Generating Strategies for Multi-Agent Pursuit-Evasion Games in Partially Observable Euclidean Space

(Extended Abstract)

Eric Raboin
University of Maryland
College Park, MD 20742
eraboin@cs.umd.edu

Ugur Kuter
Smart Information Flow Technologies
Minneapolis, MN 55401 USA
ukuter@sift.net

Dana Nau
University of Maryland
College Park, MD 20742
nau@cs.umd.edu

ABSTRACT

We present a heuristic search technique for multi-agent pursuit-evasion games in partially observable Euclidean space where a team of tracker agents attempt to minimize their uncertainty about an evasive target agent. Agents' movement and observation capabilities are restricted by polygonal obstacles, while agents' knowledge of each others' location is limited to direct observation or periodic updates from team members.

Our polynomial-time algorithm is able to generate strategies for games in continuous two-dimensional Euclidean space, an improvement over past algorithms that were only applicable to simple gridworld domains. We show experimentally that our algorithm is tolerant of interruptions in communication between agents, continuing to generate good strategies despite long periods of time where agents are unable to communicate directly. Experimental results also show that our technique generates effective strategies quickly, with decision times of less than a second for reasonably sized domains with six or more agents.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Multiagent systems*; I.2.8 [**Artificial Intelligence**]: Problem Solving, Control Methods, and Search

General Terms

Algorithms

Keywords

visibility-based pursuit-evasion, multi-agent planning, game theory

1. INTRODUCTION

Our work introduces a strategy generation technique for multi-agent pursuit-evasion games in continuous, partially observable Euclidean space. We provide a polynomial time algorithm capable of generating online strategies for a team of cooperative tracker agents that wish to pursue an evasive target. The goal of the tracker team is to minimize their uncertainty about the target's location by the end of a fixed time period. The domain may have arbitrarily shaped polygonal obstacles that limit movement as well as observability.

Minimizing uncertainty about a target's location is distinct from the goal assumed by most other pursuit-evasion formalisms. The tracker team must work to maintain visibility on the target, but also to move to strategic locations prior to visibility loss so that recovery will be possible. Previous approaches that sought to maintain visibility on the target for as long as possible [2, 3], or generate patrol strategies to find a hidden target [6, 1], may not be suited for scenarios where the target frequently passes in and out of visibility. Since we want to generate strategies quickly, this also rules out many techniques that are based on deep combinatorial search.

Prior work on this problem included a game-tree search algorithm that could generate strategies for simple gridworld domains, where time was divided into discrete time steps and agents were only permitted to move in one of four cardinal directions [4]. This previous work also assumed that agents would be in constant communication, since it generated trajectories using a heuristic method that required knowing the location of every agent on the team.

In this paper we introduce the *Limited-communication Euclidean-space Lookahead* (*LEL*) heuristic, a method for evaluating tracker strategies in games where agents can move freely in two-dimensional Euclidean space and where there may be long periods of time when communication between agents is interrupted.

Our contributions include–

- An algorithm for computing the *LEL* heuristic in two-dimensional Euclidean space with polygonal obstacles, where communication between agents may be interrupted for long periods of time.
- Complexity analysis showing that our algorithm for computing *LEL* runs in polynomial time with respect to the size of the domain and the number of agents per team.
- Experimental results showing that our algorithm quickly generates strategies for the continuous domain that are twice as effective at retaining visibility on the target when compared to a strategy that follows the shortest path to the target.

2. DEFINITIONS

We define a multi-agent, zero-sum imperfect-information game where a single *target* agent a_0 is pursued by a team of n *tracker* agents $\{a_1, a_2, \ldots a_n\}$. The goal of the tracker team is to minimize its uncertainty about the target's location by the end of the game, while the target agent is free to evade the trackers and move behind obstacles that obstruct visibility.

We assume that each agent a_i is a holonomic point robot with a fixed maximum velocity v_i. The domain can have multiple ob-

Appears in: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, Conitzer, Winikoff, Padgham, and van der Hoek (eds.), 4-8 June 2012, Valencia, Spain.

stacles, where each obstacle is a solid polygon in $\mathbb{R}^2$. We define a reachability function $R_i(L, t)$ as the set locations reachable by a_i from any location $l \in L$ by time t, and a visibility function $V_i(L)$ as the set of locations visible to a_i from any $l \in L$.

We define a target's *hidden* region as the set of locations L_{hidden} where the target could be located given the prior observations made by the tracker team. If target a_0 was last observed at location l at time t_j, then it could be located anywhere in $R_0(\{l\}, t_k - t_j)$ by time t_k. This region can be narrowed further by subtracting the locations visible to the tracker team between time t_j and t_k, accounting for the possible trajectories followed by the target during the same time period.

At the end of the game, the utility for the tracker is $-|L_{hidden}|$, while the utility for the target is $|L_{hidden}|$. The best possible outcome for the tracker team is when the exact location of the target is known, meaning L_{hidden} is empty.

3. LEL HEURISTIC

Our work introduces the *Limited-communication Euclidean-space Lookahead* (*LEL*) heuristic, a method of evaluating trajectories for the tracker team by estimating the future size of the target's *hidden* region. This estimate is based off a relaxation of the movement and observation capabilities for each of the agents, which allows the heuristic to be computed in polynomial time but still provide a reasonably good estimate of each trajectory's utility.

To compute LEL, we first determine the set of locations reachable by the target at some time t, given by $R_0(L_{hidden}, t)$. We then determine $R_i(\{l_i\}, t)$ for each tracker agent a_i given their last known location l_i. With this, we can estimate the size of the hidden region at time t by computing,

$$L_{approx}(t) = R_0(L_{hidden}, t) \setminus \bigcup_{i=1}^{n} V_i(R_i(l_i, t)) \qquad (1)$$

which is the set of all locations reachable by the target at time t, minus the set of all locations observable by the tracker. If the target has not been observed by time t, then $L_{approx}(t)$ will be a subset of the actual hidden region, since each tracker a_i will only be able to observe some of the locations in $V_i(R_i(l_i, t))$.

To evaluate a trajectory for agent a_i, we compute the sum of $|L_{approx}(t)|$ over a fixed time interval, where the initial location for a_i is set to a location along the trajectory being evaluated. We have developed a polynomial-time algorithm to do this that utilizes the *Fast Marching Method* [5], a linear-time algorithm for computing the shortest-path distances over two-dimensional Euclidean space. An example of the composite generated by this technique is shown in figure 1, where darker areas represent locations that are contained by L_{approx} for a larger portion of the time interval.

4. EXPERIMENTS

To evaluate *LEL*, we conducted experiments on 500 randomly generated domains with two-dimensional polygonal obstacles. Each trial was run for a fixed amount of time and the size of the hidden region was measured at the end of the game. For comparison purposes we evaluated both *LEL* and the *max-distance* (*MD*) heuristic, a simple hand-coded rule that instructs the tracker team to follow the "shortest-path" to the target [4].

Figure 2 shows the average size of the hidden region at the end of the game for several team sizes, where a smaller hidden region indicates greater certainty about the target's location. Teams using the *LEL* heuristic were more than twice as effective when compared to teams using the *MD* heuristic, even when the teams using the *MD* heuristic had more agents.

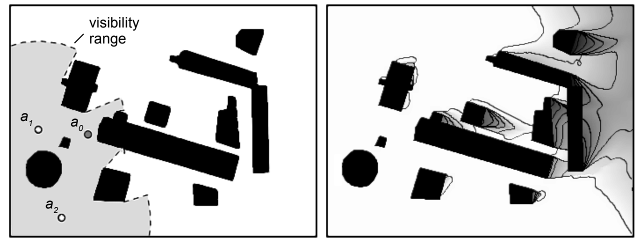

Figure 1: (left) An example game with two trackers (a_1, a_2) and one target (a_0) in their initial positions. (right) *LEL* composite for this state, darker areas indicate potential for visibility loss.

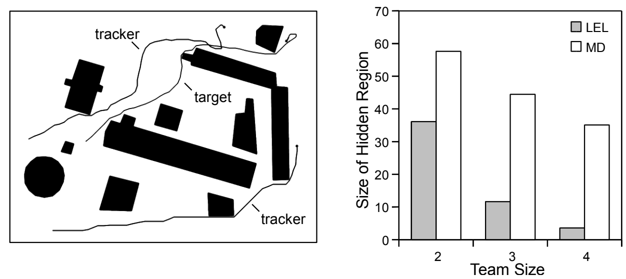

Figure 2: (left) Trajectories generated using the *LEL* heuristic in an example game, (right) Average size of the *hidden* region at the end of the game for different sized teams using the *RLA* and *MD* heuristics. Smaller values equal better performance.

5. ACKNOWLEDGEMENTS

This work was supported, in part, by DARPA and the U.S. Army Research Laboratory under contract W911NF-11-C-0037, and by a UMIACS New Research Frontiers Award. The views expressed are those of the authors and do not reflect the official policy or position of the funders.

6. REFERENCES

[1] S. LaValle, D. Lin, L. Guibas, J. Latombe, and R. Motwani. Finding an unpredictable target in a workspace with obstacles. In *Proc. International Conference on Robotics and Automation (ICRA)*, pages 737–742, 1997.

[2] T. Muppirala, S. Hutchinson, and R. Murrieta-Cid. Optimal motion strategies based on critical events to maintain visibility of a moving target. In *Proc. International Conference on Robotics and Automation (ICRA)*, pages 3826–3831, 2005.

[3] R. Murrieta, A. Sarmiento, S. Bhattacharya, and S. Hutchinson. Maintaining visibility of a moving target at a fixed distance: The case of observer bounded speed. In *Proc. International Conference on Robotics and Automation (ICRA)*, volume 1, pages 479–484, 2004.

[4] E. Raboin, D. S. Nau, U. Kuter, S. K. Gupta, and P. Svec. Strategy generation in multi-agent imperfect-information pursuit games. *Proc. International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 947–954, 2010.

[5] J. A. Sethian. A fast marching level set method for monotonically advancing fronts. In *Proc. Nat. Acad. Sci*, pages 1591–1595, 1995.

[6] I. Suzuki and M. Yamashita. Searching for a mobile intruder in a polygonal region. *SIAM Journal on Computing*, 21:863–888, 1992.