Nonlinear Generalizations of Diffusion-Based Coverage by Robotic Swarms

Karthik Elamvazhuthi and Spring Berman

Abstract— **In this paper, we generalize our diffusion-based approach [8] to achieve coverage of a bounded domain by a robotic swarm according to a target probability density that is a function of a locally measurable scalar field. We generalize this approach in two different ways. First, we show that our method can be extended in a natural way to scenarios where the robots' state space is a compact Riemannian manifold, which is the case if the robots are confined to a surface or if their configuration space is non-Euclidean due to dynamical constraints such as those present in most mechanical systems. Then, we establish the stability properties of a weighted variation of the porous media equation, a nonlinear partial differential equation (PDE). Coverage strategies based on these nonlinear PDEs have the advantage that the robots stop moving once the equilibrium probability density is reached, in contrast to our original approach [8]. We establish long-time stability properties of the target probability densities using semigroup theoretic arguments. We validate our theoretical results through stochastic simulations of a linear diffusion-based coverage strategy on a 2-dimensional sphere and numerical solutions of the weighted porous media equation on the 2-dimensional torus.**

I. INTRODUCTION

In the last two decades, there has been a considerable amount of work on multi-agent control problems. The significance of these control problems is due to the possibility of their use in multiple scenarios such as disaster response, environment monitoring, surveillance, and biomedical applications. Many initial works focused on algorithms to generate desired global behaviors from a small set of local rules, based on simplistic assumptions about agent capabilities. Over the years, these algorithms have been developed to a significant level of sophistication in terms of relaxing many of the assumptions made in earlier algorithms such as holonomy of agent motion, perfect measurements, and constant connectivity of communication networks.

One such standard assumption in many of these initial algorithms relates to the structure of the agents' state space. In most scenarios, it has been assumed that agents evolve on a Euclidean state space. However, many mechanical systems are naturally modeled on manifolds [21], [31]. Toward this end, a number of classical multi-agent algorithms have been extended to the case where agents evolve on manifolds. For example, [28] considers the problem of extending consensus algorithms to homogeneous manifolds. See also [30] for a detailed review of consensus on manifolds. Similarly, Voronoi-based coverage algorithms [7] have been extended to general Riemannian manifolds [6].

One of the goals of this paper is to relax the assumption that agents evolve on a Euclidean state space for a multi-agent coverage strategy developed by the authors in [8]. A major distinction between classical approaches to coverage [7] and the approach presented in [8] is that the latter work models the population dynamics of the swarm using a partial differential equation (PDE), which serves as a mean-field model for the swarm. The application of mean-field models to the design of control laws for multi-agent systems has seen a great amount of activity in recent years, e.g. [23], [26], [9]. In these works, the solution of the PDE is a probability density function that represents the spatiotemporal distribution of the agents. Similar work has also been done using Markov chain models that evolve either in continuous time [5], [15] or discrete time [1], [3]. The analysis in all of these works is scalable with the number of agents due the fact that it is performed on the mean-field model, which represents the swarm as a continuum. A continuum-based approach that is distinct from these approaches is proposed in [12], [27], in which the domain on which the PDE is solved represents a network of agents performing consensus-type interactions.

The diffusion-based approach in [8] can be viewed as an extension of the algorithm presented in [22] for unbounded domains. An advantage of these kinds of stochastic control laws is that they are easy to implement. The control law presented in [8] has been validated using robot experiments [20]. One of the main contributions of this paper is the extension of our coverage strategy in [8] to the case where the agents evolve on compact manifolds. The approach in [8] was based on the correspondence between stochastic differential equations (SDEs) and parabolic PDEs on Euclidean spaces. A similar correspondence also exists between SDEs and PDEs evolving on manifolds [17]. Therefore, as we show in this paper, our coverage approach naturally extends to the case where the agents evolve on finite-dimensional manifolds.

The coverage approach in [8] is based on a multiplicatively perturbed version of the classical heat equation. Hence, another problem that we investigate in this paper is whether similar control laws can be constructed using *nonlinear* diffusion models [32]. Toward this end, we analyze a weighted variation of the porous media equation. Due to the nonlinearity of this equation, the corresponding control laws require interactions between agents. These control laws drive the agents to a target spatial probability density at equilibrium, at which point they stop moving. In contrast, when the control

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

Karthik Elamvazhuthi and Spring Berman are with the School for Engineering of Matter, Transport and Energy, Arizona State University, Tempe, AZ, 85287 USA {karthikevaz, Spring.Berman}@asu.edu.

laws are defined by a linear diffusion model, the agents continue moving at equilibrium and therefore expending energy unnecessarily. This is because the coefficients of the corresponding SDE model do not tend to zero at equilibrium, even though the PDE model asymptotically converges to the target density.

II. NOTATION AND TERMINOLOGY

In this section, we define some notation and terminology that will be used throughout the paper.

We denote by $\mathcal{M}$ a finite-dimensional compact Riemannian manifold without boundary. Defining $T_{\mathbf{x}}\mathcal{M}$ as the tangent space at $\mathbf{x} \in \mathcal{M}$, we equip the manifold $\mathcal{M}$ with a Riemannian metric g, i.e., an inner product $g(\mathbf{x})$ on $T_{\mathbf{x}}\mathcal{M}$ that varies smoothly with respect to $\mathbf{x} \in \mathcal{M}$. The Riemannian volume form will be denoted by $m(\mathbf{x})$. Hence, integration with respect to this volume form will be denoted by $\int_{\mathcal{M}} \cdot \; dm(\mathbf{x})$.

We define $L^2(\mathcal{M})$ as the space of square-integrable functions with respect to the Riemannian volume form m, equipped with the norm $\|\cdot\|_2$. For each $1 \leq p < \infty$, we define $L^p(\mathcal{M})$ as the Banach space of complex-valued measurable functions over the set $\mathcal{M}$ whose absolute value raised to the p^{th} power has finite integral. We define $L^\infty(\mathcal{M})$ as the space of essentially bounded measurable functions on $\mathcal{M}$. The space $L^\infty(\mathcal{M})$ is equipped with the norm $\|z\|_\infty =$ ess $\sup_{\mathbf{x}\in\mathcal{M}}|z(\mathbf{x})|$.

For a smooth function $u : \mathcal{M} \to \mathbb{R}$, ∇u will denote the *Riemannian gradient* of the function u. Let $C_c^\infty(\mathcal{M})$ denote the set of smooth functions on $\mathcal{M}$ with compact support. We define the norm $\|\cdot\|_{H^1}$ by $\|u\|_{H^1} = \|u\|_2 + \int_{\mathcal{M}} |\nabla u(\mathbf{x})|^2 dm(\mathbf{x})$, where the term $|\nabla u(\mathbf{x})|^2$ is shorthand for $\langle \nabla u, \nabla u\rangle_g$, the norm induced by the Riemannian metric g. We will also need the weighted space $L^2_a(\mathcal{M})$, defined as follows. For a given real-valued function $a \in L^\infty(\mathcal{M})$, $L^2_a(\mathcal{M})$ refers to the set of all functions f such that $\int_{\mathcal{M}} |f(\mathbf{x})|^2 a(\mathbf{x}) dm(\mathbf{x}) < \infty$. We will always assume that the associated function a is uniformly bounded from below by a positive constant, in which case the space $L^2_a(\mathcal{M})$ is a Hilbert space with respect to the weighted inner product $\langle\cdot,\cdot\rangle_a : L^2_a(\mathcal{M}) \times L^2_a(\mathcal{M}) \to \mathbb{R}$, given by $\langle f, g\rangle_a = \int_{\mathcal{M}} f(\mathbf{x})\bar{g}(\mathbf{x})a(\mathbf{x})dm(\mathbf{x})$ for each $f, g \in L^2_a(\mathcal{M})$. We will also need the space $H^1_a(\mathcal{M}) = \big\{f \in L^2_a(\mathcal{M}) : \int_{\mathcal{M}} |\nabla(af)|^2 dm(\mathbf{x})$ for $1 \leq i \leq N\big\}$, equipped with the norm $\|f\|_{H^1_a} = \Big(\|f\|^2_a + \int_{\mathcal{M}} |\nabla(af)|^2 dm(\mathbf{x})\Big)^{1/2}$. When $a = \mathbf{1}$, where $\mathbf{1}$ is the function that takes the value 1 a.e. on $\mathcal{M}$, the spaces $L^2(\mathcal{M})$ and $H^1(\mathcal{M})$, the space of functions in $L^2(\mathcal{M})$ with square-integrable weak derivatives, coincide with the spaces $L^2_a(\mathcal{M})$ and $H^1_a(\mathcal{M})$, respectively. Let X be a Hilbert space with the norm $\|\cdot\|_X$. The space $C([0,T];X)$ consists of all continuous functions $u : [0,T] \to X$ for which $\|u\|_{C([0,T];X)} := \max_{0\leq t\leq T} \|u(t)\|_X < \infty$.

For more details regarding Riemannian manifolds, we direct the reader to standard texts such as [18], [19]. For details regarding Sobolev spaces on manifolds, see [2], [14].

III. PROBLEM MOTIVATION

In this section, we present the models that will be analyzed in this paper. All the analysis presented in this paper is based on the theory of PDEs. However, there is a strong probabilistic motivation behind these problems. Hence, we will elucidate these aspects briefly by first reviewing the work done in [8] and then elaborating on the extensions considered in this paper.

Consider a swarm of N agents deployed on a bounded domain Ω in the n-dimensional Euclidean space $\mathbb{R}^n$. The agents evolve according to a identical *reflected diffusion process*,

$$d\mathbf{X}(t) = \mathbf{v}_1(t)dt + \sqrt{2}v_2(t)d\mathbf{W}(t) + d\psi(t),$$

where $\mathbf{W}$ is the classical N-dimensional Wiener process, $\psi(t)$ is the stochastic process that confines the agents to the domain Ω, and $\mathbf{v}_1 : [0,\infty) \to \mathbb{R}^n$ and $v_2 : [0,\infty) \to \mathbb{R}$ are controllable parameters. The objective in [8] was to construct feedback control laws $\mathbf{g} : \Omega \to \mathbb{R}^N$ and $a : \Omega \to \mathbb{R}_+$ such that the closed-loop system

$$d\mathbf{X}(t) = \mathbf{g}(\mathbf{X}(t))dt + \sqrt{2a(\mathbf{X}(t))}d\mathbf{W} + d\psi(t)$$

satisfies $\lim_{t\to\infty} \mathbb{P}(\mathbf{X}(t) \in A) = \int_A f(\mathbf{x})d\mathbf{x}$ for each measurable set $A \subset \mathbb{R}^N$, where f is a given probability density function. Here, $\mathbb{P}$ denotes the probability measure induced by the process $\mathbf{X}$ on the set of sample paths. Hence, the control laws must ensure that the probability of finding an agent in an infinitesimal element $d\mathbf{x}$ is given by $f(\mathbf{x})d\mathbf{x}$. One choice of feedback laws for which this objective can be achieved is $\mathbf{g} \equiv \mathbf{0}$ and $a(\mathbf{x}) = 1/f(\mathbf{x})$. Then the probability $\mathbb{P}(\mathbf{X}(t) \in A) = \int_A \rho(\mathbf{x},t)d\mathbf{x}$, where $\rho(\mathbf{x},t)$ is the solution of the weighted diffusion equation, a linear PDE given by

$$\begin{aligned} \frac{\partial y}{\partial t} &= \Delta(a(\mathbf{x})y) \quad in \;\; \Omega \times [0,T], \\ \mathbf{n}\cdot\nabla(a(\mathbf{x})y) &= 0 \quad on \;\; \partial\Omega \times [0,T], \\ y(\mathbf{x},0) &= y_0(\mathbf{x}) \quad in \;\; \Omega, \end{aligned} \tag{1}$$

where $\mathbf{n}$ is the unit vector normal to the boundary $\partial\Omega$ of the domain Ω. Using operator theoretical arguments, it was shown in [8] that the largest eigenvalue of the operator $\Delta(a(\mathbf{x})\cdot)$ is 0. Moreover, this eigenvalue is isolated, and hence we have $\lim_{t\to\infty} \rho(\mathbf{x},t) = cf(\mathbf{x})$ in an appropriate sense for some constant c that depends only on the initial density $\rho(\mathbf{x},0) = \rho_0(\mathbf{x})$. See [8] and [10] for more rigorous formulations of the results.

Note that when $a(\mathbf{x})$ is independent of $\mathbf{x} \in \Omega$, that is, when $a(\mathbf{x})$ is constant over the domain, the weighted diffusion equation (1) reduces to the classical heat equation, for which results on long-time behavior are well-known.

In this paper, we consider the case where Ω is replaced by a general (compact) Riemannian manifold $\mathcal{M}$. For the sake of simplicity, we will assume that $\mathcal{M}$ is without a boundary. We will show that the asymptotic stability results established in [10] for the PDE (1) naturally extend to the case where the domain is a manifold. In this case, the classical Laplacian operator Δ is replaced by the Laplace-Beltrami operator, which will be denoted by Δ_g. This operator is represented in local coordinates as

$$\Delta_g(\cdot) = \frac{1}{\sqrt{|g|}} \frac{\partial}{\partial x_i} (\sqrt{|g|} g^{ij} \frac{\partial}{\partial x_j} \cdot \ ) \ , \tag{2}$$

where $|g|(\mathbf{x})$ denotes the determinant of $g(\mathbf{x})$. Hence, one of the goals of this paper will be to analyze the long-term stability properties of the weighted version of the Laplace-Beltrami operator $\Delta_g(a(\mathbf{x})\cdot)$, where $a \in L^\infty(\mathcal{M})$. From this, we will construct a stochastic process (governing the agents' motion such that the agents confined to the manifold $\mathcal{M}$ converge to a target probability density. There are indeed many analytical results on the so-called weighted Laplacian [13]. However, the operator $\Delta_g(a(\mathbf{x})\cdot)$ is generally different from this weighted Laplacian, although the two operators are equivalent in the trivial case where the coefficient $a(\mathbf{x})$ is a constant function. Note that, as in the case of a Euclidean domain, there exists a stochastic process on manifolds associated with the operator Δ_g. In fact, Δ_g generates the classical Brownian motion on manifolds [17].

Another generalization that we will consider in this paper will be the long-term behavior of the equation

$$\begin{aligned} \frac{\partial y}{\partial t} &= \Delta_g((a(\mathbf{x})y)^\alpha) \quad in \ \mathcal{M} \times [0,T], \\ y(\mathbf{x},0) &= y_0(\mathbf{x}) \quad in \ \mathcal{M}, \end{aligned} \tag{3}$$

where α is a positive integer greater than or equal to 1. For general $\alpha \neq 1$, this is a nonlinear PDE. For $a(\mathbf{x})$ that is constant in space, this is the well-known *porous media equation* [32]. The motivation behind studying this nonlinear PDE is that for $\alpha > 1$, this PDE provides an alternative control law to the one presented in [8]. In this case, the corresponding SDE is of the Mckean-Vlasov type; that is, the coefficients of the SDE depend on the probability law of the random variable itself. For $\alpha = 2$, formally, the corresponding Lagrangian particle model can be written as the ODE

$$\frac{d\mathbf{x}}{dt} = -2a(\mathbf{x})\nabla(ya(\mathbf{x})), \tag{4}$$

where y is the solution of the PDE (3). As such, for a single agent, this equation does not have a physical interpretation. However, for Euclidean domains ($\mathcal{M} = \mathbb{R}^n$), the solutions of the PDE (3) can be derived as a limit of a system of particles that are interacting through a radially symmetric potential function. See [24] for more details.

An advantage of the control law (4) is that at the equilibrium density ($\rho(\mathbf{x},t) = cf(\mathbf{x})$), its coefficients are equal to $\mathbf{0}$. This is unlike the case where the coefficients are independent of the probability law of the random variable, as in the PDE (1). Hence, the agents' velocities are zero at equilibrium when the nonlinear diffusion model (3) is used. In contrast, when the agents' control laws are derived from a linear diffusion model, the agents maintain random motion even when their density has reached equilibrium. This can be disadvantageous in applications where it would be preferable to minimize agents' motion, and hence conserve energy.

IV. ANALYSIS

A. Linear Diffusion

The proofs in this section are similar to those that we developed in [10] for a related problem in the case of a Euclidean domain. Hence, the main contribution in this section is to verify that the arguments presented in [10] also apply to the case where $\mathcal{M}$ is non-Euclidean. Toward this end, we adapt the proofs in [10] by invoking the appropriate embedding theorems, Poincaré's inequality, and the Sobolev chain rule as adapted to Sobolev functions defined on Riemannian manifolds [14].

Given $a \in L^\infty(\mathcal{M})$ such that $a \geq c$ for some positive constant c, and $\mathcal{D}(\omega_a) = H^1_a(\mathcal{M})$, we define the *sesquilinear form* $\omega_a : \mathcal{D}(\omega_a) \times \mathcal{D}(\omega_a) \to \mathbb{C}$ as

$$\omega_a(u,v) = \int_{\mathcal{M}} \nabla(a(\mathbf{x})u(\mathbf{x})) \cdot \nabla(a(\mathbf{x})\bar{v}(\mathbf{x}))dm(\mathbf{x}) \tag{5}$$

for each $u, v \in \mathcal{D}(\omega_a)$. We associate with the form ω_a an operator $A_a : \mathcal{D}(A_a) \to L^2_a(\mathcal{M})$, defined as $A_a u = v$ if $\omega_a(u,\phi) = \langle v, \phi \rangle_a$ for all $\phi \in \mathcal{D}(\omega_a)$ and for all $u \in \mathcal{D}(A_a) = \{g \in \mathcal{D}(\omega_a) : \ \exists f \in L^2_a(\mathcal{M}) \text{ s.t. } \omega_a(g,\phi) = \langle f, \phi \rangle_a \ \forall \phi \in \mathcal{D}(\omega_a)\}$.

Remark IV.1. *Note that, strictly speaking, $\omega_a(u,v)$ should be expressed as $\int_{\mathcal{M}} \langle \nabla(a(\mathbf{x})u(\mathbf{x})), \nabla(a(\mathbf{x})\bar{v}(\mathbf{x})) \rangle_g dm(\mathbf{x})$. However, for ease of presentation, we will use the abuse of notation as in the definition* (5)*.*

Lemma IV.2. *The operator $A_a : \mathcal{D}(A_a) \to L^2_a(\mathcal{M})$ is closed, densely-defined, and self-adjoint. Moreover, the operator has a purely discrete spectrum.*

Proof. Consider the associated form ω_a. This form is *closed*, i.e., the space $\mathcal{D}(\omega_a)$ equipped with the norm $\|\cdot\|_{\omega_a}$, given by $\|u\|_{\omega_a} = (\|u\|^2_a + \omega_a(u,u))^{1/2}$ for each $u \in \mathcal{D}(\omega_a)$, is complete. This is true due to the fact that the multiplication map $u \mapsto a \cdot u$ is an isomorphism from $H^1_a(\mathcal{M})$ to $H^1(\mathcal{M})$ and $H^1(\mathcal{M})$ is a Hilbert space. Moreover, the space $H^1_a(\mathcal{M})$ is dense in $L^2_a(\mathcal{M})$ [14][Theorem 2.4, Lemma 2.4] This follows from the inequality $\|au - av\|_2 \leq \|a\|_\infty \|u - v\|_2$ for each $u, v \in L^2(\mathcal{M})$, the fact that the spaces $L^2_{\mathbf{1}}(\mathcal{M}) = L^2(\mathcal{M})$ and $L^2_a(\mathcal{M})$ are isomorphic, and the fact that the space $H^1(\mathcal{M})$ is dense in $L^2(\mathcal{M})$. In addition, it follows from the definition of the form ω_a that ω_a is *symmetric*, meaning that $\omega_a(u,v) = \overline{\omega_a(v,u)}$ for each $u, v \in \mathcal{D}(\omega_a)$. The form ω_a is also *semibounded*, i.e., there exists $m \in \mathbb{R}$ such that $\omega_a(u,u) \geq m\|u\|^2_a$ for each $u \in \mathcal{D}(\omega_a)$. Hence, it follows from [29][Theorem 10.7] that the operator A_a is self-adjoint. To establish the discreteness of the spectrum of A_a, we note that $H^1(\mathcal{M})$ is compactly embedded in $L^2(\mathcal{M})$ [14][Theorem 2.9]. This implies that when $H^1_a(\mathcal{M}) = \mathcal{D}(\omega_a)$ is equipped with the norm $\|\cdot\|_{\omega_a}$, then it is also compactly embedded in $L^2_a(\mathcal{M})$. From [29][Proposition 10.6], it follows that A_a has a purely discrete spectrum. $\square$

Corollary IV.3. *Consider the PDE*

$$\begin{aligned} y_t &= \Delta_g(a(\mathbf{x})y) \qquad in \ \mathcal{M} \times [0,T], \\ y(\cdot,0) &= y^0 \qquad in \ \mathcal{M}. \end{aligned} \tag{6}$$

Let $y^0 \in L^2_a(\mathcal{M})$. Then $-A_a$ generates a semigroup of operators $(\mathcal{T}_a(t))_{t\geq 0}$ such that the unique mild solution $y \in C([0,T];L^2_a(\mathcal{M}))$ of the above PDE exists and is given by $y(\cdot,t) = \mathcal{T}_a(t)y^0$ for all $t \geq 0$. Additionally, the semigroup $(\mathcal{T}_a(t))_{t\geq 0}$ is positive, i.e., $y_0 \geq 0$ implies that $\mathcal{T}_a(t)y_0 \geq 0$ for all $t \geq 0$.

Proof. First, we note that the operator $-A_a$ is *dissipative*, i.e., $\|(\lambda+A_a)u\|_a \geq \lambda\|u\|_a$ for all $\lambda > 0$ and all $u \in \mathcal{D}(A_a)$. Next, we note that $-A_a$ is self-adjoint, and hence the adjoint operator $-A_a^*$ is dissipative as well. It follows from a corollary of the *Lumer-Phillips theorem* [11][Corollary II.3.17] that $-A_a$ generates a semigroup of operators $(\mathcal{T}_a(t))_{t\geq 0}$ that solves the PDE (6) in the mild sense.

Finally, we establish the positivity of the semigroup. Toward this end, we note that the absolute value function $|\cdot| : \mathbb{R} \to \mathbb{R}$ is Lipschitz. Hence, it follows from [14][Proposition 2.5] that $v \in H^1(\mathcal{M})$ implies that $|v| \in H^1(\mathcal{M})$ whenever v is only real-valued. This implies that if $u \in \mathcal{D}(\omega_a)$, then $|\mathrm{Re}(u)| \in \mathcal{D}(\omega_a)$, where $\mathrm{Re}(\cdot)$ denotes the real component of its argument. Then the positivity of the semigroup follows from [25][Theorem 2.7]. ☐

Next, we establish some mass-conserving properties of the semigroup generated by $-A_a$.

Lemma IV.4. *The semigroup $(\mathcal{T}_a(t))_{t\geq 0}$ has the following mass conservation property: if $y^0 \geq 0$ and $\int_{\mathcal{M}} y^0(\mathbf{x})dm(\mathbf{x}) = 1$, then $\int_{\mathcal{M}}(\mathcal{T}_a(t)y^0)(\mathbf{x})dm(\mathbf{x}) = 1$ for all $t \geq 0$.*

Proof. Let $\int_{\mathcal{M}} y^0(\mathbf{x})dm(\mathbf{x}) = 1$ with $y^0 \in L^2_a(\mathcal{M})$. Then $\int_{\mathcal{M}}(y(\mathbf{x},t) - y^0(\mathbf{x}))dm(\mathbf{x}) = -\int_{\mathcal{M}} A_a(\int_0^t y(\mathbf{x},s)ds)dm(\mathbf{x}) = -\omega_a(\int_0^t y(\mathbf{x},s)ds, 1/a) = 0$ for all $t \geq 0$. Hence, the integral-preserving property of the semigroup holds. ☐

Proposition IV.5. *For the operator $-A_a$, 0 is a simple eigenvalue with the corresponding eigenvector $f = 1/a$. Hence, if $y^0 \geq 0$ and $\int_{\mathcal{M}} y^0(\mathbf{x})dm(\mathbf{x}) = 1$, then the following estimate holds:*

$$\|\mathcal{T}_a(t)y^0 - \tilde{c}f\|_a \quad \leq \quad M_0e^{-\lambda t}\|y^0 - \tilde{c}f\|_a, \qquad (7)$$

for $\tilde{c} = 1/\int_{\mathcal{M}} f(\mathbf{x})dm(\mathbf{x})$, some positive constants M_0, λ, and all $t \geq 0$.

Proof. To establish the exponential stability estimate (7), we note that, from Poincaré's inequality on compact manifolds [14][Theorem 2.10], there exists a constant $C > 0$ such that for all $u \in H^1(\mathcal{M})$,

$$\int_{\mathcal{M}} |u(\mathbf{x}) - u_{\mathcal{M}}|^2 dm(\mathbf{x}) \;\leq\; C\int_{\mathcal{M}} |\nabla u(\mathbf{x})|^2 dm(\mathbf{x}), \quad (8)$$

where $u_{\mathcal{M}} = \frac{1}{m(\mathcal{M})}\int_{\mathcal{M}} u(\mathbf{x})dm(\mathbf{x})$. This implies that 0 is a simple eigenvalue of the Laplace-Beltrami operator $A_{\mathbf{1}}$, since the operator is self-adjoint. Since the operator A_a can be written as a composition of operators $A_{\mathbf{1}}\mathcal{M}_a$, where $\mathcal{M}_a$ is the multiplication map $u \mapsto au$ from $H^1_a(\mathcal{M})$ to $H^1(\mathcal{M})$, it follows that 0 is also a simple eigenvalue of A_a with the corresponding eigenvector $f = 1/a$. The operator A_a is positive definite, and hence its spectrum lies in the right-half of the complex plane. Then the result follows from [11][Corollary V.3.3] and the mass-preserving property of the semigroup $(\mathcal{T}_a(t))_{t\geq 0}$. ☐

B. Nonlinear Diffusion

In this section, we consider the nonlinear diffusion model (3). Here we will only summarize the main results without proofs. A detailed analysis will be presented in a forthcoming paper.

From here on, we will assume that $a \in C^\infty(\mathcal{M})$. As in previous sections, it will also be assumed that a is bounded from below by a positive constant.

We will assume that $\alpha \geq 0$. Let V^* be the dual space of $V = H^1(\mathcal{M})$. Identifying the dual of $L^2(\mathcal{M})$ with itself, we have the continuous embeddings $V \hookrightarrow L^2(\mathcal{M}) \hookrightarrow V^*$. Let F be the Riesz isomorphism from V to V^*. Then V^* is a Hilbert space with the inner product $\langle\cdot,\cdot\rangle_{V^*}$ given by $\langle u_1,u_2\rangle_V = \langle Fu_1,Fu_2\rangle_{H^1}$ for all $u_1,u_2 \in V$. Note that we are not identifying V with V^*.

Then we define the (nonlinear) operator $B_a \subset V \times V$ as

$$B_a(u) = \{-\Delta_g v;\; v(\mathbf{x}) \in \beta(a(\mathbf{x})u(\mathbf{x})),\; \text{a.e. } \mathbf{x} \in \mathcal{M}\} \qquad (9)$$

where $\beta(\cdot)$ is the subdifferential of the function $\frac{1}{m+1}|\cdot|^{m+1} : \mathbb{R} \to \mathbb{R}$. Note that here we have used the terminology $B_a \subset V \times V$ to denote a nonlinear operator, since in general such operators are multi-valued. Hence, it is typical to represent nonlinear operators using their graphs [4].

Following the arguments for the case of Dirichlet boundary conditions [4], we have the following result.

Theorem IV.6. *The operator B_a is maximal monotone. Therefore, B_a generates a nonlinear semigroup of operators $(\mathcal{S}_a(t))_{t\geq 0}$ on V. Hence, there exists a unique strong solution $\rho \in C([0,T];V)$ of the PDE (3).*

For notions of maximal monotonicity, we refer the reader to [4][Chapter 2]. This approach of establishing solutions of nonlinear PDEs using monotonicity properties of corresponding nonlinear elliptic operators is classical, but quite involved. Note that the solutions that we obtain using this approach are known to be only in $C([0,T];V)$. This is clearly a much weaker result than the one obtained for the linear diffusion model (1), in which case it was instead possible to conclude that solutions are in $C([0,T];L^2(\mathcal{M}))$ (Corollary IV.3).

Proposition IV.7. *Suppose $\rho_0 \in L^1(\mathcal{M})$. If $\int_{\mathcal{M}} \rho_0(\mathbf{x})dm(\mathbf{x}) = 1$, then the semigroup $(\mathcal{S}_a(t))_{t\geq 0}$ generated by B_a satisfies $\int_{\mathcal{M}}(\mathcal{S}_a(t)\rho_0)(\mathbf{x})dm(\mathbf{x}) = 1$ for all $t \geq 0$. Additionally, the semigroup is positive, i.e., $\rho_0 \geq 0$ implies that $S_a(t)\rho \geq 0$ for all $t \geq 0$.*

From this proposition, we can conclude the following main result on the asymptotic stability of the desired equilibrium distribution.

Theorem IV.8. *Suppose $\rho_0 \in L^1(\mathcal{M})$, $\int_{\mathcal{M}} \rho_0(\mathbf{x})dm(\mathbf{x}) = 1$, and $\rho_0 \geq 0$. Then the solution of the PDE (3), given by*

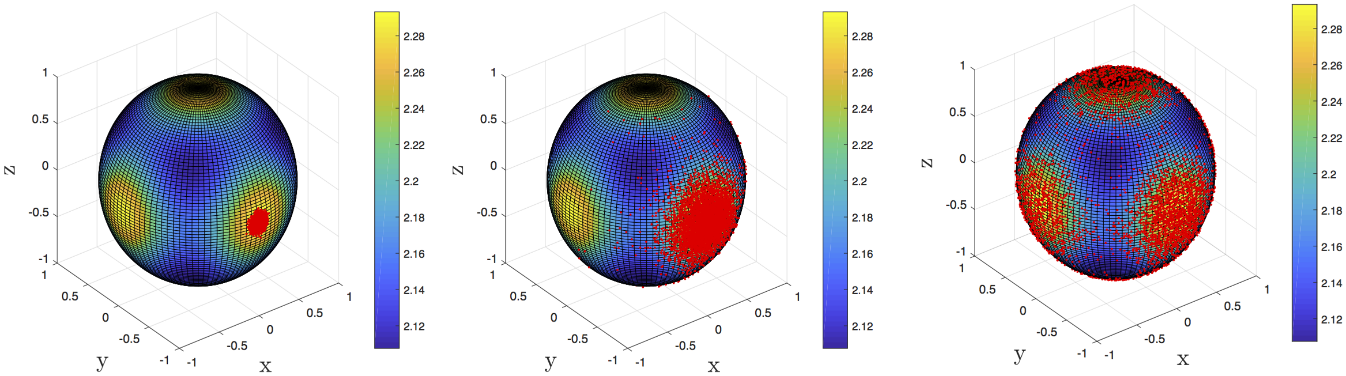

(a) $t = 0$ s

(b) $t = 2$ s

(c) $t = 10$ s

Fig. 1: Stochastic coverage of S^2 by $N = 5000$ agents (in red) at different times t, following a linear diffusion model

$\mathcal{S}_a(t)\rho_0$, *satisfies*

$$\lim_{t\to\infty} \|\mathcal{S}_a(t)\rho_0 - cf\|_V = 0, \tag{10}$$

where $c = 1/\int_{\mathcal{M}} f(\mathbf{x})dm(\mathbf{x})$ *and* $f = 1/a$.

This last result can be proven in the following way. We can directly use the functional $\Phi : V^* \to \mathbb{R} \cup \{\infty\}$ defined by

$$\Phi_a(\rho) = \begin{cases} \frac{1}{m+1}\int_{\mathcal{M}} |\rho(\mathbf{x})a(\mathbf{x})|^{\alpha+1}dm(\mathbf{x}), \\ \qquad\qquad \text{if } \rho \in L^{\alpha+1}(\mathcal{M}) \cap V^*, \\ \infty, \ \text{otherwise} \end{cases} \tag{11}$$

as a Lyapunov functional and invoke the relative compactness of the orbit of $\mathcal{S}_a(t)\rho$ to establish stability using Lasalle's invariance principle.

V. Simulation Results

In this section, we verify our theoretical results using numerical simulations.

A. Linear Diffusion

We validate the coverage approach based on linear diffusion with simulations on the two-dimensional sphere, denoted by S^2. We use the exponential coordinates [18] on the sphere to simulate the corresponding SDE. An advantage of using exponential local coordinates is that the Laplace-Beltrami operator, and hence the corresponding Brownian motion, takes a much simpler form in these coordinates. This is due to the fact that in exponential coordinates about the point $\mathbf{x} \in \mathcal{M}$, $g_{ij}(\mathbf{x}) = 0$ for $i \neq j$. Thus, in these coordinates, the SDE corresponding to the weighted diffusion equation (1) can be expressed in the highly simplified form

$$d\mathbf{X}(t) = \sqrt{2a(\mathbf{X}(t))}d\mathbf{W}(t) + O(r), \tag{12}$$

where $r = \sqrt{X_1^2 + ... + X_n^2}$.

For the scenario considered, the target density $f(\mathbf{x})$, shown in Fig. 1, is depicted on the surface of the sphere using a color density plot. Blue regions are assigned a low target density of agents, while yellow regions are assigned a high target density. The positions of $N = 5000$ agents are generated from a stochastic simulation of the SDE (12) and are superimposed on the density plot to enable comparison between the actual and target swarm distributions. As can be seen in Fig. 1, at time $t = 10$ s, the distribution of the swarm over the sphere is close to the target density.

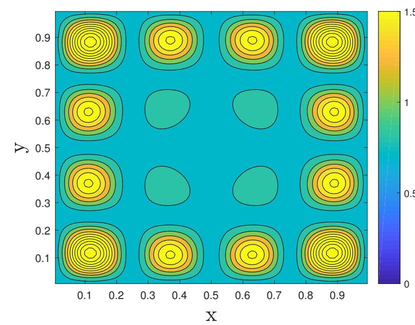

Fig. 2: Target density for the nonlinear diffusion model

B. Nonlinear Diffusion

We also validate the coverage approach based on nonlinear diffusion by numerically solving the corresponding PDE (3). We use the finite volume method [16] to solve the PDE.

The manifold $\mathcal{M}$ defined in this example is the 2-dimensional torus. Solving the PDE on the torus is equivalent to solving it on a rectangular domain with periodic boundary conditions. We set $\alpha = 2$ in the PDE. The target swarm density is shown in Fig. 2. Snapshots of the swarm density under diffusion at different times (i.e., the solution of the PDE) are shown in Fig. 3. As predicted by the previously stated asymptotic stability results, the solution is close to the target density for large enough times.

VI. Conclusions

In this paper, we presented two generalizations of a diffusion-based multi-agent coverage strategy. First, we extended the theory to general compact Riemannian manifolds without boundary. Then, we considered coverage based on

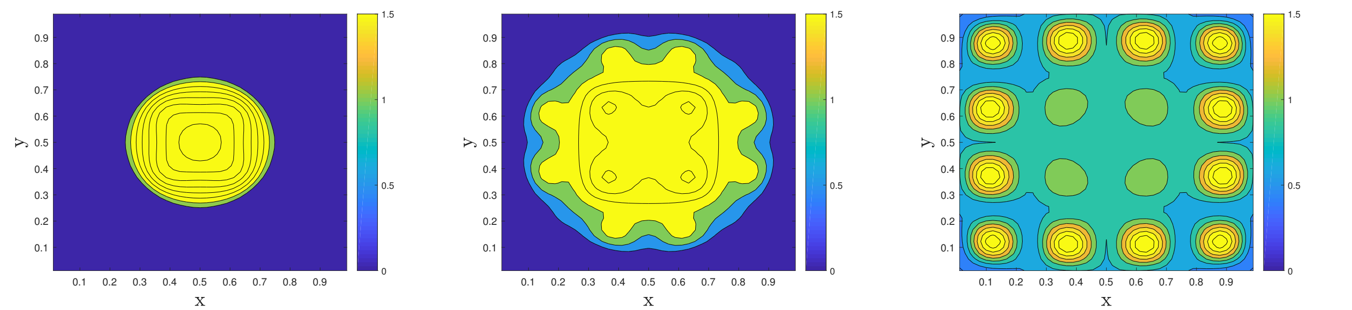

(a) $t = 2$ s (b) $t = 5$ s (c) $t = 10$ s

Fig. 3: Solution of the nonlinear diffusion model at different times t

nonlinear diffusion models. Numerical simulations verified the validity of the approaches. In future work, we will investigate the extension of this approach to manifolds with boundary. Another interesting future direction is to consider the case where each agent is nonholonomic. In addition, we will address the problem of constructing N stochastic processes, each governing the motion of an agent, such that the solution of the N–agent process converges to that of the porous media equation in a suitable sense as $N \to \infty$.

References

[1] Behçet Açikmeşe and David S Bayard. A Markov chain approach to probabilistic swarm guidance. In *American Control Conference (ACC), 2012*, pages 6300–6307. IEEE, 2012.

[2] Thierry Aubin. *Some nonlinear problems in Riemannian geometry*. Springer Science & Business Media, 2013.

[3] Saptarshi Bandyopadhyay, Soon-Jo Chung, and Fred Y Hadaegh. Probabilistic and distributed control of a large-scale swarm of autonomous agents. *IEEE Transactions on Robotics*, 33(5):1103–1123, 2017.

[4] Viorel Barbu. *Nonlinear differential equations of monotone types in Banach spaces*. Springer Science & Business Media, 2010.

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

[6] Subhrajit Bhattacharya, Robert Ghrist, and Vijay Kumar. Multi-robot coverage and exploration on Riemannian manifolds with boundaries. *The International Journal of Robotics Research*, 33(1):113–137, 2014.

[7] Jorge Cortes, Sonia Martinez, Timur Karatas, and Francesco Bullo. Coverage control for mobile sensing networks. *IEEE Transactions on Robotics and Automation*, 20(2):243–255, 2004.

[8] Karthik Elamvazhuthi, Chase Adams, and Spring Berman. Coverage and field estimation on bounded domains by diffusive swarms. In *IEEE Conf. on Decision and Control (CDC)*, pages 2867–2874. IEEE, 2016.

[9] Karthik Elamvazhuthi, Hendrik Kuiper, and Spring Berman. PDE-based optimization for stochastic mapping and coverage strategies using robotic ensembles. *Automatica*, 95:356–367, 2018.

[10] Karthik Elamvazhuthi, Hendrik Kuiper, Matthias Kawski, and Spring Berman. Bilinear controllability of a class of advection-diffusion-reaction systems. *To appear in IEEE Transactions on Automatic Control (TAC), arXiv preprint arXiv:1711.01696*, 2018.

[11] Klaus-Jochen Engel and Rainer Nagel. *One-parameter semigroups for linear evolution equations*, volume 194. Springer Science & Business Media, 2000.

[12] Paul Frihauf and Miroslav Krstic. Leader-enabled deployment onto planar curves: A PDE-based approach. *IEEE Transactions on Automatic Control*, 56(8):1791–1806, 2011.

[13] Alexander Grigoryan. *Heat kernel and analysis on manifolds*, volume 47. American Mathematical Soc., 2009.

[14] Emmanuel Hebey. *Nonlinear analysis on manifolds: Sobolev spaces and inequalities*, volume 5. American Mathematical Soc., 2000.

[15] M Ani Hsieh, Ádám Halász, Spring Berman, and Vijay Kumar. Biologically inspired redistribution of a swarm of robots among multiple sites. *Swarm Intelligence*, 2(2-4):121–141, 2008.

[16] Willem Hundsdorfer and Jan G Verwer. *Numerical solution of time-dependent advection-diffusion-reaction equations*, volume 33. Springer Science & Business Media, 2013.

[17] Nobuyuki Ikeda and Shinzo Watanabe. *Stochastic differential equations and diffusion processes*, volume 24. Elsevier, 2014.

[18] Jürgen Jost and Jèurgen Jost. *Riemannian geometry and geometric analysis*, volume 42005. Springer, 2008.

[19] John M Lee. Smooth manifolds. In *Introduction to Smooth Manifolds*, pages 1–29. Springer, 2003.

[20] Hanjun Li, Chunhan Feng, Henry Ehrhard, Yijun Shen, Bernardo Cobos, Fangbo Zhang, Karthik Elamvazhuthi, Spring Berman, Matt Haberland, and Andrea L Bertozzi. Decentralized stochastic control of robotic swarm density: Theory, simulation, and experiment. In *IEEE/RSJ Conf. on Intelligent Robots and Systems (IROS)*, pages 4341–4347, 2017.

[21] Jerrold E Marsden and Tudor S Ratiu. *Introduction to mechanics and symmetry: a basic exposition of classical mechanical systems*, volume 17. Springer Science & Business Media, 2013.

[22] Alexandre R Mesquita, João P Hespanha, and Karl Åström. Optimotaxis: A stochastic multi-agent optimization procedure with point measurements. In *International Workshop on Hybrid Systems: Computation and Control*, pages 358–371. Springer, 2008.

[23] Dejan Milutinovic and Pedro Lima. Modeling and optimal centralized control of a large-size robotic population. *IEEE Transactions on Robotics*, 22(6):1280–1285, 2006.

[24] Karl Oelschläger. Large systems of interacting particles and the porous medium equation. *Journal of Differential Equations*, 88(2):294–346, 1990.

[25] El-Maati Ouhabaz. *Analysis of heat equations on domains (LMS-31)*. Princeton University Press, 2009.

[26] Amanda Prorok, Nikolaus Correll, and Alcherio Martinoli. Multi-level spatial modeling for stochastic distributed robotic systems. *International Journal of Robotics Research*, 30(5):574–589, 2011.

[27] Jie Qi, Rafael Vazquez, and Miroslav Krstic. Multi-agent deployment in 3-D via PDE control. *IEEE Transactions on Automatic Control*, 60(4):891–906, 2015.

[28] Alain Sarlette and Rodolphe Sepulchre. Consensus optimization on manifolds. *SIAM Journal on Control and Optimization*, 48(1):56–76, 2009.

[29] Konrad Schmüdgen. *Unbounded self-adjoint operators on Hilbert space*, volume 265. Springer Science & Business Media, 2012.

[30] Rodolphe Sepulchre. Consensus on nonlinear spaces. *Annual Reviews in Control*, 35(1):56–64, 2011.

[31] Mark W Spong, Seth Hutchinson, and Mathukumalli Vidyasagar. *Robot modeling and control*. John Wiley and Sons, Inc., 2005.

[32] Juan Luis Vázquez. *The porous medium equation: mathematical theory*. Oxford University Press, 2007.