Automatica 50 (2014) 149–154

Contents lists available at ScienceDirect

Automatica

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

Brief paper

Distributed optimal control for multi-agent trajectory optimization[☆]

Greg Foderaro[a], Silvia Ferrari[a,1], Thomas A. Wettergren[b]

[a] Duke University, Durham, NC, United States

[b] Naval Undersea Warfare Center, Newport, RI, United States

ARTICLE INFO

Article history:
Received 30 May 2012
Received in revised form
15 July 2013
Accepted 28 August 2013
Available online 20 November 2013

Keywords:
Distributed control
Optimal control
Robot navigation
Multilevel control
Multiscale dynamical systems

ABSTRACT

This paper presents a novel optimal control problem, referred to as distributed optimal control, that is applicable to multiscale dynamical systems comprised of numerous interacting agents. The system performance is represented by an integral cost function of the macroscopic state that is optimized subject to a hyperbolic partial differential equation known as the advection equation. The microscopic control laws are derived from the optimal macroscopic description using a potential function approach. The optimality conditions of the distributed optimal control problem are first derived analytically and, then, demonstrated numerically through a multi-agent trajectory optimization problem.

1. Introduction

Many complex systems ranging from renewable resources (Sanchirico & Wilen, 2005) to very large scale robotic (VLSR) systems (Reif & Wang, 1999) can be described as multiscale dynamical systems comprised of many interactive agents. On small spatial and temporal scales, the dynamics of every agent can be described by a small system of ordinary differential equations (ODEs), referred to as the microscopic or detailed equation. On larger spatial and temporal scales, the agents' dynamics and interactions give rise to macroscopic coherent behaviors, or coarse dynamics, that can be modeled by partial differential equations (PDEs) (Kevrekidis et al., 2003). In many cases, the macroscopic PDE model can be derived by mapping the microscopic states of the agents to a macroscopic description using an appropriate restriction operator, such as the distribution of the agents or its lower-order moments (Kevrekidis et al., 2003).

This paper presents a distributed optimal control (DOC) problem formulation and optimality conditions applicable to a class of multiscale dynamical systems in which the restriction operator is the distribution of the agents, and the macroscopic dynamics are given by a PDE known as the advection equation. The DOC approach is demonstrated by solving a trajectory optimization problem in which a large number of unicycle robots must travel from an initial to a final macroscopic state, in the presence of obstacles. It was recently shown that optimizing the trajectories of N agents in an obstacle-populated environment is polynomial-space-hard (PSPACE-hard) in N (Hopcroft, Schwartz, & Sharir, 1984). A problem is considered PSPACE-hard if every problem in its class is at least as difficult as any problem solvable in polynomial space (PSPACE). The class of PSPACE problems contains many problems for which no efficient solutions are known. Therefore, a PSPACE-hard problem is generally considered to be computationally intractable for large N, as it would require exponential deterministic time in the worse case (Rich, 2008).

Several approaches have been proposed for tackling the control of VLRS systems, and avoid complexity issues for large N (Cheah, Hou, & Slotine, 2009). These approaches include prioritized planning techniques (Thrun, Bennewitz, & Burgard, 2002), and path-coordination methods (LaValle & Hutchinson, 1998), which first plan the agents' trajectories independently, and then adjust

[☆] This research was funded by the ONR Code 321. The material in this paper was partially presented at the 49th IEEE Conference on Decision and Control (CDC), December 15–17, 2010, Atlanta, Georgia, USA. This paper was recommended for publication in revised form by Associate Editor Kok Lay Teo under the direction of Editor Ian R. Petersen.

E-mail addresses: greg.foderaro@duke.edu (G. Foderaro), sferrari@duke.edu (S. Ferrari), t.a.wettergren@ieee.org (T.A. Wettergren).

[1] Tel.: +1 919 660 5310; fax: +1 919 660 8963.

0005-1098 © 2013 The Authors. Published by Elsevier Ltd. Open access under CC BY-NC-ND license.
http://dx.doi.org/10.1016/j.automatica.2013.09.014

the microscopic control laws to avoid mutual collisions. Behavior-based control methods seek feasible solutions by programming a set of simple behaviors for each agent, and by showing that the agents' interactions give rise to a macroscopic behavior, such as dispersion (Reif & Wang, 1999). Swarm-intelligence methods, such as foraging and schooling (Gazi & Passino, 2004), view each agent as an interchangeable unit subject to local objectives and constraints through which the swarm can converge to a range of pre-defined distributions.

The DOC approach presented in this paper does not rely on decoupling the agents' dynamics, or on specifying the agents' distribution *a priori*. Instead, DOC optimizes the macroscopic performance of the system subject to agent dynamics that are coupled via the objective function, and relies on the macroscopic evolution equation and restriction operator that characterize the multiscale system to reduce the computational complexity of the optimal control problem. As a result the computation required is significantly reduced compared to classical optimal control, and the trajectories of cooperative agents can be computed over large spatial and time scales without sacrificing optimality or completeness. The DOC optimality conditions are derived using calculus of variations, and validated using numerical solutions obtained via a direct optimization method. Simulations are presented to illustrate the performance of the DOC approach on a trajectory optimization problem involving hundreds of agents, and multiple cooperative objectives.

2. Problem formulation and assumptions

This paper considers the problem of computing the optimal state and control trajectories for a multiscale dynamical system comprised of N dynamical systems, referred to as agents, that can each be described by a small system of ODEs, referred to as the detailed equation,

$$\dot{\mathbf{x}}_i(t) = \mathbf{f}[\mathbf{x}_i(t), \mathbf{u}_i(t), t], \qquad \mathbf{x}_i(T_0) = \mathbf{x}_{i_0}, \quad i = 1, \ldots, N \tag{1}$$

where $\mathbf{x}_i \in \mathcal{X} \subset \mathbb{R}^n$ and $\mathbf{u}_i \in \mathcal{U} \subset \mathbb{R}^m$ denote the microscopic state and control of the ith agent, respectively, $\mathbf{x}_{i_0}$ is the initial value of the microscopic state, $\mathcal{X}$ denotes the microscopic state space, and $\mathcal{U}$ denotes the space of admissible microscopic controls. On larger spatial and temporal scales, the interactions of the N agents give rise to macroscopic coherent behaviors, or coarse dynamics, that are modeled by PDEs. The macroscopic state of the multiscale system, denoted by $\mathbf{X} \in \mathbb{R}^l$, consists of $l < n$ variables that capture the macroscopic system dynamics and performance, such as lower-order moments of the microscopically-evolving agent distribution (Kevrekidis et al., 2003).

From the agent distribution, it is possible to determine a restriction operator $\wp_{\mathbf{x}_i}$ that maps the microscopic states to the macroscopic description (Kevrekidis et al., 2003). Since $\mathbf{x}_i$ is a time-varying continuous vector, $\wp_{\mathbf{x}_i}$ is a time-varying probability density function (PDF), $\wp_{\mathbf{x}_i} : \mathcal{X} \times \mathbb{R} \to \mathbb{R}$, such that $X = \wp_{\mathbf{x}_i}(\mathbf{x}_i, t)$, and $l = 1$. Then, for any agent i, the probability of event $\mathbf{x}_i \in B$ at time t is,

$$P(\mathbf{x}_i \in B, t) = \int_B \wp_{\mathbf{x}_i}(\mathbf{x}_i, t) d\mathbf{x}_i, \tag{2}$$

for any subset $B \subset \mathcal{X}$, where $\wp_{\mathbf{x}_i}$ is a nonnegative function that satisfies the normalization property,

$$\int_{\mathcal{X}} \wp_{\mathbf{x}_i}(\mathbf{x}_i, t) d\mathbf{x}_i = 1 \tag{3}$$

and is abbreviated to $\wp$ in the remainder of this paper. For example, if $\mathbf{x}_i$ is the position of agent i at time t, the agent can be viewed as a fluid particle in the Lagrangian approach, and $\wp(\mathbf{x}_i, t)$ can be viewed as the forward PDF of particle position (Pope, 2000). Furthermore, $N\wp(\mathbf{x}_i, t)$ represents the density of agents in $\mathcal{X}$.

The macroscopic system performance is a function of the agent distribution and control, and it can be expressed as an integral cost function of $\wp$ and $\mathbf{u}_i$,

$$J = \phi[\wp(\mathbf{x}_i, T_f)] + \int_{T_0}^{T_f} \int_{\mathcal{X}} \mathscr{L}[\wp(\mathbf{x}_i, t), \mathbf{u}_i(t), t] d\mathbf{x}_i dt \tag{4}$$

where $\mathscr{L}$ is the Lagrangian, and ϕ is the terminal cost. DOC seeks to determine the macroscopic state and microscopic control trajectories that minimize J over a (large) time interval $(T_0, T_f]$, subject to the coarse dynamics, the normalization condition (3), and state constraints.

Through state constraints, it is possible to guarantee that, at any time $t \in (T_0, T_f]$, $\mathbf{x}_i \in \mathcal{X}$ for all i, and, thus, agents in $\mathcal{X}$ are never created nor destroyed. The PDE that governs the motion of a conserved, scalar quantity, such as a PDF, as it is advected by a known velocity field is a hyperbolic PDE known as the advection equation (Boyd, 2001). Based on the advection equation, when $\wp$ is advected by the velocity field $\mathbf{v}_i = \dot{\mathbf{x}}_i \in \mathbb{R}^n$, known from the detailed Eq. (1), the evolution of $\wp$ can be derived from the continuity equation and Gauss' theorem. It can be shown that the time-rate of change of $\wp$ can be written in terms of the divergence of the vector $(\wp\mathbf{v}_i)$, as shown by the advection equation,

$$\frac{\partial \wp}{\partial t} = -\nabla \cdot \{\wp(\mathbf{x}_i, t)\, \mathbf{v}_i(t)\} \tag{5}$$

$$= -\nabla \cdot \{\wp(\mathbf{x}_i, t)\, \mathbf{f}[\mathbf{x}_i, \mathbf{u}_i, t]\} \tag{6}$$

where, the gradient ∇ denotes a row vector of partial derivatives with respect to the elements of $\mathbf{x}_i$, $(\cdot)$ denotes the dot product, and the divergence is written as the dot product between $(\wp\mathbf{v}_i)$ and the gradient ∇. The reader is referred to Boyd (2001) for a detailed derivation of the advection equation. Assuming the initial agent distribution is a known PDF g_0, the macroscopic evolution equation (6) is subject to the following initial and boundary conditions,

$$\wp(\mathbf{x}_i, T_0) = g_0(\mathbf{x}_i), \quad \forall \mathbf{x}_i \in \mathcal{X} \tag{7}$$

$$\wp(\mathbf{x}_i, t) = 0, \quad \forall \mathbf{x}_i \in \partial\mathcal{X}, \ \forall t \in (T_0, T_f] \tag{8}$$

where $\partial\mathcal{X}$ denotes the boundary of $\mathcal{X}$, such that agents remain in the interior of $\mathcal{X}$ at all times. Additionally, $\wp$ must obey the normalization condition (3), and the state constraint

$$\wp(\mathbf{x}_i, t) = 0, \quad \forall \mathbf{x}_i \notin \mathcal{X}, \ \forall t \in (T_0, T_f]. \tag{9}$$

Then, the DOC problem consists of finding the optimal agent distribution, $\wp^*$, and microscopic controls, $\mathbf{u}_i^*$, that minimize the macroscopic cost function (4) subject to the dynamic constraint (6), the normalization condition (3), the initial and boundary conditions (7)–(8), and the state constraint (9). Since the DOC problem does not obey the classical optimal control formulation (Stengel, 1986), new optimality conditions are derived in the next section, and then they are validated numerically in Section 5 through a multi-agent trajectory optimization problem presented in Section 4.

3. DOC optimality conditions

The necessary conditions for optimality are derived by using calculus of variations to determine the agent distribution and control laws that minimize the integral cost function (4). Since the optimization of (4) is subject to a set of dynamic and equality constraints, the integral to be minimized is found by adjoining the dynamic constraints to (4) using a Lagrange multiplier (Fox, 1987). By this approach, necessary conditions for optimality are found from the first-order effects of control variations that must be zero at all times for the integral cost to be stationary. Then, higher-order sensitivity to control variations can be tested to discriminate between cases in which the integral is a minimum, a maximum, or is neither (Fox, 1987).

From the distributive property of the dot product and by change of sign, the advection equation (6) is rewritten as the time-varying equality constraint,

$$\frac{\partial \wp}{\partial t} + (\nabla \wp) \cdot \mathbf{f} + \wp(\nabla \cdot \mathbf{f}) = 0 \tag{10}$$

where the functions' arguments are omitted for brevity. Since (10) is a dynamic constraint that must be satisfied at all times, a time-varying Lagrange multiplier, $\lambda = \lambda(\mathbf{x}_i, t)$, is used to adjoin the equality constraint (10) to the integral cost (4). Then, the augmented cost function,

$$J_A = \phi[\wp(\mathbf{x}_i, T_f)] + \int_{T_0}^{T_f} \int_{\mathcal{X}} \left\{ \mathscr{L}(\wp, \mathbf{u}_i, t) + \lambda \left[\frac{\partial \wp}{\partial t} + (\nabla \wp) \cdot \mathbf{f} + \wp(\nabla \cdot \mathbf{f}) \right] \right\} d\mathbf{x}_i dt \tag{11}$$

is to be minimized with respect to the functional forms of the time-varying agent distribution $\wp$ and control $\mathbf{u}_i$, and subject to the equality constraints (3), (7)–(9).

The integrand of (11) must satisfy stationarity conditions throughout $(T_0,\ T_f]$ in order for J_A to be stationary (Fox, 1987). This is proven by introducing the Hamiltonian,

$$\mathscr{H}[\wp(\cdot), \mathbf{u}_i(\cdot), \lambda(\cdot), t] \equiv \mathscr{L}(\cdot) + \lambda[(\nabla \wp) \cdot \mathbf{f} + \wp(\nabla \cdot \mathbf{f})] \tag{12}$$

which is a function of the agent distribution, the control, and the Lagrange multiplier, and is analogous to the Hamiltonian from Pontryagin's minimum principle (Stengel, 1986). The augmented cost function (11) is then rewritten in terms of the Hamiltonian and simplified using integration by parts,

$$\begin{aligned} J_A &= \phi + \int_{T_0}^{T_f} \int_{\mathcal{X}} \left(\mathscr{H} + \lambda \frac{\partial \wp}{\partial t} \right) d\mathbf{x}_i dt \\ &= \phi + \int_{\mathcal{X}} [\lambda(\mathbf{x}_i, T_f)\wp(\mathbf{x}_i, T_f) - \lambda(\mathbf{x}_i, T_0)\wp(\mathbf{x}_i, T_0)] d\mathbf{x}_i \\ &\quad + \int_{T_0}^{T_f} \int_{\mathcal{X}} \left(\mathscr{H} - \wp \frac{\partial \lambda}{\partial t} \right) d\mathbf{x}_i dt. \end{aligned} \tag{13}$$

By the fundamental theorem of calculus of variations (Fox, 1987), an integral with fixed end points, T_0 and T_f, is stationary for weak variations if the first order effect of variations in the function to be optimized is zero throughout $(T_0, T_f]$. Thus, for J_A to be stationary, the first-order effect of control variations $\delta \mathbf{u}_i(t)$ on (13) must be zero for all $t \in (T_0, T_f]$. By the causality of the macroscopic dynamic equation (6), control perturbations lead to perturbations in $\wp$, and thus the first variation of J_A is

$$\begin{aligned} \delta J_A &= \int_{T_0}^{T_f} \int_{\mathcal{X}} \left[\left(\frac{\partial \mathscr{H}}{\partial \wp} - \dot{\lambda} \right) \delta \wp(\delta \mathbf{u}_i) + \frac{\partial \mathscr{H}}{\partial \mathbf{u}_i} \delta \mathbf{u}_i \right] d\mathbf{x}_i dt \\ &\quad + \left(\frac{\partial \phi}{\partial \wp} + \int_{\mathcal{X}} \lambda d\mathbf{x}_i \right) \delta \wp(\delta \mathbf{u}_i) \Big|_{t=T_f} \\ &\quad - \int_{\mathcal{X}} \lambda d\mathbf{x}_i \delta \wp(\delta \mathbf{u}_i) \Big|_{t=T_0}. \end{aligned}$$

For an extremum, we must have $\delta J_A = 0$ for all $\delta \wp$, $\delta \mathbf{u}_i$, and the variations from $\delta \wp$ and $\delta \mathbf{u}_i$ must independently vanish along the optimal solution curve. Since it can be assumed that the initial control has no effect on the initial state conditions, the equations,

$$\dot{\lambda} = \frac{\partial \mathscr{H}}{\partial \wp} = \frac{\partial \mathscr{L}}{\partial \wp} + \lambda(\nabla \cdot \mathbf{f}) \tag{14}$$

and,

$$0 = \frac{\partial \mathscr{H}}{\partial \mathbf{u}_i} = \frac{\partial \mathscr{L}}{\partial \mathbf{u}_i} + \lambda \left[(\nabla \wp) \frac{\partial \mathbf{f}}{\partial \mathbf{u}_i} + \wp \frac{\partial}{\partial \mathbf{u}_i} (\nabla \cdot \mathbf{f}) \right] \tag{15}$$

must be satisfied for $T_0 \le t \le T_f$, subject to the terminal conditions

$$\int_{\mathcal{X}} \lambda(\mathbf{x}_i, T_f) d\mathbf{x}_i = - \frac{\partial \phi}{\partial \wp} \Big|_{t=T_f}. \tag{16}$$

Eqs. (14)–(16) constitute necessary conditions for optimality for the DOC problem in Section 2. Thus, the optimal agent distribution $\wp^*$ must satisfy (14)–(16) along with the normalization condition (3), the initial and boundary conditions (7)–(8), and the state constraint (9). If these conditions are satisfied, the extremals can be tested using higher-order variations to verify that they lead to a minimum of the augmented cost function J_A in (11). In particular, sufficient conditions for optimality could be derived from the second-order derivatives of the Hamiltonian (12) with respect to $\mathbf{u}_i$, or Hessian matrix that is positive definite for a convex Hamiltonian. In this paper, we consider admissible solutions of (14)–(16) to be optimal if perturbations at any $t \in (T_0,\ T_f]$ only increase the value of J_A.

The microscopic control laws are determined from the optimal macroscopic description $\wp^*$ by defining an attractive potential that pulls the agents towards $\wp^*$. Since $\wp^*$ is a time-varying distribution, the potential function is defined as a quadratic function of the error between $\wp^*$ and the estimated agent distribution, $\hat{\wp}$, at time $(t + \delta t)$:

$$U \triangleq \frac{1}{2} [\hat{\wp}(\mathbf{x}_i, t + \delta t) - \wp^*(\mathbf{x}_i, t + \delta t)]^2. \tag{17}$$

The time interval δt is a small time constant that is chosen to prevent the agents from lagging behind $\wp^*$. The estimate $\hat{\wp}(\mathbf{x}_i, t + \delta t)$ is computed by stepping the advection equation (6) forward in time by an interval δt from $\hat{\wp}(\mathbf{x}_i, t)$, and $\hat{\wp}(\mathbf{x}_i, t)$ is computed via kernel density estimation from the agents' positions at time t (Simonoff, 1996). Then, a microscopic control law that minimizes (17) is obtained from the negative gradient of U, based on the detailed Eq. (1), such that $\mathbf{u}_i^* = \mathbf{c}[\wp^*(\mathbf{x}_i, t + \delta t)]$.

4. Multi-agent trajectory optimization

The DOC problem and optimality conditions presented in the previous section are demonstrated through a multi-agent trajectory optimization problem. Consider a system of N cooperative unicycle robots traveling through an obstacle-populated compact space $\mathcal{W} \subset \mathbb{R}^2$, referred to as the *workspace*, and occupied by M obstacles $\mathcal{B}_1, \ldots, \mathcal{B}_M$, where $\mathcal{B}_j \subset \mathcal{W}$. The dynamics of each robot are described by the nonlinear unicycle model,

$$\dot{x}_i = v_i \cos \theta_i \qquad \dot{y}_i = v_i \sin \theta_i \qquad \dot{\theta}_i = \omega_i \tag{18}$$

where $\mathbf{q}_i = [x_i\, y_i\, \theta_i]^T$ is the configuration of agent i, which contains the xy-coordinates, x_i and y_i, and heading angle, θ_i, with $i = 1, \ldots, N$. The microscopic control vector of agent i is $\mathbf{u}_i = [v_i\, \omega_i]^T$, where v_i and ω_i are the linear and angular velocities, respectively.

The macroscopic state of the system is described by the time-varying PDF, or restriction operator, $\wp : \mathcal{X} \times \mathbb{R} \to \mathbb{R}$, such that the probability of $\mathbf{x}_i = [x_i\, y_i]^T$ is given by (2), in terms of $\wp$. It follows that $\mathcal{W} = \mathcal{X}$, and $\wp$ can be regarded as the density of agents in $\mathcal{W}$ at time $t \in (T_0,\ T_f]$. Given an initial distribution $g_0(\mathbf{x}_i)$, the agents must travel in $\mathcal{W}$ to meet a goal distribution $g(\mathbf{x}_i)$, while avoiding obstacles, and minimizing energy consumption. The goal distribution is assumed to be time-invariant, and all M obstacles' positions and geometries are assumed known without error. This section shows that all of these trajectory optimization objectives can be expressed in terms of the PDF, $\wp$, to be optimized.

A measure of the difference between $\wp$ and the goal distribution, g, is given by the instantaneous Kullback–Leibler (KL) divergence at time t,

$$D(\wp \parallel g) = \int_{\mathcal{X}} \wp(\mathbf{x}_i, t) \log_2 \frac{\wp(\mathbf{x}_i, t)}{g(\mathbf{x}_i)} d\mathbf{x}_i \tag{19}$$

where, by definition, the support set of $\wp$ is contained by the support set of g, and the value $0 \log_2(0/0)$ is replaced with 0 for continuity (Cover & Thomas, 1991). Although the KL divergence is not a true distance function because it is not symmetric, it is a suitable objective function because its value increases when the difference between $\wp$ and g increases, and vice versa. Also, the KL divergence of $\wp$ and g is zero when the two distributions are equal.

A repulsive potential U_{rep} can be generated from the obstacles' geometries $\mathcal{B}_1, \ldots, \mathcal{B}_M$ in $\mathcal{W}$, as shown in Latombe (1991). Then, the obstacle avoidance objective can be represented by the product $\wp U_{rep}$. The energy consumption is modeled as a quadratic function of the control. The DOC cost function to be minimized is,

$$J = \int_{T_0}^{T_f} \left[w_d\, D(\wp \parallel g) + \int_{\mathcal{W}} \left(w_r\, \wp\, U_{rep} + w_e\, \mathbf{u}_i^T \mathbf{R} \mathbf{u}_i \right) d\mathbf{x}_i \right] dt \tag{20}$$

where, $\mathbf{R}$ is a diagonal positive-definite matrix. The scalar weights w_d, w_r, and w_e can be chosen by the user or from a Pareto optimization curve, and represent the desired tradeoff between the three competing objectives. By this formulation of the cost function, the KL divergence of $\wp$ and g is minimized throughout $(T_0,\ T_f]$.

The solution of the DOC problem can be approached by a parametrization technique that approximates the function to be optimized by a weighted linear combination of basis functions (Loxton, Teo, & Rehbock, 2008; Wang, Gui, Teo, Loxton, & Yang, 2009). Finite Gaussian mixture models are commonly used to provide parametric approximations of PDFs. Thus, in this paper, the agent distribution is approximated by a mixture model comprised of z components with Gaussian PDFs $f_1, \ldots, f_z$, and corresponding mixing proportions (or weights) $w_1, \ldots, w_z$. The n-dimensional multivariate Gaussian PDF,

$$f_j(\mathbf{x}_i, t) = \frac{e^{\{-(1/2)[\mathbf{x}_i - \boldsymbol{\mu}_j(t)]^T \boldsymbol{\Sigma}_j(t)^{-1} [\mathbf{x}_i - \boldsymbol{\mu}_j(t)]\}}}{(2\pi)^{n/2} |\boldsymbol{\Sigma}_j(t)|^{1/2}} \tag{21}$$

is referred to as the component density of the mixture, and is characterized by a time-varying mean vector $\boldsymbol{\mu}_j \in \mathbb{R}^n$, and a time-varying covariance matrix $\boldsymbol{\Sigma}_j \in \mathbb{R}^{n \times n}$, with $j = 1, \ldots, z$. We assume that, at any $t \in (T_0, T_f]$, the agent distribution can be approximated as follows,

$$\wp(\mathbf{x}_i, t) \approx \sum_{j=1}^{z} w_j(t) f_j(\mathbf{x}_i, t) \tag{22}$$

where, $0 \le w_j \le 1$ for any j, and $\sum_{j=1}^{z} w_j = 1$ (McLachlan, 2000). In this paper, it is assumed that z is fixed, and chosen by the user. Then, an approximately-optimal agent distribution $\wp^*$ can be obtained by determining the optimal trajectories of the mixture model parameters, i.e., $\boldsymbol{\mu}_j^*$, $\boldsymbol{\Sigma}_j^*$, and w_j^*, for $j = 1, \ldots, z$.

In addition to satisfying the DOC constraints and optimality conditions, the mixture model parameters must be determined such that the component densities $f_1, \ldots, f_z$ are nonnegative and obey the normalization condition for all $t \in (T_0, T_f]$. This is accomplished by discretizing the continuous DOC problem in space and time, about a finite set of collocation points in $\mathcal{X} \times (T_0, T_f]$. Let $\Delta\mathbf{x}$ and Δt denote constant space and time discretization intervals, respectively, that, to guarantee numerical stability, are chosen according to the Courant–Friedrichs–Lewy condition (Tannehille, Anderson, & Pletcher, 1997). Then, by formulating the discretized DOC problem as a finite dimensional NLP, the optimal mixture model parameters can be computed via sequential quadratic programming (SQP) (Bertsekas, 2007). The details of the DOC numerical solution will be provided in a separate paper.

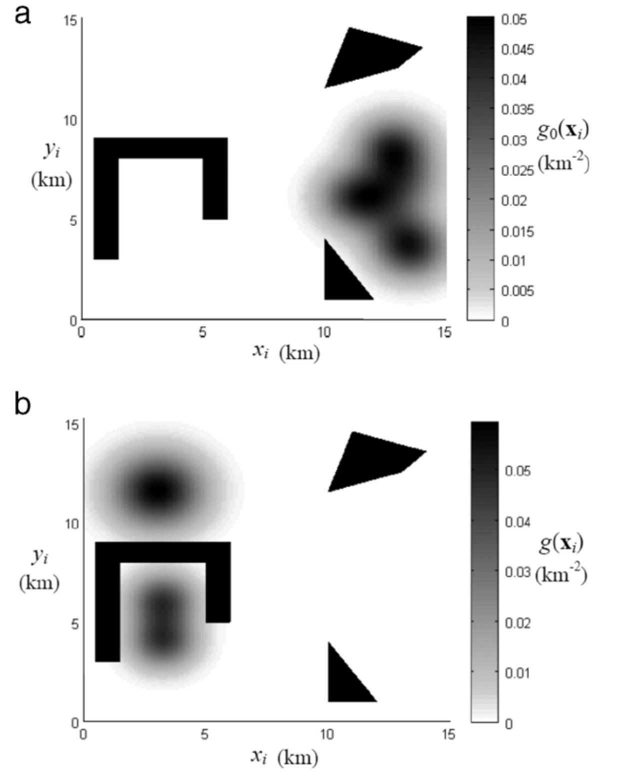

Fig. 1. Initial (a) and goal (b) agent distributions for a workspace with three obstacles (solid black).

Once an optimal agent distribution $\wp^*$ is obtained from the DOC problem (18)–(20), the microscopic control laws are obtained from the negative gradient of the potential function in (17). For robots described by the unicycle model (18), the microscopic control law is,

$$\mathbf{u}_i = [v_c \quad Q(\hat{\theta}_i, -\nabla U)]^T \tag{23}$$

where,

$$Q(\cdot) = \{a(\hat{\theta}_i) - a[\Theta(-\nabla U)]\}\text{sgn}\{a[\Theta(-\nabla U)] - a(\hat{\theta}_i)\}$$

is the minimum differential between the agent's actual heading angle $\hat{\theta}_i$ and the desired heading angle $\Theta(-\nabla U)$, v_c is the agent's speed, $\text{sgn}(\cdot)$ is the sign function, and $a(\cdot)$ is an angle wrapping function (Latombe, 1991).

5. Simulation results

The DOC solution of the multi-agent trajectory optimization problem presented in the previous section is illustrated through an example in which $N = 500$ agents with unicycle dynamics (18) must travel from the initial distribution, g_0, to the goal distribution, g, plotted in Fig. 1, during a time interval $(0,\ 22]$ h. The initial microscopic states $\mathbf{x}_{i_0}$ are determined by sampling g_0. Subsequently, the agents must travel in a workspace $\mathcal{W} = [0,\ L] \times [0,\ L]$, with $L = 15$ km, and three obstacles plotted in solid black in Fig. 1. All agents are assumed to have a linear velocity $v_i = 0.7$ km/h, and an angular velocity $\omega_i \in [-\omega_{\max},\ +\omega_{\max}]$, where $\omega_{\max} = 0.52$ Rad/s. The cost function weights, $w_d = 20$, $w_r = 0.1$, and $w_e = 1$, are chosen based on the units and relative magnitudes of the three navigation objectives.

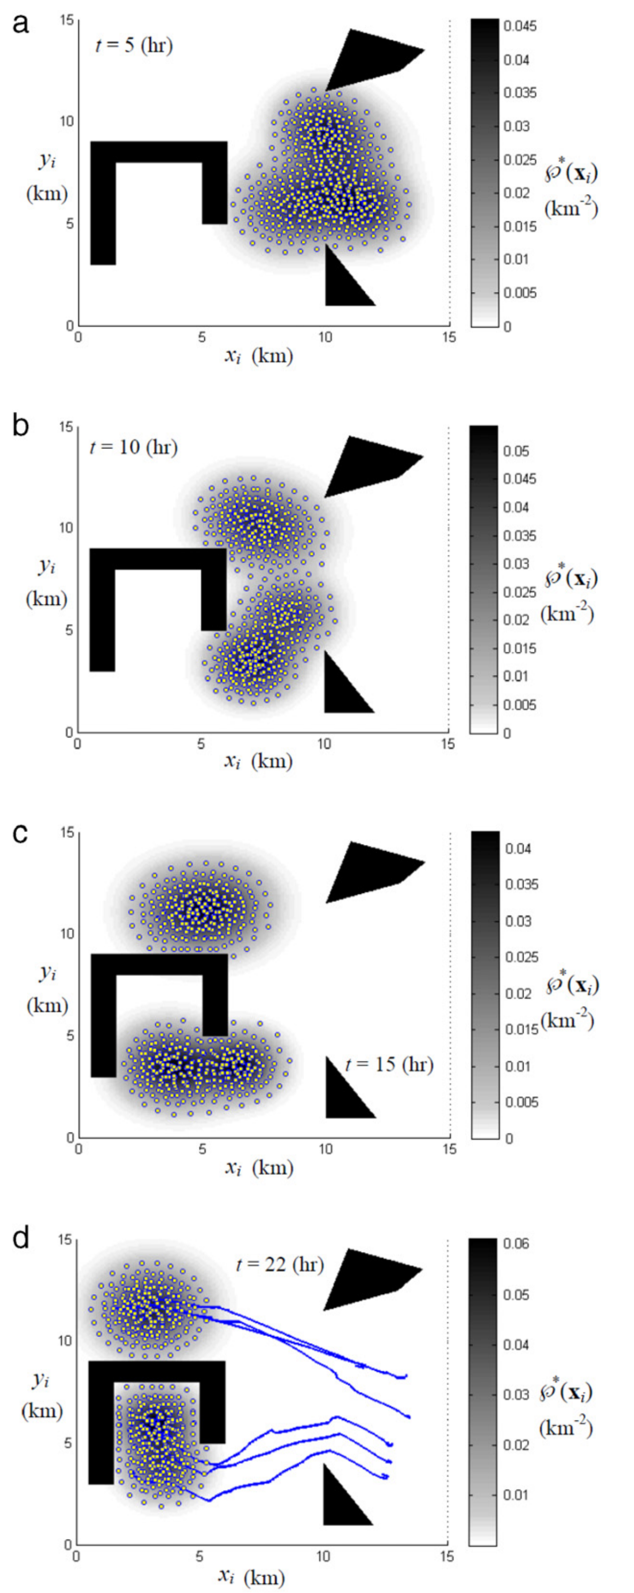

Fig. 2. Optimal evolution of agent distribution and microscopic state (yellow circles) for $N = 500$ microscopic agents, at four instants in time. State trajectories (blue lines) of 6 randomly chosen agents are plotted in (d).

The computational complexity of the optimization performed in this example is of the order of the dominant computation of the algorithm's quadratic program (QP) subproblem, which is a QR decomposition using Householder Triangularization (Powell, 1977). This leads to a complexity of $O(z^2XK^3)$ that does not grow

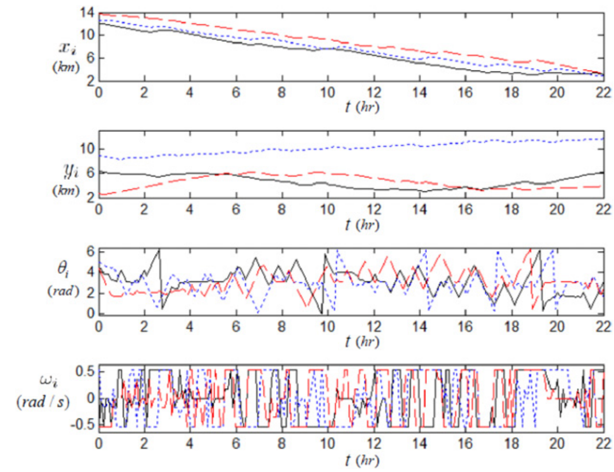

Fig. 3. Microscopic state and control histories for 3 agents randomly chosen from the example in Fig. 2.

with the number of agents N. The number of mixture components, $z = 6$, is chosen to obtain the best tradeoff between accuracy and computational complexity. Time is discretized in intervals of $\Delta t = 1$ h, such that $K = 22$, and the state space is discretized using $X = 900$ collocation points. As a result, the optimal agent distribution could be computed in several hours on a Core-2 Duo CPU 2.13-GHz computer with 8-GB RAM, while the corresponding classical optimal control problem for $N = 500$ was found to be intractable on the same machine.

The optimal agent distribution, $\wp^*$, and the values of the agents' microscopic state variables, $\mathbf{x}_i$, are plotted in Fig. 2 at four sample moments in time, $t = 5$ h (a), $t = 10$ h (b), $t = 15$ h (c), and $t = 22$ h (d). The evolution of the microscopic state, $\mathbf{x}_i$, is simulated by integrating the closed-loop detailed Eq. (18) numerically for all i, using a time interval $\delta t = 3$ s. At every time step of the numerical integration, the feedback control law is evaluated according to (23), from the attractive potential (17) defined in terms of the optimal distribution $\wp^*$. The time-histories of the DOC microscopic state and control for three, randomly-chosen agents are plotted in Fig. 3, and the state trajectories of six, randomly-chosen agents are plotted in Fig. 2(d). The results show that, as specified by the cost function (20), over time $\wp^*$ meets the goal distribution g, while agents also avoid obstacles in $\mathcal{W}$, and minimize energy consumption.

The optimal agent distributions obtained via SQP are also used to show that any perturbations from the optimal mixture model parameters increase the error in the optimality conditions derived in Section 3. Fig. 4 shows the effects of perturbations in the covariances of two mixture components at $t = 21$ h, for the optimal distribution in Fig. 2. Here, the jth component's covariance is modified such that $\Sigma_j = \Sigma_j^* + c_j I_2$, where c_j is the perturbation parameter varied in Fig. 4, and e_1 and e_2 denote the mean-squared errors for the optimality conditions (14) and (15), respectively. These results are representative of an extensive set of simulations in which the means, covariances, and component weights were perturbed from optimal at various times. In all cases, the optimality conditions were validated numerically by showing that e_1 and e_2 were at a minimum for the mixture model parameters $\boldsymbol{\zeta}^*$ computed via SQP.

6. Conclusion

This paper presents a novel DOC problem formulation that extends the capabilities of classical optimal control to multiscale dynamical systems. The DOC optimality conditions are derived analytically for the case in which the macroscopic description

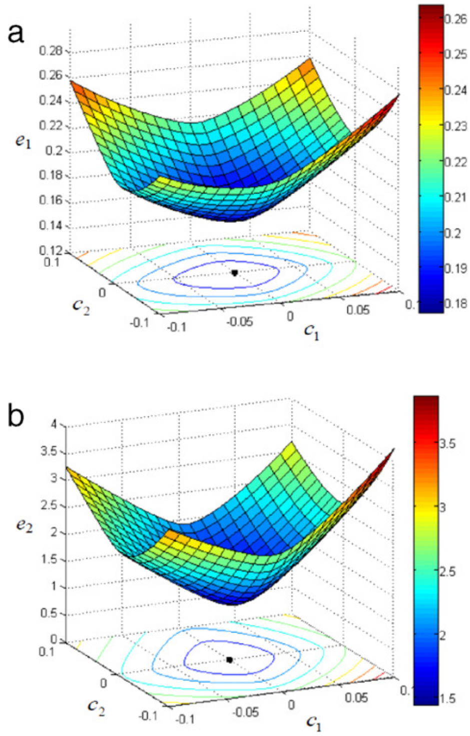

Fig. 4. Numerical error for optimality conditions (14) (a) and (15) (b) as a function of covariance perturbation parameters.

is characterized by the agent distribution, and the macroscopic dynamics are modeled by the advection equation. Numerical simulations are used to validate the optimality conditions, and to demonstrate the effectiveness of the approach for optimizing the trajectories of many unicycle robots that must travel through an obstacle-populated environment.

References

Bertsekas, D. P. (2007). Nonlinear programming. In *Athena scientific*. MA: Belmont.

Boyd, J. P. (2001). *Chebyshev and Fourier spectral methods* (2nd ed.). New York, NY: Dover.

Cheah, C. C., Hou, S. P., & Slotine, J. J. E. (2009). Region-based shaped control for a swarm of robots. *Automatica*, *45*, 2406–2411.

Cover, T. M., & Thomas, J. A. (1991). *Elements of information theory*. John Wiley and Sons, Inc.

Fox, C. (1987). *An introduction to the calculus of variations*. New York: Dover Publications, Inc.

Gazi, V., & Passino, K. M. (2004). Stability analysis of social foraging swarms. *IEEE Transactions on Systems, Man and Cybernetics*, *34*(1), 539–557.

Hopcroft, J. E., Schwartz, J. T., & Sharir, M. (1984). On the complexity of motion planning for multiple independent objects; pspace-hardness of the "warehouseman's problem". *The International Journal of Robotics Research*, *3*(4), 76–88.

Kevrekidis, I. G., Gear, C. W., Hyman, J. M., Kevrekidis, P. G., Runborg, O., & Theodoropoulos, C. (2003). Equation-free, coarse-grained multiscale computation: enabling microscopic simulators to perform system-level analysis. *Communications in Mathematical Sciences*, *1*(4), 715–762.

Latombe, J. C. (1991). *Robot motion planning*. Kluwer Academic Publishers.

LaValle, S. M., & Hutchinson, S. A. (1998). Optimal motion planning for multiple robots having independent goals. *IEEE Transactions on Robotics and Automation*, *14*, 912–925.

Loxton, R. C., Teo, K. L., & Rehbock, V. (2008). Optimal control problems with multiple characteristic time points in the objective and constraints. *Automatica*, *44*(11), 2923–2929.

McLachlan, G. (2000). *Finite mixture models*. New York, NY: Wiley Interscience.

Pope, Stephen B. (2000). *Turbulent flows*. New York, NY: Cambridge University Press.

Powell, M. J. D. (1977). *Numerical analysis, lecture notes in mathematics*, *A fast algorithms for nonlinearly constrained optimization calculations*. Berlin, Heidelberg: Springer.

Reif, J. H., & Wang, H. (1999). Social potential fields: a distributed behavioral control for autonomous robots. *Robotics and Autonomous Systems*, *27*(3), 171–194.

Rich, E. (2008). *Automata, computability and complexity*. Pearson Education, Inc.

Sanchirico, J., & Wilen, J. (2005). Optimal spatial management of renewable resources: matching policy scope to ecosystem scale. *Journal of Environmental Economics and Management*, *50*, 23–46.

Simonoff, J. S. (1996). *Smoothing methods in statistics*. Springer.

Stengel, R. F. (1986). *Optimal control and estimation*. Dover Publications, Inc.

Tannehille, J. C., Anderson, D. A., & Pletcher, R. H. (1997). *Computational fluid mechanics and heat transfer*. Taylor and Francis.

Thrun, S., Bennewitz, M., & Burgard, W. (2002). Finding and optimizing solvable priority schemes for decoupled path planning techniquues for teams of mobile robots. *Robotics and Autonomous Systems*, *41*(2), 89–99.

Wang, L. Y., Gui, W. H., Teo, K. L., Loxton, R., & Yang, C. H. (2009). Time delayed optimal control problems with multiple characteristic time points: computation and industrial applications. *Journal of Industrial and Management Optimization*, *5*(4), 705–718.

Greg Foderaro received the B.S. degree in Mechanical Engineering from Clemson University in 2009 and is pursuing a Ph.D. degree at Duke University. His research interests are in underwater sensor networks, robot path planning, multiscale dynamical systems, pursuit-evasion games, and spiking neural networks. He is a student member of the IEEE.

Silvia Ferrari is Paul Ruffin Scarborough Associate Professor of Engineering at Duke University, where she directs the Laboratory for Intelligent Systems and Controls (LISC). Her principal research interests include robust adaptive control of aircraft, learning and approximate dynamic programming, and optimal control of mobile sensor networks. She received the B.S. degree from Embry-Riddle Aeronautical University and the M.A. and Ph.D. degrees from Princeton University. She is a senior member of the IEEE, and a member of ASME, SPIE, and AIAA. She is the recipient of the ONR young investigator award (2004), the NSF CAREER award (2005), and the Presidential Early Career Award for Scientists and Engineers (PECASE) award (2006).

Thomas A. Wettergren received the B.S. degree in electrical engineering and the Ph.D. degree in applied mathematics, both from Rensselaer Polytechnic Institute in Troy, NY. He works at the Naval Undersea Warfare Center in Newport, RI, where he currently serves as the US Navy Senior Technologist for Operational and Information Science. He also is an Adjunct Professor of Mechanical Engineering at Pennsylvania State University. His research interests are in development of new analytical and computational methods for the mathematical modeling and control of multi-agent undersea sensing systems. He has been the recipient of the NAVSEA Scientist of the Year, the Assistant Secretary of the Navy Top Scientist of the Year, and the IEEE-USA Harry Diamond awards. He is a senior member of the IEEE and a member of SIAM.