Decentralized Control of Multi-Agent Systems using Local Density Feedback

Shiba Biswal, Karthik Elamvazhuthi, and Spring Berman

***Abstract*—In this paper, we stabilize a discrete-time Markov process evolving on a compact subset of $\mathbb{R}^d$ to an arbitrary target distribution that has an $L^\infty(\cdot)$ density and does not necessarily have a connected support on the state space. We address this problem by stabilizing the corresponding Kolmogorov forward equation, the *mean-field model* of the system, using a density-dependent transition kernel as the control parameter. Our main application of interest is controlling the distribution of a multi-agent system in which each agent evolves according to this discrete-time Markov process. To prevent agent state transitions at the equilibrium distribution, which would potentially waste energy, we show that the Markov process can be constructed in such a way that the operator that pushes forward measures is the identity at the target distribution. In order to achieve this, the transition kernel is defined as a function of the current agent distribution, resulting in a nonlinear Markov process. Moreover, we design the transition kernel to be *decentralized* in the sense that it depends only on the local density measured by each agent. We prove the existence of such a decentralized control law that globally stabilizes the target distribution. Further, to implement our control approach on a finite N-agent system, we smoothen the mean-field dynamics via the process of *mollification*. We validate our control law with numerical simulations of multi-agent systems with different population sizes. We observe that as N increases, the agent distribution in the N-agent simulations converges to the solution of the mean-field model, and the number of agent state transitions at equilibrium decreases to zero.**

I. Introduction

In this work, we address the problem of stabilizing a multi-agent system evolving on a compact, connected subset of $\mathbb{R}^d$ to a target distribution. This problem can arise, for example, in the allocation of nodes in an electric power grid [4] or a wireless network [38], or the redistribution of an ensemble of agents such as a swarm of robots (e.g. [1], [19]). Our main application is a variety of multi-agent applications that require task reallocation or spatial redistribution, such as environmental monitoring, surveillance, disaster response, and autonomous construction. We consider groups of agents that all follow the same dynamics and control policies, which are independent of the agents' identities. We assume that each agent can obtain local measurements of the agent population but do not require inter-agent communication.

Instead of specifying the spatiotemporal evolution of each individual agent, a *microscopic approach* to agent control, we will design agent control laws using a fluid approximation of the multi-agent system, called the *macroscopic* or *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 determined by the *Kolmogorov forward equation* corresponding to the Markov process. Mean-field models have a long history; their origin is in statistical physics, in particular for modeling dense gases [37], and they have been applied in epidemiology [3], game theory [5], and in the modeling of computer networks [10], electric power system loads [33], and energy demand networks [14].

We will address the problem of stabilizing the mean-field model using the *transition kernel* as the control parameter. In contrast to commonly used graph-theoretic approaches for controlling multi-agent systems [12], [36], control approaches based on mean-field models scale well to very large numbers of agents. Moreover, a range of tools are available to analyze and control mean-field dynamical models, which have the advantage of linearity in the absence of agent interactions.

In this paper, we design stochastic agent control laws using the mean-field model of the agent population dynamics. Essentially, we construct an ergodic Markov process such that its stationary distribution is a target $L^\infty(\cdot)$ function. Moreover, we require the agents to stop transitioning at this target distribution. This condition requires that the operator acting on densities be the identity operator at the target distribution. We consider multi-agent systems in which the agents have specified nonlinear discrete-time dynamics. This paper builds on our work in [6], wherein we constructed a Markov process that can be stabilized to probability distributions that have continuous densities. Moreover, in contrast to this paper, explicit agent dynamics were not specified in [6].

We detail our **contribution** in the following three points:

1. We design the transition kernel to stabilize the mean-field model to target measures that have $L^\infty(\cdot)$ densities, a larger class of measures than we previously considered in [6]–[8]. In [7], [8], we considered measures that have $L^\infty(\cdot)$ densities that are strictly positive a.e. (almost everywhere) on the domain. These works generalize the Perron-Frobenius theorem [29], which does not guarantee stabilization to distributions on discrete state spaces that are not strongly connected (see [1] for discrete-time Markov chains and [22] for continuous-time Markov chains), to Markov processes on continuous state spaces. However, in this paper, we are able to stabilize the mean-field model to distributions that are not supported everywhere on the domain by using a control law that is density-dependent.

A similar measure control problem is addressed in [13],

This work was supported by the Arizona State University Global Security Initiative.

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.

in which the authors consider an optimal control problem that drives a linear system evolving on $\mathbb{R}^d$, for a fixed initial condition on the measure, to target Gaussian measures. One of the earlier works on the topic of distribution control is [30], in which reducing the difference between distributions is posed as a minimization problem.

2. The convergence of a Markov process to an equilibrium distribution does not necessarily imply that the agents evolving according to the process also converge to equilibrium states. In fact, agents may continue to transition between states, which can cause them to waste energy. To prevent agents from continuing to switch between states at the equilibrium distribution, we construct the Markov process such that its forward operator, which pushes forward measures, is the identity operator at the desired equilibrium. This results in a time-dependent transition kernel that is a function of the distribution and gives rise to a *nonlinear Markov process*. Such stochastic processes, called *density-dependent population processes*, are used to model the dynamics of logistic growth, epidemics, and chemical reaction networks [25]. Preliminary work in this direction for discrete-time Markov chains evolving on finite graphs has been done in [34].

3. Since we establish that the transition kernel must depend on the distribution, our third goal is to construct the kernel to have a decentralized structure. A kernel with this structure corresponds to agent control policies that require each agent to estimate the population only in its local neighborhood. Toward this end, we construct a kernel for the mean-field model that is defined pointwise; that is, it is a function of the value of the distribution at the current state. We proved the existence of such feedback control laws in the case of continuous-time Markov chains evolving on finite graphs in [21], [24]. A similar problem is addressed in [35], which develops a decentralized control approach by *a priori* restricting the controller to have a decentralized structure. Another related work [17] designs a centralized controller and uses estimation algorithms to determine the entire agent distribution in a decentralized manner.

Our approach of analyzing the stability of a dynamical system from a measure-theoretic point of view is quite classical [32]. This approach is also used extensively in the context of mean-field games [28], optimal transport theory [39], and mean-field control [27]. In [7], we present a review of significant works that have influenced research on the stabilization of Markov processes.

II. Notation

In this section, we present notation that will be used throughout the paper. We define $\bar{\mathbb{R}}_+ := [0, \infty)$ and $\mathbb{R}_+ := (0, \infty)$. Similarly, we define $\bar{\mathbb{Z}}_+$ as the set of all non-negative integers and $\mathbb{Z}_+$ as the set of all positive integers. Given a dimension $d \geq 1$, the closed ball in $\mathbb{R}^d$ of radius δ centered at x will be denoted by $B_\delta(x)$. For an arbitrary set A, the symbol $|A|$ will refer to the cardinality of A.

We denote the state space by $\Omega \subseteq \mathbb{R}^d$, a measurable compact set. The Borel sigma algebra on Ω, corresponding to the standard topology on $\mathbb{R}^d$, is denoted by $\mathcal{B}(\Omega)$. The set of admissible control inputs and its corresponding Borel sigma algebra will be denoted by U and $\mathcal{B}(U)$, respectively. 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 space of probability measures on Ω and U by $\mathcal{P}(\Omega)$. 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 $f : \mathbb{R}^n \rightarrow \mathbb{R}$ such that $d\nu = f dm$; this function is called the *Radon-Nikodym derivative* of ν with respect to m [26].

For a measure space $(\mathcal{X}, \nu)$, we define $L^p(\mathcal{X}, \nu)$, where $p \in [1, \infty)$, as the space $\{f : \mathcal{X} \rightarrow \mathbb{R} : f \text{ is measurable and } \|f\|_p < \infty\}$, where $\|f\|_p = (\int |f|^p d\nu)^{1/p}$. In addition, we define $L^\infty(\mathcal{X}, \nu) = \{f : \mathcal{X} \rightarrow \mathbb{R} : f \text{ is measurable and } \|f\|_\infty < \infty\}$, where $\|f\|_\infty = \operatorname{ess\,sup}_{x \in \mathcal{X}} |f(x)|$. Where it is understood, the measure will be dropped from the notation of L^p spaces. $C(\mathcal{X})$ is the space of continuous functions on $\mathcal{X}$. For a function $f : \mathcal{X} \rightarrow \mathbb{R}$, the *support* of f is the closure of the set of points in $\mathcal{X}$ where f is nonzero. The characteristic function over a set A will be denoted as $\chi_A(\cdot)$. The Dirac measure concentrated a point x is denoted as δ_x, where $\delta_x(A) = 1$ if $x \in A$ and $\delta_x(A) = 0$ otherwise.

For measurable spaces $\mathcal{X}$ and $\mathcal{Y}$ with sigma algebras $\mathcal{M}$ and $\mathcal{N}$, respectively, a *transition kernel* or *Markov kernel* is a map $\mathcal{T} : \mathcal{X} \times \mathcal{N} \rightarrow [0, 1]$, where $\mathcal{T}(x, \cdot)$ is a measure on $\mathcal{Y}$ for each fixed $x \in \mathcal{X}$ and $\mathcal{T}(\cdot, E)$ is a Borel measurable function on $\mathcal{X}$ for each fixed $E \in \mathcal{N}$. The transition kernel $\mathcal{T}$ induces an operator $T : \mathcal{P}(\mathcal{X}) \rightarrow \mathcal{P}(\mathcal{Y})$ 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 $\mathcal{Y}$. 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 $\mathcal{T}(x, du) = h(x, u)du$. The density $h : \mathcal{X} \times \mathcal{Y} \rightarrow \bar{\mathbb{R}}$ will also be called the *kernel function* of the transition kernel $\mathcal{T}$.

For a function $F : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}^d$, we define F_x as the map from $\mathcal{Y} \rightarrow \mathbb{R}^d$ when $x \in \mathcal{X}$ is fixed, and similarly, we define F_y as the map from $\mathcal{X} \rightarrow \mathbb{R}^d$ when $y \in \mathcal{Y}$ is fixed.

III. Problem Formulation

We now state the problem addressed in this paper. Consider a system of N agents that evolve in discrete time on the set $\Omega \subset \mathbb{R}^d$. We assume that the agents are *identity-free*, by which we mean that they evolve independently of one another and according to the same dynamics.

We suppose that the dynamics of each agent $k \in \{1, \dots, N\}$ is governed by a continuous map $F : \Omega \times U \rightarrow \mathbb{R}^d$. We assume that F is *non-singular*, which means that for all $E \in \mathcal{B}(\Omega)$, $m(F_u^{-1}(E)) = 0$ and $m(F_x^{-1}(E)) = 0$ whenever $m(E) = 0$. We also assume that $F(x, 0) = x$. The dynamics of each agent is then given by the following nonlinear discrete-time control system:

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

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

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

Suppose that every agent $k \in \{1, \ldots, N\}$ uses the same control law $u_n = v(x_n^k)$ at each time n; that is, the control law is independent of the agent identity k. In this case, by the *law of large numbers*, when $N \to \infty$, the empirical distribution $\frac{1}{N}\sum_{k=1}^{N}\delta_{\xi_n^k}$ converges to a deterministic quantity $\mu_n \in \mathcal{P}(\Omega)$, which evolves according to the following forward equation,

$$\mu_{n+1} = F^{\#}(\cdot, u_n)\mu_n, \quad \mu_0 \in \mathcal{P}(\Omega), \tag{2}$$

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

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

for each $E \in \mathcal{B}(\Omega)$. Since we are interested in the problem of stabilizing system (2) to a given target measure μ^d with density f^d, we must determine whether there exists a sequence of feedback laws u_n such that starting from any initial measure, the system (2) converges to μ^d. In general, this problem cannot be solved using deterministic feedback laws, as was shown in [23]. Therefore, in the next section we will construct a *stochastic* feedback law using a state-to-control transition kernel $K : \Omega \times \mathcal{B}(U) \to [0, 1]$. On a continuous state space, the transition kernel plays the role of the transition probability matrix on a discrete state space. That is, given that an agent is at state $x \in \Omega$, it chooses a subset of control inputs $W \subset U$ with probability $K(x, W)$. We note that deterministic control laws $v : \Omega \to U$ are a special type of stochastic control law in that $K(x, du) = \delta_{v(x)}$; that is, given the state x, the probability of choosing the control $v(x)$ is 1.

To achieve our goal of redistributing agents over Ω, we will construct K to be a function of the distribution μ; this dependence will be denoted as K_μ. For $\mu \in \mathcal{P}(\Omega)$, the transition kernel K_μ induces a nonlinear forward Kolmogorov operator $P_\mu : \mathcal{P}(\Omega) \to \mathcal{P}(\Omega)$, defined as

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

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

$$\mu_{n+1} = P\mu_n, \quad \mu_0 \in \mathcal{P}(\Omega). \tag{3}$$

Hence, taking the mean-field limit of the empirical distribution enables us to treat the N-agent system as a continuum, as described in the introduction.

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

$$\mathbb{P}(\xi_{n+1}^k \in E | \xi_n^k = x) = Q_\mu(x, E), \tag{4}$$

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

$$(P_\mu\mu)(E) = \int_{\Omega} Q_\mu(x, E)d\mu(x). \tag{5}$$

In this paper, instead of arbitrary measures in $\mathcal{P}(\Omega)$, we will consider those measures that have L^1 *densities* (derivatives with respect to m). By restricting P to this subset of $\mathcal{P}(\Omega)$, for $f \in L^1(\Omega)$ we can define a nonlinear operator $\widetilde{P}_f$ on $L^1(\Omega)$; the exact construction will be carried out in the next section. Then (3) can be rewritten as

$$f_{n+1} = \widetilde{P}_{f_n} f_n, \quad f_0 \in L^1(\Omega). \tag{6}$$

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

Problem III.1. *Let $\widetilde{P}_{f_n}$ be the forward operator induced by the operator P defined in* (5)*. Given a target distribution $\mu^d \in \mathcal{P}(\Omega)$ with density $f^d \in L^{\infty}(\Omega)$ and a non-singular continuous map $F : \Omega \times U \to \mathbb{R}^d$, determine whether there exists a transition kernel $Q_\mu : \Omega \times \mathcal{B}(\Omega) \to [0, 1]$ such that:*

1) *Equation* (6) *satisfies $\lim_{n\to\infty} \widetilde{P}_{f_n} \circ \ldots \circ \widetilde{P}_{f_1} \circ \widetilde{P}_{f_0} f_0 = f^d$ for all initial measures $f_0 \in L^1(\Omega)$, and*
2) *$\widetilde{P}_{f^d} = I$, where I is the identity operator.*

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

Having proven the existence of such an operator $\widetilde{P}_f$, in Section V we will introduce the system of N agents that evolve according to the N-agent Markov process that is an approximation of the mean-field model (3). Since P (via Q) can be constructed such that μ^d is an equilibrium of the system

(3), we observe that in simulations of the corresponding N-agent system, presented in Section VI, the empirical distribution $\frac{1}{N}\sum_{k=1}^{N}\delta_{\xi_n^k}$ converges to an empirical distribution that approximates f^d as $n \to \infty$.

IV. Stability Result

In this section, an operator $\widetilde{P}$ that solves Problem III.1 will be constructed. As stated in Problem III.1, $f^d \in L^\infty(\Omega)$ is the density of the target measure. In our previous work [7], we assumed that f^d is supported m almost everywhere on Ω; in this paper, we relax this assumption. The cost of this generality comes at the price of working with a nonlinear operator $\widetilde{P}$, which is also necessary to ensure that agent transitions between states stop once the equilibrium distribution is reached.

We begin by stating our **assumptions**.

1. We assume that Ω is a *path connected*, compact subset of $\mathbb{R}^d$. Path connectedness of Ω means that any two points $x, y \in \Omega$ can be connected by a *path* in Ω, which is a continuous map $p : [0,1] \to \Omega$ with $p(0) = x$, $p(1) = y$.

2. We also require Ω to satisfy the *cone condition* (Definition 4.6, [2]), which ensures that the boundary of Ω is regular enough. A domain $\mathcal{D}$ is said to satisfy the cone condition if there exists a cone $\mathcal{C}$ such that each $x \in \Omega$ is the vertex of a 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.

3. We require the measures in $\mathcal{P}(\Omega)$ to be absolutely continuous w.r.t m. This implies that its derivative (density) w.r.t. m must be in $L^1(\Omega)$.

4. The density f^d of the target distribution is in $L^\infty(\Omega)$.

5. As mentioned in Section III, we require F to be non-singular. Additionally, F must satisfy *Lusin's property* (see page 5).

6. Lastly, for the system (1) to be controllable, we need the following local controllability condition.

Definition IV.1. *The system* (1) *is said to be one-step locally controllable if there exists $r > 0$ such that, for every $x \in \Omega$, $B_r(x) \cap \Omega \subseteq F(x, U)$.*

In other words, a system is one-step locally controllable if from any given state $x \in \Omega$, the system can reach every other state (in Ω) that is within a radius r. From here on, we will consider r to be fixed as per this definition.

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

$$a_f(x) = \begin{cases} \frac{f(x)-f^d(x)}{f(x)} \text{ for } m\text{-a.e. } x \text{ if } f(x) - f^d(x) > 0; \\ 0 \quad \text{otherwise.} \end{cases} \tag{7}$$

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

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

$$k(x,u) \begin{cases} > 0 \text{ for } m\text{-a.e. } x \in \Omega, u \in U \text{ st. } F(x,u) \in \Omega; \\ = 0 \text{ otherwise;} \end{cases} \tag{8}$$

$$\int_U k(x,u)du = 1 \text{ for } m\text{-a.e. } x \in \Omega. \tag{9}$$

Before we proceed, we must determine whether we can construct a measurable $k \in L^\infty(\Omega \times U, m \times m)$ that satisfies these properties. We note that due to the first condition (8), the integral in the second condition (9) is computed over the set $U_x := F_x^{-1}(\Omega)$. This integral can therefore be expressed as $\int_U k(x,u)\chi_{U_x}(u)du = 1$. Since F_x is continuous, the set U_x is measurable for each x. The following lemma can be used to construct such a measurable k.

Lemma IV.2. *For $\forall x \in \Omega$, we have the following results:*

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

Proof. Result (1) is proved in Proposition V.2 of [7]. Proving this result requires the domain Ω to have a 'smooth enough' boundary, which is ensured by the cone condition.

To prove results (2) and (3), let $G = \{(x,u) \in \Omega \times U : F(x,u) \in \Omega\}$. G is Borel measurable because F is continuous in both variables. Since χ_G is a Borel measurable function, the *Tonelli theorem* [26] implies that $(\chi_G)_x$ is Borel measurable for each $x \in \Omega$. Since $(\chi_G)_x(u) = \chi_{U_x}(u)$, we have that $\chi_{U_x}(u)$ is a measurable function in both variables, proving result (2). Then, by the Tonelli theorem, we have that $x \mapsto \int_U (\chi_G)_x du$ is Borel measurable. Since $(\chi_G)_x(u) = \chi_G(x,u)$, we have that $\int_U (\chi_G)_x du = m(F_x^{-1}(\Omega)) = m(U_x)$. That is, $x \mapsto m(U_x)$ is Borel measurable. □

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

Next, we define a transition kernel that depends on the current distribution μ, $K_\mu : \Omega \times \mathcal{B}(U) \to [0,1]$. For $W \in \mathcal{B}(U)$,

$$K_\mu(x,W) = K_\mu(x, W \cap U_x) = K_\mu^1 + K_\mu^2, \quad \text{where} \tag{10}$$
$$K_\mu^1 = a_{f_\mu}(x)\int_W k(x,u)du,$$
$$K_\mu^2 = (1 - a_{f_\mu}(x))\delta_0(W).$$

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

Remark IV.3. *We note that the local controllability assumption, Definition IV.1, implies that there exists a measurable control $V(x) \in U$ such that $F(x, V(x)) = x$ (see Proposition V.13 in [7]). Therefore, the condition $F(x,0) = x$ is not restrictive. However, we impose this condition here for the sake of simplicity and note that we can extend our results even when this condition is not satisfied following the steps in [7].*

The proof of the next result follows from Lemma IV.2.

Lemma IV.4. *The kernel K_μ is well-defined. That is, $K_\mu(\cdot, W)$ is a measurable function on Ω for each fixed $W \in \mathcal{B}(U)$ and $K_\mu(x, \cdot)$ is a probability measure on U for each fixed $x \in \Omega$.*

Using K_μ, we define a closed-loop kernel $Q_\mu : \Omega \times \mathcal{B}(\Omega) \to [0, 1]$. The subscript μ in Q_μ is included to indicate that Q_μ is a function of μ, due to the dependence of K_μ on μ. For $E \in \mathcal{B}(\Omega)$,

$$
\begin{aligned}
Q_\mu(x,E) &= \int_U \chi_E(F(x,u))K_\mu(x,du) && (11)\\
&= a_{f_\mu}(x)\int_U \chi_E(F(x,u))k(x,u)du \ + \\
&\quad (1-a_{f_\mu}(x))\int_U \chi_E(F(x,u))d\delta_0 \\
&= Q^1_\mu + Q^2_\mu, \text{ where} && (12)\\
Q^1_\mu &= a_{f_\mu}(x)\int_U \chi_E(F(x,u))k(x,u)du,\\
Q^2_\mu &= (1-a_{f_\mu}(x))\delta_x(E).
\end{aligned}
$$

For the next result, we require F to satisfy Lusin's property [9] in both the x and u variables, which in simple terms means that F maps sets of measure zero to sets of measure zero. For $x \in \Omega$ fixed, we say that $F_x : (U, m) \to (\mathbb{R}^d, m)$ satisfies Lusin's property if $m(F_x(W)) = 0$ for every $W \in U$ with $m(W) = 0$. Lusin's property for F_u has a similar definition.

Lemma IV.5. *The kernel Q_μ is well-defined; that is, $Q_\mu(\cdot, E)$ is a measurable function on Ω for each $E \in \mathcal{B}(\Omega)$ and $Q_\mu(x, \cdot)$ is a probability measure on Ω for each $x \in \Omega$. Further, if F satisfies Lusin's property, then Q^1_μ is regular.*

Proof. The proof that Q_μ is well-defined is similar to the proof that K is well-defined (Lemma IV.4). To prove that Q^1_μ is regular, we first require that $Q^1_\mu(x, \cdot) \ll m$ for every x. Indeed, if $E \in \mathcal{B}(\Omega)$ is such that $m(E) = 0$, then due to the non-singularity of F with respect to both variables x and u, we have that $(m \times m)(F^{-1}(E)) = 0$. Therefore, for $x \in \Omega, u \in U$, we have that $\chi_E(F(x, u)) = \chi_{F^{-1}(E)}(x, u) = 0$ in the integral that defines Q^1_μ . Hence, $Q^1_\mu(x, E) = 0$.

The proof that Q^1_μ has a kernel function $q \in L^\infty(\Omega \times \Omega)$ is given in Proposition IV.4 of [7]. □

Next, we define a nonlinear operator $P_\mu : \mathcal{P}(\Omega) \to \mathcal{P}(\Omega)$ in terms of Q_μ as follows:

$$
\begin{aligned}
(P_\mu\mu)(E) &= \int_\Omega Q_\mu(x,E)d\mu(x) && (13)\\
&= \int_\Omega\int_U \chi_E(F(x,u))K_\mu(x,du)d\mu(x)\\
&= \int_\Omega\int_U a_{f_\mu}(x)\chi_E(F(x,u))k(x,u)du \ d\mu(x) \ +\\
&\quad \int_\Omega (1-a_{f_\mu}(x))\delta_x(E)d\mu(x)\\
&= \int_\Omega\int_U a_{f_\mu}(x)\chi_E(F(x,u))k(x,u)du \ d\mu(x) \ +\\
&\quad \int_E (1-a_{f_\mu}(x))d\mu(x). && (14)
\end{aligned}
$$

Due to the properties of Q_μ (Lemma IV.5), we immediately have the following lemma.

Lemma IV.6. *Operator P_μ preserves $\mathcal{P}(\Omega)$, and furthermore, it preserves absolutely continuous measures.*

By restricting P_μ to those measures that are absolutely continuous w.r.t m, that is, measures that have L^1 densities, we can define $\widetilde{P}_{f_\mu} : L^1(\Omega) \to L^1(\Omega)$. The next few steps will be toward this effort. We note that the operator $\widetilde{P}_{f_\mu}$ is a function of f_μ through the dependence of Q_μ on μ.

Lemma IV.6 implies that $P_\mu\mu \ll m$; let $\widetilde{P}_{f_\mu} f_\mu$ be the density, that is, for $E \in \mathcal{B}(\Omega)$,

$$(P_\mu\mu)(E) = \int_E (\widetilde{P}_{f_\mu} f_\mu)(y)dy.$$

We note that since Q^1_μ is regular, there must exist a function $q \in L^\infty(\Omega \times \Omega)$. Therefore, from (12), we have that

$$Q^1_\mu(x, E) = \int_E q(x, y)dy = \int_U \chi_E(F(x, u))K^1_\mu(x, du).$$

Using this expression, (14) can be rewritten as follows. For $E \in \mathcal{B}(\Omega)$,

$$
\begin{aligned}
(P_\mu\mu)(E) &= \int_\Omega\int_E a_{f_\mu}(x)q(x,y)dy \ f_\mu(x)dx \ +\\
&\int_E (1-a_{f_\mu}(x))f_\mu(x)dx = \int_E (\widetilde{P}_{f_\mu} f_\mu)(y)dy.
\end{aligned}
$$

Applying Fubini's theorem [26] to the equation above, we obtain an expression for an operator $\widetilde{P}$ defined on $L^1(\Omega)$ as follows. For $f \in L^1(\Omega)$,

$$
\begin{aligned}
\widetilde{P}_f &= \widetilde{P}^1_f + \widetilde{P}^2_f, \text{ where} && (15)\\
(\widetilde{P}^1_f f)(y) &= \int_\Omega a_f(x)q(x,y)f(x)dx,\\
(\widetilde{P}^2_f f)(y) &= (1-a_f(y))f(y).
\end{aligned}
$$

The operator $\widetilde{P}_f$ preserves $L^1(\Omega)$, stated in the proposition below.

Proposition IV.7. *For $f \in L^2(\Omega) \subset L^1(\Omega)$ we have that*

1) *$\widetilde{P}_f \in L^1(\Omega)$, i.e. $\widetilde{P}_f$ is well-defined. Moreover, $\widetilde{P}_f$ preserves probability densities; in other words, it is a Markov operator [32].*
2) *In fact, $\widetilde{P}_f : L^2(\Omega) \to L^2(\Omega)$ is well-defined.*

The second result above is a consequence of Proposition II.4.7 of [15], detailed in Proposition IV.4 of [7]. We will require this result in Section V.

Note that for each fixed f, $\widetilde{P}_f$ is a linear operator. The next result will be used in Section V. In the following result, $\mathbb{B}(L^2(\Omega))$ stands for the set of bounded linear operators on $L^2(\Omega)$.

Lemma IV.8. *The map from $L^2(\Omega) \to \mathbb{B}(L^2(\Omega))$, defined as $f \mapsto \widetilde{P}_f$, is uniformly bounded; that is, for every $f \in L^2(\Omega)$, $\|\widetilde{P}_f\| \le C$ for some $C > 0$. Moreover, this result also holds true for $\widetilde{P}_f$ as an operator on $L^1(\Omega)$.*

Proof. This follows from the fact that $\widetilde{P}_f$ depends on f through the a_f function, which is in L^∞ for any $f \in L^1(\Omega)$

or $L^2(\Omega)$. An application of Theorem 6.18 in [26] then proves the result for $\widetilde{P}_f^1$. The result holds true trivially for $\widetilde{P}_f^2$, since it is a *multiplication operator*. □

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

Next, we will show that f^d is a globally asymptotically stable equilibrium of system (6).

Theorem IV.9. *For the system* (6)*, f^d is globally asymptotically stable in the $L^1(\Omega, m)$ norm.*

For the proof of this theorem, we need the following lemma.

Lemma IV.10. *At any given time n, if $y \in \Omega$ is such that $f_n(y) \leq f^d(y)$, then $f_{n+1}(y)$ monotonically increases with n.*

Proof. Consider the case when for some $y \in \Omega$, $f_n(y) > f^d(y)$. Then, it follows that $a_{f_n}(y) > 0$. Expression (15) then becomes:

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

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

$$\begin{aligned} f_{n+1}(y) \geq f_n(y) > f^d(y); \\ f_n(y) > f_{n+1}(y) \geq f^d(y). \end{aligned} \tag{16}$$

Consequently, it is not possible that $f_{n+1}(y) < f^d(y)$ for any value of n. Next, consider the case when $y \in \Omega$ is such that $f_n(y) \leq f^d(y)$. In this case, $a_{f_n}(y) = 0$. Expression (15) then reduces to:

$$\widetilde{P}_{f_n} f_n(y) = \int_\Omega a_{f_n}(x) k(x,y) f_n(x) dx + f_n(y). \tag{17}$$

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

$$\begin{aligned} f_{n+1}(y) \geq f^d(y) > f_n(y); \\ f^d(y) > f_{n+1}(y) \geq f_n(y). \end{aligned} \tag{18}$$

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

Define the sets

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

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

We can now state the proof of Theorem IV.9. The proof employs an argument by contradiction that if the density f_n converges to a function other than f^d, then the measure μ_n is pushed from sets where its density f_n is greater than f^d to sets where $f_n < f^d$. This is straightforward to conclude from the definitions of the transition kernels K_μ and Q_μ; however, to prove the convergence of f_n to f^d, it is necessary to precisely quantify the measure that is pushed during each time step, which is computed in the proof.

Proof of Theorem IV.9. By construction, for any n we have that $\Omega = E_n^1 \sqcup E_n^2 \sqcup E_n^3$; moreover, these sets do not intersect one another. We also have that each f_n is a probability density on Ω, and hence must integrate to 1 over $E_n^1 \sqcup E_n^2 \sqcup E_n^3$. By definition, on E_n^2, $f_n = f^d$. Therefore, to prove this result, it is sufficient to show that on the set E_n^1, $\|f_n - f^d\|_1 \to 0$ as $n \to \infty$, as this would imply that on Ω, $\|f_n - f^d\|_1 \to 0$ as $n \to \infty$.

On E_n^1, by (18), we have that $f_{n+1} \geq f_n$, and hence that $f^d - f_n \geq f^d - f_{n+1}$. Set $F_n = (f^d - f_n)^+$, where for an arbitrary function $h : \mathbb{R}^d \to \mathbb{R}$, h^+ denotes the positive part of h. Then F_n is monotonically decreasing on Ω. The sequence $(F_n)_n$ is bounded, and monotonically decreasing, which implies that F_n converges pointwise to a function, say g. By the *monotone convergence theorem* [26], we then have that $\int_\Omega F_n \to \int_\Omega g$. If $g = 0$, then we have our result. If $g \neq 0$, then since f_n is a probability density on Ω, $\int_\Omega F_n \to \int_\Omega g$ implies that $\int_\Omega (f_n - f^d)^+ \not\to 0$. We will next prove by contradiction that g is in fact 0.

We suppose that $g \neq 0$. Let $\int_\Omega g \geq \gamma$, where $\gamma > 0$. Define $S = \{x \in \Omega : g(x) > 0\}$. We note that the definition of S is independent of time. Given the conditions in (16) and (18), it follows that $E_n^1 \supset E_{n+1}^1$ for all n. Due to the convergence of F_n to g, we must have that for all n, $S \subset E_n^1$. Moreover, $\lim_{n\to\infty} m(E_n^1) \to m(S)$. Note that,

$$\int_S f^d(x) - f_n(x) dx \geq \int_S g(x) dx > \gamma. \tag{19}$$

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

$$\begin{aligned} m(B(x_j)) = \; & m(B(x_j) \cap S) + m(B(x_j) \cap (E_n^1 \setminus S)) \\ & + m(B(x_j) \cap (E_n^1)^c). \end{aligned} \tag{20}$$

Let $m(B(x_j) \cap S) \geq \epsilon_0$, for some $\epsilon_0 > 0$. If $m(B(x_j) \cap (E_n^1)^c) = 0$ at the current time n, then we look for a large enough time $T \in \mathbb{Z}_+$ such that $m(E_T^1 \setminus S)) \leq \epsilon_1 << \epsilon_0$. At times $n \geq T$, (20) shows that $m(B(x_j) \cap (E_n^1)^c) > 0$, ensuring the existence of at least one ball from the cover that has intersections of positive measure with both S and $(E_n^1)^c$.

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

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

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

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

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

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

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

Consequently, for each t_n we have,

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

which implies that

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

As $n \to \infty$, the integral on the right-hand side of the equation above tends to $-\infty$, contradicting the fact that this integral is an upper bound on the integral of g over S, as per (19). Thus, we must have that $\int_{N_1\cap(E_n^1)^c} f_n - f^d \to 0$ as $n \to \infty$.

We will now use an induction argument to show that $\int_{(E_n^1)^c} f_n - f^d \to 0$. We have just shown that this was true for the neighborhood of S given by $N_1 \cap (E_n^1)^c$. We assume that $\int_{N_k\cap(E_n^1)^c} f_n - f^d \to 0$ for some $k > 1$. We will prove that this also holds true for $N_{k+1} \cap (E_n^1)^c$. Suppose that it is not true; then, $\int_{N_{k+1}\cap(E_n^1)^c} f_n - f^d \geq \delta_1$ *i.o* for some $\delta_1 > 0$. Again, denote the sequence of times when this happens by $\{t_n\}_n$. By construction, N_{k+1} does not intersect S; however, N_{k+1} may intersect E_n^1 (possibly a subset of N_k), to which it can push measure. We now demonstrate that N_{k+1} pushes most of its measure to $N_k \cap (E_n^1)^c$. We have established that for any $n \geq T$, $m(N_k \cap E_n^1) \leq m(E_n^1 \setminus S) \leq \epsilon_1$, which is arbitrarily small. Hence, $m(N_k \cap E_n^1)$ must be arbitrarily small, and therefore $m(N_k \cap (E_n^1)^c)$ must have positive measure. Consequently, we have that

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

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

V. The N-Agent System

In this section, we will define the microscopic description of the system, i.e., the model of individual agents' state transitions, and study how it relates to the macroscopic or mean-field model (3). The following mathematical definitions are adapted from [16].

Consider a population of N agents evolving on the state space Ω. Let the state of each agent k at time n be given by the random variable $\xi_n^k \in \Omega,\ k = \{1,\dots,N\}$. Each agent transitions between states on Ω according to the transition kernel Q_μ defined in (12). The N-agent system can therefore be described as a Markov chain $\xi_n = (\xi_n^1,\dots,\xi_n^N)$ with state space Ω^N. To a measure $\nu \in \mathcal{P}(\Omega)$, we associate a measure $\nu^{\otimes N} = \nu \times \ldots \times \nu \in \mathcal{P}(\Omega^N)$. The empirical measure $m(x)$ associated with the point $x = (x^1,\dots,x^N) \in \Omega^N$, where each entry x^k is the state of agent k, is given by a normalized sum of Dirac measures associated with each agent,

$$m(x) = \frac{1}{N}\sum_{k=1}^N \delta_{x^k}. \tag{21}$$

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

$$\begin{aligned}\mathbf{P}(\xi_0 \in dx) &= \mu_0^{\otimes N}(dx),\\ \mathbf{P}(\xi_n \in dx|\xi_{n-1} = z) &= (P_m m(z))^{\otimes N}(dx),\end{aligned} \tag{22}$$

where $dx = dx^1 \times \ldots \times dx^N$ and P_m is as defined in (14). However, unlike (14) m does not have a density, therefore m must be converted to an absolutely continuous measure so that the operator P_m makes sense, we will do so later in this section. At time $n = 0$, the N-agent system can be modeled as N independent random variables $\xi_0^1,\dots,\xi_0^N$ with common distribution μ_0. At time $n \geq 1$, define $\mu_n^N := m(\xi_n)$. Then, μ_{n+1}^N is evaluated as

$$\mu_{n+1}^N = P_m m(\xi_n). \tag{23}$$

Thus, from the equation above, at time n the N-agent system is modeled as N random variables $\xi_n^1,\dots,\xi_n^N$ that are conditional on ξ_{n-1} and distributed according to $P_m m(\xi_{n-1})$. The agents' states are therefore not independent of one another; their distribution is dependent on the system configuration at time $n-1$. Although the evolution of each agent's state is not Markovian, the distribution of the N-agent system evolves according to an *interacting* Markov chain. At time $n = 0$, $\mu_0^N \to \mu_0$ as $N \to \infty$. At times $n \geq 1$, due to the aforementioned interaction between agents, the *law of large numbers* does not apply. Thus, another method must be used to establish the limit $\mu_n^N \overset{N\to\infty}{\longrightarrow} \mu_n$, where μ_n evolves according to (3). This limit is called the *mean-field limit*. The work [16] proved this limit for systems of the form (3) in which the right-hand side is continuous. In [31], this limit is referred to as the *dynamic law of large numbers*; it is proven for Markov processes whose evolution is governed by a partial differential equation (PDE). A comprehensive survey of mean-field approximations of both discrete-time and continuous-time Markov chains is given in [11].

Since the empirical measure m is a sum of Dirac measures, it is not absolutely continuous with respect to the Lebesgue measure. We will "mollify" the Dirac measures in order to be able to use results from the previous section and to apply the operators $\widetilde{P}_f$ and P_μ defined in (14) and (15), respectively, to absolutely continuous measures. Mathematically, this means that the measure m is convolved with a smooth function $\phi : \mathbb{R}^d \to \mathbb{R}$, a *mollifier*, to obtain a smooth function (density). The convolution of m and ϕ is carried out as

$$\phi * m = \int_\Omega \phi(x) dm = \frac{1}{N} \sum_{i=1}^{N} \phi(x - x^k). \tag{24}$$

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

$$\phi(x) = \begin{cases} e^{-\left(\frac{1}{1-\|x\|^2}\right)}, & x \in (-1,1), \\ 0, & \text{otherwise.} \end{cases} \tag{25}$$

To change the support of ϕ, we define a function ϕ_h on $\mathbb{R}^d$ for some $h > 0$ as follows [26]:

$$\phi_h(x) = h^{-d} \phi\left(\frac{x}{h}\right). \tag{26}$$

Note that $\int \phi_h = 1$, which is independent of h. Moreover, the "mass" of ϕ_h becomes concentrated at the origin as $h \to 0$; that is, ϕ_h tends to a Dirac measure as $h \to 0$. Figure 1 shows a visualization of two bump functions with $h = 0.1$ and one with $h = 0.05$. Since the integral of all bump functions is 1, to compensate for the decrease in h, the peak of the bump function with $h = 0.05$ is significantly higher than the peaks of the functions with $h = 0.1$.

The introduction of the mollifier also has implications for the implementation of the N-agent system in practice. For an agent with state x, given a distribution f, the transition kernel K_μ in (10) is defined such that it requires pointwise evaluation of the function $f(x)$ in the term $a_f = (f - f^d)/f$ from (7). However, to evaluate the density ϕ at its state x using (24), the agent must know the states x_i of all other agents whose states are within a distance h of its own. For example, if the agents' states are their positions in space, mollification of the empirical measure implies that each agent must estimate the density ϕ in (24) based on its relative distance to all agents that are located within a neighborhood whose size is determined by the parameter h. As $h \to 0$, this neighborhood shrinks, and the density tends to the Dirac measure, which is singular. Note that here, we assume ideal sensing in order to simplify the analysis. More realistic descriptions of the N-agent system should include models of sensor noise.

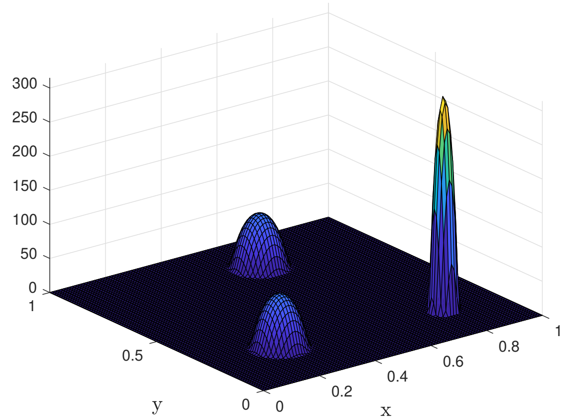

Fig. 1. Visualization of two bump functions ϕ_h with $h = 0.1$ and one bump function with $h = 0.05$.

In order to derive the macroscopic (mean-field) model from the microscopic description of the system, i.e. the dynamics of N individual agents, one typically needs to take the mean-field limit, as described earlier in this section. Since we have introduced the mollifier, a second limit needs to be proven as well. Both limits are defined below.

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

$$\mu_{n+1}^h = P_{\phi_h * \mu_n^h} \mu_n^h, \quad \mu_0^h \in \mathcal{P}(\Omega). \tag{27}$$

Due to the introduction of the mollifier, we expect the N-agent system (23) to converge to the system above, which is different from (3). That is,

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

This limit is usually proven in the *weak topology* and can be established for discrete-time systems using results from [16]. Applying these results requires proving that the right-hand side of (3) is continuous in the weak topology, which is significantly challenging for our system. Thus, we will reserve this investigation for future work.

2) $h \to 0$: The second limit proves that the solution of (27) converges to the solution of (3); that is, for all $n \in \mathbb{Z}_+$,

$$\mu_n^h \to \mu_n \quad \text{as} \quad h \to 0. \tag{28}$$

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

A. The limit as $h \to 0$

We prove the limit (28) for a dense subset of $L^1(\Omega)$; specifically, we consider distributions $\mu \in \mathcal{P}(\Omega)$ that have $L^2(\Omega)$ densities. Moreover, we require f^d to be bounded from below a.e. on Ω.

Let $\mu_0 \ll m$ with density $f_0 \in L^2(\Omega)$. In Proposition IV.7, we proved that $\widetilde{P}_f$ preserves $L^2(\Omega)$; that is, $f_n = \widetilde{P}_{f_{n-1}} \circ \ldots \circ$

$\widetilde{P}_{f_1} \circ \widetilde{P}_{f_0} f_0 \in L^2(\Omega)$ for $n \in \mathbb{Z}_+$ Therefore, system (27) can be rewritten on $L^2(\Omega)$ as

$$f_{n+1}^h = \widetilde{P}_{\phi_h * f_n^h} f_n^h, \quad f_0^h \in L^2(\Omega). \tag{29}$$

Since $m(\Omega) < \infty$, $L^2(\Omega) \subset L^1(\Omega)$, and therefore we will consider system (6) to be a system on $L^2(\Omega)$ instead of $L^1(\Omega)$. We will show that solutions of the above system converge to those of (6) in the $L^1(\Omega)$ norm.

Theorem V.1. *Suppose the initial condition f_0 be in $L^2(\Omega)$. Let f_n^h and f_n be solutions of* (29) *and* (6)*, respectively. If f^d is bounded from below a.e. on Ω, then*

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

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

To prove this result, we need the following proposition, whose proof is given in the Appendix.

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

1) *For $f \in L^2(\Omega)$,*

$$\|\widetilde{P}_{\phi_h * f} g - \widetilde{P}_f g\|_1 \xrightarrow{h \to 0} 0$$

2) *If $f_i \xrightarrow{i \to \infty} f$ in the $L^1(\Omega)$ norm, then*

$$\|\widetilde{P}_{f_i} g - \widetilde{P}_f g\|_1 \xrightarrow{i \to \infty} 0$$

We can now prove Theorem V.1.

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

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

Then, by statement (1) of Proposition V.2, $\|f_1^h - f_1\|_1 = \|\widetilde{P}_{\phi_h * f_0} f_0 - \widetilde{P}_{f_0} f_0\|_1 \to 0$ as $h \to 0$. Assume that this is true for some $n > 1$; i.e., $\|f_n^h - f_n\|_1 \to 0$ as $h \to 0$. We will show that this limit holds true for $n+1$ using the following computation:

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

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

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

This follows from the fact that the inner limit tends to $\widetilde{P}_{f_n^h}$ by statement (1) of Proposition V.2, and the outer limit $\lim_{f_n^h \to f_n} \widetilde{P}_{f_n^h}$ tends to $\widetilde{P}_{f_n}$ by statement (2) of Proposition V.2. Therefore, the bracket $\widetilde{P}_{\phi_h * f_n^h} - \widetilde{P}_{f_n}$ in the first term tends to 0 as $h \to 0$, and hence we have our result. □

VI. Simulations

In this section, we present numerical solutions of the mean-field model (3) and simulations of the corresponding N-agent system. We provide verification via these simulations that as $N \to \infty$, the simulations of the N-agent system (stochastic simulations) approach the solution of the deterministic system (3).

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

$$\begin{aligned}
x_{n+1}^1 &= x_n^1 + u_n^1 \cos(u_n^2), \\
x_{n+1}^2 &= x_n^2 + u_n^1 \sin(u_n^2).
\end{aligned} \tag{30}$$

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

To simulate the mean-field model (6), we need to discretize both Ω and U. The set Ω is partitioned into $n_x \in \mathbb{Z}_+$ sets, $\widetilde{\Omega} = \{\Omega_1, \ldots, \Omega_{n_x}\}$, where $\Omega = \cup_{i=1}^{n_x} \Omega_i$ and the sets Ω_i have intersections of zero Lebesgue measure. The set of control inputs U is approximated as a set of $n_u \in \mathbb{Z}_+$ discrete elements, $\widetilde{U} = \{v_1, \ldots, v_{n_u}\}$, where $v_i \in U$ for each i. Define index sets $\mathcal{I} = \{1, \ldots, n_x\}$ and $\mathcal{J} = \{1, \ldots, n_u\}$. Using these definitions, 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 state of this Markov chain to be i. We use a modified version of *Ulam's method* [18] to construct this approximation. In the uncontrolled setting, Ulam's method is a classical technique for constructing approximations of the *pushforward map* (Perron-Frobenius operators) induced by dynamical systems. Let p_{ij}^l denote the probability of the system state being in the set Ω_j in the next time step, given that the system state is uniformly randomly distributed over the set Ω_i and the selected control input is v_l. To obtain p_{ij}^l via the modified Ulam's method, we assume that a fixed number of agents, say M, are uniformly distributed over Ω_i. For each agent $m \in \{1, \ldots, M\}$ with state $x_m \in \Omega_i$, we compute $F(x_m, v_l)$. Then, we define the transition probabilities of the approximating controlled Markov chain as follows:

$$p_{ij}^l = \frac{\left|\{y \in \Omega_j : x_m = F_l^{-1}(y),\ m = 1, \ldots, M\}\right|}{\left|\{y \in \Omega : x_m = F_l^{-1}(y),\ m = 1, \ldots, M\}\right|},$$

where $F_l(\cdot) = F(\cdot, v_l)$. We next define an equivalent of the state-to-control transition kernel K. Let $\tilde{k}_{il}$ be the probability of choosing the control variable v_l, given that the system state x_m is in Ω_i. We set $\tilde{k}_{il} > 0$ if for some m, $F(x_m, v_l) \in \Omega$, while ensuring that $\tilde{k}_{il}$ is a probability.

We now define the discretization of the mean-field model (3). Let $\mu \in \mathcal{P}(\widetilde{\Omega})$ and $j \in \mathcal{I}$, and let μ^d be the discretization of f^d on $\tilde{\Omega}$. Let $\mathbf{P} \in \mathbb{R}^{n_x \times n_x}$ be the discretization of the operator P_μ defined in (14). Then the discretization of system (3) is given by:

$$\mu_{n+1} = \mathbf{P}\mu_n, \tag{31}$$
$$\mathbf{P}\mu(j) = \sum_{i\in\mathcal{I}} a_\mu(i) \sum_{l\in\mathcal{J}} \tilde{k}_{il}\; p^l_{ij}\mu(i) + (1 - a_\mu(j))\mu(j),$$

where $a_\mu(i) = (\mu(i) - \mu^d(i))/\mu(i)$ if $\mu(i) - \mu^d(i) > 0$, and $a_\mu(i) = 0$ otherwise. Figure 3 shows snapshots of the simulation of this system at several times n.

Algorithm 1 Simulation of N agents

1: **Input**: $\Omega, U, k, F, N, f^d, h, T_f$
2: **Initialize** $n = 0, a^k = 0, x^k_0 \in \Omega$ for all $k = 1, \dots, N$
3: **while** $n \le T_f$ **do**
4:     **for** $k = 1 : N$ **do**
5:         $y = x^k_n$ ▷ Current location of agent k
6:         $s = 0$
7:         **for** all $j \in \mathcal{N}(k)$ **do**
▷ $\mathcal{N}(k) :=$ {agents within distance h of k}
8:             $z = x^j_n$
9:             $s = s + \mathrm{PHI}(y, z, h)$
10:         **end for**
11:         $f_n(y) = \frac{1}{|\mathcal{N}(k)|} s$
12:         **if** $f_n(y) > f^d(y)$ **then**
13:             $a^k = \frac{f_n(y) - f^d(y)}{f_n(y)}$
14:         **end if**
15:         **if** $a^k > 0$ **then**
16:             Draw v uniformly from $(0, 1)$
17:             **if** $v \le a^k$ **then**
18:                 Draw $u \sim k(y, \cdot)$ from U
19:                 $y = F(y, u)$
20:             **end if**
21:         **end if**
22:         $x^k_{n+1} = y$
23:     **end for**
24:     $n = n + 1$
25: **end while**

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

Algorithm 1 presents the pseudocode that simulates the evolution of agents over a domain Ω with a control set U, until a specified final time T_f. An agent considers another agent to be its neighbor if their relative distance is less than h, the parameter of the bump function ϕ_h described in the previous section. We denote the set of neighbors of agent k at any given time by $\mathcal{N}(k)$. At every time step, each agent computes the value of the bump function based on the relative distances of its neighbors. Note that C in Line 31 is a normalizing constant which is chosen to ensure that Φ is a probability density. Figures 4-7 show snapshots of the N-agent simulation for agent population sizes of N =100, 500, and 1000, with $h = 0.1$ in the first three figures and 0.05 in the last figure.

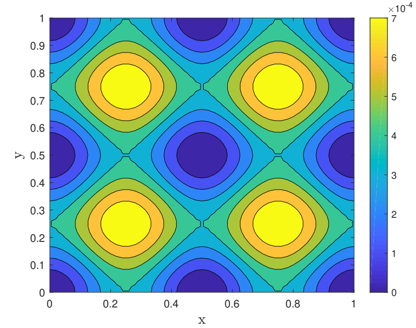

Fig. 2. Target distribution f^d

We first investigate the effect of increasing N while keeping h fixed on the time evolution of the simulated N-agent system. Figure 3 shows that as time n increases, the mean-field model indeed converges asymptotically to the target distribution in Fig. 2. We observe that the convergence slows down significantly after time $n = 500$. Following our discussion in Section V, we expect the stochastic simulations of the N-agent system to converge to the discretization of system (27) in the limit $N \to \infty$. Although system (27) is different from the system (3), note that the solutions of the two systems (3) and (27) converge in the limit $h \to 0$. The snapshots in Figs. 4-6 show that as the population size N is increased with a fixed value of h, the agent distribution in the N-agent simulation approaches the solution of system (31), plotted in Fig. 3. In all three figures, the agent distribution converges to a discrete approximation of the continuous target distribution.

Next, we study the effect of N on the frequency of agent state transitions. For each of the N-agent simulations shown in Figs. 4-7, Figs. 8a-8d plot the time evolution of the 2-norm of five randomly selected agents' states. Figs. 8a-8c show that the agents' frequency of state transitions significantly decreases with increasing N; the agents eventually stop transitioning between states (i.e., stop moving) for both $N = 500$ and $N = 1000$. This trend can be attributed to our approximation of a continuous distribution by a discrete function representing the state of the N-agent system. For low values of N, the resulting coarse discretization of f^d might yield an operator $\widetilde{P}_{\phi_h * \mu^N_n}$ that is not a sufficiently accurate approximation of $\widetilde{P}_{f^d} = I$, the condition that stops agent state transitions. Higher values of N produce a finer discretization of f^d, which improves the accuracy of the approximation of $\widetilde{P}_{f^d} = I$. This validates our claim that control policies designed for the mean-field model can be implemented on a population of individual agents to achieve a target distribution, as long as this population is sufficiently large.

Finally, we investigate the effect of changing h while keeping N fixed. Similar to the agent distribution in Fig. 6 ($h = 0.1$), the agent distribution in Fig. 7 ($h = 0.05$) approaches the solution of (31) shown in Fig. 3 as N increases. However, the relative closeness of the distributions in Figs. 6

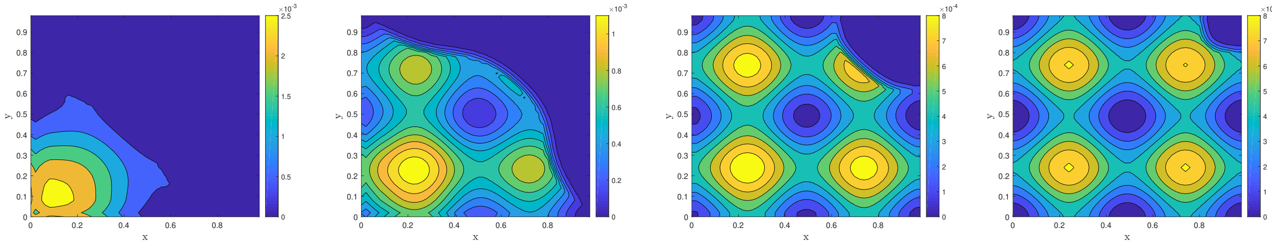

(a) $n = 50$ (b) $n = 200$ (c) $n = 500$ (d) $n = 2000$

Fig. 3. Snapshots of the simulation of system (31) at several times n.

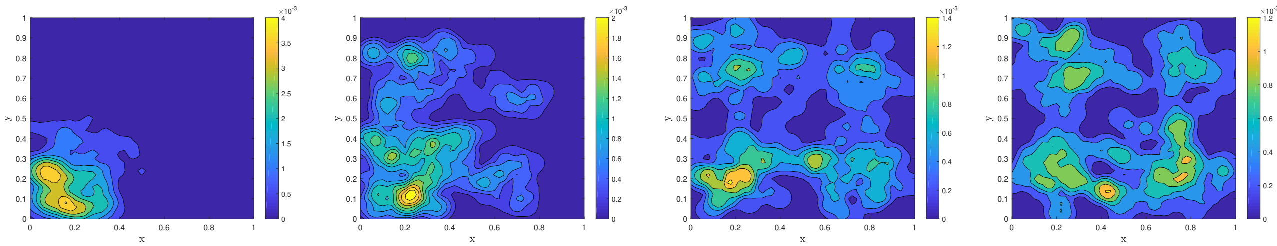

(a) $n = 50$ (b) $n = 200$ (c) $n = 500$ (d) $n = 2000$

Fig. 4. Snapshots of a stochastic simulation of $N = 100$ agents, with $h = 0.1$, at several times n.

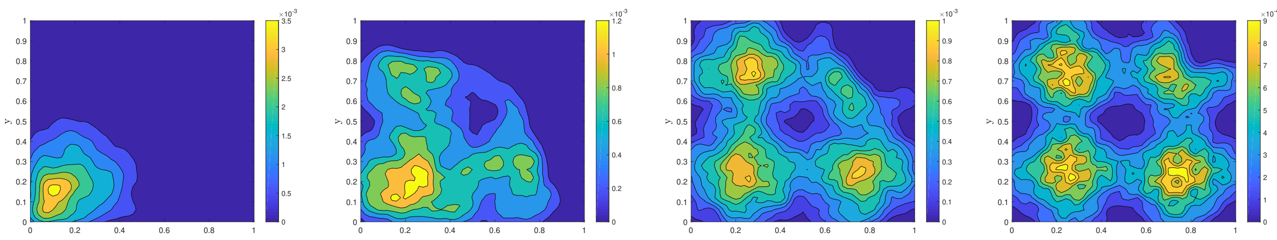

(a) $n = 50$ (b) $n = 200$ (c) $n = 500$ (d) $n = 2000$

Fig. 5. Snapshots of a stochastic simulation of $N = 500$ agents, with $h = 0.1$, at several times n.

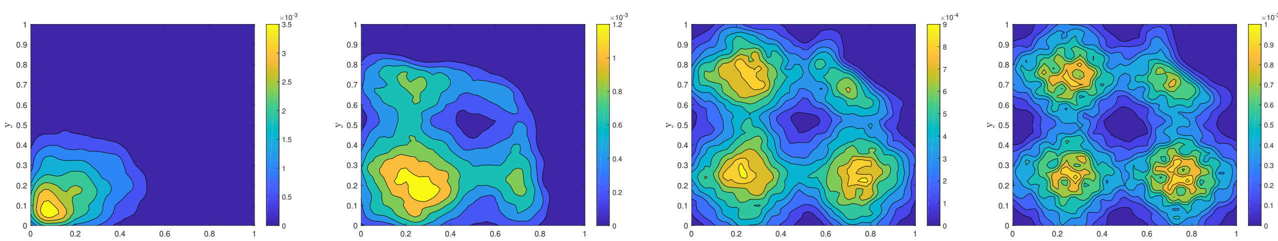

(a) $n = 50$ (b) $n = 200$ (c) Density at time $n = 500$ (d) $n = 2000$

Fig. 6. Snapshots of a stochastic simulation of $N = 1000$ agents, with $h = 0.1$, at several times n.

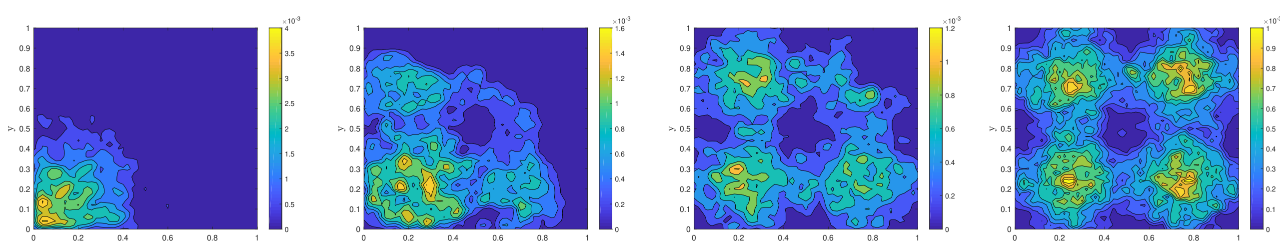

(a) $n = 50$ (b) $n = 200$ (c) $n = 500$ (d) $n = 2000$

Fig. 7. Snapshots of a stochastic simulation of $N = 1000$ agents, with $h = 0.05$, at several times n.

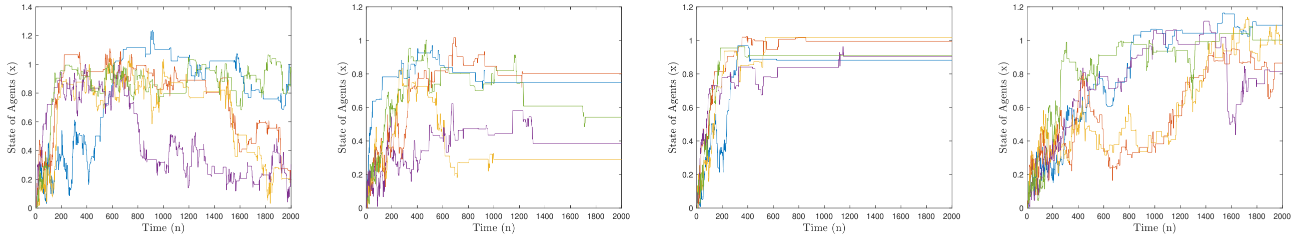

(a) $N = 100$, $h = 0.1$ (b) $N = 500$, $h = 0.1$ (c) $N = 1000$, $h = 0.1$ (d) $N = 1000$, $h = 0.05$

Fig. 8. Time evolution of the 2-norm of five randomly selected agents' states in each of the N-agent simulations.

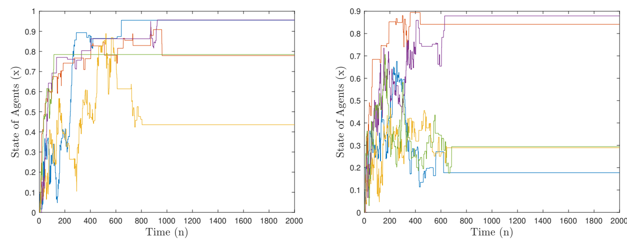

(a) $N = 100$, $h = 0.25$ (b) $N = 500$, $h = 0.2$

Fig. 9. Time evolution of the 2-norm of five randomly selected agents' states in two N-agent simulations with different values of N and h (snapshots of corresponding stochastic simulations not shown).

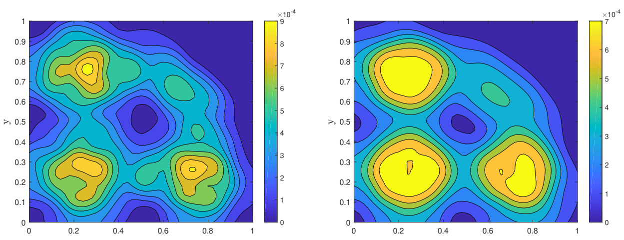

(a) $h = 0.2$ (b) $h = 0.25$

Fig. 10. Snapshots at time $n = 2000$ of stochastic simulations of $N = 1000$ agents with different values of h.

and 7 to the distribution in Fig. 3 is not apparent from the figures. This can be explained by noting that in this case, we are holding N constant and decreasing h, thereby reversing the order of limits that we considered in Section V. There is no mathematical guarantee that the limits commute, and hence, we do not necessarily expect that with reduced h, the N-agent simulations will more closely approach the solution of (31). Moreover, a lower value of h for a fixed N yields a smaller neighborhood in which each agent evaluates the local density, which can produce a less accurate approximation of $\widetilde{P}_{f^d} = I$. As explained previously, this can result in persistent agent state transitions, which are evident in the simulation of $N = 1000$ agents when h is reduced from 0.1 (Fig. 8c) to 0.05 (Fig. 8d). Increasing h, on the other hand, can result in the eventual cessation of agent transitions in smaller agent populations N. This is demonstrated in Figs. 9a-9b, which show that when h is increased from 0.2 to 0.25, the population of N agents stops transitioning for a lower value of N. The snapshots of stochastic simulations for $N = 1000$ at time $n = 2000$ in Figs. 6d, 7d, 10a, and 10b demonstrate that the agent distribution becomes smoother as h is increased, due to the smoothening effect of the mollification.

VII. Conclusion

In this paper, we have used a discrete-time mean-field model describing the state dynamics of a multi-agent system to design decentralized state-feedback agent control laws that drive the agents asymptotically to a target state distribution. To implement the control laws, the agents only require knowledge of the local agent density; for example, the density of agents within their sensing range. The mean-field model considered here is the forward Kolmogorov equation of a discrete-time Markov process that can be stabilized to an arbitrary distribution that has $L^\infty(\cdot)$ derivatives. Moreover, the Markov process can be constructed such that its forward operator is the identity operator at the desired distribution. This prevents agents from switching between states once the equilibrium distribution is reached. Although stability and convergence results were proven for the mean-field model, simulations of the corresponding N-agent system demonstrate that for relatively small numbers of agents ($N \gtrsim 500$), the agents indeed redistribute themselves to the target distribution and thereafter cease switching between states. Our use of density-dependent feedback control laws enables us to specify a more general class of target distributions than in our prior works [7], [8], in which we considered only open-loop control laws. In the future, we would like to establish the mean-field limit of the system considered in this paper, as well as extend our results to swarms of agents governed by M-step controllable dynamical models (where $M > 1$).

Appendix A
Proof of Proposition V.2

Here, we prove the two convergence results stated in the proposition.

(1) Let $f \in L^2(\Omega)$. Then, $\phi_h * f \in L^2(\Omega)$. By Theorem 8.14 of [26], $\phi_h * f \xrightarrow{h \to 0} f$ in the L^2 norm. To prove convergence of $\widetilde{P}_{\phi_h * f}$ to $\widetilde{P}_f$ as operators on $L^1(\Omega)$, choose $g \in L^2(\Omega)$ (since Ω has finite measure, $g \in L^1(\Omega)$), and compute the following:

$$\begin{aligned} &\|\widetilde{P}_{\phi_h * f} g - \widetilde{P}_f g\|_1 \\ &= \int_\Omega \left| \widetilde{P}_{\phi_h * f} g(y) - \widetilde{P}_f g(y) \right| dy. \end{aligned} \tag{32}$$

Recall that according to (15), $\widetilde{P}_f = \widetilde{P}_f^1 + \widetilde{P}_f^2$. We will now evaluate the integral (32) in terms of the two operators $\widetilde{P}_f^1$ and $\widetilde{P}_f^2$.

In (32), the component of the integrand that depends on $\widetilde{P}^1_f$ is given by:

$$\widetilde{P}^1_{\phi_h * f}g(y) - \widetilde{P}^1_f g(y) = \int_\Omega A(x,y)dx, \text{ where}$$
$$A(x,y) = a_{\phi_h * f}(x)q(x,y)g(x) - a_f(x)q(x,y)g(x).$$

We now define the following sets:

$$\begin{aligned} E_1 &= \{x \in \Omega : \phi_h * f(x) > f^d(x)\}, \\ E_2 &= \{x \in \Omega : \phi_h * f(x) \leq f^d(x)\}, \\ E_3 &= \{x \in \Omega : f(x) > f^d(x)\}, \\ E_4 &= \{x \in \Omega : f(x) \leq f^d(x)\}. \end{aligned}$$

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

$$\begin{aligned} \int\limits_{E_1 \cap E_3} A \ &\leq \ \|q\|_\infty \int_{E_1 \cap E_3} a_{\phi_h * f} g - a_f g \qquad (33) \\ &= \ \|q\|_\infty \int_{E_1 \cap E_3} \frac{\phi_h * f - f^d}{\phi_h * f} g - \frac{f - f^d}{f} g \\ &= \ \|q\|_\infty \int_{E_1 \cap E_3} \frac{\phi_h * f - f}{\phi_h * f} \frac{f^d}{f} g. \end{aligned}$$

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

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

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

$$\begin{aligned} \int\limits_{E_2 \cap E_3} A \ &\leq \ \|q\|_\infty \int\limits_{E_2 \cap E_3} -\frac{f - f^d}{f} g \qquad (34) \\ &\leq \ \|q\|_\infty \left\|\frac{f - f^d}{f}\right\|_\infty \|g\|_2 \ m(E_2 \cap E_3). \end{aligned}$$

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

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

Note that,

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

Continuing the computation from above,

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

By Proposition 2.29 of [26], since $\phi_h * f - f \to f$ in the L^2 norm as $h \to 0$, then $\phi_h * f - f \to f$ in measure; that is, $m(\{\phi_h * f - f \leq \delta\}) \to 0$ as $h \to 0$ for every $\delta > 0$. Therefore, we must have that $m(E_2 \cap E_3) \to 0$ as $h \to 0$, and consequently, the integral of A over $E_2 \cap E_3$ must converge to 0. Now, consider the integral of A over S_3:

$$\begin{aligned} \int\limits_{E_1 \cap E_4} A \ &\leq \ \|q\|_\infty \int_{E_1 \cap E_4} \frac{\phi_h * f - f}{\phi_h * f} g \qquad (35) \\ &\leq \ \|q\|_\infty \left\|\frac{1}{\phi_h * f}\right\|_\infty \|g\|_2 \ \|\phi_h * f - f\|_2 \ m(E_1 \cap E_4). \end{aligned}$$

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

$$\int_{E_2 \cap E_4} A = \int_{E_2 \cap E_4} a_{\phi_h * f} g - a_f g = 0. \qquad (36)$$

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

Returning to the integral (32), the component of the integrand that depends on $\widetilde{P}^2_f$ is given by:

$$\begin{aligned} &\widetilde{P}^2_{\phi_h * f} g(y) - \widetilde{P}^2_f g(y) \\ &= (1 - a_{\phi_h * f}(y))\, g(y) - (1 - a_f(y))\, g(y) \\ &= a_f(y) g(y) - a_{\phi_h * f}(y) g(y) := B(y). \qquad (37) \end{aligned}$$

This term is equal to the integrand of each of the four integrals considered in (33)-(36). Since we showed that each of these integrands tends to 0 as $h \to 0$, we must have that $B(y) \to 0$ as well.

We can now evaluate (32) as

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

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

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

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] Hakan Andersson and Tom Britton. *Stochastic epidemic models and their statistical analysis*, volume 151. Springer Science & Business Media, 2012.

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

[5] Michel Benaïm and Jörgen W Weibull. Deterministic approximation of stochastic evolution in games. *Econometrica*, 71(3):873–903, 2003.

[6] Shiba Biswal, Karthik Elamvazhuthi, and Spring Berman. Stabilization of multi-agent systems to target distributions using local interactions. *IFAC-PapersOnLine*, 54(9):60–65, 2021. International Symposium on Mathematical Theory of Networks and Systems (MTNS) 2020.

[7] Shiba Biswal, Karthik Elamvazhuthi, and Spring Berman. Stabilization of nonlinear discrete-time systems to target measures using stochastic feedback laws. *IEEE Transactions on Automatic Control*, 66(5):1957–1972, 2021.

[8] Shiba Biswal, Kathik Elamvazhuthi, and Spring Berman. Fastest mixing Markov chain on a compact manifold. In *2019 IEEE 58th Conference on Decision and Control (CDC)*, pages 547–554, 2019.

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

[10] Charles Bordenave, David McDonald, and Alexandre Proutiere. A particle system in interaction with a rapidly varying environment: Mean field limits and applications. *arXiv preprint math/0701363*, 2007.

[11] Luca Bortolussi, Jane Hillston, Diego Latella, and Mieke Massink. Continuous approximation of collective system behaviour: A tutorial. *Performance Evaluation*, 70(5):317–349, 2013.

[12] Francesco Bullo, Jorge Cortes, and Sonia Martinez. *Distributed control of robotic networks: a mathematical approach to motion coordination algorithms*, volume 27. Princeton University Press, 2009.

[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] Yue Chen, Md Umar Hashmi, Joel Mathias, Ana Bušić, and Sean Meyn. Distributed control design for balancing the grid using flexible loads. In *Energy Markets and Responsive Grids*, pages 383–411. Springer, 2018.

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

[16] Pierre Del Moral et al. Measure-valued processes and interacting particle systems. application to nonlinear filtering problems. *The Annals of Applied Probability*, 8(2):438–495, 1998.

[17] Nazlı Demir, Utku Eren, and Behçet Açıkmeşe. Decentralized probabilistic density control of autonomous swarms with safety constraints. *Autonomous Robots*, 39(4):537–554, 2015.

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

[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, 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, Shiba Biswal, and Spring Berman. Mean-field stabilization of robotic swarms to probability distributions with disconnected supports. In *2018 Annual American Control Conference (ACC)*, pages 885–892. IEEE, 2018.

[22] Karthik Elamvazhuthi, Shiba Biswal, and Spring Berman. Controllability and decentralized stabilization of the Kolmogorov forward equation for Markov chains. *Automatica*, 124:109351, 2021.

[23] 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.

[24] Karthik Elamvazhuthi, Matthias Kawski, Shiba Biswal, Vaibhav Deshmukh, and Spring Berman. Mean-field controllability and decentralized stabilization of Markov chains. In *2017 IEEE Conference on Decision and Control (CDC)*, pages 3131–3137, 2017.

[25] Stewart N Ethier and Thomas G Kurtz. *Markov processes: characterization and convergence*, volume 282. John Wiley & Sons, 2009.

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

[27] Massimo Fornasier, Benedetto Piccoli, and Francesco Rossi. Mean-field sparse optimal control. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 372(2028):20130400, 2014.

[28] Diogo A Gomes, Joana Mohr, and Rafael Rigao Souza. Discrete time, finite state space mean field games. *Journal de mathématiques pures et appliquées*, 93(3):308–328, 2010.

[29] Roger A Horn and Charles R Johnson. *Matrix analysis*. Cambridge university press, 2012.

[30] Miroslav Kárnỳ. Towards fully probabilistic control design. *Automatica*, 32(12):1719–1722, 1996.

[31] Vassili N Kolokoltsov. *Nonlinear Markov processes and kinetic equations*, volume 182. Cambridge University Press, 2010.

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

[33] Roland Malhamé and Chee-Yee Chong. On the statistical properties of a cyclic diffusion process arising in the modeling of thermostat-controlled electric power system loads. *SIAM Journal on Applied Mathematics*, 48(2):465–480, 1988.

[34] T William Mather and M Ani Hsieh. Distributed robot ensemble control for deployment to multiple sites. *Robotics: Science and Systems VII*, 2011.

[35] T William Mather and M Ani Hsieh. Synthesis and analysis of distributed ensemble control strategies for allocation to multiple tasks. *Robotica*, 32(02):177–192, 2014.

[36] Mehran Mesbahi and Magnus Egerstedt. *Graph theoretic methods in multiagent networks*, volume 33. Princeton University Press, 2010.

[37] Alain-Sol Sznitman. Topics in propagation of chaos. In *Ecole d'été de probabilités de Saint-Flour XIX—1989*, pages 165–251. Springer, 1991.

[38] Hamidou Tembine et al. Energy-constrained mean field games in wireless networks. *Strategic Behavior and the Environment*, 4(2):187–211, 2014.

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

Shiba Biswal received the Ph.D and M.S. degree in mechanical engineering and the M.A. degree in mathematics from Arizona State University (ASU), Tempe, AZ, in 2020, 2015 and 2019, respectively. 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 multi-agent 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.