Optimal Control of Stochastic Coverage Strategies for Robotic Swarms

Karthik Elamvazhuthi[1] and Spring Berman[1]

Abstract— **This paper addresses a trajectory planning and task allocation problem for a swarm of resource-constrained robots that cannot localize or communicate with each other and that exhibit stochasticity in their motion and task-switching policies. We model the population dynamics of the robotic swarm as a set of advection-diffusion-reaction partial differential equations (PDEs), a linear parabolic PDE model that is bilinear in the robots' velocity and task-switching rates. These parameters constitute a set of time-dependent control variables that can be optimized and broadcast to the robots prior to their deployment. The planning and allocation problem can then be formulated as a PDE-constrained optimization problem, which we solve using techniques from optimal control. Simulations of a commercial pollination scenario validate the ability of our control approach to drive a robotic swarm to achieve predefined spatial distributions of activity over a closed domain, which may contain obstacles.**

I. INTRODUCTION

Swarm robotic systems comprised of hundreds or thousands of inexpensive, relatively expendable platforms have the potential to collectively perform tasks on large spatial and temporal scales quickly, robustly, adaptively, and autonomously. The production and deployment of robotic swarms is approaching feasibility due to recent advances in computing, sensing, actuation, power, control, and 3D printing technologies [1]. In the last few years, the miniaturization of these technologies has led to many novel platforms for swarm robotic applications, including micro aerial vehicles (MAVs) and nano air vehicles (NAVs) [2], [3] for tasks such as exploration, mapping, environmental monitoring, chemical source localization, search-and-rescue, surveillance, and reconnaissance. At even smaller scales, micro-nano systems composed of nanoparticles, DNA machines, synthetic bacteria, and magnetic materials are currently being developed for micro object manipulation and biomedical applications, including molecular imaging, drug and gene delivery, therapeutics, and diagnostics [4], [5], [6], [7].

While the technology to create robotic swarms is progressing, open problems remain in the control of such systems for successful operation with theoretical guarantees on the collective performance. It may be impractical or impossible for robots to use maps, global position information, or inter-robot communication when they move through unstructured, unpredictable, and potentially hostile environments, such as during disaster response operations and intelligence-surveillance-reconnaissance missions. However, it remains a challenge to deploy autonomous robots that can perform tasks in environments where prior data is unavailable and the GPS signal and radio communications are limited or unreliable [8], [9]. In addition, the robot control policies may need to be implemented on platforms with extreme constraints on sensing, computing, and communication, including micro- and nanoscale robots and MAVs/NAVs such as the RoboBee, a recently developed insect-scale flapping-wing MAV [3]. Control frameworks for robotic swarms must be *scalable* with the robot population size and *adaptable* to changes in target collective behaviors, environmental disturbances, and robot failures and errors. The control policies must accommodate *non-deterministic behaviors* that arise in autonomous systems [10]. Stochasticity can arise from noise due to sensor and actuator errors; inherent randomness in robot encounters with each other and with environmental features; and, for nanorobots, the effects of Brownian motion and chemical interactions at scales below tens of micrometers [4].

Despite a wide range of research on multi-robot systems, little work in the field has addressed the synthesis of provably correct robot controllers for *both robot motion and task allocation simultaneously* in the realistic scenarios that we have described above. Various approaches to multi-robot task allocation have, however, received considerable attention in the literature. Our proposed methodology addresses the problem of the single-task robot, multi-robot task (ST-MR) problem [11], in which robots can execute at most one task at a time and tasks require multiple robots. *Stochastic* approaches to the ST-MR problem for robotic swarms have been developed in which tasks are executed at random times by unidentified robots and an allocation emerges from the collective swarm activity. These approaches include threshold-based algorithms [12], [13], [14], inspired by division of labor mechanisms in social insects, and works that optimize the stochastic robot task-switching rates using *non-spatial* macroscopic models that describe the time evolution of the robot population in each state [15], [16], [17], [18]. Our proposed approach employs *spatial* macroscopic swarm models similar to those used in [19], [20], which are based on the Fokker-Planck partial differential equation; however, unlike our approach, these works do not address the problem of controller optimization. Optimal control of a spatial swarm model is considered in [21] for the specific purpose of directing the swarm to a desired location.

In [22], [23], we presented a rigorous methodology for optimizing the motion and task switching policies of a spatially inhomogeneous robotic swarm to achieve target coverage

This work was supported in part by National Science Foundation (NSF) award no. CMMI-1436960.

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

objectives. Our approach incorporates both individual-level (*microscopic*) and population-level (*macroscopic*) models that describe the robots' roles, task transitions, and motion. The macroscopic model takes the form of a set of advection-diffusion-reaction partial differential equations (PDEs) that govern the swarm population dynamics. The parameters of the PDE model are optimized to produce a specified spatial distribution of activity over a domain by a swarm of robots executing the corresponding control policies. We validated the control framework through simulations of an MAV swarm in a commercial pollination scenario where the robots must achieve a target density field of flower visits over several rows of crops.

This paper employs the control framework and pollination application that we previously described in [22], [23]. Whereas our prior work used stochastic optimization methods to compute the parameters of the PDE model, in this paper we present an approach to computing these parameters using *optimal control techniques*. Optimal control guarantees local optimality and converges to the same solution each time, if it converges. This approach computes the control parameters much more quickly than our previously implemented stochastic optimization method, making it possible to use for real-time control.

Our model and the control parameters form a *bilinear control system*, in which the control acts multiplicatively with the state [24]. There has been a considerable amount of work on bilinear control systems in the finite-dimensional case, and a number of works in the infinite-dimensional case of PDEs has also been carried out. Several studies have been conducted on the problem of optimal control of bilinear systems. The stabilization problem in the PDE case has also received some attention. This includes several controllability studies of bilinear control systems [25] like the wave equation and reaction-diffusion models [26]. A long-standing problem in this field has been the controllability of quantum systems [27]. Closely related to such systems is the concept of ensemble control, which has been exploited to achieve global control of multi-agent systems [28]. An optimal control approach using space- and time-dependent velocity fields [29] has been used to address a multi-agent trajectory planning problem. The problem that we address is a variant of the concurrent planning and task allocation problem addressed in [30]. In contrast to the work in [30], the dimensionality of our robot controller state space is independent of the number of robots, and the controls are broadcast to the entire swarm at once rather than communicated to each individual robot.

II. PROBLEM STATEMENT

We consider a scenario in which a swarm of robotic bees must pollinate several rows of crops. We aim to design robot control policies that produce a uniform density of flower visits along crop rows, and that can achieve any ratio between numbers of flower visits at plants in different rows. In contrast to our previous work in [22], [23], we consider environments that are bounded rather than unbounded and that may contain obstacles.

A. Robot capabilities

The robots would have sufficient power to undertake brief flights that originate from a location called the *hive*, and they would return to the hive to recharge. A computer at the hive can serve as the supervisory agent in our architecture. The computer calculates the parameters of the robot motion and task transitions for a specified pollination objective and transmits these parameters to the robots when they are docked at the hive for charging and uploading data. During a flight, the robots are assumed to be capable of recognizing a flower that is very close by, distinguishing between different types of flowers, flying to a flower, and hovering briefly while obtaining pollen from the flower using an appropriate appendage. Each robot is equipped with a compass and thus can fly with a specified heading. We also assume that robots can detect obstacles within their local sensing range and adjust their flight path to avoid collision. Notably, the robots are not assumed to have localization capabilities, since it is infeasible to use GPS sensors on highly power-constrained platforms.

B. Robot controller

Each member of a swarm of N robots performs the following actions during a flight. Upon deploying from the hive, each robot flies with a time-dependent velocity $\mathbf{v}(t) \in \mathbb{R}^2$. Concurrently with this deterministic motion, the robot exhibits random movement that arises from inherent noise due to sensor and actuator errors. We assume that the flowers are distributed densely enough such that a robot can always detect at least one flower in its sensing range when it flies over plants. While a robot is flying over a row with flowers of type j, it decides with a time-dependent probability per unit time $k_j(t)$ to pause at a flower in its sensing range and hover for pollination. The robot resumes flying with a fixed probability per unit time k_f, which determines the time taken to pollinate. The optimal control approach described in section IV computes the parameters $\mathbf{v}(t)$ and $k_j(t)$ prior to the robots' flight.

III. MODELS OF THE COVERAGE SCENARIO

A. Microscopic Model

The microscopic model is used to simulate the individual robots' motion and probabilistic decisions that are produced by the robot controller in section II-B. We model a robot's changes in state as a Chemical Reaction Network (CRN) in which the species are F, a flying robot; H_j, a robot that is hovering over a flower of type j; and V_j, an instance of a robot visit to a flower of type j. The reactions are:

$$F \xrightarrow{k_j(t)} H_j + V_j \tag{1}$$

$$H_j \xrightarrow{k_f} F \tag{2}$$

A robot i has position $\mathbf{x}_i(t) = [x_i(t)\ y_i(t)]^T$ at time t. The deterministic motion of each flying robot is governed by

the time-dependent velocity field $\mathbf{v}(t) = [v_x(t)\ v_y(t)]^T$. The robot's random movement is modeled as a Brownian motion that drives diffusion with an associated diffusion coefficient D, which we assume that we can characterize. We model the displacement of the robot over each timestep Δt using the standard-form Langevin equation [31],

$$\mathbf{x}_i(t+\Delta t) - \mathbf{x}_i(t) = \mathbf{v}(t)\Delta t + (2D\Delta t)^{1/2}\ \mathbf{Z}(t), \qquad (3)$$

where $\mathbf{Z} \in \mathbb{R}^2$ is a vector of independent, normally distributed random variables with zero mean and unit variance. When a robot encounters an obstacle or a wall, it avoids a collision by flying according to a specular reflection from the boundary.

B. Macroscopic Model

We can describe the time evolution of the expected spatial distribution of the swarm with a *macroscopic model* consisting of a set of advection-diffusion-reaction (ADR) partial differential equations [32]. The states of this model are the population density fields $y_1(\mathbf{x},t)$ of flying robots, $y_2(\mathbf{x},t)$ of hovering robots, and $y_3(\mathbf{x},t)$ of flower visit events. The velocity field $\mathbf{v}(t)$ and transition rates $k_j(t)$ are time-dependent control parameters. The model is defined over a bounded domain, $\Omega \subset \mathbb{R}^2$, with Lipschitz continuous boundary $\partial\Omega$. We define $Q = \Omega \times (0,T)$ and $\Sigma = \partial\Omega \times (0,T)$ for some fixed final time T. The vector $\mathbf{n} \in \mathbb{R}^2$ is the outward normal to $\partial\Omega$. There are n_f types of flowers, and the function $H_i : \Omega \to \{0,1\}$ is a spatially-dependent coefficient that models the presence ($H_i(\mathbf{x}) = 1$) or absence ($H_i(\mathbf{x}) = 0$) of flowers of type i at point $\mathbf{x}$ in the domain.

Given these definitions, the macroscopic model of the pollination scenario is defined as:

$$\begin{aligned}
\frac{\partial y_1}{\partial t} &= \nabla \cdot (D\nabla y_1 - \mathbf{v}(t)y_1) - \sum_{i=1}^{n_f} k_i H_i y_1 + k_f y_2 \ \ in\ \ Q,\\
\frac{\partial y_2}{\partial t} &= \sum_{i=1}^{n_f} k_i H_i y_1 - k_f y_2 \ \ in\ \ Q,\\
\frac{\partial y_3}{\partial t} &= \sum_{i=1}^{n_f} k_i H_i y_1, \qquad (4)
\end{aligned}$$

with the no-flux boundary conditions

$$D\partial_{\mathbf{n}} y_1 - \mathbf{v}(t)y_1 \cdot \mathbf{n} = 0 \ \ on\ \ \Sigma. \qquad (5)$$

Initially, the flying robots are distributed according to a Gaussian density centered at a point $\mathbf{x}_0$, and there are no hovering robots or visits in the domain.

We numerically solve the macroscopic model with an explicit finite-volume method in which the advection term is solved using the Lax-Wendroff method, implemented with a superbee flux limiter to prevent spurious oscillations [33]. We numerically integrate the model using the method of lines and solve the resulting ordinary differential equations using the ode45 function in MATLAB.

IV. Optimal Control of Robot Behaviors

Our algorithm for computing optimal control policies is based on the well-known gradient descent method. Methods of optimal control help to reduce the amount of computation that is required to compute the gradient of the objective functional with respect to the control, subject to constraints in the form of differential equations. This is done using the adjoint state equation. In the case of finite-dimensional systems, the adjoint/co-state equation can be derived using the Hamiltonian and Pontryagin's maximum principle. However, in the infinite-dimensional case, a general maximum principle does not exist and the existence of the Hamiltonian has been proved only for a limited class of systems. We instead derive the directional derivatives of the control-to-state mapping and use the generalized chain rule of differentiation of composite mappings in Banach spaces as is common in the literature [34], [35]. For this purpose, we must define the spaces that are relevant for our analysis. This will aid us in establishing that our candidate for the derivative of the control-to-state mapping is indeed the derivative. As proven in [36], an optimal control exists for this problem.

A. Preliminaries

We define the Sobolev space $H^1(\Omega) = \{z \in L^2(\Omega) : \frac{\partial z}{\partial x_1} \in L^2(\Omega), \frac{\partial z}{\partial x_2} \in L^2(\Omega)\}$ [37]. Here, the spatial derivative is to be understood as a weak derivative defined in the distributional sense. We equip the space with the usual Sobolev norm, $\|y\|_{H^1(\Omega)} = (\|y\|^2_{L^2(\Omega)} + \sum_{i=1}^2 \|\frac{\partial y}{\partial x_i}\|^2_{L^2(\Omega)})^{1/2}$. The space $H^1(\Omega)^*$ is the dual space of $H^1(\Omega)$, the space of bounded linear functionals on $H^1(\Omega)$ through the inner product of $L^2(\Omega)$. Furthermore, let $X = H^1(\Omega) \times L^2(\Omega)^n$. Then we have $X^* := H^1(\Omega)^* \times L^2(\Omega)^n$. X and X^* inherit their topology from their factors.

We consider the system Equation (4) in the following general form for notational convenience:

$$\begin{aligned}
\frac{\partial y}{\partial t} &= Ay + \sum_{i=1}^{m+2} u_i B_i y + f \ \ in\ \ Q,\\
\vec{n} \cdot (\nabla y_1 - \vec{u}_b y_1) &= g \ \ in\ \ \Sigma,\\
y(0) &= y_0, \qquad (6)
\end{aligned}$$

where A and $B_i : L^2(0,T;X) \to L^2(0,T;L^2(\Omega)^{1+n})$ are operators, $f \in F = L^2(0,T;L^2(\Omega)^{1+n})$, $g \in G = L^2(0,T;L^2(\partial\Omega)$, and $y_0 \in L^2(\Omega)^{1+n}$. We define $A_g : L^2(0,T;X) \to L^2(0,T;X^*)$ as the variational form of the operator A. We understand a function $y \in Y = L^2(0,T;X)$, with $y_t \in Y^* = L^2(0,T;X^*)$, to be a weak solution of the system provided that:

$$\langle \frac{\partial y}{\partial t}, \phi \rangle_{Y^*,Y} = \langle A_g y, \phi \rangle_{Y^*,Y} + \sum_{i=1}^{m+2} \langle u_i B_i y, \phi \rangle_F + \langle f, \phi \rangle_F \qquad (7)$$

for all $\phi \in L^2(0,T;X)$.

For $n = 2$, we have the following form for the operators

A and B_i:

$$A = \begin{bmatrix} \nabla^2 & k_f & 0 \\ 0 & -k_f & 0 \\ 0 & 0 & 0 \end{bmatrix} \quad B_1 = \begin{bmatrix} -\frac{\partial}{\partial x_1} & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$
$$B_2 = \begin{bmatrix} -\frac{\partial}{\partial x_2} & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$
$$B_i = \begin{bmatrix} -H_{i-2} & 0 & 0 \\ H_{i-2} & 0 & 0 \\ H_{i-2} & 0 & 0 \end{bmatrix} \qquad 3 \le i \le k_f + 2 \tag{8}$$

We know that $y \in L^2(0,T;X)$ and $y_t \in L^2(0,T;X^*)$ imply that $y \in C([0,T];L^2(\Omega)^{1+n})$. Here, $C([0,T];Y)$ is equipped with the norm $\|y\|_{C([0,T];Y)} = \sup_{t\in[0,T]} \|y(t)\|_Y$. The boundary conditions are equipped with A_g in the variational formulation using Green's theorem as,

$$A_g = \begin{bmatrix} M_g & k_f & 0 \\ 0 & -k_f & 0 \\ 0 & 0 & 0 \end{bmatrix} \tag{9}$$

Here, $M_g : L^2(0,T;V) \to L^2(0,T;V^*)$ is the Laplacian in the variational form and is defined as,

$$\langle M_g y, \phi \rangle_{V^*,V} = -\langle D\nabla y, \nabla\phi \rangle_{L^2(\Omega)} + \int_{\partial\Omega} (g + \vec{n}\cdot\vec{b}y)\phi dx \tag{10}$$

The solution of Equation (4) corresponds to A_0 and $f = 0$. The controls are $\vec{v} = (u_1, u_2) = \vec{u}_b$, $u_i = k_{i-2}$ for $3 \le i \le m+2$ and $m = n_f$.

B. The Optimal Control Problem

Our optimal control problem can be framed as follows:

$$\min_{(y,u)\in Y\times U_{ad}} J(y,u) = \frac{1}{2}\|Wy(\cdot,T) - y_\Omega\|^2_{L^2(\Omega)^{1+n}} + \frac{\lambda}{2}\|u\|^2_{L^2(0,T)^m} \tag{11}$$

subject to Equation (6) for $f = 0$ and $g = 0$. Here, y_Ω is the target spatial distribution of robot activity, $Y = C([0,T], L^2(\Omega)^{1+n})$, and

$$U_{ad} = \{u \in L^2(0,T)^{m+2};\ u_i^{min} \le u_i \le u_i^{max}\ \ a.e.\ \ in\ \ (0,T)\}$$

is the set of admissible control inputs. Note that, due to the essential bounds on u, we have that $u \in L^\infty(0,T)^{m+2}$. Additionally, we take $W \in \mathscr{L}(L^2(\Omega)^{m+2})$. W is typically a weighting function that weights the relative significance of minimizing the distance between different states and their targets.

Lemma 4.1: Given $f \in L^2(0,T;L^2(\Omega))^{1+n}$, $g \in L^2(0,T;L^2(\partial\Omega))$, and the initial condition $y_0 \in L^2(\Omega)^{1+n}$, a unique solution exists for the problem in Equation (6). We have the following estimate for the unique solution y in $C([0,T];L^2(\Omega)^{1+n})$:

$$\|y\|_{C([0,T];L^2(\Omega)^{1+n})} + \|y\|_{L^2(0,T;X)} \le K(\|y_0\|_{L^2(\Omega)^{1+n}} + \|f\|_{L^2(0,T;L^2(\Omega))} + \|g\|_{L^2(0,T;L^2(\partial\Omega))}, \tag{12}$$

where K depends only on Ω, $\max_{1\le i\le m+2} |u_i^{max}|$, $\max_{1\le i\le m+2} |u_i^{min}|$, and $\max_{1\le i\le m+2} |b_i|$.

Proof: See [36]. ∎

Now we introduce the control-to-state mapping, Ξ: $U_{ad} \to Y$, which maps a control, u, to y, the corresponding solution defined through Equation (6) for $f = 0$. This will help us in studying the differentiability of Ξ in U_{ad} and consequently the differentiability of the objective functional, J.

Proposition 4.2: The mapping Ξ is Gateaux differentiable at every $u \in U_{ad}$, and its Gateaux derivative, $\Xi'(u) : U_{ad} \to Y$, evaluated at $h \in U_{ad}$, i.e. $\Xi'(u)h$, is given by the solution of the following equation:

$$\begin{aligned} &\frac{\partial w}{\partial t} = Aw + \sum_{i=1}^{m+2} u_i B_i w + \sum_{i=1}^{m+2} h_i B_i y \\ &\vec{n}\cdot(\nabla w_1 - \vec{u_b}\cdot w_1) = \vec{n}\cdot(\vec{h_b} y_1) \\ &w(0) = 0. \end{aligned} \tag{13}$$

Proof: We define $y_\varepsilon = \Xi(u + \varepsilon h)$. We show that $y_\varepsilon \to y$ as $\varepsilon \to 0$. Define $g = y_\varepsilon - y$. Then we have,

$$\begin{aligned} &\frac{\partial g}{\partial t} = Ag + \sum_{i=1}^{m+2} (u_i + \varepsilon h) B_i g + \varepsilon \sum_{i=1}^{m+2} h_i B_i y \\ &\vec{n}\cdot(\nabla g_1 - (\vec{u_b} + \varepsilon\vec{h_b})\cdot g_1) = \vec{n}\cdot(\varepsilon\vec{h_b} y_1) \\ &g(0) = 0. \end{aligned} \tag{14}$$

For ε sufficiently small, $u + \varepsilon h \in U_{ad}$. Thus, it follows from Lemma 4.1 that

$$\|g\|_{C([0,T];L^2(\Omega)^{1+n})} \le C(\|\varepsilon \sum_{i=1}^{m+2} h_i B_i y\|_{L^2(Q)^{1+n}} + \|\varepsilon y_1\|_{L^2(\Omega)})$$

and so

$$\|g\|_{C([0,T];L^2(\Omega)^{1+n})} \le \varepsilon K(\|y\|_X + \|y_1\|_{L^2(\Omega)}),$$

where K is a constant.

Hence, $y_\varepsilon \to y$ as $\varepsilon \to 0$. Next, we define $z = g/\varepsilon - w$. Then, it is required to prove that $z \to 0$ as $\varepsilon \to 0$. From the definition of z, we get

$$\begin{aligned} &\frac{\partial z}{\partial t} = Az + \sum_{i=1}^{m+2} u_i B_i z + \sum_{i=1}^{m+2} h_i B_i g \\ &\vec{n}\cdot(\nabla z_1 - \vec{u_b}\cdot z_1) = \vec{n}\cdot(\vec{h_b} g_1) \\ &z(0) = 0. \end{aligned} \tag{15}$$

Invoking Lemma 4.1, since $g \to 0$, we infer that $z \to 0$ as $\varepsilon \to 0$ and hence, $g/\varepsilon \to w$. ∎

We now consider the reduced problem,

$$\min_{u\in U_{ad}} \hat{J}(u) := J(\Xi(u), u) \tag{16}$$

Let $A^\#$ and $B_i^\#$ be the formal adjoints of operators A and B_i respectively.

Theorem 4.3: The reduced objective functional $\hat{J}$ is differentiable in the Gateaux sense, and the derivative has the form

$$\langle \hat{J}'(u), h\rangle_{L^2(0,T)^{m+2}} = \int_0^T \langle \vec{n}\cdot(\vec{h_b}p_1), y_1\rangle_{L^2(\partial\Omega)} + \int_0^T \langle \sum_{i=1}^{m+2} h_i B_i y, p\rangle_{L^2(\Omega)^{1+n}} + \lambda\langle u, h\rangle_{L^2(0,T)^{m+2}}, \quad (17)$$

where p is the solution of the backward-in-time adjoint equation,

$$\begin{aligned} -\frac{\partial p}{\partial t} &= A^{\#}p + \sum_{i=1}^{m+2} u_i B_i^{\#} p \\ \vec{n}\cdot\nabla p_1 &= 0 \\ p(T) &= W^*(Wy(\cdot,T) - y_\Omega). \end{aligned} \quad (18)$$

Proof: We use the generalized chain rule of differentiation of operators in Banach spaces to prove the above result.

Consider $G : C([0,T];L^2(\Omega)^{1+n}) \to L^2(\Omega)^{1+n}$, which maps the state to its final value. This linear continuous mapping is well-defined for functions in the domain $C([0,T];L^2(\Omega)^{1+n})$ due to continuity in time over a compact set.

Using the chain rule of differentiation [34], [38] (since $\hat{J}$ is Frechet differentiable and Ξ is Gaeteaux differentiable), the Gateaux derivative of $\hat{J}$ is given by

$$\langle \hat{J}'(u), h\rangle = \langle J_y(y,u), \Xi'(u)h\rangle + \langle J_u(y,u), u\rangle, \quad (19)$$

which is equal to

$$\langle \hat{J}'(u), h\rangle = \langle G^*W^*(WGy - y_\Omega), w\rangle + \lambda\langle u, h\rangle. \quad (20)$$

Thus we have,

$$\langle \hat{J}'(u), h\rangle = \langle W^*(WGy - y_\Omega), Gw\rangle + \lambda\langle u, h\rangle. \quad (21)$$

Then,

$$\langle \hat{J}'(u), h\rangle = \langle p(\cdot,T), w(\cdot,T)\rangle + \lambda\langle u, h\rangle.$$

Consider the term $\langle p(\cdot,T), w(.,T)\rangle$. Using integration by parts in time, we find that this term is:

$$\int_0^T \langle \frac{\partial p}{\partial t}, w\rangle + \int_0^T \langle p, \frac{\partial w}{\partial t}\rangle + \langle p(0), w(0)\rangle$$

and hence is equal to:

$$\int_0^T \langle \frac{\partial p}{\partial t}, w\rangle + \int_0^T \langle p, A_0 w + \sum_{i=1}^{m+2} u_i B_i w + \sum_{i=1}^{m+2} h_i B_i y\rangle.$$

Let us now define the formal adjoints of these operators,

$$A_0^{\#} = \begin{bmatrix} M_0^{\#} & 0 & 0 \\ k_f & -k_f & 0 \\ 0 & 0 & 0 \end{bmatrix}, \quad (22)$$

such that $M_0^{\#} : L^2(0,T;V) \to L^2(0,T;V^*)$ is given by

$$\langle M_0^{\#} y, \phi\rangle_{V^*,V} = -\langle D\nabla y, \nabla\phi\rangle_{L^2(\Omega)}. \quad (23)$$

Equation (18) has a solution in the weak sense, and

$$-\langle \frac{\partial p}{\partial t}, \phi\rangle = \langle A_0^{\#} p, \phi\rangle + \sum_{i=1}^{m+2} \langle u_i B_i^{\#} p, \phi\rangle \quad (24)$$

for all $\phi \in L^2(0,T;X)$ The previous step can be written as,

$$\int_0^T \langle \frac{\partial p}{\partial t}, w\rangle + \int_0^T \langle A_0^{\#} p + \sum_{i=1}^{m} u_i B_i^{\#} p, w\rangle + \int_0^T \vec{n}\cdot(\vec{h_b}py) + \int_0^T \langle p, \sum_{i=1}^{m+2} h_i B_i y\rangle.$$

It follows that

$$\langle p(\cdot,T), w(.,T)\rangle = \int_0^T \int_{\partial\Omega} \vec{n}\cdot(\vec{h_b}p_1 y_1) + \int_0^T \langle p, \sum_{i=1}^{m} h_i B_i y\rangle,$$

and hence we have our result. ∎

The adjoint state equation for the system defined in Equation (4) with respect to the objective functional, J, is therefore given by:

$$\begin{aligned} -\frac{\partial p_1}{\partial t} &= \nabla\cdot(D\nabla p_1 + \mathbf{v}(t)p_1) + \sum_{i=1}^{n_f} k_i H_i(-p_1 + p_2 + p_3) \ \ in \ \ Q, \\ -\frac{\partial p_2}{\partial t} &= k_f p_1 - k_f p_2 \ \ in \ \ Q, \\ -\frac{\partial p_3}{\partial t} &= 0 \ \ in \ \ Q, \end{aligned} \quad (25)$$

with the Neumann boundary conditions

$$\vec{n}\cdot\nabla p_1 = 0 \ \ on \ \ \Sigma \quad (26)$$

and final time condition

$$p(T) = W^*(Wy(\cdot,T) - y_\Omega). \quad (27)$$

V. SIMULATED POLLINATION SCENARIO

A. Simulation Setup

We developed microscopic and macroscopic models of scenarios in which a swarm of robots is tasked to achieve a specified spatial distribution of flower visits over five crop rows. We considered four different scenarios. We computed optimal control parameters of the macroscopic model to achieve two types of target spatial distributions of visits over the field: one in which visits were required along all five crop rows (*Objective 1*), and another in which they were required only along two of the crop rows (*Objective 2*). For both objectives, we simulated an environment with and without obstacles to investigate the effect of the boundary conditions on the optimized robot control policies.

For each scenario, we simulated 1000 robots in a normalized domain of size 1 m × 1 m. We set $k_f = 0.2$ to define an expected pollination time of $k_f^{-1} = 5$ s, a realistic value for certain bee species [22]. In the optimization, the robot speed was bounded between -10 and 10 m/s, and the transition rates k_j were bounded between 0 and 1.25 s^{-1}. The microscopic model was simulated over a grid of 21×21 cells. To account for numerical diffusion, the partial differential equation was simulated over a finer grid of 51×51 cells.

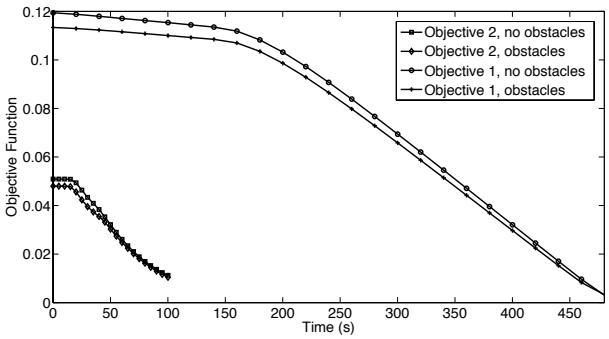

Fig. 1. Objective function over time for all four scenarios

B. Results

Figure 1 shows that in all four scenarios, our optimal control approach successfully minimizes the objective function, driving it nearly to zero in the time span allotted for the simulation. The resulting optimized parameters over time are plotted in Figure 2.

The top plot of Figure 2 corresponds to the Objective 1 case, in which crops rows 2 and 4 were assigned twice as high a target density of flower visits as rows 1, 3, and 5. The robots start at the bottom of the field in this case and have 480 s to achieve the target density. The robot speed is kept almost at zero throughout the optimization run; the robots' motion is dominated by diffusion, and after the robots diffuse over the entire domain (at 150 s), the transition rates are increased to approximately constant levels. The transition rate k_2, implemented when a robot is over row 2 or 4, is driven to about the twice the value of k_1, implemented for rows 1, 3, and 5, which results in twice as many flower visits in rows 2 and 4.

The bottom plot of Figure 2 corresponds to the Objective 2 case, in which the target visit density is set to zero in rows 1, 2, and 3 and to a nonzero value in rows 4 and 5 (the rightmost two rows). The time allotted for the task is 100 s, a shorter time than for Objective 1. In this case, the robots start at the left end of the field, and their optimized speed in the positive x direction is kept high to drive them quickly to the right side of the field. The transition rate k_1 increases as the robots slow down in the x direction, causing them to focus the bulk of their flower visits on the rightmost two rows. The robots aggregate against the right boundary as they encounter it and start moving with a leftward velocity component around 37 s, which directs them again over the rows to be pollinated.

Figure 3 through Figure 6 compare snapshots of the microscopic simulations (left columns) and macroscopic model numerical solutions (right columns) for each scenario. The corresponding snapshots of the two models are similar in each scenario, which validates the ability of our macroscopic model to predict the behavior of an ensemble of individual robots. Although the obstacles block robots and force them to fly along altered paths, the resulting distributions of flower visits still approximately conform to the target distributions.

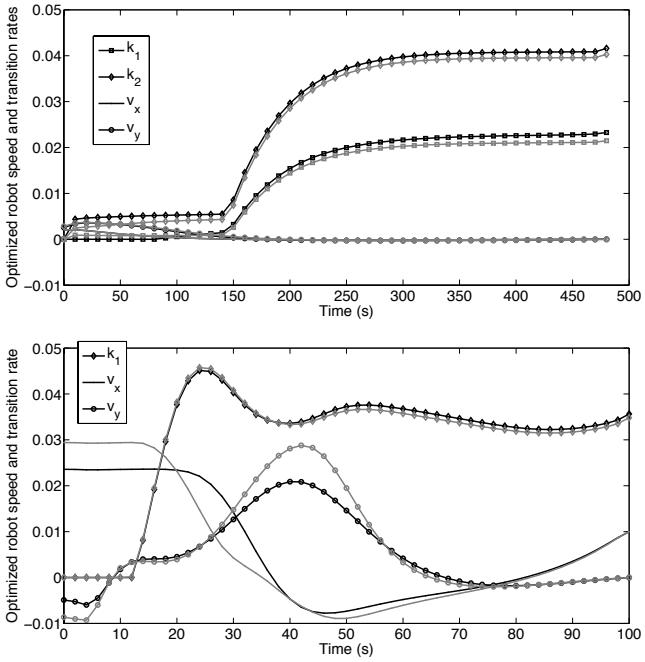

Fig. 2. Optimized robot parameters for Objective 1 (top) and Objective 2 (bottom). Gray plots show parameters for environments with obstacles, and black plots show parameters for environments without obstacles.

VI. CONCLUSIONS

In this work, we have presented an optimal control approach to achieving target spatial distributions of robot activity over a domain when the robots have minimal capabilities: local sensing, heading information, and no global position information or communication. In our simulated scenarios, we have assumed that the environment is known beforehand and that there are no environmental disturbances or robot failures. In future work, we will consider scenarios in which the robots must identify the distributions of features of interest from their sensor data before executing their tasks (for instance, mapping regions where flowers are blooming in order to have information on where they should pollinate). The robot control strategies should also be able to incorporate feedback from the robots in order to fulfill the coverage task in the presence of unknown environmental disturbances, such as wind in the pollination scenario. For such scenarios, we must identify types of observers with minimal measurement costs that will provide sufficiently rich state reconstruction to enable real-time control in a broadcast control framework. As is often the case for infinite-dimensional systems, exact observability will not be possible unless all agents communicate back their state estimates. Some related work on industrial processes with similar controls has been done in [39]. Finally, we can expand the types of control schemes that we consider to include ones with inter-agent interactions, which is a common feature in PDE models of natural coordinated behaviors such as flocking, schooling, and taxis [40], [41].

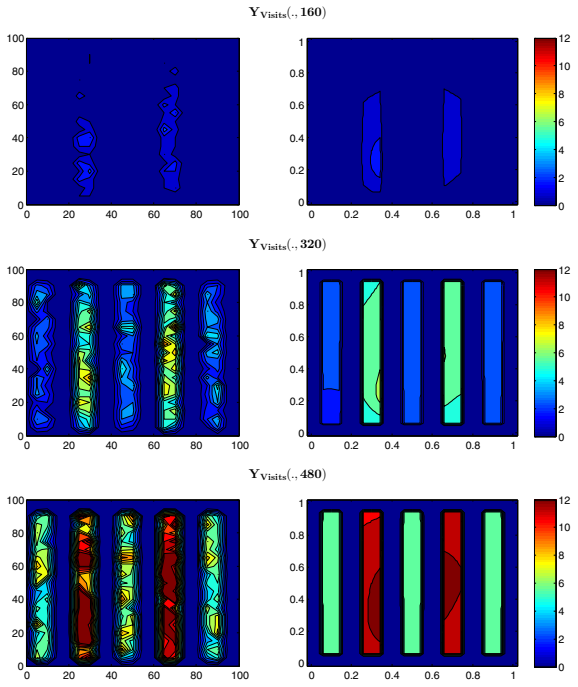

Fig. 3. Distribution of flower visits at three times in the microscopic (left) and macroscopic (right) models with parameters optimized for Objective 1, no obstacles

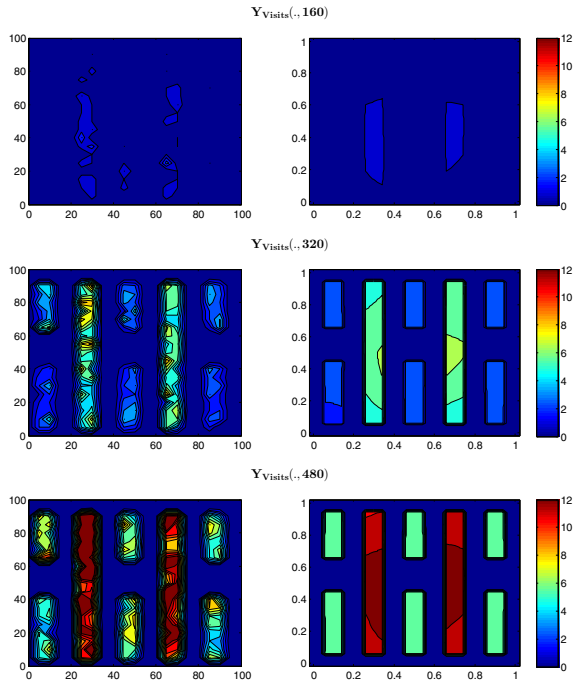

Fig. 4. Distribution of flower visits at three times in the microscopic (left) and macroscopic (right) models with parameters optimized for Objective 1, with obstacles

REFERENCES

[1] R. Bogue, "3D printing: the dawn of a new era in manufacturing?" *Assembly Automation*, vol. 33, no. 4, pp. 307–311, 2013.

[2] V. Kumar and N. Michael, "Opportunities and challenges with autonomous micro aerial vehicles," *Int'l. Journal of Robotics Research*, vol. 31, no. 11, pp. 1279–1291, 2012.

[3] K. Y. Ma, P. Chirarattananon, S. B. Fuller, and R. J. Wood, "Controlled flight of a biologically inspired, insect-scale robot," *Science*, vol. 340, no. 6132, pp. 603–607, 2013.

[4] E. Diller and M. Sitti, "Micro-scale mobile robotics," *Foundations and Trends in Robotics*, vol. 2, pp. 143–259, 2013.

[5] S. Tong, E. J. Fine, Y. Lin, T. J. Cradick, and G. Bao, "Nanomedicine: Tiny particles and machines give huge gains," *Annals of Biomedical Engineering*, pp. 1–17, 2013.

[6] C. Mavroidis and A. Ferreira, "Nanorobotics: Past, present, and future," in *Nanorobotics*, C. Mavroidis and A. Ferreira, Eds. Springer New York, 2013, pp. 3–27.

[7] S. Hauert, S. Berman, R. Nagpal, and S. N. Bhatia, "A computational framework for identifying design guidelines to increase the penetration of targeted nanoparticles into tumors," *Nano Today*, vol. 8, no. 6, pp. 566–576, 2013.

[8] E. Olson, J. Strom, R. Morton, A. Richardson, P. Ranganathan, R. Goeddel, M. Bulic, J. Crossman, and B. Marinier, "Progress towards multi-robot reconnaissance and the MAGIC 2010 competition," *Journal of Field Robotics*, vol. 29, no. 5, pp. 762–792, Sept. 2012.

[9] "Unmanned Systems Integrated Roadmap, FY 2013-2038," http://www.defense.gov/pubs/DOD-USRM-2013.pdf, reference Number 14-S-0553.

[10] "A Roadmap for U.S. Robotics: From Internet to Robotics, 2013 Edition," http://robotics-vo.us/sites/default/files/2013$%$20Robotics$%$20Roadmap-rs.pdf, Robotics Virtual Organization.

[11] G. A. Korsah, "Exploring bounded optimal coordination for heterogeneous teams with cross-schedule dependencies," Ph.D. dissertation, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, 2011.

[12] M. J. B. Krieger, J.-B. Billeter, and L. Keller, "Ant-like task allocation and recruitment in cooperative robots," *Nature*, vol. 406, pp. 992–995, 2000.

[13] W. Agassounon and A. Martinoli, "Efficiency and robustness of threshold-based distributed allocation algorithms in multi-agent systems," in *Proc. First Int'l. Joint Conf. on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2002, pp. 1090–1097.

[14] T. H. Labella, M. Dorigo, and J.-L. Deneubourg, "Division of labor in a group of robots inspired by ants' foraging behavior," *ACM Trans. Auton. Adapt. Syst.*, vol. 1, no. 1, pp. 4–25, 2006.

[15] N. Correll, "Parameter estimation and optimal control of swarm-robotic systems: A case study in distributed task allocation," in *Int'l. Conf. Robotics and Automation (ICRA)*, 2008, pp. 3302–3307.

[16] W. Liu and A. F. T. Winfield, "Modeling and optimization of adaptive foraging in swarm robotic systems," *Int'l. J. of Robotics Research*, vol. 29, no. 14, pp. 1743–1760, 2010.

[17] S. Berman, Á. Halász, M. A. Hsieh, and V. Kumar, "Optimized stochastic policies for task allocation in swarms of robots," *IEEE Trans. on Robotics*, vol. 25, no. 4, pp. 927–937, Aug. 2009.

[18] L. Odhner and H. Asada, "Stochastic recruitment control of large ensemble systems with limited feedback," *ASME J. Dyn. Sys. Meas. Control*, vol. 132, no. 4, pp. 041 008–1–041 008–9, 2010.

[19] H. Hamann and H. Wörn, "A framework of space-time continuous models for algorithm design in swarm robotics," *Swarm Intelligence*, vol. 2, no. 2-4, pp. 209–239, 2008.

[20] A. Prorok, N. Correll, and A. Martinoli, "Multi-level spatial modeling for stochastic distributed robotic systems," *Int'l. J. of Robotics Research*, vol. 30, no. 5, pp. 574–589, 2011.

[21] D. Milutinovic and P. Lima, "Modeling and optimal centralized control

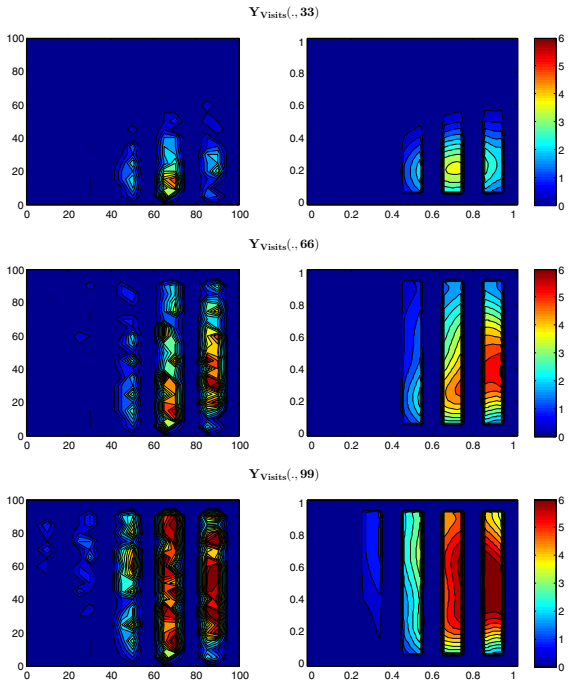

Fig. 5. Distribution of flower visits at three times in the microscopic (left) and macroscopic (right) models with parameters optimized for Objective 2, no obstacles

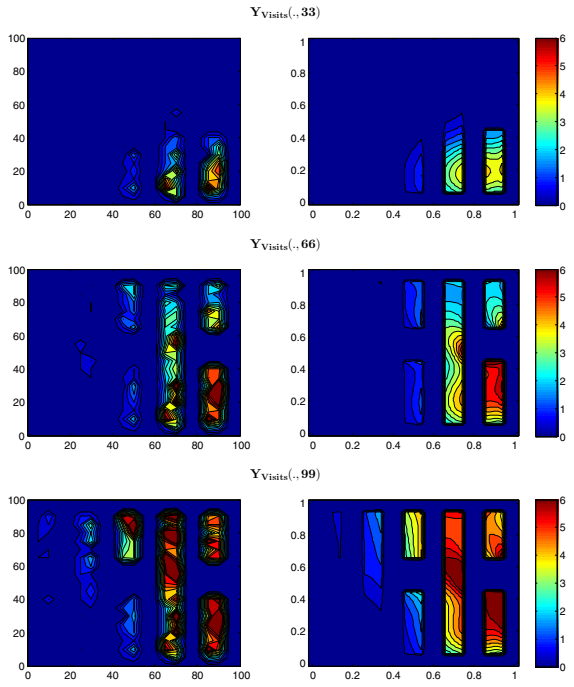

Fig. 6. Distribution of flower visits at three times in the microscopic (left) and macroscopic (right) models with parameters optimized for Objective 2, with obstacles

of a large-size robotic population," *IEEE Trans. on Robotics*, vol. 22, no. 6, pp. 1280–1285, 2006.

[22] S. Berman, R. Nagpal, and Á. Halász, "Optimization of stochastic strategies for spatially inhomogeneous robot swarms: A case study in commercial pollination," in *Int'l. Conf. Intelligent Robots and Systems (IROS)*, 2011, pp. 3923–3930.

[23] K. Dantu, S. Berman, B. Kate, and R. Nagpal, "A comparison of deterministic and stochastic approaches for allocating spatially dependent tasks in micro-aerial vehicle collectives," in *Int'l. Conf. Intelligent Robots and Systems (IROS)*, 2012, pp. 793–800.

[24] D. Elliott, *Bilinear control systems*. Springer, 2009.

[25] J. M. Ball, J. E. Marsden, and M. Slemrod, "Controllability for distributed bilinear systems," *SIAM Journal on Control and Optimization*, vol. 20, no. 4, pp. 575–597, 1982.

[26] A. Y. Khapalov, *Controllability of Partial Differential Equations Governed by Multiplicative Controls*. Berlin Heidelberg: Springer-Verlag, 2010, vol. 1995, Lecture Notes in Mathematics.

[27] G. Turinici, "Beyond bilinear controllability: applications to quantum control," in *Control of Coupled Partial Differential Equations*. Springer, 2007, pp. 293–309.

[28] A. Becker, C. Onyuksel, and T. Bretl, "Feedback control of many differential-drive robots with uniform control inputs," in *Int'l. Conf. on Intelligent Robots and Systems (IROS)*, 2012, pp. 2256–2262.

[29] G. Foderaro, "A distributed optimal control approach for multi-agent trajectory optimization," Ph.D. dissertation, Duke University, 2013.

[30] M. Turpin, N. Michael, and V. Kumar, "Capt: Concurrent assignment and planning of trajectories for multiple robots," *Int'l. Journal of Robotics Research*, vol. 33, no. 1, pp. 98–112, 2014.

[31] D. T. Gillespie, "The chemical Langevin equation," *J. Chem. Phys.*, vol. 113, no. 1, pp. 297–306, 2000.

[32] S. Berman, V. Kumar, and R. Nagpal, "Design of control policies for spatially inhomogeneous robot swarms with application to commercial pollination," in *Int'l. Conf. Robotics and Automation (ICRA)*, 2011, pp. 378–385.

[33] R. J. Leveque, *Finite-Volume Methods for Hyperbolic Problems*. Cambridge Univ. Press, 2004.

[34] F. Tröltzsch, *Optimal Control of Partial Differential Equations: Theory, Methods, and Applications*. American Mathematical Society, 2010, vol. 112, Graduate Studies in Mathematics.

[35] A. Belmiloudi, *Stabilization, Optimal and Robust Control: Theory and Applications in Biological and Physical Sciences*. London: Springer-Verlag, 2008.

[36] K. Elamvazhuthi, "A variational approach to planning, allocation and mapping in robot swarms using infinite dimensional models," Master's thesis, Arizona State University, 2014.

[37] L. C. Evans, *Partial Differential Equations*, 2nd ed. American Mathematical Society, 2010, vol. 19, Graduate Studies in Mathematics.

[38] M. Hinze, R. Pinnau, M. Ulbrich, and S. Ulbrich, *Optimization with PDE Constraints*. New York: Springer, 2009.

[39] D. Vries, K. J. Keesman, and H. Zwart, "A Luenberger observer for an infinite dimensional bilinear system: a UV disinfection example," in *Proc. of the 3rd Int'l. Federation of Automatic Control (IFAC) Symposium on System, Structure and Control*, 2007.

[40] M. R. D'Orsogna, Y.-L. Chuang, A. L. Bertozzi, and L. S. Chayes, "Self-propelled particles with soft-core interactions: patterns, stability, and collapse," *Physical Review Letters*, vol. 96, no. 10, p. 104302, 2006.

[41] J. D. Murray, *Mathematical Biology I: An Introduction*, 3rd ed. Springer, 2007, vol. 17, Interdisciplinary Applied Mathematics.