Automatica 50 (2014) 2936–2942

Contents lists available at ScienceDirect

Automatica

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

Brief paper

Supervised coverage control of multi-agent systems☆

Gökhan M. Atınç[1], Dušan M. Stipanović, Petros G. Voulgaris

Coordinated Science Laboratory, University of Illinois, 1308 West Main Street, Urbana, IL, 61801, USA

ARTICLE INFO

Article history:
Received 5 August 2013
Received in revised form
1 June 2014
Accepted 8 July 2014
Available online 23 October 2014

Keywords:
Coverage control
Collision avoidance
Multi-agent systems
Cooperative control
Autonomous mobile robots

ABSTRACT

In this paper, we consider a dynamic coverage problem for multi-agent systems, where the main objective of a group of mobile agents is to explore a given compact region. We propose a novel control scheme, where we introduce a supervisor that assists a group of agents with the centralized coverage control law and the global trajectory tracking control law. The coverage control law ensures the coverage task is done until the agents end up in local minima, and when they do, the global trajectory tracking control law ensures that the agents are deployed to uncovered regions. Our control scheme is designed to be decoupled such that only one control law is active at a given time. In addition to the coverage objective, we design control laws for coverage agents to avoid collisions and maintain proximity to a supervisor. Moreover, we utilize feedback linearization to use the proposed control scheme for coverage control of kinematic unicycle agents. We validate our approach via numerical simulations.

1. Introduction

In this work, we are concerned with the coverage control problem for multi-agent systems. In general, by coverage control, we refer to the static coverage control, and the dynamic coverage control. The static coverage problem can be traced back to locational optimization problems where the main objective is the optimal placement of sensors to cover a given domain. There are several works that deal with the static coverage problem; in Cortés, Martinez, Karatas, and Bullo (2004), the authors propose dynamic versions of the Lloyd algorithm to iteratively calculate control laws that drive agents to centroids of the Voronoi regions in a Voronoi partitioning of the domain. Similar approaches have been proposed in Cortés (2010), Cortés, Martinez, and Bullo (2004), Gao, Cortés, and Bullo (2008), Kwok and Martìnez (2010), Laventall and Cortés (2009), Schwager, Rus, and Slotine (2009) and Schwager, Slotine, and Rus (2011). In Li and Cassandras (2005), the authors propose control laws to maximize the detection probability of random events in a given domain while explicitly considering the communication cost in optimization. We refer to the aforementioned coverage problem as static because the agents come to a final configuration to accomplish the coverage objective.

The dynamic coverage problem is fundamentally different than the static one. The objective in dynamic coverage control is to develop control strategies for a group of mobile agents with limited range sensors such that the coverage level reaches a desired value at every point. We refer to it as dynamic because the agents move around to explore the given area instead of converging to a final configuration; thus, mobility is the key feature of the dynamic coverage problem. There has been a number of works in this area in the past few years. In Hussein and Stipanović (2006), a coverage error function for formulating a dynamic coverage problem is proposed. Agents switch to global control laws in order to relocate to uncovered regions when they get stuck in local minima of the error function. A discussion on the stability of the switching system is not provided in this work. Other works that build on the same approach include Hussein and Stipanović (2007a,b), Hussein, Stipanović, and Wang (2007), Song, Feng, Fan, and Wang (2011), Stipanović, Valicka, Tomlin, and Bewley (2013) and Wang and Hussein (2010). The switching behavior of the discontinuous control law is present in all of these works. In Hubel et al. (2008), the authors consider information decay in the environment. Although this approach provides a more realistic framework for the coverage problem, the switching behavior of the control law is still an issue. Moreover, only the coverage objective is considered in the work. Other papers that deal with the coverage control problem are Franco, Paesa, López-Nicolás, Sagüés, and Llorente (2012) and Franco, López-Nicolás, Sagüés, and Llorente (2013). Franco et al.

☆ This work has been supported by Qatar National Research Fund under NPRP Grant 4-536-2-199, the AFOSR grant FA95501210193 and the ARO grant W911NF-10-1-0044. The material in this paper was partially presented at the 52nd IEEE Conference on Decision and Control (CDC), December 10–13, 2013, Florence, Italy. This paper was recommended for publication in revised form by Associate Editor Riccardo Scattolini under the direction of Editor Frank Allgöwer.

E-mail addresses: gatinc2@illinois.edu (G.M. Atınç), dusan@illinois.edu (D.M. Stipanović), voulgari@illinois.edu (P.G. Voulgaris).

[1] Tel.: +1 217 220 2865; fax: +1 217 244 5705.

http://dx.doi.org/10.1016/j.automatica.2014.10.023
0005-1098/

(2012) propose a control scheme where the control laws constantly weigh between coverage and global deployment strategies, which rely on a hierarchical grid decomposition. In Franco et al. (2013), authors propose a coverage scheme with variable gains for handling the problem of persistent dynamic coverage. The stability analysis of the system, collision avoidance and proximity maintenance are not present in neither of these works.

In Haumann, Breitenmoser, Willert, Listmann, and Siegwart (2011) and Haumann, Listmann, and Willert (2010), coverage problem is converted into a problem of finding optimal directions that maximize the number of unexplored points that lie in front of each agent. Authors rightfully argue that collision avoidance is a natural outcome of the Voronoi partitioning. That being said, the partitioning needs to be updated continuously which is computationally expensive. Moreover, the proposed solution adds another optimization problem onto the partitioning problem, thus further increasing the computational complexity.

In this work, we propose a novel supervised dynamic coverage scheme for wheeled mobile robots that combines (i) coverage control law that ensures that agents accumulate sensory information until they become stationary, (ii) trajectory tracking control law that is designed to deploy idle agents to unexplored regions, (iii) avoidance control law that guarantees inter-agent collision avoidance and (iv) proximity control law that ensures that robots maintain their proximity to a supervisor.

A favorable feature of our approach is that, the computational burden is less compared to the static coverage control design; Voronoi partitioning of a given domain may be costly if the number of agents is large. The disadvantage of our supervised coverage scheme is that, it is centralized whereas the static coverage approaches are decentralized. Due to the nature of the dynamic coverage problem, the global coverage level of the whole domain has to be tracked; to do this, we explicitly include a supervisor. The global deployment laws can easily be generated/regenerated since the supervisor knows the positions of the coverage agents. Thus, by including a supervisor in our scheme, we allow the coverage agents to be deployed to uncovered regions in a systematic way at the cost of imposing centralized communication. That being said, we can convert our scheme into a decentralized one by considering multiple teams of agents, each with their own supervisor; the members of a team would communicate with their own supervisor only, and the supervisors would exchange the coverage level information.

One of the novelties of our approach lies in the structure of the control laws. Although the control laws proposed in Hussein and Stipanović (2006), Hussein and Stipanović (2007a,b), Hussein et al. (2007), Song et al. (2011), Stipanović et al. (2013) and Wang and Hussein (2010) are similar to the control scheme of our work, these control laws depict switching behavior, thus they are discontinuous. Moreover, a rigorous stability analysis of the switching behavior is not present. In contrast, the control signals in our work are designed to be continuous. In order to ensure the continuity of the control signals, we incorporate the trajectory tracking objective into the coverage error function and impose conditions so that the control signals smoothly transition between different operation modes. We utilize gradient-based control laws for collision avoidance and proximity maintenance which do not require any partitioning of the domain, thus reducing the computational complexity of the problem. Additionally, our proposed control law is such that the two components are decoupled, meaning that they are not active simultaneously. This decoupling allows us to break down the control problem into simpler subproblems that we analyze separately for different modes and provide the complete stability analysis for the overall coverage problem, including collision avoidance and proximity maintenance, which is, to the best of our knowledge, unique within the framework of dynamic coverage control approaches.

The decoupled structure of the proposed control law renders our scheme suitable for gain-scheduling type of approaches; the supervisor may select different gains for different modes in advance based on the value of the coverage error while maintaining the continuity of the signals. In addition, it allows for an asynchronous implementation of the supervised scheme in which the agents can simultaneously be in different modes. Finally, by constructing a scheme in which agents transition between different objectives while maintaining the continuity of signals, we propose a novel control design framework that can be extended to general control problems.

The rest of the paper is organized as follows. In Section 2, we discuss different elements of the coverage problem. In Section 3, we provide stability results for the supervised coverage control problem. We discuss supervised coverage control for wheeled mobile robots in Section 4. Subsequently, we present numerical simulations in Section 5. We provide concluding remarks in Section 6.

2. Elements of the dynamic coverage problem

In the following subsections, we describe the building blocks of our problem. We denote a compact domain to be covered by $\mathcal{D} \subset \mathbb{R}^2$.

2.1. Smooth transition signal

In our development, we use a continuously differentiable function $\gamma(t)$ that we will utilize later in designing the control laws. As long as a Boolean condition is "false", $\gamma(t) \equiv 1$. When the condition becomes "true", $\gamma(t)$ decreases smoothly from 1 to 0. As long as another Boolean condition is "false", $\gamma(t) \equiv 0$. When the second condition becomes "true", $\gamma(t)$ increases smoothly from 0 to 1. Any function that behaves in the described way can be utilized as $\gamma(t)$.

We want to distinguish between the regions where $\gamma(t)$ changes behavior; i.e., **Mode 1:** $\gamma(t) \equiv 1$, **Mode 2:** $0 \leq \gamma(t) \leq 1$, $\dot{\gamma}(t) \leq 0$, **Mode 3:** $\gamma(t) \equiv 0$ and **Mode 4:** $0 \leq \gamma(t) \leq 1, 0 \leq \dot{\gamma}(t)$.

2.2. Sensing and accumulated information

Along the lines of Hussein and Stipanović (2006), Hussein and Stipanović (2007a,b), Hussein et al. (2007) and Stipanović et al. (2013), we utilize the following sensor model:

$$S_i(t, \|p_i - \tilde{p}\|^2) = \gamma(t)\tilde{S}_i(\|p_i - \tilde{p}\|^2), \\ \triangleq \gamma(t)\frac{M_{cov_i}}{R^4_{cov_i}} \max\{0, R^2_{cov_i} - \|p_i - \tilde{p}\|^2\}^2, \quad (1)$$

where $p_i \in \mathbb{R}^2$ is the position of an agent, $\tilde{p} \in \mathcal{D}$, M_{cov_i} is the maximum sensing level and R_{cov_i} is the sensing region. Using the sensor model (1), we formulate the accumulated sensory information in the following way:

$$Q(t, \mathbf{P}) = C^{**} - C^{**}e^{-k^*A}, \quad (2)$$

where $\mathbf{P} = \begin{bmatrix} p_1^T & \cdots & p_N^T \end{bmatrix}^T$ and k^* is a design variable. Also, $A(t, \tilde{p}) = \int_0^t \sum_{i=1}^N S_i(t, \|p_i(\tau) - \tilde{p}\|^2)d\tau$. One may think of C^* as a proxy for quantifying how well a certain area is explored. If the coverage level at a particular region is C^*, we consider that region to be sufficiently explored. In order to ensure that the desired coverage level C^* can be exactly attained at any point in a given domain, we design $Q(t, \mathbf{P})$ such that its horizontal asymptote, C^{**}, is at a level that is slightly greater than C^*.