***(Jentzsch-Perron)****[26] 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 V.8. *If Ω is path connected and satisfies the cone condition, and F is a C^1 function and satisfies Lusin's property, then 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,

$$\widetilde{P} = (\widetilde{S} - I)\varepsilon D + I, \quad 0 < \varepsilon << 1 \tag{19}$$

Remark V.9. *For ε small enough, $\widetilde{P}$ is a positive operator.*

Remark V.10. *The transformation* (19) *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 [19] and to stabilize a class of hybrid-switching diffusions to target invariant measures in [22].*

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 (11). Instead, we will show that P can be represented as (9) 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 (19), we evaluate the right-hand side of this equation:

$$\int_E \int_\Omega q(x,y)a(x)f_\mu(x)dydx + \int_E (1 - a(x))f_\mu(x)dx, \tag{20}$$

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 (19), we will assume that $\widehat{Q}$ is of the following form,

$$\widehat{Q}(x, E) = \int_E q(x,y)a(x)dy + (1 - a(x))\delta_x(E). \tag{21}$$

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:

$$\int_\Omega \int_E q(x,y)a(x)f_\mu(x)dydx + \int_\Omega (1 - a(x))f_\mu(x)\delta_x(E)dx$$
$$= \int_\Omega \int_E q(x,y)a(x)f_\mu(x)dydx + \int_E (1 - a(x))f_\mu(x)dx.$$

By applying Fubini's theorem to the first term, we observe that the above expression is exactly equal to (20). Hence, $\widehat{Q}$ is indeed the Markov kernel of P, and P can be represented in a form similar to (9) as shown below. For all $\mu \in \mathcal{P}(\Omega)$ and all $E \in \mathcal{B}(\Omega)$,

$$(P\mu)(E) = \int_\Omega \widehat{Q}(x, E)d\mu(x) \tag{22}$$
$$= \int_\Omega \int_U \chi_E(F(x,u))\widehat{K}(x, du)d\mu(x), \tag{23}$$

where $\widehat{K}$ is the state-to-control kernel of P. The existence of $\widehat{K}$ will be ensured by Theorem V.13 below. As we will show in the next two theorems, Theorems V.11 and V.12, this constructed P is our solution to Problem III.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 V.11. *The operator $\widetilde{P}$ defined in* (19) *has* 1 *as its largest eigenvalue, and this eigenvalue is algebraically simple and isolated (i.e., is not a limit point.).*

Step 7: For (4) 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 V.12. *For all ε small enough, if f_π, f_d are bounded from below, then $\widetilde{P}$ in* (19) *is primitive.*

Step 8: Finally, we prove that $\widehat{K}$, the state-to-control kernel of P in (23), is well-defined.

Theorem V.13. *Let system* (4) *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* (21) *is the Markov kernel of the corresponding closed-loop system, and hence the equalities in* (22)-(23) *hold true.*

VI. Formulation of the Optimization Problem

In this section, we present a solution to a relaxed version of Problem III.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 (4). In this section, we investigate whether we can pose an optimization problem to search for such an operator P such that the system (4) converges exponentially fast to the equilibrium μ_d. The spectral gap of P will determine the rate of convergence of system (4); the larger the gap, the faster the convergence. Recall the assumptions on μ_d as stated in Section V: $\mu_d \ll m$, 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 [10], 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 $\widehat{P} = M_{f_d}^{-1}\widetilde{P}M_{f_d}$ in (15), which has the same spectrum as $\widetilde{P}$. The advantage here is that $\widehat{P}$ is bistochastic, as proved in (17), 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 $\mathrm{span}(\mathbf{1})$ corresponding to its largest eigenvalue 1:

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

where $\mathrm{Proj}_{(\cdot)}$ 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 $\mathrm{Proj}_{\mathbf{1}}(v) = \frac{\langle v, \mathbf{1}\rangle}{\|\mathbf{1}\|_2^2}\mathbf{1}$, and the projection of v onto $\mathbf{1}^\perp$ is $\mathrm{Proj}_{\mathbf{1}^\perp} = I - \mathrm{Proj}_{\mathbf{1}}$. Therefore, we have

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

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 (5) is defined in constraint (26) below. The relationship between $\widehat{P}$ and $\widetilde{P}$ is enforced as constraint (25) in the optimization problem, defined as follows:

$$\min_K \quad \left\|\widehat{P}(K) \circ \mathrm{Proj}_{\mathbf{1}^\perp}\right\| \tag{24}$$

subject to

$$\widehat{P} = M_{f_d}^{-1}\widetilde{P}M_{f_d}, \tag{25}$$

$$(P\mu)(A) = \int_\Omega \int_U \chi_A(F(x,u))K(x,du)d\mu \tag{26}$$
$$\forall A \in \mathcal{B}(\Omega),\ \mu \in \mathcal{P}(\Omega),$$

$$K(x,E) \geq 0 \quad \forall x \in \Omega,\ \forall E \in \mathcal{B}(U), \tag{27}$$

$$\int_\Omega K(x,U)dx = 1 \quad \forall x \in \Omega, \tag{28}$$

$$Q_K(x,A) = \int_U \chi_A(F(x,u))K(x,du) \quad \forall A \in \mathcal{B}(\Omega), \tag{29}$$

$$\int_\Omega f_d(y)Q_K(x,dy) = f_d(x) \quad \forall x \in \Omega. \tag{30}$$

The constraints (27)-(28) ensure that K is indeed a Markov kernel. Constraint (29) defines the closed-loop transition kernel Q_K in terms of K, and constraint (30) ensures that f_d is the stationary distribution of $\widetilde{P}$.

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.

VII. Numerical Optimization

In this section, we present a numerical approach to solving the optimization problem (24)-(30). Our approach can be applied to control systems of the form (1), 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,

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

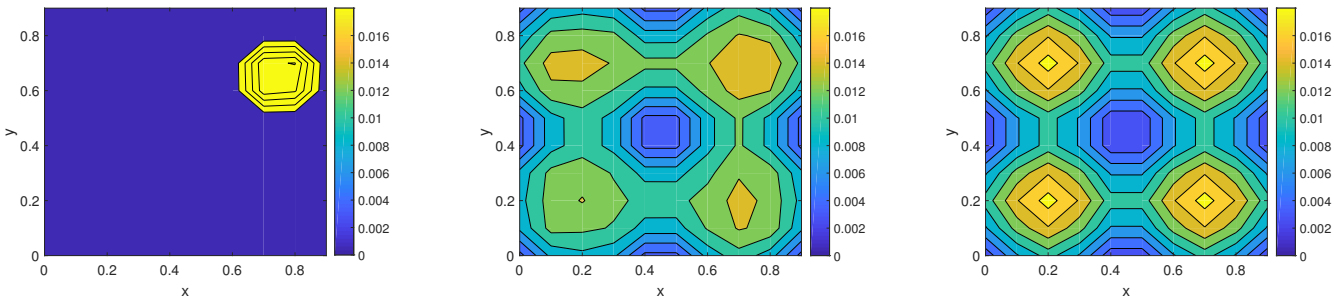

(a) $n = 2$ (b) $n = 7$ (c) $n = 15$

Fig. 1. Simulation of the additive system (39) in Example 1 at three times n.

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* [17] 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^l_{ij} 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:

$$p^l_{ij} = \frac{m(\Omega_i \cap F_l^{-1}(\Omega_j))}{m(\Omega_i)}, \tag{31}$$

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 (5), in matrix form. Then $\mathbf{P}$ is given by:

$$(\mathbf{P}\mu)(j) = \sum_{i\in\mathcal{I}} \sum_{l\in\mathcal{J}} \tilde{k}_{il}\ p^l_{ij}\mu(i). \tag{32}$$

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 (24)-(30). We define a bistochastic matrix $\widehat{\mathbf{P}}$ according to (15). This equation is enforced as constraint (34) in the quadratic program, defined as follows:

$$\min_{\mathbf{K}} \left\| \widehat{\mathbf{P}} - \frac{\mathbf{1}\mathbf{1}^T}{n_x} \right\| \tag{33}$$

subject to

$$\widehat{\mathbf{P}} = \mathbf{M}_d^{-1}\mathbf{P}\mathbf{M}_d, \tag{34}$$

$$(\mathbf{P}\mu)(j) = \sum_{i\in\mathcal{I}} \sum_{l\in\mathcal{J}} \tilde{k}_{il}\ p^l_{ij}\mu(i) \quad \forall j \in \mathcal{I},\ \forall \mu \in \mathcal{P}(\widetilde{\Omega}), \tag{35}$$

$$\tilde{k}_{il} \geq 0 \quad \forall i \in \mathcal{I},\ \forall l \in \mathcal{J}, \tag{36}$$

$$\mathbf{K}\mathbf{1} = \mathbf{1}, \tag{37}$$

$$\mathbf{P}\mu_d = \mu_d. \tag{38}$$

The constraint (35) above is written from (32), where p^l_{ij}, $i, j \in \mathcal{I},\ l \in \mathcal{J}$, is obtained via Ulam's method as per (31). Note that $\mathbf{1}$ in (33) is a vector in $\mathbb{R}^{n_x}$. We observe that this problem is convex and similar to the optimization problem solved in [10].

VIII. Simulation Results

In this section, we apply the numerical optimization procedure to two control systems of the form (1) evolving in $\mathbb{R}^2$. To solve the optimization problem (33)-(38), we used CVX, a MATLAB package for solving convex programs [25]. Since the optimization problem is a quadratic program, it becomes computationally intractable for very fine discretizations of the domain Ω. Therefore, in the examples in Subsections VIII-A and VIII-B, we use a relatively coarse discretization. In Subsection VIII-C, we solve a feasibility problem for which a finer discretization is possible. In all three cases, the quadratic program (33)-(38) was solved to obtain a state-to-control transition probability matrix $\mathbf{K}$. Defining $\mathbf{P}$ from the resulting $\mathbf{K}$ according to (35), we simulated system (4) 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.

A. Example 1: Additive Model

We first consider a linear additive vector field F in system (1):

$$x_{n+1} = x_n + u_n, \tag{39}$$

where $x_n = [x^1_n\ x^2_n]^T \in \Omega$ 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 $U = [-1,1]^2$. 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 1a-1c show snapshots of the simulation of system (4) at three times. Figure 4 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.

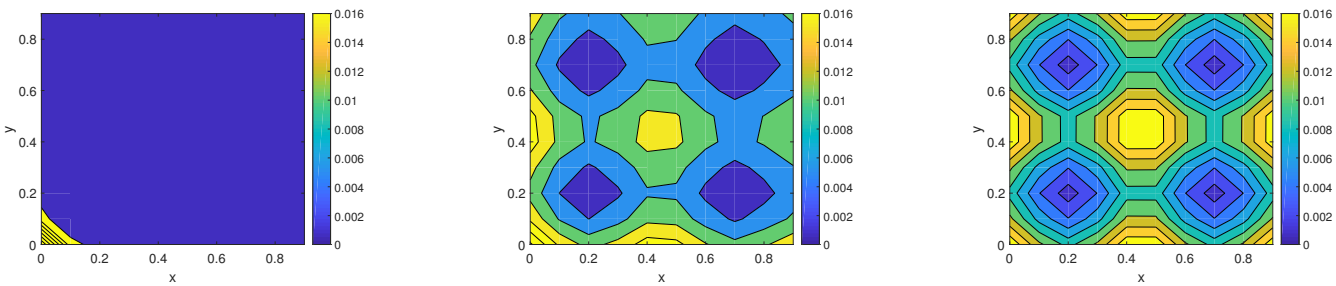

(a) $n = 2$ (b) $n = 7$ (c) $n = 15$

Fig. 2. Simulation of the unicycle system (40) in Example 2 at three times n.

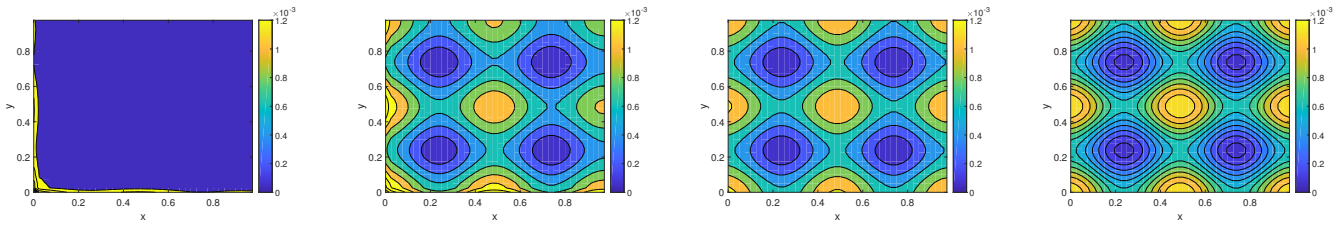

(a) $n = 5$ (b) $n = 10$ (c) $n = 15$ (d) $n = 20$

Fig. 3. Simulation of the unicycle system (40) in Example 3 at four times n.

B. Example 2: Unicycle Model

We next consider a nonlinear vector field F in system (1) that represents a unicycle model:

$$\begin{aligned} x^1_{n+1} &= x^1_n + u^1_n \cos(u^2_n), \\ x^2_{n+1} &= x^2_n + u^1_n \sin(u^2_n). \end{aligned} \tag{40}$$

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 2a-2c show snapshots of the simulation of system (4) at three times, and Figure 4 plots the natural logarithm of the error metric γ_n over time. Again, the measure converges asymptotically to the target measure.

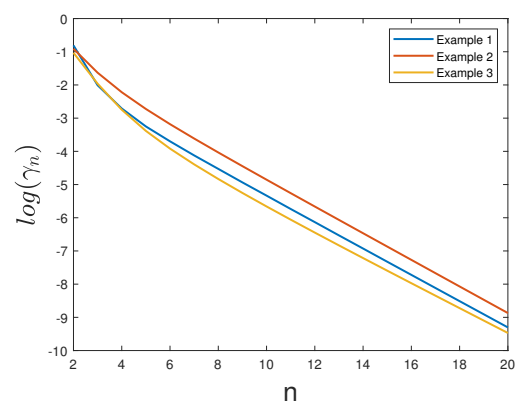

Fig. 4. The time evolution of the natural logarithm of the error between the simulated and target measures in each example.

C. Example 3: Feasibility Problem

In the two previous cases, the optimization problem (33)-(38) 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 (4), we solve the feasibility problem. This entails searching for any matrix $\mathbf{K}$ such that the steady-state distribution of system (4) is μ_d; i.e., any $\mathbf{K}$ that satisfies the constraints (34)-(38). This serves to demonstrate that the feasibility problem can be solved for larger grid sizes than $n_x = 100$. The simulated system (1) is defined as the unicycle model (40) 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 3a-3d and 4 show snapshots of the simulation of system (4) 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 4 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 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 VI. 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.

IX. Conclusion

In this paper, we presented an approach to stabilizing a class of discrete-time nonlinear systems that evolve on a compact subset of $\mathbb{R}^d$ to target probability measures that are absolutely continuous with respect to the Lebesgue measure, have L^∞ derivatives, and are positive almost everywhere on this compact set. We also presented a method to optimize the convergence rate of the system to this target measure and validated our method in simulation. Possible directions for future work are to extend our stabilization result to arbitrary discrete-time nonlinear systems that are n-step controllable and to prove the existence of an optimal solution to the optimization problem that we have posed. In addition, we plan to extend our results to continuous-time nonlinear systems evolving on continuous state spaces.

Appendix A

Proof of Lemma IV.2. 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 (5), 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 *(i)* of Definition IV.1 follows from property (6), and accordingly the integrand in (5) is non-negative. Condition *(ii)* 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 $\|\bar{P}\|_1 = 1$. $\square$

Proof of Proposition IV.4. The proof will be divided into the following key steps.

1. Prove that the closed-loop kernel Q defined in (10) 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 [14], on $L^2(\Omega \times \Omega, m \times m)$.
3. Apply Proposition II.4.7 of [14] 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 (5), we obtain:

$$\begin{aligned}(P\delta_z)(E) &= \int_\Omega \int_U \chi_E(F(x,u))\ k(x,u)\ du\ d\delta_z\\ &= \int_U \chi_E(F(z,u))\ k(z,u)\ du \qquad (41)\\ &\leq \|k\|_\infty \int_U \chi_E(F(z,u))\ du \qquad (42)\\ &= \|k\|_\infty\ m(F_z^{-1}(E)). \qquad (43)\end{aligned}$$

The equality (41) follows from *Fubini's theorem* [24]. 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, [9]) Since the change of variables theorem can only be applied to an open set, we restrict F to $\mathrm{int}(U)$ (i.e., $F|_{\mathrm{int}(U)}$). 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 (42) 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,

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

Combining this result with (42), 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,

$$\begin{aligned}(P\mu)(E) &= \int_E (\widetilde{P}f_\mu)(x)dx = \int_\Omega Q(x,E)f_\mu(x)dx \qquad (44)\\ &\leq \int_\Omega \|k\|_\infty c_1 m(E) f_\mu(x)dx = Cm(E),\end{aligned}$$

where C is a constant. The second equality follows from (9), and the inequality follows from computations above. Hence, we have achieved a uniform bound on $(P\mu)(E)$, which by Lemma IV.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 [4], 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 IV.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 (44), we have

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

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 [14], given by $q \in L^\infty(\Omega \times \Omega, m \times m)$, we can apply Proposition II.4.7 of [14] to obtain our result, namely, that $\widetilde{P}$ is well-defined, bounded, and compact. □

Proof of Proposition IV.5. We will use Theorem 6.18 of [24] 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 (13), we have

$$\int_\Omega q(x,y) dy = \int_\Omega \frac{q(x,y)}{f_\pi(y)} f_\pi(y) dy = \int_\Omega \frac{q(x,y)}{f_\pi(y)} d\pi(y) = 1. \tag{45}$$

Next, we evaluate the integral $\int_\Omega \frac{q(x,y)}{f_\pi(y)} d\pi(x)$. Using property (14), 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 [24] 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 $r(\widehat{P}) \leq 1$. □

Proof of Proposition V.2. 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 [9], as was done in the proof of Proposition IV.4, to obtain the following lower bound on $m(F_x^{-1}(B_x))$:

$$\begin{aligned} \int_U \chi_{B_x}(F(x,u)) du &\geq c_2 \int_U \chi_{B_x}(F(x,u)) |\det(D_u F_x)| du \\ &= c_2 \int_{F_x(U)} \chi_{B_x}(y) dy = c_2 \; m(B_x \cap F_x(U)) \\ &= c_2 \; m(B_x) \geq m(\mathcal{C}_B) \end{aligned} \tag{46}$$

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* [24] implies that $(\chi_G)_x$ is Borel measurable for each $x \in \Omega$, and therefore 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, 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. This concludes the proof. □

A. Proof of Proposition V.3

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,

$$(A\mu)(W) = \int_\Omega K(x, W) d\mu(x), \; W \in \mathcal{B}(U) \tag{47}$$

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 A.1. *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 IV.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 V.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 IV.3, the operator $\bar{A}$ takes $L^1(\Omega, m)$ to $L^\infty(U, m)$. □

Proof of Proposition V.3. The proof follows from an application of Theorem 1.3 of [4] in combination with the approach used in Proposition IV.4. □

B. Proof of Proposition V.4

Before presenting the proof, we need the following characterization of ideals on a finite-dimensional measure space from [18]. 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 $E \in \mathcal{M}$:

$$I_E := \{g : E \subseteq \{g = 0\}\} \tag{48}$$

Proof of Proposition V.4. 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 (48) 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

$$Q(x, E) = (S\delta_x)(E) = \int_U \chi_E(F(x,u)) k(x,u) du = 0 \quad \text{for } \mu\text{-a.e. } x \in E^c \tag{49}$$

Next, we examine the two possible ways that the integral in (49) 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 5, $B_r(\cdot)$ 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 5.

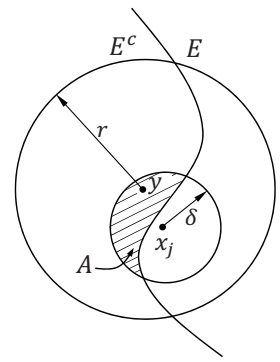

Fig. 5. Illustration of the subset A (shaded region) used in the proof of Proposition V.4.

Returning to the integral (49), 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 $V := F_y^{-1}(B_\delta(x_j) \cap E)$ and restricting the domain of integration in (49) 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 (49), and hence $\widetilde{S}$ is indeed irreducible. □

Proof of Proposition V.6. (1) 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 V.4) 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 $z \in \mathcal{N}(x)$,

$$(S\delta_z)(B_\epsilon(x)) \geq c m(B_\epsilon(x)). \tag{50}$$

First, we will assume that (50) 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 (50) 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 V.1.

From (18), 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 (46); hence, we have omitted a few steps here.

$$\begin{aligned}(S\delta_z)(B_\epsilon) &= \int_U \chi_{B_\epsilon}(F(z, u))k(z, u)du \quad \text{(from (41))} \\ &\geq c_3 \int_U \chi_{B_\epsilon(x)}(F(z, u))|\det D_u F_z|du \\ &= c_3 \int_{F_z(U)} \chi_{B_\epsilon(x)}(y)dy = c_3 m(B_\epsilon(x))\end{aligned}$$

This shows that for every $x \in \Omega$, $(S\delta_z)(B_\epsilon(x)) \geq c_3 m(B_\epsilon(x))$. This proves that (50) holds true, and thus it is indeed true that $f_\pi^{-1} \in L^\infty(\Omega, m)$.

(2) The proof follows from the discussion in Section IV and Proposition IV.5. □

C. Proof of Proposition V.11

We require the following definitions from [28] 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 $AR_T(z) := A(T - zI)^{-1}$ 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 Proposition V.11. 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, [28]), 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 V.6 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 [28], 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 [28], 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 [41], 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 [32] 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 [41] 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.

Second, we show that 1 is the spectral radius and an isolated eigenvalue of $\widetilde{P}$. As per Proposition 4.1 of [41], 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 [32] 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 [41] 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. □

D. Proof of Proposition V.12

To prove that $\widetilde{P}$ is primitive, we require the following theorem from [26]. 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 $\mathcal{X}$ [41]. 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 Proposition V.12. 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 IV.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 (19), 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 V.9). 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. □

Proof of Proposition V.13. 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 [5]. Since system (4) 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 [5], 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 $W \in \mathcal{B}(U)$,

$$\widehat{K}(x, W) = a(x)K(x, W) + (1 - a(x))\delta_{v(x)}(W), \qquad (51)$$

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 (51) 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 $E \in \mathcal{B}(\Omega)$:

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

We obtained the first term in the integral (†) by using (41) 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. □

References

[1] Behçet Açıkmeşe and David S Bayard. Markov chain approach to probabilistic guidance for swarms of autonomous agents. *Asian Journal of Control*, 17(4):1105–1124, 2015.

[2] Robert A Adams and John JF Fournier. *Sobolev spaces*, volume 140. Elsevier, 2003.

[3] Andrei Agrachev and Paul Lee. Optimal transportation under nonholonomic constraints. *Transactions of the American Mathematical Society*, 361(11):6019–6047, 2009.

[4] Wolfgang Arendt and Alexander V Bukhvalov. Integral representations of resolvents and semigroups. In *Forum Mathematicum*, volume 6, pages 111–136. Walter de Gruyter, Berlin/New York, 1994.

[5] Jean-Pierre Aubin and Hélène Frankowska. *Set-valued analysis*. Springer Science & Business Media, 2009.

[6] Fabio Bagagiolo and Dario Bauso. Mean-field games and dynamic demand management in power grids. *Dynamic Games and Applications*, 4(2):155–176, 2014.

[7] Alain Bensoussan, Jens Frehse, Phillip Yam, et al. *Mean field games and mean field type control theory*, volume 101. Springer, 2013.

[8] Spring Berman, Ádám Halász, M Ani Hsieh, and Vijay Kumar. Optimized stochastic policies for task allocation in swarms of robots. *IEEE Transactions on Robotics*, 25(4):927–937, 2009.

[9] Vladimir I Bogachev. *Measure theory*, volume 1. Springer Science & Business Media, 2007.

[10] Stephen Boyd, Persi Diaconis, and Lin Xiao. Fastest mixing Markov chain on a graph. *SIAM Review*, 46(4):667–689, 2004.

[11] René Carmona and François Delarue. *Probabilistic Theory of Mean Field Games with Applications I-II*. Springer, 2018.

[12] Yongxin Chen, Tryphon T Georgiou, and Michele Pavon. Optimal transport over a linear dynamical system. *IEEE Transactions on Automatic Control*, 62(5):2137–2152, 2017.

[13] Yongxin Chen, Tryphon T Georgiou, and Michele Pavon. Steering the distribution of agents in mean-field games system. *Journal of Optimization Theory and Applications*, 179(1):332–357, 2018.

[14] John B Conway. *A course in functional analysis*, volume 96. Springer Science & Business Media, 2013.

[15] Apurba Kumar Das, Arvind U Raghunathan, and Umesh Vaidya. Transfer operator-based approach for optimal stabilization of stochastic systems. In *2017 American Control Conference (ACC)*, pages 1759–1764. IEEE, 2017.

[16] Vaibhav Deshmukh, Karthik Elamvazhuthi, Shiba Biswal, Zahi Kakish, and Spring 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.

[17] Jiu Ding and Aihui Zhou. *Statistical properties of deterministic systems*. Springer Science & Business Media, 2010.

[18] Tanja Eisner, Bálint Farkas, Markus Haase, and Rainer Nagel. *Operator theoretic aspects of ergodic theory*, volume 272. Springer, 2015.

[19] Karthik Elamvazhuthi and Spring Berman. Nonlinear generalizations of diffusion-based coverage by robotic swarms. In *2018 IEEE Conference on Decision and Control (CDC)*, pages 1341–1346, Dec. 2018.

[20] Karthik Elamvazhuthi and Spring Berman. Mean-field models in swarm robotics: A survey. *Bioinspiration & Biomimetics*, 15(1):015001, Oct. 2019.

[21] Karthik Elamvazhuthi, Piyush Grover, and Spring Berman. Optimal transport over deterministic discrete-time nonlinear systems using stochastic feedback laws. *IEEE Control Systems Letters*, 3(1):168–173, 2019.

[22] Karthik Elamvazhuthi, Hendrik Kuiper, Matthias Kawski, and Spring Berman. Bilinear controllability of a class of advection-diffusion-reaction systems. *IEEE Transactions on Automatic Control*, 64(6):2282–2297, 2018.

[23] Alessio Figalli and Ludovic Rifford. Mass transportation on sub-Riemannian manifolds. *Geometric and functional analysis*, 20(1):124–159, 2010.

[24] Gerald B Folland. *Real analysis: modern techniques and their applications*. John Wiley & Sons, 2013.

[25] Michael Grant and Stephen Boyd. CVX: Matlab software for disciplined convex programming, version 2.1. http://cvxr.com/cvx, March 2014.

[26] JJ Grobler. Spectral theory in Banach lattices. In *Operator theory in function spaces and Banach lattices*, pages 133–172. Springer, 1995.

[27] Ahed Hindawi, J-B Pomet, and Ludovic Rifford. Mass transportation with LQ cost functions. *Acta Applicandae Mathematicae*, 113(2):215–229, 2011.

[28] Peter D Hislop and Israel Michael Sigal. *Introduction to spectral theory: With applications to Schrödinger operators*, volume 113. Springer Science & Business Media, 2012.

[29] S. M. Huck and J. Lygeros. Stochastic localization of sources with convergence guarantees. In *2013 European Control Conference (ECC)*, pages 602–607, July 2013.

[30] Andrzej Lasota and Michael C Mackey. *Chaos, fractals, and noise: stochastic aspects of dynamics*, volume 97. Springer Science & Business Media, 2013.

[31] Jean-Michel Lasry and Pierre-Louis Lions. Mean field games. *Japanese Journal of Mathematics*, 2(1):229–260, 2007.

[32] Ivo Marek. Frobenius theory of positive operators: Comparison theorems and applications. *SIAM Journal on Applied Mathematics*, 19(3):607–628, 1970.

[33] Alexandre R Mesquita and Joao P Hespanha. Jump control of probability densities with applications to autonomous vehicle motion. *IEEE Transactions on Automatic Control*, 57(10):2588–2598, 2012.

[34] Alessio Porretta. On the planning problem for the mean field games system. *Dynamic Games and Applications*, 4(2):231–256, 2014.

[35] Arvind Raghunathan and Umesh Vaidya. Optimal stabilization using Lyapunov measures. *IEEE Transactions on Automatic Control*, 59(5):1316–1321, 2014.

[36] Ludovic Rifford. *Sub-Riemannian geometry and optimal transport*. Springer, 2014.

[37] Gareth O Roberts, Jeffrey S Rosenthal, et al. Geometric ergodicity and hybrid Markov chains. *Electronic Communications in Probability*, 2(2):13–25, 1997.

[38] Gareth O Roberts and Richard L Tweedie. Geometric L^2 and L^1 convergence are equivalent for reversible Markov chains. *Journal of Applied Probability*, 38(A):37–41, 2001.

[39] Naci Saldi, Tamer Basar, and Maxim Raginsky. Markov–Nash equilibria in mean-field games with discounted cost. *SIAM Journal on Control and Optimization*, 56(6):4256–4287, 2018.

[40] Filippo Santambrogio. Optimal transport for applied mathematicians. *Birkäuser, NY*, pages 99–102, 2015.

[41] Helmut H Schaefer. Banach lattices. In *Banach Lattices and Positive Operators*, pages 46–153. Springer, 1974.

[42] René L Schilling. *Measures, integrals and martingales*. Cambridge University Press, 2017.

[43] Hamidou Tembine. Energy-constrained mean field games in wireless networks. *Strategic Behavior and the Environment*, 4(1):99–123, 2014.

[44] Umesh Vaidya, Prashant G Mehta, and Uday V Shanbhag. Nonlinear stabilization via control Lyapunov measure. *IEEE Transactions on Automatic Control*, 55(6):1314–1328, 2010.

[45] Cédric Villani. *Topics in optimal transportation*, volume 58. American Mathematical Society, 2003.

Shiba Biswal received the M.S. and Ph.D. degrees in mechanical engineering and the M.A. degree in mathematics from Arizona State University (ASU), Tempe, AZ, in 2015, 2020, and 2019, respectively. Currently, she is a Hedrick Assistant Adjunct Professor at the Department of Mathematics, University of California, Los Angeles. Her research interests include nonlinear control theory, probability theory, stochastic processes, and stochastic hybrid systems.

Karthik Elamvazhuthi received the B.Tech. degree in mechatronics engineering from the Manipal Institute of Technology, Manipal, India in 2011. He received the M.S. and Ph.D. degrees in mechanical engineering from Arizona State University, Tempe, AZ, in 2014 and 2019, respectively. Currently, he is a CAM Assistant Adjunct Professor at the Department of Mathematics, University of California, Los Angeles. His research interests include modeling and control of robotic swarms using methods from partial differential equations and stochastic processes.

Spring Berman (M'07) received the M.S.E. and Ph.D. degrees in mechanical engineering and applied mechanics from the University of Pennsylvania in 2008 and 2010, respectively. From 2010 to 2012, she was a postdoctoral researcher in computer science at Harvard University. Since 2012, she has been a professor of mechanical and aerospace engineering at Arizona State University. Her research focuses on the analysis of behaviors in biological and engineered collectives and the synthesis of control strategies for robotic swarms. She received the 2014 DARPA Young Faculty Award and the 2016 ONR Young Investigator Award.