Automatica 95 (2018) 356–367

Contents lists available at ScienceDirect

Automatica

journal homepage: www.elsevier.com/locate/automatica

PDE-based optimization for stochastic mapping and coverage strategies using robotic ensembles[☆]

Karthik Elamvazhuthi [a], Hendrik Kuiper [b], Spring Berman [a,*]

[a] School for Engineering of Matter, Transport and Energy, Arizona State University, Tempe, AZ, USA
[b] School of Mathematical and Statistical Sciences, Arizona State University, Tempe, AZ, USA

ARTICLE INFO

Article history:
Received 30 December 2016
Received in revised form 28 March 2018
Accepted 29 March 2018

Keywords:
Swarm robotics
Autonomous mobile robots
Optimal control
Bilinear control systems
Partial differential equations
Distributed-parameter systems
Stochastic systems
Decentralized systems

ABSTRACT

This paper presents a novel partial differential equation (PDE)-based framework for controlling an ensemble of robots, which have limited sensing and actuation capabilities and exhibit stochastic behaviors, to perform mapping and coverage tasks. We model the ensemble population dynamics as an advection–diffusion–reaction PDE model and formulate the mapping and coverage tasks as identification and control problems for this model. In the mapping task, robots are deployed over a closed domain to gather data, which is unlocalized and independent of robot identities, for reconstructing the unknown spatial distribution of a region of interest. We frame this task as a convex optimization problem whose solution represents the region as a spatially-dependent coefficient in the PDE model. We then consider a coverage problem in which the robots must perform a desired activity at a programmable probability rate to achieve a target spatial distribution of activity over the reconstructed region of interest. We formulate this task as an optimal control problem in which the PDE model is expressed as a bilinear control system, with the robots' coverage activity rate and velocity field defined as the control inputs. We validate our approach with simulations of a combined mapping and coverage scenario in two environments with three target coverage distributions.

1. Introduction

Over the past few decades, partial differential equation (PDE) models of multi-agent systems have been used extensively in mathematical biology to analyze collective behaviors such as chemotaxis, flocking, schooling, predator–prey interactions, and pattern formation (Okubo, 1986). Many of these models are linear or nonlinear advection–diffusion type PDEs, which describe the spatiotemporal evolution of probability densities of agents. Mathematical tools such as bifurcation analysis, optimization, and control theory can be applied to these continuum *macroscopic models* to make qualitative and quantitative predictions about the system behavior. Typically, each PDE model corresponds to a discrete *microscopic model* that captures the stochastic and deterministic actions of individual agents. While these microscopic models are more accurate descriptions of the agents' behavior, the macroscopic models enable tractable analysis for large agent numbers.

Recently, this work has motivated the use of similar types of PDEs to model and control the spatiotemporal dynamics of very large collectives, or *swarms* (Francesca & Birattari, 2016), of small, resource-constrained robots (e.g., Karydis & Kumar, 2016; Sitti et al., 2015) that are currently being developed for applications such as environmental monitoring, exploration, surveillance, disaster response, and biomedical procedures. PDEs have been used to characterize the distributions of chemotactic robots in a diffusive fluid environment (Galstyan, Hogg, & Lerman, 2005), miniature robots inspecting a model of jet turbine blades (Prorok, Correll, & Martinoli, 2011), and honeybee-inspired agents that aggregate at the optimal value of a scalar field (Correll & Hamann, 2015). The parameters of these PDE models can be mapped to control inputs that drive the robots' motion and probability rates of switching between states or tasks, and the collective behavior of the robots follows the PDE model prediction in expectation. Several works have exploited this correspondence to control the spatial distribution of an ensemble (Foderaro, Ferrari, & Wettergren, 2014; Milutinovic & Lima, 2006). These control approaches can be viewed as extensions of stochastic task allocation schemes based on nonspatial rate equation models (Berman, Halász, Ani Hsieh, & Kumar, 2009; Correll & Martinoli, 2006; Lerman, Jones, Galstyan,

☆ This work was supported by National Science Foundation (NSF) award no. CMMI-1436960. The material in this paper was partially presented at the IEEE Robotics and Automation Society Conference (ICRA), May 25–30, 2015, Seattle, Washington, USA. This paper was recommended for publication in revised form by Associate Editor Thomas Meurer under the direction of Editor Miroslav Krstic.

* Corresponding author.
E-mail addresses: karthikevaz@asu.edu (K. Elamvazhuthi), kuiper@asu.edu (H. Kuiper), spring.berman@asu.edu (S. Berman).

https://doi.org/10.1016/j.automatica.2018.06.007
0005-1098/

& Matarić, 2006; Martinoli, Easton, & Agassounon, 2004). Other applications of continuum population dynamical models to multi-agent control include optimized confinement strategies (Haque, Rahmani, Egerstedt, & Yezzi, 2014), consensus using the theory of mean field controls (Nourian, Caines, Malhame, & Huang, 2013), controlled flocking (Piccoli, Rossi, & Trélat, 2014) that includes non-parallel motions (Han & Ge, 2015), and pattern generation in the presence of obstacles (Pimenta, Michael, Mesquita, Pereira, & Kumar, 2008). There has also been some recent work on using PDEs to model Laplacian network dynamics of agents for formation control; see Elamvazhuthi and Berman (2014), Frihauf and Krstic (2011), Meurer and Krstic (2011) and references therein.

We apply this PDE-based modeling framework to develop a control approach for allocating tasks among an ensemble of robots. In our scenarios, a *task* is defined as a desired activity that a robot performs in a certain spatial region of the environment. The tasks can be performed in parallel, and multiple robots can be simultaneously allocated to each task. While various deterministic approaches have been developed for multi-robot task allocation, including centralized and decentralized market-based techniques (Choi, Brunet, & How, 2009; Dias, Zlot, Kalra, & Stentz, 2006) and centralized methods for optimal task assignment and trajectory planning (Adler, de Berg, Halperin, & Solovey, 2015; Turpin, Michael, & Kumar, 2014), their computation and/or communication requirements do not scale well to very large numbers of robots and tasks. In contrast to these works, we develop a *stochastic* approach in which tasks are performed at random times by unidentified robots with limited computing capabilities and no global localization. The robots can receive information that is broadcast from a central supervisor, although they cannot communicate among themselves. Such limitations will be common in swarm robotic platforms, e.g. micro aerial vehicles (Karydis & Kumar, 2016) and microrobots (Sitti et al., 2015), and in scenarios where the robots operate in GPS-denied environments where communication is impractical or unreliable. In our proposed approach, a task allocation emerges from the collective ensemble activity.

We first consider a *mapping* problem in which the objective is to estimate a scalar spatial field from unlocalized data obtained by the robots. We then define a *coverage* problem in which the ensemble must produce a target spatial density of activity over a region of interest, which may be estimated in the *mapping* problem. For this problem, we express the PDE model as a *bilinear control system* (Ball, Marsden, & Slemrod, 1982) and formulate an optimal control problem that computes the control inputs. Since we do not assume that agents are capable of global localization or estimation of the local agent population density, we frame the coverage problem as an open-loop control problem that does not require feedback on agent positions or densities. We follow the variational approach described in Tröltzsch (2010) for optimal control of the PDE model. While there has been some prior work on bilinear optimal control of systems of PDEs (Annunziato & Borzì, 2014; Ou & Schuster, 2010), these works do not address the types of PDEs that we consider. An optimal control problem for a bilinear parabolic PDE was formulated in Ou and Schuster (2010) with the diffusion coefficient as the control. In Annunziato and Borzì (2014), bilinear control of a class of advection–reaction systems was considered; unlike our PDE models, these systems did not include diffusion.

The mapping and controller synthesis approaches described in this paper require a central supervisor with the computational capabilities necessary to solve the associated optimization problems. Despite this centralized component, the approaches are scalable with the number of agents in the ensemble since each agent executes the same controllers with the same control variables, which are preprogrammed or broadcast by the supervisor. In our coverage strategy, there are only three control variables to be computed; in contrast, the most naive approach to controlling an ensemble of N agents moving in d dimensions would require the computation of N^d control inputs.

We first presented our *coverage* approach in Elamvazhuthi and Berman (2015), where we introduced a similar optimal control problem, derived the gradient of the objective functional with respect to the control parameters, and used a gradient descent algorithm to compute the optimal control. This paper provides a complete analysis of our approach in Elamvazhuthi and Berman (2015) by investigating the well-posedness of the PDE model and the optimal control problem. The theory of weak solutions that we use to establish the well-posedness of the PDE model is classical (Evans, 1998). However, to the best of our knowledge, there have been no prior results on well-posedness that can be directly applied to our model, which is a system of PDEs in which diffusion is present only in one of the species, the control variables are time-dependent, and a zero-flux boundary condition is imposed on the boundary of a Lipschitz domain. In this paper, we prove the existence and uniqueness of solutions of our PDE model by deriving suitable energy estimates for the solutions. We also use these derived energy estimates to ensure that the computation of the gradient, performed using the adjoint equation approach, is well-posed. Moreover, we prove the existence of an optimal control for the problem using standard compactness arguments adapted to the PDE control setting (Tröltzsch, 2010). In addition to this analysis, our formulation of the *mapping* problem in the same framework is a novel contribution of this paper; in Elamvazhuthi and Berman (2015) it was assumed that the environment is known beforehand.

The paper is organized as follows. Section 2 describes the robot capabilities and their programmed behaviors during the mapping and coverage assignments, and Section 3 defines the microscopic and macroscopic models of the ensemble and its activity during each assignment. Section 4 defines key mathematical terminology that is used in Sections 5 and 6 to formulate and analyze the mapping and coverage objectives, respectively, as optimization problems that incorporate the macroscopic models. We validate our approach in Section 7 with simulations in which a region of interest must first be mapped and then covered with a target distribution of robot activity, and we conclude in Section 8.

2. Task descriptions and assumptions

We consider a scenario in which (1) a small number of agents must map a region of interest in an unknown, bounded environment, which we refer to as the *mapping assignment*, and then (2) a larger ensemble of agents must produce a target spatial distribution of activity within the mapped regions, which we call the *coverage assignment*. For instance, this activity could consist of sensor measurements, or as in our previous work (Elamvazhuthi & Berman, 2015), contacts with flowers to effect crop pollination. The mapping and coverage assignments will be formulated in a decoupled manner by posing them as two separate optimization problems in terms of their associated mean-field PDE models. We will then demonstrate through numerical simulations that these two problems can be solved sequentially in order to achieve the desired coverage objective.

2.1. Robot capabilities

We assume that the agents lack global localization, inter-agent communication, and prior data about the environment. Each agent is equipped with local sensing capabilities, allowing it to detect and distinguish between different types of regions within its sensing range, and a compass, enabling it to move with a specified heading. Additionally, the agents have sufficient memory to store the times at which they record observations of regions of interest. Similarly,