Stabilization of Multi-Agent Systems to Target Distributions using Local Interactions *

Shiba Biswal* Karthik Elamvazhuthi* Spring Berman**

* Department of Mathematics, University of California, Los Angeles,

CA 90095, USA (e-mail: {sbiswal,karthikevaz}@math.ucla.edu)

** School for Engineering of Matter, Transport and Energy, Arizona

State University, Tempe, AZ, 85281 USA

(e-mail:spring.berman@asu.edu).

Abstract: In this paper, we construct a mean-field discrete-time Markov process evolving on a compact subset of \mathbb{R}^{d} can be stabilized to an arbitrary target distribution that has a continuous density. This density, unlike in our previous works, need not have a connected support on the state space. Our main application of interest is characterizing the distribution of a multi-agent system that evolves according to a discrete-time Markov process. Even if the Markov process converges to an equilibrium distribution, the agents may continue to switch between states, potentially wasting energy. In order to prevent this unnecessary switching, we show that the Markov process can be constructed in such a way that the operator that pushes forward measures is the identity operator at the target measure. The challenge in the stability analysis of the system arises from the fact that the transition kernel is a function of the current distribution, resulting in a nonlinear Markov process. Moreover, we aim to design the transition kernel, which is the feedback control law for the Markov process, to be *decentralized* in the sense that it depends on the local density of agents. We prove by construction that there exists a control law that is decentralized and globally stabilizes the desired measure. In order to implement this control law, the individual agents must estimate the local population density. We validate our control law with numerical simulations of multi-agent systems with different population sizes. We observe that the number of agent state transitions at equilibrium significantly decreases as the population size increases.

Keywords: Large Scale Systems, Stochastic Modeling and Stochastic Systems Theory, Operator Theoretic Methods in Systems Theory, Robotics

1. INTRODUCTION

In this paper, we address the problem of stabilizing a discrete-time Markov process evolving on a compact, connected subset of \mathbb{R}^d to a desired distribution. Our application of interest is controlling the distribution of a multi-agent system in which the agents evolve according to the Markov process that we consider. The time evolution of the distribution is given by the *Kolmogorov forward equation*. In our previous works, Biswal et al. (2019a,b), we considered a similar scenario in which we stabilized measures that have an L^{∞} density. In this paper, we consider measures that have *continuous* densities, a much smaller class of measures than we previously considered. However, unlike in our earlier works, the distribution's support need not be strictly positive almost everywhere. The reason for this will be made clear shortly.

There are numerous well-established methods for control of multi-agent systems, many of which are described in Bullo et al. (2009); Lewis et al. (2013); Mesbahi and Egerstedt (2010). However, many of these control approaches lack scalability with respect to very large agent populations. When all agents follow the same control laws and these control laws are independent of agents' identities, an alternative approach is to apply control techniques to a fluid approximation of the swarm in the form of a mean-field model (Elamvazhuthi and Berman, 2019). 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. In the absence of agent interactions, the mean-field model often has the advantage of being linear, or at least simpler to analyze than the dynamics of a population of finite number of agents. Therefore, in this paper, we focus on the problem of stabilizing the mean-field model of the system using the corresponding transition kernel, which determines the state transition probabilities of the agents, as the control parameter. We demonstrate that the control law designed for the mean-field model enables a control approach that scales well with the agent population size.

Our goals in this paper are threefold. First, we aim to design a Markov process that can stabilize any probability measure that has a continuous density and is not

^{*} This work was supported by the Arizona State University Global Security Initiative.

necessarily positive everywhere on the domain. The significance of the second property stems from the fact that in general, discrete-time Markov chains cannot be stabilized to distributions that do not have connected supports, as shown in Elamvazhuthi et al. (2017). 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, our second goal is to construct a Markov process such that its forward operator, which pushes forward measures, is the identity operator at equilibrium. This results in a time-dependent transition kernel that is a function of the distribution and gives rise to a nonlinear Markov process. Since we establish that the kernel must be dependent on the distribution, our third goal is to construct the kernel to have a decentralized structure. A kernel with this structure corresponds to control laws that require each agent to estimate the population only in its local neighborhood, rather than obtain feedback on the entire agent distribution. Toward this end, we construct an explicit 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. This is the reason for working with measures that have continuous densities. We proved the existence of such feedback laws in the case of continuous-time Markov chains evolving on finite graphs in Elamvazhuthi et al. (2017, 2018). A similar problem is addressed in Mather and Hsieh (2014), which develops a decentralized control approach by a priori restricting the controller to have a decentralized structure, in order to avoid requiring agents to have information on the entire distribution. Another related work is Demir et al. (2015), which 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 (Lasota and Mackey, 2013), and it is also used extensively in the context of mean-field games (Gomes et al., 2010), optimal transport theory (Villani, 2003), and mean-field control (Fornasier et al., 2014). We present a review of significant works that have influenced research on the stabilization of Markov processes in Biswal et al. (2019b).

2. NOTATION

In this section, we present notation that will be used throughout the paper. We define $\mathbb{R}_+ := [0, \infty)$, and $\mathbb{R}_+ := (0, \infty)$. Similarly, we define \mathbb{Z}_+ as the set of all nonnegative 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 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 . We denote the spaces of probability measures on Ω 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 $f : \mathbb{R}^n \to \mathbb{R}$ such that $d\nu = f dm$; this function is called the *density* or the derivative of ν with respect to m, Folland (2013).

For a measure space (\mathcal{X}, ν) , we define $L^p(\mathcal{X}, \nu)$, where $p \in [1, \infty)$, as the space $\{f : \mathcal{X} \to \mathbb{R} : f \text{ is measurable and } \|f\|_p < \infty\}$, where $\|f\|_p = (\int |f|^p d\nu)^{1/p}$. In addition, we define $L^{\infty}(\mathcal{X}, \nu) = \{f : \mathcal{X} \to \mathbb{R} : f \text{ is measurable and } \|f\|_{\infty} < \infty\}$, where $\|f\|_{\infty} = \text{ess sup}_{x \in \mathcal{X}} |f(x)|$. $C(\mathcal{X})$ is the space of continuous functions on \mathcal{X} . We define $\mathcal{C}(\mathcal{X}) \subset L^1(\mathcal{X})$ to be a set of continuous functions that satisfy $\int_{\mathcal{X}} f = 1$. We endow $\mathcal{C}(\mathcal{X})$ with the $\|\cdot\|_1$ norm.

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 $\sup_x \frac{||Tx||_{\mathcal{Y}}}{||x||_{\mathcal{X}}}$.

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} ,

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

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 $\mathcal{T}(x, du) = h(x, u)du$. The function $h : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}_+$ will be called the *kernel* function of the transition kernel \mathcal{T} .

3. PROBLEM FORMULATION

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 *identityfree*, that is, that they evolve independently of one another and according to the same dynamics. Let ξ_n^k denote the state of each agent $k \in \{1, \ldots, N\}$ in Ω at time $n \in \mathbb{Z}_+$. 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 transition kernel 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 probability density f^d . Since we assume that the agents are identity-free, we will define the transition kernel 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 . Since our goal is to control the distribution of the N-agent system, we consider the mean-field limit of $\frac{1}{N} \sum_{k=1}^{N} \delta_{\xi_n^k}$ as $N \to \infty$. Therefore, we will consider the following discrete-time flow on the space of densities $C(\Omega)$:

$$f_{n+1} = \widetilde{P}_{f_n} f_n, \quad n = 0, 1, 2 \dots$$

$$f_0 \in \mathcal{C}(\Omega)$$
(1)

where $\widetilde{P}_{f_n} : \mathcal{C}(\Omega) \to \mathcal{C}(\Omega)$, defined next, is called the forward operator. \widetilde{P}_{f_n} is a nonlinear operator that depends on the current density f_n , the subscript f_n emphasizes this dependence.

Let $f \in \mathcal{C}(\Omega)$ be the density of a measure $\mu \in \mathcal{P}(\Omega)$. Then, we will define \tilde{P}_f via a transition kernel $K_{\mu} : \Omega \times \mathcal{B}(\Omega) \to [0,1]$ that depends on the measure μ . Assume K_{μ} is regular, that is if $K_{\mu}(x, \cdot) = q_f(x, y)dy$, then $q_f \in L^{\infty}(\Omega \times \Omega, m \times m)$. We define \tilde{P}_f as

$$(\widetilde{P}_f f)(y) = \int_{\Omega} q_f(x, y) f(x) d(x).$$
(2)

To ensure that \widetilde{P}_f preserves probability densities, we impose the following property on q_f :

$$q_f(x,y) \begin{cases} \ge 0, \text{ for } m\text{-a.e. } y \in \Omega \\ = 0, \text{ otherwise} \end{cases}$$
(3)

$$\int_{\Omega} q_f(x, y) dy = 1, \text{ for } m\text{-a.e. } x \in \Omega.$$
(4)

These properties ensure that K_{μ} is stochastic.

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

Problem 1 Given $f^d \in \mathcal{C}(\Omega)$, determine whether there exists a transition kernel $K_{\mu} : \Omega \times \mathcal{B}(\Omega) \to [0, 1]$ such that (1) 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 densities $f_0 \in \mathcal{C}(\Omega)$, and furthermore, $\widetilde{P}_{f^d} = I$, where the forward operator \widetilde{P}_{f^d} is defined in (2).

The condition $\widetilde{P}_{f^d} = I$, the identity operator, is to ensure that the agent's trajectory, at the desired equilibrium distribution f^d , remains static. This condition leads to the nonlinearity of the operator \widetilde{P}_f .

4. MAIN RESULT

In this section, an operator \tilde{P} that solves Problem 1 will be constructed. To begin, we state the assumptions. Suppose $f^d \in \mathcal{C}(\Omega)$ be the desired distribution in Problem 1. We note that f^d need not be supported *m* almost everywhere on Ω , unlike our work in Biswal et al. (2019b). We also assume that Ω is a *path connected*, *compact* subset of \mathbb{R}^d . We also require Ω to satisfy the *cone condition* (Biswal et al., 2019b), which ensures that the boundary of Ω is regular enough.

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

$$a_f(x) = \begin{cases} \frac{f(x) - f^d(x)}{f(x)} & \text{if } f(x) - f^d(x) > 0\\ 0 & \text{otherwise.} \end{cases}$$
(5)

We note that a_f is continuous on Ω . Moreover, the supremum of a_f is 1.

Let r > 0 be fixed. To simplify notation, denote $B_r(x) \cap \Omega$ as B_x . Let $k : \Omega \times \Omega \to \mathbb{R}$ be a function in $L^{\infty}(\Omega \times$ $\Omega, m \times m$). Further, to simplify analysis we will assume in this paper that k is uniformly distributed over the ball B_x , that is, for $x \in \Omega$ and $y \in B_x$, $k(x, y) = \frac{1}{m(B_x)} \chi_{B_x}$. Therefore, k satisfies

$$\int_{\Omega} k(x,y) \chi_{B_x}(y) dy = 1.$$
(6)

The question of measurability of a k that satisfies the above property, is answered in our extension of this paper Biswal et al. (2021). Let $A \in \mathcal{B}(\Omega)$. Define transition kernel $K_{\mu} : \Omega \times \mathcal{B}(\Omega) \to [0, 1]$ as

$$K_{\mu}(x,A) = K_{\mu}(x,B_{x}\cap A)$$

= $a_{f_{\mu}}(x) \int_{A} k(x,y) \chi_{B_{x}}(y) dy + (1 - a_{f_{\mu}}(x)) \delta_{x}(A).$ (7)

The kernel is defined such that the corresponding Markov chain stays at x with probability $1 - a_{f_{\mu}}(x)$ and moves to a state in the set A with probability $a_{f_{\mu}}(x)$, and when it moves, the distribution is given by the density k(x, y).

Using K_{μ} , we define an operator P_{μ} that acts on measures in $\mathcal{P}(\Omega)$ as

$$(P_{\mu}\mu)(A) = \int_{\Omega} K_{\mu}(x, A)d\mu(x) = \int_{\Omega} K_{\mu}(x, B_{x} \cap A)d\mu(x)$$
$$= \int_{\Omega} \int_{A} a_{f_{\mu}}(x)\chi_{B_{x}}(y)k(x, y)dyd\mu(x) +$$
$$\int_{\Omega} (1 - a_{f_{\mu}}(x))\delta_{x}(A)d\mu(x)$$
$$= \int_{\Omega} \int_{A} a_{f_{\mu}}(x)\chi_{B_{x}}(y)k(x, y)dyd\mu(x) +$$
$$\int_{A} (1 - a_{f_{\mu}}(x))d\mu(x).$$
(8)

The subscript μ in P_{μ} , emphasizes the fact that P_{μ} is a nonlinear operator that depends on the measure it is acting upon.

We note how a Dirac measure δ_z , for some $z \in \Omega$, behaves under the action of P_{μ} in successive time steps. Since $k(x, \cdot) = \frac{1}{m(B_x)} \chi_{B_x}$, we have $P_{\mu} \delta_z(A) = \frac{m(B_z \cap A)}{m(B_z)}$. That is, an application of P_{μ} to the Dirac measure results in a measure that is supported on the ball of radius r centered at z. The operator P_{μ} so defined therefore has a 'spreading effect'.

We note a few important properties of the transition kernel K_{μ} and the operator P_{μ} in the following lemma. The proof is omitted as it is follows straightforwardly from the definition.

Lemma 1. K_{μ} is a well-defined Markov kernel, that is, it is a measure on Ω in the second variable and a measurable function on Ω in the first. Further, $P_{\mu} : \mathcal{P}(\Omega) \to \mathcal{P}(\Omega)$ is a well-defined operator, it preserves absolutely continuous measures.

To obtain an operator that acts on functions (densities), specifically, $C(\Omega)$ instead of measures $\mathcal{P}(\Omega)$, we carry out the following computations. In view of the lemma above, we must have that $P_{\mu}\mu$ is absolutely continuous w.r.t m, suppose $\widetilde{P}_{f_{\mu}}f_{\mu}$ be its density. That is, for $A \in \mathcal{B}(\Omega)$

$$(P_{\mu}\mu)(A) = \int_{A} (\widetilde{P}_{f_{\mu}}f_{\mu})(y)dy \tag{9}$$

Applying Fubini's theorem (Folland (2013)) to (9) and equating to (8) we obtain

$$\begin{split} \int_A (\widetilde{P}_{f_\mu} f_\mu)(y) dy &= \int_A \int_\Omega a_{f_\mu}(x) \chi_{B_x}(y) k(x,y) dy d\mu(x) + \\ &\int_A (1 - a_{f_\mu}(x)) d\mu(x). \end{split}$$

From the above equation we can write an expression for \widetilde{P}_f as follows. For $f \in C(\Omega)$ we define \widetilde{P}_f on $\mathcal{C}(\Omega)$ as

$$P_f f(y) = P_f^1 f + P_f^2 f, \text{ where}$$

$$(\widetilde{P}_f^1 f)(y) = \int_{\Omega} a_f(x) k(x, y) \chi_{B_x}(y) f(x) dx +$$

$$(\widetilde{P}_f^2 f)(y) = (1 - a_f(y)) f(y)$$

$$(10)$$

Next, we prove that \widetilde{P} indeed preserves $\mathcal{C}(\Omega)$.

Lemma 2. $\widetilde{P}: \mathcal{C}(\Omega) \to \mathcal{C}(\Omega)$ is well-defined.

Proof. Let $f \in \mathcal{C}(\Omega)$. Owing to the property of k, it is easy to see that \widetilde{P}_f preserves functions that integrate to 1. Next, we will prove that $\widetilde{P}_f f$ is continuous. We note that the non-integral term $\widetilde{P}_f^2 f$ in (10) is trivially continuous, hence, we will prove continuity of $\widetilde{P}_f^1 f$. This integral term is non-zero if $a_f(y), a_f(z) > 0$, which we assume to be true. In order to prove continuity, we invoke the basic ε - δ argument, that is, we will prove that for every $\varepsilon > 0$, there exists a $\delta > 0$, such that for all $y, z \in \Omega$, whenever $|y-z| \leq \delta$, $|\widetilde{P}f(y) - \widetilde{P}_f^1 f(z)| \leq \varepsilon$. Fix $\varepsilon > 0$. For $y, z \in \Omega$, $|\widetilde{P}_f^1 f(y) - \widetilde{P}_f^1 f(z)|$ evaluates to:

$$\int_{\Omega} (a_f(x)k(x,y)\chi_{B_x}(y)f(x) - a_f(x)k(x,z)\chi_{B_x}(z)f(x)) dx$$

$$\leq \|k\|_{\infty} \int_{\Omega} (f(x) - f^d(x))(\chi_{B_x}(y) - \chi_{B_x}(z))dx$$

$$\leq \|k\|_{\infty} \|f - f^d\|_{sup} \int_{\Omega} (\chi_{B_x}(y) - \chi_{B_x}(z))dx. \tag{11}$$
We have that

We have that,

$$\int_{\Omega} \chi_{B_x}(y) dx = m \left(\bigcup_{x \in \Omega: y \in B_x} B_x \right) = m(B_r(y) \cap \Omega),$$

and similarly for the point z. Therefore,

$$\int_{\Omega} \chi_{B_x}(y) - \chi_{B_x}(y) dx = m(B_r(y) \cap \Omega) - m(B_r(z) \cap \Omega)$$

= $m(B_r(y) \cap \Omega \setminus B_r(z)) - m(B_r(z) \cap \Omega \setminus B_r(y))$
 $\leq m(B_r(y) \setminus B_r(z)) - m(B_r(z) \setminus B_r(y))$

It is easy to see that right hand side is bounded from above by a function of $\delta = |y - z|$. That is, if $|\tilde{P}_f^1 f(y) - \tilde{P}_f^1 f(z)| \leq \varepsilon$, then $\varepsilon = g(\delta)$, where g is a bijective function that can be determined by geometry not included here for the sake of brevity. Choosing $\delta = g^{-1}(\varepsilon)/\|k\|_{\infty} \|f - f^d\|_{sup}$ completes the continuity argument.

The operator \tilde{P}_f trivially satisfies $\tilde{P}_{f^d} f^d = f^d$. Moreover, it satisfies $\tilde{P}_{f^d} = I$, in order to stop the agents from transitioning between states at the target density f^d . Next, we will show that f^d is a globally asymptotically stable equilibrium of system (1).

Theorem 3. For the system (1), f^d is globally asymptotically stable in the $L^1(\Omega, m)$ norm.

Here only a sketch of the proof is provided, the theorem will be proved for more general f^d in the extension of this paper Biswal et al. (2021).

Proof. (*sketch*) To start, we make the following observation. Consider the case when $y \in \Omega$ is such that it satisfies $f_n(y) > f^d(y)$. Then $a_f(y) > 0$. From the expression (10) we obtain,

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

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

$$\begin{aligned}
f_{n+1}(y) &\geq f_n(y) > f^d(y) \\
f_n(y) &> f_{n+1}(y) \geq f^d(y).
\end{aligned} (12)$$

That is, it is not possible for the function $f_{n+1}(y) < f^d(y)$, for any *n*. Next, consider the case when $y \in \Omega$ is such that $f_n(y) \leq f^d(y)$. In this situation $a_f(y) = 0$. Expression (10) then reduces to

$$\widetilde{P}f_n(y) = \int_{\Omega} a_{f_n}(x)k(x,y)\chi_{B_x}(y)f_n(x)dx + f_n(y).$$
(13)

Similar to the previous case, attributing to the fact that the first term is non-negative, only one of the following must be true:

$$\begin{aligned}
f_{n+1}(y) &\geq f^d(y) > f_n(y) \\
f^d(y) &> f_{n+1}(y) \geq f_n(y).
\end{aligned} \tag{14}$$

That is, in this case f_{n+1} monotonically increases.

For any given n, define the sets $E_n^1 = \{y \in \Omega : f_n(y) < f^d(y)\}$, $E_n^2 = \{y \in \Omega : f_n(y) = f^d(y)\}$ and $E_n^3 = \{y \in \Omega : f_n(y) > f^d(y)\}$. By construction, $\Omega = E_n^1 \sqcup E_n^2 \sqcup E_n^3$; moreover, these sets do not intersect one another. Since each f_n is a probability density on Ω , it 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 (14), we have that $f_{n+1} \ge f_n$, and hence $f^d - f_n \ge f^d - f_{n+1}$. 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 . Moreover, since $(F_n)_n$ is bounded, this implies that F_n converges pointwise to a function, say g. By the monotone convergence theorem, we then have that $\int_{\Omega} F_n \to \int_{\Omega} g$. If g = 0, then we have our result. We will next prove by contradiction that g is in fact 0.

Let $g \neq 0$, and let $\int_{\Omega} g \geq \gamma > 0$. Define $S = \{x \in \Omega : g(x) > 0\}$. We note that the definition of S is independent of time. Given the conditions in (12) and (14), 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, $\lim_{n\to\infty} m(E_n^1) \to m(S)$. Fix n, in order to obtain a contradiction, the proof involves proving that the measure that is pushed from $(E_n^1)^c$ to S is greater than γ contradicting the claim.

5. SIMULATIONS

In this section, we present stochastic simulations of the system (1). Consider a population of N agents evolving



Fig. 1. Visualization of the bump functions



Fig. 2. Target distribution

on a state space Ω . Denote the state of each agent k at time n by $X_k^N(n) = x_k$, $k = \{1, \ldots, N\}$. The empirical measure μ_n^N at time n is given by a normalized sum of Dirac measures associated with each agent,

$$\mu_n^N = \frac{1}{N} \sum_{k=1}^N \delta_{x_k}.$$
 (15)

The empirical measure, being a sum of Dirac measures, does not have a density; that is, it is not absolutely continuous with respect to the Lebesgue measure. Therefore, we cannot apply \tilde{P} or P in (8), (10) to this quantity. In order to be able to apply \tilde{P} on this empirical measure, we "mollify" the Dirac measures. Mathematically, this means that the measure μ^N is convolved with a smooth function $\phi : \mathbb{R}^d \to \mathbb{R}$, a *mollifier*. The convolution of μ^N and ϕ is given by

$$\phi * \mu^N = \int_{\Omega} \phi(x, y) d\mu^N = \sum_{i=1}^N \phi(x - x_i).$$
 (16)

Therefore, the result of this convolution is a sum of smooth functions, which is smooth. Loosely speaking, each Dirac measure is replaced by a smooth function. We can now apply \tilde{P} or P to the right-hand side of this equation. We note a few properties of a mollifier ϕ on \mathbb{R}^d (Folland, 2013). For some h > 0, define

$$\phi_h(x) = h^{-d} \phi\left(\frac{x}{h}\right). \tag{17}$$

We observe that $\int \phi_h$ is independent of h. Moreover, the "mass" of ϕ_h becomes concentrated at the origin as $h \to 0$. In our simulations, we have chosen the standard *bump* function with a compact support:

$$\phi(x) = \begin{cases} e^{-\left(\frac{1}{1-\|x\|^2}\right)}, & x \in (-1,1) \\ 0, & \text{otherwise.} \end{cases}$$
(18)

Figure 1 shows a visualization of the bump functions of size h = 0.1; each Dirac measure is now replaced by

a smooth approximating bump function. The definition of \widetilde{P}/P is pointwise description or strictly local. The transition kernel K (7) is such that it requires a parameter a_f (5) which is positive at some $x \in \Omega$, only if the density at x is higher than the target density's value at x. However, this description is not feasible practically. The mollification procedure changes this local description to a non-local description. Therefore, from an agent's perspective, the mollification of the empirical measure implies that each agent estimates the density ϕ in (16) based on the relative distance between itself and its neighbors that are within a radius h. Therefore, as $h \to 0$, the bump function tends to the Dirac measure and the interactions tend to being local.

We now present stochastic simulations of agents evolving on $\Omega \subset \mathbb{R}^2$. In the example below, Ω is defined as the unit square $[0,1] \times [0,1]$. The target distribution is shown in Fig. 2; it is set to be $\mu_d = \sin^2(2\pi x^1) + \sin^2(2\pi x^2)$, where $[x^1 \ x^2]^T \in \Omega$. Agent population sizes of 100, 500, and 1000 are simulated. The bump parameter h is set to 0.1. As explained in the previous paragraph, h determines the area over which the agents interact; that is, each agent considers another agent to be its neighbor as long as their relative distance is within h. The radius r of the ball over which $k(x, \cdot)$ is positive is set to 0.1. The initial conditions of the agent states were randomly chosen.

For each population size, Figs. 3, 4, and 5 show snapshots of the time evolution of the agent distribution, as well as the trajectories of five randomly selected agents. In the agent trajectory plot, norm of the state is plotted against time. We observe that as the population size increases, the agent distribution becomes closer to the target distribution, and the agents' frequency of switching between states significantly decreases. This can be attributed to the fact that we are approximating a continuous function, the distribution, with a finite number of points, which represent the agent positions. In summary, we have shown that control policies designed using the mean-field model can be implemented on a population of individual agents, as long as this population is sufficiently large.

6. CONCLUSION

In this paper, we constructed a Markov process that converges to an arbitrary distribution that has continuous derivatives. This distribution need not be strictly positive everywhere on the domain. Moreover, the Markov process can be constructed such that agent transitions between states at the equilibrium distribution is prevented. Although the results were proven for the mean-field model, we showed via stochastic simulations that for populations of a few hundred agents, the agent distribution converges fairly closely to the target distribution at equilibrium.

REFERENCES

- Biswal, S., Elamvazhuthi, K., and Berman., S. (2019a). Fastest mixing Markov chain on a compact manifold. 2019 IEEE Conference on Decision and Control (CDC), Nice, France.
- Biswal, S., Elamvazhuthi, K., and Berman, S. (2019b). Stabilization of nonlinear discretetime systems to target measures using stochastic



feedback laws. In revision for *IEEE Transactions* on Automatic Control. Preprint available at: https://www.researchgate.net/publication/330103622. Biswal, S., Elamvazhuthi, K., and Berman, S. (2021).

Fig. 5. Stochastic simulation for 1000 agents at different times n.

(b) Density at time n = 250

(a) Initial density at time n = 0

- Diswai, S., Elamvazhuthi, K., and Berman, S. (2021). Decentralized control of multi-agent systems using local density feedback. Conditionally accepted to *IEEE Transactions on Automatic Control.*
- Bullo, F., Cortes, J., and Martinez, S. (2009). Distributed control of robotic networks: a mathematical approach to motion coordination algorithms, volume 27. Princeton University Press.
- Demir, N., Eren, U., and Açıkmeşe, B. (2015). Decentralized probabilistic density control of autonomous swarms with safety constraints. *Autonomous Robots*, 39(4), 537– 554.
- Elamvazhuthi, K. and Berman, S. (2019). Mean-field models in swarm robotics: A survey. *Bioinspiration & Biomimetics*, 15(1), 015001.
- Elamvazhuthi, K., Biswal, S., and Berman, S. (2018). Mean-field stabilization of robotic swarms to probability distributions with disconnected supports. In 2018 Annual American Control Conference (ACC), 885–892. IEEE.
- Elamvazhuthi, K., Kawski, M., Biswal, S., Deshmukh, V., and Berman, S. (2017). Mean-field controllability and decentralized stabilization of Markov chains. In 2017 IEEE Conference on Decision and Control (CDC),

3131-3137.

(c) Density at time n = 500

Folland, G.B. (2013). Real analysis: modern techniques and their applications. John Wiley & Sons.

(d) Agent states

- Fornasier, M., Piccoli, B., and Rossi, F. (2014). Meanfield sparse optimal control. *Philosophical Transactions* of the Royal Society A: Mathematical, Physical and Engineering Sciences, 372(2028), 20130400.
- Gomes, D.A., Mohr, J., and Souza, R.R. (2010). Discrete time, finite state space mean field games. Journal de mathématiques pures et appliquées, 93(3), 308–328.
- Lasota, A. and Mackey, M.C. (2013). Chaos, fractals, and noise: stochastic aspects of dynamics, volume 97. Springer Science & Business Media.
- Lewis, F.L., Zhang, H., Hengster-Movric, K., and Das, A. (2013). Cooperative control of multi-agent systems: optimal and adaptive design approaches. Springer Science & Business Media.
- Mather, T.W. and Hsieh, M.A. (2014). Synthesis and analysis of distributed ensemble control strategies for allocation to multiple tasks. *Robotica*, 32(02), 177–192.
- Mesbahi, M. and Egerstedt, M. (2010). *Graph theoretic methods in multiagent networks*, volume 33. Princeton University Press.
- Villani, C. (2003). Topics in optimal transportation, volume 58. American Mathematical Society.