Automatica 72 (2016) 64–72

Contents lists available at ScienceDirect

Automatica

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

Brief paper

Cooperative pursuit with Voronoi partitions☆

Zhengyuan Zhou [a], Wei Zhang [b], Jerry Ding [c], Haomiao Huang [d], Dušan M. Stipanović [e], Claire J. Tomlin [c]

[a] Department of Electrical Engineering, Stanford University, Stanford, CA, 94305, USA
[b] Department of Electrical and Computer Engineering, Ohio State University, Columbus, OH, 43210, USA
[c] Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA, 94720, USA
[d] Kuna Systems, CA, USA
[e] Control and Decision Group, Coordinated Science Laboratory, University of Illinois, Urbana, IL, 61801, USA

ARTICLE INFO

Article history:
Received 26 January 2014
Received in revised form
6 February 2016
Accepted 21 April 2016

Keywords:
Pursuit–evasion games
Voronoi
Cooperative pursuit

ABSTRACT

This work considers a pursuit–evasion game in which a number of pursuers are attempting to capture a single evader. Cooperation among multiple agents can be difficult to achieve, as it may require the selection of actions in the joint input space of all agents. This work presents a decentralized, real-time algorithm for cooperative pursuit of a single evader by multiple pursuers in bounded, simply-connected planar domains. The algorithm is based on minimizing the area of the generalized Voronoi partition of the evader. The pursuers share state information but compute their inputs independently. No assumptions are made about the evader's control strategies other than requiring the evader control inputs to conform to a speed limit. Proof of guaranteed capture is shown when the domain is convex and the players' motion models are kinematic. Simulation results are presented showing the efficiency and effectiveness of this strategy.

1. Introduction

This paper studies a multi-agent pursuit–evasion game, with a number of pursuers attempting to capture a single evader in a simply connected planar region. The pursuers' speeds are equal to or greater than that of the evader, and the objective is to find a successful cooperative strategy for the pursuers. Finding cooperation strategies among multiple agents in adversarial games can be challenging, as computing solutions over the joint state space of multiple agents can greatly increase computational complexity. The class of pursuit–evasion games considered here can, in a general theoretical setting, be treated as a multi-agent differential game. The solution to such a problem can, in principle, be obtained by solving the corresponding Hamilton–Jacobi–Isaacs (HJI) partial differential equation (PDE) (Başar & Olsder, 1999; Evans & Souganidis, 1984; Isaacs, 1967; Mitchell, Bayen, & Tomlin, 2005). In particular, one can define the game through a value function representing the time-to-capture, with the evader attempting to maximize this function and the pursuers attempting to minimize this function. Under certain technical conditions, the value of the game can be characterized as the solution to an HJI equation, which can be in turn used to synthesize optimal controls for the pursuers to minimize time-to-capture. Solutions to HJI equations are typically found either using the method of characteristics (Başar & Olsder, 1999; Isaacs, 1967), in which optimal trajectories are found by integrating backward from a known terminal condition, or via numerical approximation of the value function on a grid of the continuous state space (Ding, Sprinkle, Shankar Sastry, & Tomlin, 2008; Falcone & Ferretti, 2002; Huang, Ding, Zhang, & Tomlin, 2011; Mitchell et al., 2005).

The practical usage of the differential game approach, however, faces several computational challenges. While the characteristic solutions are useful in understanding optimal solutions qualitatively, they require backward integration from terminal configurations, which can make it difficult to generate strategies when only

☆ This work has been supported in part by NSF under CPS:ActionWebs (CNS-931843), by ONR under the HUNT (N0014-08-0696) and SMARTS (N00014-09-1-1051) MURIs and by grant N00014-12-1-0609, by AFOSR under the CHASE MURI (FA9550-10-1-0567). The material in this paper was presented at the 50th IEEE Conference on Decision and Control and European Control Conference, December 12–15, 2011, Orlando, FL, USA. This paper was recommended for publication in revised form by Associate Editor Hideaki Ishii under the direction of Editor Ian R. Petersen.

E-mail addresses: zyzhou@stanford.edu (Z. Zhou), zhang@ece.osu.edu (W. Zhang), jerryding@berkeley.edu (J. Ding), haomiao@stanford.edu (H. Huang), dusan@illinois.edu (D.M. Stipanović), tomlin@eecs.berkeley.edu (C.J. Tomlin).

http://dx.doi.org/10.1016/j.automatica.2016.05.007

the initial configurations of the agents are known. On the other hand, HJI computation on grids suffers from the curse of dimensionality: computing solutions to HJI equations is computationally infeasible for scenarios with a large number of agents, as the grid required for approximating the value function grows exponentially in the dimensions of the joint configuration space.

For certain games and game configurations, it is possible to construct strategies for the agents geometrically. For example, pure-distance pursuit, in which a pursuer minimizes the instantaneous distance to the evader, has been shown to be the optimal pursuit strategy for certain zero-sum differential games in open environments (Isaacs, 1967), as well as a choice of strategy which guarantees capture in simply-connected regions (Alexander, Bishop, & Ghrist, 2006). In other cases, strategies based upon a geometric argument have also been found for coordinating groups of pursuers in open, unbounded spaces (Kopparty & Ravishankar, 2005). These methods are computationally efficient in generating control strategies, but are often limited to relatively simple game environments with no obstacles and typically assume homogeneous player speed.

In addition to games in continuous time and continuous spaces, research in discrete, turn-based games played on graphs have shown that three pursuers are sufficient and sometimes necessary to capture any evader in a planar graph (Aigner & Fromme, 1984; Parsons, 1978). These results for discrete games have led to strategies for a class of continuous games known as visibility-based pursuit–evasion (Gerkey, Thrun, & Gordon, 2006; Guibas, Latombe, Lavalle, Lin, & Motwani, 1997; LaValle & Hinrichsen, 2001). The graph-based analysis has also recently inspired results for continuous games showing that three pursuers are also sufficient and sometimes necessary to capture an evader with equal speed in bounded polygonal domains with obstacles (Bhadauria, Klein, Isler, & Suri, 2012). Similar to the graph and visibility pursuit strategies, they operate on the principle of successively reducing the game domain into a single simply-connected region by having pursuers block the evader from portions of the game space. However, to the best of our knowledge, currently implementable versions of these strategies do not yet exist. For all these methods, the game domain reduction requires searching over a large set of discrete actions, limiting the size of problems that can be practically solved.

In this paper, we present a decentralized pursuit strategy based on the Voronoi decomposition of the game domain with respect to agent positions, where the pursuers cooperatively minimize the area of the evader's Voronoi partition. Each pursuer influences the evader's partition only through the shared Voronoi boundary. Thus, each pursuer's input decouples from that of the other pursuers and can be computed independently. However, their inputs are coupled through the Voronoi partition, giving rise to cooperation among the pursuers. The pursuit algorithm is decentralized in the sense that the pursuers compute their control actions independently given the agent positions, which is the only shared information. This approach allows computation to take place in the low dimensional configuration space of individual agents instead of the high dimensional joint state space of all agents (as is the case in the HJI computation), thus enabling real-time implementation. We mention that the body of work described in Ames et al. (2014) and Bhadauria et al. (2012) is relevant to the multi-agent pursuit–evasion problems discussed in this paper, and is exciting as it provides formal proofs for the existence of pursuit strategies that guarantee capture, with provable bounds on the time-to-capture. On the other hand, the practical computation of such ideal strategies is still a subject of ongoing investigations. Our approach can be viewed as complementary to these theoretical contributions, in that we derive computable solution strategies for the particular problem scenario of convex environments and equal speeds. The insights obtained from this scenario may provide the foundation for the construction and implementation of more general strategies.

1.1. Our contributions

Our contributions are threefold. First, we propose a Voronoi partition based pursuit strategy and show that, under the assumptions of convex environments, kinematic agent dynamics, and equal speeds of all players, this strategy results in guaranteed capture of the evader in finite time. Some elements of the results for convex domains and equal speeds were first presented in Huang, Zhang, Ding, Stipanović, and Tomlin (2011). This work elaborates upon the previous results, gives an equivalent characterization of the control input that is more intuitive and easily implementable (Theorem 2) and provides a simpler proof of guaranteed capture based on a construction of an energy function (Theorem 3).

Second, we generalize the Voronoi pursuit strategy to non-convex game domains and unequal agent speeds, in particular when the pursuers are faster than the evader, and show how to apply a modified fast marching method (FMM) (Zhou, Takei, Huang, & Tomlin, 2012) to quickly compute generalized Voronoi partitions and pursuer inputs. This approach provides a scalable, computationally efficient, and easily implementable algorithm for computing cooperative pursuer inputs in a multi-pursuer scenario.

Third, extensive simulation studies are carried out to show the effectiveness of the proposed strategy and compare its performance and computational complexity with several representative methods in the literature. In particular, it is shown that the algorithm encourages effective cooperative pursuit among multiple agents, resulting in superior performance to techniques such as the pure-pursuit, in which the pursuers attempt to minimize the instantaneous distance to the evader, and comparable performance to the optimal pursuit strategy based on HJI calculation, which is computational intensive and therefore not feasible beyond a single-pursuer–single-evader scenario.

2. The cooperative pursuit problem

Consider a multi-agent pursuit–evasion game involving N pursuers and a single evader, taking place in an open, simply connected region Ω in $\mathbb{R}^2$. Let $x_e \in \mathbb{R}^2$ be the position of the evader and $x_p^i \in \mathbb{R}^2$ be the position of pursuer i. The equations of motion are

$$\begin{aligned} \dot{x}_e &= d, \qquad x_e(0) = x_e^0, \\ \dot{x}_p^i &= u_i, \qquad x_p^i(0) = x_p^{i,0}, \quad i = 1, \ldots, N, \end{aligned} \tag{1}$$

where d and u_i are the velocity control inputs of the evader and pursuers, respectively, and $x_e^0, x_p^{i,0} \in \Omega$ are the initial evader and pursuer positions. The respective agent inputs are constrained to lie within sets $U_i \subset \mathbb{R}^2$ for the pursuers and $D \subset \mathbb{R}^2$ for the evader. In this paper, U_i and D are assumed to be the following:

$$D = \{d \mid \|d\| \le v_{e,\max}\}, \qquad U_i = \{u_i \mid \|u_i\| \le v_{i,\max}\}, \tag{2}$$

where $v_{e,\max}$ and $v_{i,\max}$ are the maximum speeds of the evader and the pursuers, respectively, and $\|\cdot\|$ denotes the Euclidean norm in $\mathbb{R}^2$. The motions of the evader and pursuers, as described by Eq. (1), are also constrained to lie within the region Ω, i.e.,

$$x_e(t),\ x_p^i(t) \in \Omega, \quad \forall t \ge 0. \tag{3}$$

Any velocity input $d(t)$ or $u_i(t)$ satisfying the constraints (2) and (3) is called an admissible input for the evader or pursuer i, respectively.

The goal of the pursuers is to capture the evader by having at least one of the pursuers come within a distance $r_c > 0$ of the evader. To achieve this capture condition, each pursuer selects control inputs using a pursuit strategy $\mu_i(x_e, x_p^1, \ldots, x_p^N)$ based upon observations of the evader and pursuer positions at each time