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/$ – see front matter © 2013 The Authors. Published by Elsevier Ltd. All rights reserved.
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).