Proceedings of the American Control Conference
Anchorage, AK May 8-10, 2002

A study in cooperative control: The RoboFlag Drill

Matthew G. Earl and Raffaello D'Andrea
Cornell University
{mge1,rd29}@cornell.edu

Abstract

We take an optimization approach, inspired by [1], toward synthesizing control strategies for cooperative muti-agent systems. We successfully apply this approach to a subproblem derived from a new test-bed for studies in cooperative systems called the RoboFlag competition.

1 Introduction and results

We are interested in using collections of agents in a cooperative manner to autonomously perform complex tasks. This is not a new idea given that there exist autonomous multi-agent systems today that work fairly well. Cornell's RoboCup team [2] is a good example. The system consists of five autonomous vehicles that compete against another team of five vehicles in a game of soccer on a field the size of a ping-pong table. When you watch the Cornell team compete it is fairly obvious that the robots are cooperating to achieve a common goal—to win the game. This is most apparent when the robots pass the ball to one another to avoid defending robots and ultimately to get a clear shot on goal. To achieve a cooperative strategy the system is decomposed hierarchically into several levels. The lower level components, such as trajectory generation, have been rigorously considered in [2]. High level components, such as offensive plays for passing the ball to achieve a clear shot on goal, have been developed using human intuition [2] since the problem is extremely difficult.

RoboCup is cutting edge technology but, in general, you do not see cooperative multi-agent systems in practice. However, there is a demand for autonomous systems to take care of the dull, dirty, and dangerous tasks that arise in life and, furthermore, we feel that autonomous multi-agent systems will have interesting advantages in scalability, redundancy, robustness, among other advantages, over single-agent systems. The reason multi-agent systems are not used in practice is in part due to two important issues. First, for many applications, the control strategy for such systems must meet performance and safety criteria which currently is hard to guarantee. And second, developing strategies for large complex systems is extremely difficult. Therefore, there is a need for systematic analysis and synthesis tools. Numerous researchers from computer science, controls, and the operations research communities have contributed valuable work toward this end—see the papers [3, 4, 5, 6] and the references therein; in fact, an interdisciplinary approach seems ideal. In this paper we present first steps toward strategy synthesis, via an optimization approach inspired by [1], within the context of a new test-bed for studies in cooperative control.

The RoboCup competition is a valuable means to study multi-agent systems but, due to the small field and fast pace of the game, complex strategies are usually not optimal. Due to this the second author has proposed RoboFlag as a companion competition in hopes to provide a fertile test-bed for advances in cooperative control technology involving complex strategies. The RoboFlag competition is described in detail on the second authors web page; here we give a brief description to give a flavor for the game. The competition is between two teams of eight robots appease and is similar to the well known games "capture the flag" and "paintball." The objective is to infiltrate your opponents territory, get their flag, and bring it to your home base while, at the same time, trying to thwart your opponent from capturing your own flag. While pursuing the objective there are a number of other factors to worry about. Some parts of the field are off limits, moving obstacles and golf balls shot by the opponents must be avoided, and each robot has a limited amount of fuel. This game is complicated and it would be a daunting task to develop winning strategies at this point in time.

As a first step we consider a vastly simplified version of the competition which we call the RoboFlag Drill. The drill involves two teams of robots, the attackers and the defenders, on a playing field with a region at its center called the defense zone. The attackers are drones directed toward the center of the defense zone along a straight line path at a constant velocity. The objective for the defenders is to thwart the attackers from entering the defense zone by intercepting each attacker before it enters the zone. Once an attacker enters the defense zone or is intercepted by a defender it remains stationary for the remainder of the drill. While pursuing its objective defenders must avoid collisions with other defenders and stationary obstacles as well as avoid entering the defense zone which is off limits to defending robots.

To intelligently perform the drill the defenders need a strategy that is consistent with the objective and ideally optimal with respect to the objective. We assume the strategy is implemented with a centralized controller with perfect knowledge of the system, perfect access to all states, and with the ability to transmit control signals to the defenders instantaneously. The controller needs to figure out the inputs to provide each defending robot so that the objective is achieved. To achieve this we pose the problem as an optimization. We seek a set of control inputs that minimize the number of attackers that enter the defense zone over the duration of the drill and, in addition, is consistent with the system dynamics (robot dynamics) and the constraints (no collisions, etc.).

Inspired by [1] we model the RoboFlag Drill as a system composed of continuous and discrete states (a hybrid system) with linear dynamics subject to inequality constraints and logical (if-then-else) rules. In this formulation robot dynamics are modeled with second order differential equations as in [2]. We convert this model to mixed logical dynamical (MLD) form [1] with the help of the software package HYSDEL [7]. With the system dynamics and constraints in MLD form we pose our objective precisely. Following the procedure in [1], we convert the optimization problem to a mixed integer linear program (MILP) and solve using the commercially available solver CPLEX [8]. Figure 1 shows three defenders successfully thwarting eight attacking robots from entering the defense zone. In [9] similar ideas, developed independently, are used for aircraft trajectory planning.

For the simplified problem at hand we have successfully generated a strategy to minimize the objective, but the procedure is computationally intensive and assumes a static environment. However, by exploiting problem specific structure and imposing a control hierarchy we feel that the optimization can be performed repeatedly throughout the duration of the drill (via model predictive control) and thus generate near optimal strategies for dynamically changing and uncertain environments.

Acknowledgments

We thank Richard Murray and Jason Hickey from Caltech and Carla Gomes, Bart Selman, and Pritam Ganguly from Cornell for helpful discussions.

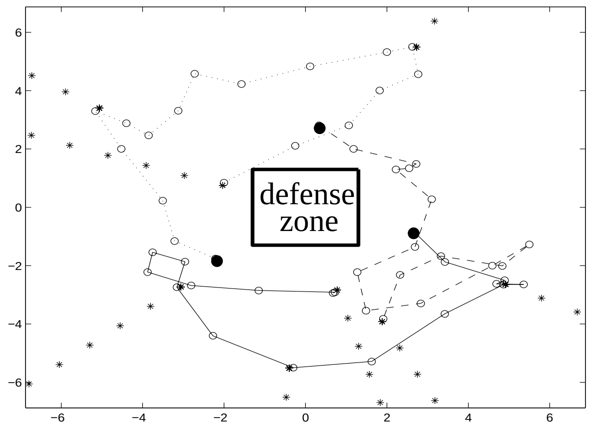

Figure 1: Three defending robots (circles) successfully defend the zone from eight attacking robots (stars). Filled circles denote the initial defender locations.

References

[1] A. Bemporad and M. Morari. Control of systems integrating logic, dynamics, and constraints. *Automatica*, vol. 35, pp. 407-428, Mar. 1999.

[2] D'Andrea, R., Kalmár-Nagy, T., Ganguly, P., Babish, M. The Cornell Robocup Team. In Kraetzschmar, G., Stone P., Balch T. editors: *Robot Soccer WorldCup IV, Lecture Notes in Artificial Intelligence.* Springer, 2001.

[3] R. Alur, A. Das, J. Esposito, R. Fierro, Y. Hur, G. Grudic, V. Kumar, I. Lee, J. P. Ostrowski, G. Pappas, J. Southall, J. Spletzer, and C. Taylor. A framework and architecture for multirobot coordination. *Experimental Robotics: LNCS Series.* Springer-Verlag, 2001.

[4] M. Bowling and M. Veloso. Multiagent learning using a variable learning rate. *Artificial Intelligence.* 2002, to appear.

[5] M. J. Mataric. Learning in Behavior-Based Multi-Robot Systems: Policies, Models, and Other Agents. In *Cognitive Systems Research, special issue on Multi-disciplinary studies of multi-agent learning.* Ron Sun, ed., 2(1), Apr 2001, 81-93.

[6] P. Stone. *Layered Learning in Multiagent Systems: A Winning Approach to Robotic Soccer.* MIT Press, 2000.

[7] F.D. Torrisi, A. Bemporad, and D. Mignone. HYSDEL—A language for describing hybrid systems. Technical Report AUT00-03, ETH Zurich, 2000. `http://control.ethz.ch/ hybrid/hysdel`

[8] `http://www.ilog.com/products/cplex/`

[9] A. Richards and J. P. How. Aircraft Trajectory Planning With Collision Avoidance Using Mixed Integer Linear Programming. In *Proc. American Control Conf.*, 2002.