

Scalable and Practical Pursuit-Evasion with Networked Robots

Marcos A. M. Vieira · Ramesh Govindan · Gaurav S. Sukhatme

Received: date / Accepted: date

Abstract In this paper, we consider the design and implementation of practical pursuit-evasion games with networked robots, where a communication network provides sensing-at-a-distance as well as a communication backbone that enables tighter coordination between pursuers. We first develop, using the theory of zero-sum

This material is based in part upon work supported by the National Science Foundation under Grants No. CCF-0820230, CNS-0540420, CNS-0325875 and CCR-0120778 and a gift from the Okawa Foundation. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author (s) and do not necessarily reflect the views of the National Science Foundation. Marcos Vieira was supported in part by Grant 2229/03-0 from CAPES, Brazil.

Marcos A. M. Vieira
University of Southern California
3710 S. McClintock Avenue
Ronald Tutor Hall (RTH) 418
Los Angeles, CA 90089
Tel.: +1-213-8215627
Fax: +1-213-7407285
E-mail: mvieira@usc.edu

Ramesh Govindan
University of Southern California
MC 2905
Ronald Tutor Hall (RTH412)
3710 S. McClintock Ave
Los Angeles, CA 90089-2905
Tel: +1-213-7404509
Fax: +1-213-7407285
E-mail: ramesh@usc.edu

Gaurav S. Sukhatme
University of Southern California
MC 2905
Ronald Tutor Hall (RTH 405)
3710 South McClintock Avenue
Los Angeles, California 90089-2905
Tel.:+1-213-7400218
Fax: +1-213-8215696
E-mail: gaurav@usc.edu

games, an algorithm that computes the minimal completion time strategy for pursuit-evasion when pursuers and evaders have same speed, and when all players make optimal decisions based on complete knowledge. Then, we extend this algorithm to when evader are significantly faster than pursuers. Unfortunately, these algorithms do not scale beyond a small number of robots. To overcome this problem, we design and implement a partition algorithm where pursuers capture evaders by decomposing the game into multiple multi-pursuer single-evader games. We show that the partition algorithm terminates, has bounded capture time, is robust, and is scalable in the number of robots. We then describe the design of a real-world mobile robot-based pursuit evasion game. We validate our algorithms by experiments in a moderate-scale testbed in a challenging office environment. Overall, our work illustrates an innovative interplay between robotics and communication.

Keywords Networked Robots · Pursuit-Evasion Game · Wireless Sensor Network

1 Introduction

We are motivated by practical problems in security and monitoring for large, structured, spaces (e.g., to ensure the integrity of a large building or complex). The problem we focus on is pursuit-evasion wherein robots must pursue and catch evaders.

In Pursuit-Evasion Games (PEGs), multiple robots (the pursuers) collectively determine the location of one or more evaders, and try to corral them. The game terminates when every evader has been corralled by one or more robots.

Since the definition of the discrete pursuit-evasion game by Parsons [31], PEGs have received significant attention [3][27][33][16][37]. The main difference between existing approaches and our approach is that we play the multi-pursuer multi-evader game with full knowledge of the game. Pursuers make use of an embedded sensor network to have complete knowledge of the game.

The main contributions of our work are as follows. First, we describe an optimal capture time algorithm (Algorithm 1) to minimize the time to capture of a multi-pursuer multi-evader game, even with optimal adversarial evader behavior. Second, we extend this algorithm to the case (Algorithm 2) when the evaders are infinitely faster than the pursuers. Although these algorithms are optimal, they are exponential in the number of players. Thus, our third contribution is a scalable algorithm (Algorithm 3) which guarantees performance close to optimal. Finally, we validate our algorithms by experiments on a moderate-scale real-world testbed.

Several versions of PEG exist. In certain frameworks, it is acceptable to merely “sight” an evader for it to be “located” [16], in others, a precise coordinate must be reported [29]. Other formulations insist on a certain speed of convergence with fewer constraints on accuracy [3,37]. Finally, formulations vary depending on whether the multi-robot control algorithm is required to have provably correct behavior, whether the number of evaders is known a priori, and whether they are malicious or benign. Each variation of the problem brings with it a different set of challenges, and several of these variations have been solved to varying degrees. In section 8, we provide a detailed review of the literature.

We focus on PEGs in bounded, spatially complex environments similar to today’s office environments. In such environments, we can assume imperfect geometric regularity (e.g., the presence of corridors and 90-degree turns, but possibly a regular placement of doorways or elevator exits). However, it is increasingly true that such environments are well provisioned with wireless communication capability, and that many such environments will likely have dense embedded sensing (for surveillance or environmental control). This environmental embedded sensing can provide, for the players, full knowledge of the game.

We consider the class of PEGs played on a discrete graph [3][27][33]. Specifically, we use a topological map of the environment, whose nodes correspond to coarse-grained regions and whose links connect neighboring regions [20,22,38]. Discrete graph-based games are acceptable for many uses of pursuit-evasion (e.g., surveillance, finding survivors). Of course, the actual games execute in a continuous space, but the discrete graph

is used to define our game model and to localize the participants.

The pursuers make use of the resources of an embedded network. This network provides sensing, communication and computational resources to the robots: *Sensing*: In the traditional setting the only sensing resources available to robots are the ones they carry. In our setting robots, via communication, effectively sense-at-a-distance. Networked sensing can provide proprioceptive as well as exteroceptive sensing (e.g. network nodes can track robot pose as well as environmental features) and potentially provides sensing redundancy. *Communication*: In addition to effectively increasing the sensing range of robots, the presence of the embedded network effectively extends the inter-robot communication range and improves the quality of inter-robot communication.

Computation: The sensor network could potentially be viewed as a computational resource by the robots — this is particularly true since sensor networks are no longer confined to large numbers of very simple networked processors. This means that the robot planning and coordination algorithms could run on one (or more) of the more capable nodes in the network, while the robots’ onboard computation is restricted to low-level control loops and sensor filtering.

In this paper, we consider a version of the game in which p pursuers collectively attempt to capture r evaders (Section 2). We are interested in the convergence time of the game (i.e., the minimum number of steps for the pursuers to capture the evaders). We start by presenting the first provably minimal convergence time algorithm for multiple pursuers and multiple evaders (Section 3). In our formulation, all players have complete knowledge of the state of other players, and the evaders also pursue an optimal evasion strategy. Based on the optimal policy that pursuers should use in order to capture evaders with the minimum number of steps, we design an assignment algorithm that optimally decomposes the game (Section 4).

Unfortunately, the computation required for this optimal algorithm does not scale beyond a small number of participants (less than 10, with today’s technology). To provide a practical and scalable system, we consider an approximate strategy: we decompose the multi-player game into multiple multi-pursuer single-evader games. We prove that our approximation terminates, has bounded captured time, is robust, and is scalable in the number of robots, being suited for practical applications. Clearly, convergence time depends on the topology as well as on the number of pursuers and evaders: we explore these dependencies for a few canonical topologies (Section 6). Finally, we present re-

sults from running an implementation of our algorithm on a physical robot testbed (Section 7). The experiments confirm the feasibility of our algorithm, even on a worst-case initial configuration.

2 Assumptions, Terminology and Definitions

In this section, we start by stating the sensing and communication assumptions for our PEGs, then discuss the class of games we are interested in. We then lay down some terminology, and formally define the objective of our PEGs. This sets the stage to present the optimal capture time algorithm and also a scalable near-optimal algorithm for PEG, which are discussed in the next sections.

We focus on PEGs in bounded, spatially complex, environments similar to today’s office environments. In such environments, we can assume imperfect geometric regularity (e.g., the presence of corridors and 90-degree turns, but possibly a regular placement of doorways or elevator exits). Because such environments are obstructed, they present limited line-of-sight visibility. However, it is increasingly true that such environments are well provisioned with wireless communication capability, and that many such environments will likely have dense embedded sensing (for surveillance or environmental control).

In this paper, we assume such *network-assisted* environments; these environments provide sensing-at-a-distance to circumvent line-of-sight limitations. Moreover, they provide a network communication capability that enables much tighter coordination than would have been possible otherwise. More specifically, the network a) contains sensors that are able to *approximately* localize all participants, and b) provides a communication backbone that enables participants to exchange game state.

Based on these assumptions, we consider the class of PEGs in which all participants have complete (but possibly imprecise) knowledge of the positions of all participants. Furthermore, our PEGs are played on a discrete space: we model the environment as a topological map, whose nodes correspond to coarse-grained regions and whose links connect neighboring regions [20, 22]. Our discrete-space assumption is acceptable for many uses of pursuit-evasion (e.g., surveillance, finding survivors) and we argue that more precise capture can be implemented with the addition of a simple proximity sensor (e.g., a low resolution camera) if necessary. We are interested in the class of games where enough pursuers (we make this more precise in the next section) exist to guarantee termination. Initially, we also assume that pursuers and evaders move at the same speed (more

precisely, we assume that they move exactly one hop in the topology at each time step). We later relax this assumption and consider even the case where evader can move faster than pursuers.

Within this framework, we are interested in the optimal strategy that pursuers and evaders should play, where our measure of optimality is the capture time (defined below). Before we discuss this, we lay down some terminology.

Let $G = (V, L)$ be a finite connected undirected graph with V vertices and L links or edges. There are two sets of players called pursuers P and evaders E . Initially, P and E occupy some vertices of G . In describing the algorithm, we assume that time is discrete and increments at steps of 1; in our implementation, of course, we make no such assumption. At each time step, all pursuers and evaders are given the positions of all participants. Both teams play a game on G according to the following rule. At each step, each pursuer chooses a neighboring vertex of G to move to, the evaders do the same. They then move to the corresponding vertex in G , as defined in [28], and repeat the previous step. The team of pursuers P wins if it “captures” all evaders. If an evader can avoid capture indefinitely, then the evader team wins the game. In the literature [7], the necessary number of pursuers to capture an evader in a graph G is denoted by $(c(G))$.

Let $p=|P|$, $r=|E|$, $v=|V|$. Let P_i be the current position of the i^{th} pursuer and E_i be the current position of the i^{th} evader. The tuple $a = \langle P_0, \dots, P_p, E_0, \dots, E_e \rangle$ represents the current position of all participants. We define a boolean variable $T(turn)$ to denote if it is the pursuers’ turn to move or not (recall that, in our algorithm, pursuers and evaders alternate at each time step¹). We say that the tuple $\langle a, T \rangle$ encodes the *state* s of a game.

In general, a game can have $2 * v^{p+r}$ states, because each pursuer and evader can be at one of v positions, and for each configuration of pursuers and evaders, there are 2 turns (evader and pursuer moves). Each game can be represented by a sequence of transitions through this state space. Each pursuer and evader executes a deterministic algorithm (called its *policy*) for determining, given the current state, what the next move should be. We call ρ the pursuer policy and ϵ the evader policy. Since we consider deterministic policies, if in a particular game, a state is repeated, the game will not terminate and the evaders win.

We can now define our game as follows:

Input Coarse estimated positions of p robots and r

¹ This assumption enables us to analyze our algorithm. Of course, in real world experiments, it is difficult to ensure this synchrony.

evaders in a bounded environment E .

Output Motion commands for p robots.

Goal Minimize the capture time of the evaders.

Restriction No motion model for the evaders available to pursuers.

The game terminates when a *capture state* is reached. In a capture state, at least one pursuer occupies the same vertex in which an evader resides. There exists a different definition of termination: if, during the evolution of the game, a pursuer reaches an evader’s position, the evader exits the game. It is easy to see that our definition results in a game that is strictly harder than this variant. When evaders exit the game, the remaining pursuers can always, and more quickly, capture the remaining set of evaders. Indeed, there exist cases where our game might not terminate (a 2-pursuer 2-evader game, in some topologies) but this variant will.

3 Optimal Strategy

In defining the game, we have avoided mention of the particular strategy that the pursuers and evaders use. In this section, we discuss the optimal strategy.

3.1 Equipotent players

In this subsection, we assume that pursuers and evaders move at the same speed (more precisely, we assume that they move exactly one hop in the topology at each time step).

In our game, we have assumed that both pursuers and evaders have complete information about the positions of all players. Consider now the *optimal strategy* for the pursuers and evaders: for the former, to capture the evaders in the shortest time, and for the latter, to avoid capture for the longest possible time. To formalize this intuition, we turn to zero-sum games.

Pursuit-Evasion is a zero-sum game since the pursuers’ gain or loss is exactly balanced by the losses or gains of the evader. The evader’s goal is to escape as long as possible whereas the pursuers have to capture the evaders as fast as possible. Zero-sum games have been extensively studied in the game theory literature, and our solution models a PEG as a zero-sum game that uses the minimax algorithm [35]. This algorithm minimizes the maximum possible loss for each player in the game.

To describe this algorithm, consider first that the evolution of any PEG can be represented by a game graph, a directed graph with possible cycles. The start state (as defined by the starting configuration of the pursuers and evaders) has a directed edge from itself to

all possible next states that the pursuers can make from the start state. (In our game, we assume that pursuers and evaders move alternatively). In turn, from each of these states, there is a directed edge to all possible next states resulting from evaders’ moves from that state. The graph can thus be recursively defined. In general, a game is a traversal on this graph. If this traversal ends in a capture state, the pursuers win the game. However, it is also possible for the traversal to repeat states: such a traversal will result in a non-terminating game and the evaders win.

Suppose now, that in the game graph, we assign to each state S , a cost function C .

- When an evader moves, $C(s)$ denotes the maximum distance from state s to a capture state, and
- When a pursuer moves, $C(s)$ represents the minimum distance from state s to a capture state.

In our game, we consider the following policies:

- The pursuers’ policy ρ is: choose that neighboring state in the game graph which has the smallest $C(s)$. Intuitively, this moves the game at each step as close as possible to a capture state.
- The evaders’ policy ϵ is: choose that neighboring state in the game graph which has the largest $C(s)$. Intuitively, this moves the game at each step as far away as possible from a capture state. Thus, the evader is truly adversarial.

In what follows, we first show how to compute the game graph efficiently. Then, we prove that these policies are optimal from the pursuer’s perspective: if the game terminates, they reach a capture state in the shortest possible number of moves.

We construct the game graph by generating all states and all possible transitions between states. For each state, we can calculate the cost to reach a capture state using a bottom-up approach. Using this, pursuers and evaders follow the strategy discussed above.

We construct the game graph by generating all states and all possible transitions between states. Initially, all states have cost ∞ , except the initial set of capture states which have cost 0 (Algorithm 1).

Let us define F_0 as a set of capture states, in other words, the states where at least one pursuer occupies each vertex in which an evader resides. These states have cost value 0 since no move is necessary for the termination of the game. Now define F_1 as the set of states which can reach one of the capture states in one pursuer move. F_1 consists of both states where it is the pursuer’s turn to move, and states where it is the evader’s turn to move. A state where it is the evader’s turn to move belongs to F_1 where irrespective of the next move made by the evader, the pursuer can still

reach the capture state in a single step. Similarly, we can define inductively the F_{i+1} set of states as the states from which the pursuers only needs $i + 1$ steps to terminate the game, irrespective of evader's movements. For any state $s = (a, 1)$, when it is the pursuer's turn to move, the cost $C(s) = \text{minimum}(C(s')) + 1$ where s' is any state that the pursuer can transition to from s . For any state $s = (a, 0)$, when it is the evader's turn to move, the cost $C(s) = \text{maximum}(C(s'))$, where s' is any state that the evader can transition to from s .

For a pursuer, the goal is to reach a capture state, i.e. minimize C . For an evader, if $C(s)$ is finite, where s is the current state, then irrespective of what the evader does, it will be captured within $C(s)$ pursuer's move. The evader should then choose a transition a to a' which maximizes the time of capture. If any state has a cost ∞ , the number of pursuers is not sufficient to capture the evaders, and it is necessary to add more pursuers to the game.

Algorithm 1 computes the state transition diagram, and the associated costs. Pursuers and evaders can use this to implement the optimal strategy ρ and ϵ : during a pursuer turn, ρ implies selecting the neighboring state with the smallest $C(s)$ and during an evader turn, ϵ implies choosing that neighboring state that has the largest $C(s)$. The algorithm loops while we add more states to F_i . Eventually, when there are no more states to be added, the algorithm terminates. Since each state can only be added once, the algorithm terminates.

To prove the optimality of our formulation, it suffices to prove the following theorem:

Theorem 1 *For any given topological graph G , if $l(s)$ is the minimal number of steps required to reach a capture state for any state s , then $C(s) = l(s)$.*

Proof We will prove this by induction on i , the length of the optimal number of steps till capture.

Basis: If $i = 0$, the game is over ($l(s) = 0$) without requiring a state transition. This means that there must be at least one pursuer in each node occupied by an evader (by our definition of capture). This is exactly our definition of a capture state (lines 5-12 Algorithm 1) for which $C(s) = 0$.

Inductive Step: Assume that, for all states s for which $l(s) = i$, $C(s) = i$. We now try to show that for any state s' for which $l(s') = i + 1$, $C(s') = i + 1$.

Proving by contradiction, assume exists a state s for which $l(s) = i + 1$ but $C(s) \neq i + 1$. We divide into two cases: (1) $C(s) < i + 1$ or (2) $C(s) > i + 1$.

Case 1: This case results into a contradiction because $l(s)$ is optimal. $C(s)$, which is the number of steps our game would take, cannot be shorter than this.

Algorithm 1 Algorithm for computing the game graph.

```

1: {Initialization}
2: Generate all states
3: Generate all possible transitions
4:
5: for all state  $s$  do
6:   if  $s$  is a capture state then
7:     add  $s$  to  $F_0$ 
8:      $C(s) \leftarrow 0$  {cost function}
9:   else
10:     $C(s) \leftarrow \infty$ 
11:   end if
12: end for
13:  $i \leftarrow 0$ 
14: repeat
15:    $i \leftarrow i + 1$ 
16:    $change \leftarrow false$ 
17:    $U \leftarrow$  set of all unmarked states that have a transition to a marked state.
18:   for all  $s$  in  $U$  do
19:     if  $s$  is a pursuer move then
20:       if  $s$  has at least one transition to a marked state then
21:         add  $s$  to  $F_i$  {mark  $s$ }
22:          $C(s) \leftarrow \min$  (over all marked neighbors) + 1
           {count this move}
23:         add transition to  $\rho$ 
24:          $change \leftarrow true$ 
25:       end if
26:     else
27:       {evader move}
28:       if all transition from  $s$  reach a marked state then
29:         add  $s$  to  $F_i$  {mark  $s$ }
30:          $C(s) \leftarrow \max$  (over all marked neighbors)
31:         add transition to  $\epsilon$ 
32:          $change \leftarrow true$ 
33:       end if
34:     end if
35:   end for
36: until not  $change$ 

```

Case 2: In state s , there are two possibilities: either it is the pursuers' turn to move, or the evaders'.

Case 2.1: Suppose it is the pursuers' turn. Since we have $l(s) = i + 1$, there is a transition from state s to some state s' , where $l(s') = i$. We also have $C(s') = i$ by the induction hypothesis. From line 22 in Algorithm 1, $C(s)$ is the minimum of all $C(s'') + 1$, where s'' is a neighboring state of s . Remember, s' is a neighboring state of s . Thus, $C(s) \leq i + 1$, but we have already discussed that $C(s)$ cannot be less than $i + 1$. So, $C(s) = i + 1$, a contradiction.

Case 2.2: Suppose it is the evaders' turn. For every possible move from s to some s' , s' is already a capture state (by construction in line 28 of Algorithm 1). $\forall s' C(s') \leq i + 1$ since the algorithm only considers marked states in line 28, and, at that stage of the algorithm, $C(s') \leq i + 1$ for all marked states. From

line 30, $C(s)$ will be the maximum of $C(s')$. Thus, $C(s) \leq i + 1$, again a contradiction.

Our solution is complete because for all state s we calculate $C(s)$, which gives the optimal length for all possible state s .

Theorem 1 proves algorithm 1 is optimal because for any possible game configuration, it is possible to reach a capture state (game is over) in the minimum number of steps. Note that this proof of optimality holds for any graph G : intuitively, the proof works because it abstracts any graph G by the quantity $c(G)$, which is the minimal number of pursuers required to ensure eventual capture on G .

For a graph with v vertices, the space complexity is $O(v^{p+r})$ since we need to store all states. For each state, we might need to visit all other state, thus, the time complexity is $O(v^{2*(p+r)})$.

In practice, to play the game, we first pre-compute a complete state transition diagram offline. This is pre-loaded on all the robots, and each pursuer or evader makes a decentralized local state transition decision, given the current state (the positions of all the robots), to calculate next state.

3.2 Superior Evader

We now consider the case when every evader is a *superior* evader, i.e., is faster than any pursuer [40]. Specifically, while pursuers are limited (in each time step) to moving exactly one hop on the topological graph, a superior evader can move to any node from its current position, except those nodes to whom paths are blocked by a pursuer.

To solve this new problem, it turns out that, for this variant as well, it is possible to compute the same game graph model as in previous section. The total number of states in the game is the same as before, and only one class of transitions is different. For each of the game state where is the evaders' turn to move, there are now more potential neighbors since the evader can move more than one step in a single turn. This will increase the number of transitions of the game.

To compute the game graph, we only need to add a pre-step, which is presented in Algorithm 2. For all the game states, it calculates all possible nodes an evader can reach without visiting a node where there is a pursuer. The implementation to detect if exists such a path consists of deleting the nodes where there is a pursuer from the topological graph and using Depth-first search or Breadth-first search to traverse the graph. Note that Line 3 of Algorithm 2 creates a new graph S which removes, from G , nodes where the pursuer is located (be-

Algorithm 2 Algorithm for computing the game graph with Superior Evaders.

```

1: for all state  $s$  do
2:   if  $s$  is an evader move then
3:     Create graph  $S$  by copying topological graph  $G$  and
       removing nodes where pursuers are located in state  $s$ .
4:     for all evader in  $E$  do
5:       Depth-first Search( $S, E_i$ )
6:     end for
7:     add transitions from  $s$  to all game states where evaders
       were able to visit(reach).
8:   end if
9: end for

```

cause the evader should avoid these nodes). This makes it easier, in line 5, to calculate all possible nodes an evader can reach without visiting a node where there is a pursuer. We then create a transition on the game graph for all possible visited nodes in the topological graph.

The overall algorithm (detecting if the game ends, what is the best policy) does not change: the capture states are defined as before, and the optimal pursuer and evader strategies are the same as before. More precisely, to calculate the policy for pursuers to capture superior evader, we still use the Algorithm 1 but with the game graph computed in Algorithm 2.

Thus, Theorem 1 still holds because given any state s , if the optimal path l takes s , where $l(s) = i$ then $C(s) = i \forall s, \forall i$.

In the superior evader case, there are $O(v^{p+r})$ states. The Depth-first search complexity is $O(|V| + |L|)$ for each evader. In the worse case, we might need to add v^r transitions for each state. Thus, the overall complexity of this pre-step is $O((v^{p+r}) * (r * v + v^r))$. Assuming $r < v$, the overall complexity for the entire game is still $O(v^{2*(p+r)})$.

Unfortunately, regardless of whether the evaders are equipotent or superior to pursuers, the computational cost of enumerating the state transition diagram is not practical; for instance, using a modern desktop, we have been unable to compute the robot's strategy for 9 robots.

This motivates the work presented in the next section, where we present a scalable algorithm that partition a PEG into multiple multi-pursuer, single-evader games. We have evaluated our scalable algorithm for equipotent participants, but expect our results to be similar for the superior evader case.

4 Partition Strategy

In our partitioning strategy, the pursuers divide the evaders amongst themselves and play $c(G)-1$ sub-games

(recall that $c(G)$ is the minimum number of pursuers required to guarantee termination on a graph G). In each of these sub-games, the pursuers and the evader each play the optimal strategy discussed in Section 3, and the goal is still to minimize the time to capture the evader. The key insight is that computing the state transition diagram for these partitioned games is computationally feasible. This section discusses an assignment algorithm that allocates $c(G)$ pursuers to each evader.

Our assignment algorithm assumes $p \geq c(G) * r$, i.e., that the number of pursuers is at least that required to ensure that $c(G)$ pursuers can be assigned to each evader. The inputs to the assignment algorithm include the graph G and the initial positions of all pursuers and evaders. The assignment algorithm outputs for each pursuer which evader to pursue (and eventually capture).

We model the assignment problem of r teams and r evaders as a matching problem. Consider the bipartite graph $G = (T, E, L)$. The set T contains nodes, each of which represents a *team* of pursuers: each team represents a distinct combination of pursuer robots. Each node in the set E represents a single evader. Finally, each edge $l = (u, v)$ from the set L represents the assignment of team u to evader v . Each edge l has a cost c_{uv} which is the time to capture evader v with team u . It is possible to compute c_{uv} , the expected time to capture evader v by team u , by evaluating a $c(G) - 1$ game. Realize that we can pre-compute this cost and we only need to verify one game to know the cost of every position. We represent the cost c_{uv} in a matrix C .

The goal is to find the assignment that minimizes the maximum time to capture all evaders.

An assignment can be represented as a matrix $X = [x_{ij}]$ where x_{ij} is equal to 1 if team i is assigned to evader j , and equal to 0 otherwise. The assignment problem is written formally as follows:

$$\begin{aligned} \min \quad & \max_{i,j=1,\dots,N} c_{ij} x_{ij} \\ \text{subject to} \quad & \sum_{j=1}^n x_{ij} = 1 \quad \forall i \in N \\ & \sum_{i=1}^n x_{ij} = 1 \quad \forall j \in N \\ & x_{ij} \in \{0, 1\} \quad \forall (i, j) \in N. \end{aligned}$$

This problem is called *linear bottleneck assignment* (LBAP) [19] and can be solved optimally by any of the polynomial-time algorithms based on network flow theory [9].

The assignment algorithm described above is summarized in Algorithm 3.

Algorithm 3 Assignment of pursuers to evaders

- 1: Compute a cost metric for a $c(G) - 1$ game.
 - 2: Generate all possible team configurations, with r teams which has $c(G)$ pursuers and one pursuer does not belong to more than one team.
 - 3: **for all** feasible team configuration **do**
 - 4: calculate matching cost as a LBAP.
 - 5: update assignment with minimum max game cost.
 - 6: **end for**
 - 7: **return** min max assignment.
-

Algorithm 3 assigns $c(g) * r$ pursuers to r evaders. If $p > c(g) * r$, the unassigned pursuers chose randomly which evader to capture. This may modify the minimum capture time of all evaders, and adds robustness to algorithm 3 in case of robot failure.

Algorithm 3 solves the global optimization problem described above even with the restriction that a pursuer can participate in at most one team. Would a simple greedy algorithm have worked? A greedy assignment strategy can result in a non-optimal terminating game with extremely large, but finite, capture time. Consider a $2 - 2$ game where the $c(G) = 1$ and the cost matrix $C = \begin{bmatrix} 1 & 2 \\ 3 & h \end{bmatrix}$. The optimal assignment that minimizes the time to capture of all evaders for this matrix is pursuer 1 to evader 2 and pursuer 2 to evader 1, which gives max time to capture $c = 3$. The greedy assignment would instead assign pursuer 1 to evader 1 and pursuer 2 to evader 2, with cost $c = h$, where h is a large finite number. Figure 1 illustrates such a game. In the greedy strategy, evader 2 could “escape” into the part of the graph denoted by dots: pursuer 2 will eventually be able to catch evader 2 (because $c(G)$ is 1), but after many steps.

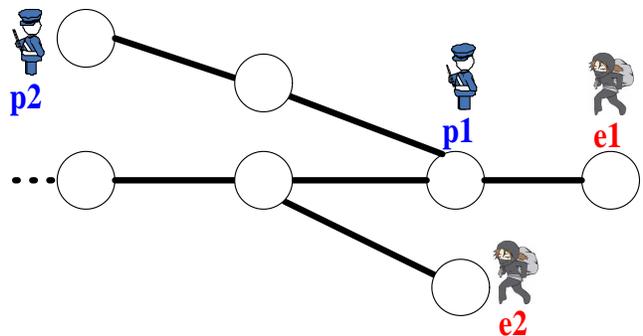


Fig. 1 An instance of a game which will terminate with huge capture time if pursuers play a greedy strategy.

Each pursuer robot runs the assignment algorithm locally. The inputs to the algorithm are the current positions of all pursuers and evaders; this information is obtained from the network (and may, of course, be inex-

act because of sensing noise). Each pursuer executes the assignment only once, and sticks with the assignment until the game terminates.

4.1 Complexity

The number of evaluated team configurations x can be modelled as the number of ways to put p distinct balls in r identical boxes. Each distinct ball represents a pursuer and each identical box represents an evader. The boxes are identical because we do not need to determine the identity of the evaders, the matching will determine this. For instance, if $c(G) = 3$, $r = 2$, $p = 6$, configuration $\langle p_1 p_2 p_3 \rangle \langle p_4 p_5 p_6 \rangle$ is equivalent to $\langle p_4 p_5 p_6 \rangle \langle p_1 p_2 p_3 \rangle$ in our matching input. Moreover, each box should have $c(G)$ balls.

The number of ways to put p distinct balls in r identical boxes is to first imagine the boxes as distinct and then determine what we over-counted. The number of distributions of a set of p distinct balls into a set of r distinct boxes if each box is to hold a specified number of balls is $\binom{p}{p_1, p_2, \dots, p_r} = \frac{p!}{\prod_{i=1}^r p_i!}$ where $p_1 + p_2 + \dots + p_r = p$. In our case, since all $p_i = c(G)$, we have $\frac{p!}{(c(G)!)^r}$. Due the fact that the boxes are identical, we over-counted $r!$. Hence, the number of configuration x we have is $x = \frac{p!}{(c(G)!)^{r*(r!)}$.

Given that $n! \simeq n^n$ (Stirling's approximation), we have $x \simeq \frac{p^p}{(c(G))^{c(G)*r*(r!)}}$. If $p = c(G) * r$, we have the number of evaluated configuration $x \simeq r^{p-r}$. The complexity of the Linear Bottleneck Assignment Problem (LBAP) with N nodes is $O(N^2)$ [32]. We evaluate x configurations, which has r nodes. Thus, the complexity of this part of algorithm 3 is $O(r^2 x)$. We also need to evaluate a $c(G) - 1$ game to know the cost metric, which has complexity $O(|V|^{2(c(G)+1)})$ (from Section 3). Thus, our overall complexity is $O(r^2 x + |V|^{2(c(G)+1)})$, where $x = \frac{p!}{(c(G)!)^{r*(r!)}} \simeq r^{p-r}$ instead of previous $O(|V|^{2(p+r)})$. Since $r \ll |V|$, this algorithm is significantly more computationally efficient.

4.2 Properties

Here we enumerate the properties of the partition strategy.

Termination: Lemma 1 guarantees game will terminate, assuming no robot fails.

Lemma 1 *Algorithm 3 guarantees that the $p - r$ game will terminate.*

Proof Since we decompose the $p - r$ game into r parallel $c(G) - 1$ sub-games, and each of these games is guaran-

teed to terminate by Theorem 1, the overall $p - r$ game terminates.

Optimality of partitioning: Algorithm 3 optimally partitions the evaders across the pursuers, subject to the cost metric. More precisely, there is no better way in partition the game into $c(G)-1$ games. This property follows from the optimality of the LBAP assignment algorithm.

Bounded capture time: The completion time for the $p - r$ game is the maximum completion time across the $c(G) - 1$ sub-game. However, is it still an open question how far off from the optimal completion time this partition strategy is.

Scalability: Algorithm 3 scales better than the optimal strategy, since its scaling is dominated by r^{p-r} , while the optimal scales as $|V|^{p+r}$, and usually $r \ll |V|$.

Robustness: If $p > c(G) * r$, algorithm 3 can assign extra pursuers to evaders to ensure robustness to robot failures.

To summarize, algorithm 3 works as follows. In a decentralized manner, each robot pre-computes a $c(G) - 1$ state transition diagram. Then, after being informed by the network of the positions of all robots, each robot runs the assignment algorithm described above once. Thereafter, using the pre-computed state transition diagram and network localization updates, all robots continuously play the game until all evaders are captured.

5 Design

In this section, we discuss the hardware and network testbed which form the basis for our pursuit-evasion experiments. We then describe, in some detail, our PEG software design and implementation. This sets the stage for our system evaluation, which is discussed in the next section.

5.1 Platform

The Robot Platform. We use a commoditized robotics platform and made minimal modifications to it using commercial off-the-shelf products. Our platform consists of an iRobot Create and a small embedded computer mounted on top of it (Figure 2).

The Create, a differential drive robot, has a round chassis of 33 cm diameter. The robot essentially has two kinds of sensors. First, a pair of tactile sensors that, together with a bumper, can help determine whether a robot hits an obstacle and the angle at which it does so. Second, a suite of infrared (IR) sensors: the bumper

contains an IR wall sensor on the right and an omnidirectional IR receiver in the top, and four additional IR sensors mounted underneath the bumper facing down. *We do not add additional sensing hardware to the Create.*

The embedded computer, the Ebox 3854, is an 800 MHz embedded PC with 256MB shared DDR memory, and supports a 1280x1024 VGA interface, one 10/100 LAN, and USB, mini PCI and compact flash sockets. The Ebox contains a 4GB compact flash (CF) card. Finally, we chose the EMP-8602 mini-PCI 802.11 a/b/g wireless card to provide network connectivity. This choice of platform is primarily designed to enable ease of programming (the Ebox runs standard Linux), and has enough ports to support various connectivity options. The embedded computer is powered by the Create’s battery through a DC-DC power adapter (picoPSU).

The embedded computer is powered by the Create’s battery. Since the two devices have incompatible voltage, we use a commercial DC-DC power adapter called the picoPSU. To use the adapter we had to fabricate custom connectors to the adapter, and to reconfigure the pins on the adapter itself, the details of which are beyond the scope of this paper.

The embedded computer runs Linux Fedora Core 6 as the operating system. For sensing and control, we developed a Create driver for Player [13], using which we are able to move the robot, turn on/off LEDs, read the bumpers, buttons and IR sensors. We set the nominal speed to 0.2 m/s.



Fig. 2 Team of Creates which depicts our robot platform - the iRobot Create with an Ebox

The Network. The robots use the network shown in Figure 3. This network is deployed above the false ceiling on one floor of a large office building and consists of two tiers: an upper tier containing 6 Stargates (embedded computers running Linux), and a lower-tier containing 56 TelosB motes (tiny commercial sensor nodes). The sensor nodes run an 8MHz Texas Instru-

ments MSP430 microcontroller, have 10KB RAM and a 2.4 GHz IEEE 802.15.4 Chipcon Wireless Transceiver with a nominal bit rate of 250 Kbps.

The network provides the robots two capabilities. The upper-tier nodes form a wireless communication backbone that can be used for inter-robot communication and coordination. In the absence of this backbone, communication between the robots can be significantly delayed as a result of transient network outages caused by robot mobility. The presence of a static network infrastructure ensures that robots can quickly communicate with each other, and thereby coordinate more effectively. The lower-tier nodes provide the capability of a virtual position sensor, and can sense the position of pursuers and evaders. We describe the details of this capability below.

One might, of course, envision other ways in which these two capabilities might have been provided. For the purposes of this paper, the exact form of the network is immaterial: we make a plausible argument that many built environments in the future are likely to have wireless communication support and a programmable sensor network that could be used for approximate positioning. We now describe in turn the various software components in our system.

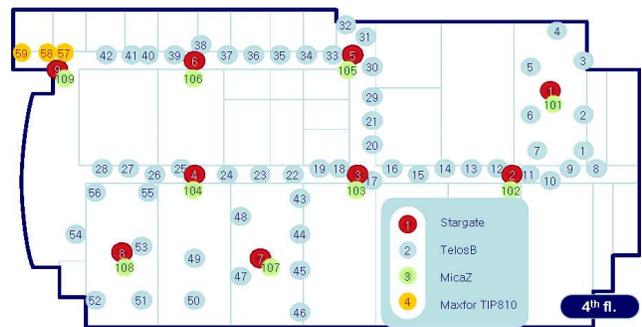


Fig. 3 Layout of the two-tier network testbed which is composed of Stargates (upper-tier) and wireless sensor motes (lower-tier)

5.2 Network-Assisted Localization

Our robots do not have extra hardware such as sonar or lasers to determine robot poses. In our system, the network estimates robot position. As we have discussed above, we use a topological map as a representation of the environment. The goal of our network-assisted localization subsystem is to approximately place each robot at a node in this topological map.

This subsystem uses the signal strength of periodic “beacon” messages emitted roughly once every second

(we add some randomization to the interval to reduce collisions) by the robots themselves.

The second-tier nodes receive these beacons, and report all beacons whose signal strength is above a certain threshold. Tenet [14], a readily available open-source software package for programming wireless sensor networks, is the software that collects the beacon signal strength. A centralized robot location server then applies a voting scheme on a sliding window of reports (the width of the window is 7) to generate a location estimate. Since radio propagation characteristics can vary significantly even with small displacements, we use a simple heuristic filter to smooth the estimate. This filter essentially updates a robot’s current position only after n consecutive reports are received. Clearly, this trades off increased latency for a smoother evolution of the estimate. Figure 4 motivates the need for smoothing the estimate, and shows the performance of the filter on our testbed.

Even so, our localization algorithm can have errors. To quantify the performance of our network-assisted localization scheme, we conducted an experiment where we measured the location error for each of 1500 beacon reports (Figure 5). Our technique is correct at least 95% of the time. All errors are at 1-hop errors: that is, the actual and estimated locations were adjacent nodes on the topological map. These errors occur when the robot is close to the boundary between these two adjacent nodes. We argue that the impact of this localization error on the game is negligible: in practice, a proximity sensor could have detected robots in adjacent nodes.

Our architecture also allows for a decentralized localization system. Instead of retrieving location information from the server, this information can be flooded on the upper-tier network, of which the robots are a part. This flooding cost is minimal, since we only need to transmit the position of the players, which requires $(p + r)(\lceil \log_2 v \rceil)$ bits. We have experimented with this version as well, and the performance of the system is comparable to the centralized system, as shown in Section 7.

One can imagine a variety of localization methods being used by the robots. Map-based localization using either a sonar, laser, camera (or a combination of these) could be used (IR on the Create is too short range to be useful for anything other than proximity sensing - we exploit it in the wall-following behavior described below). However, as discussed above, we choose to work with simple robots and do not augment them with sophisticated sensing. Given that we operate in an indoor bounded environment, one could perhaps conceive of doing localization by filtering odometry and incorporating corrections based on proximity to certain landmarks

(e.g., detectable by the IR sensor and/or bumpers, or better still the Create’s recharging station). The Create’s odometry is poor. Thus we do not currently follow this path, without completely ruling it out in the future.

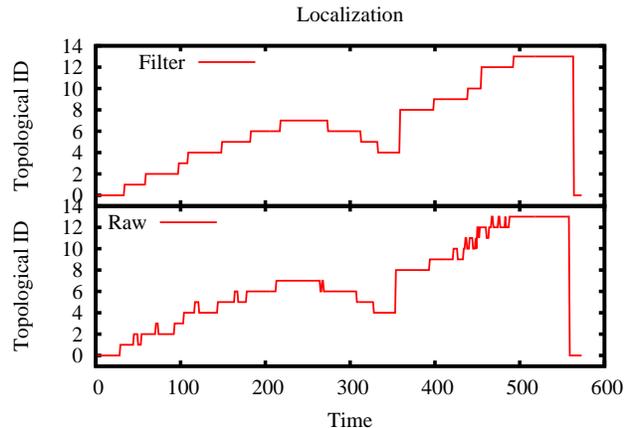


Fig. 4 Network-assisted localization estimates: unfiltered estimates (bottom) are noisy, filtered estimates (top) eliminate noise.

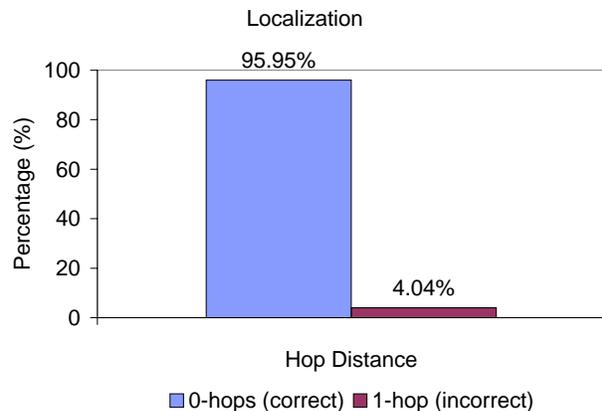


Fig. 5 Localization Accuracy

5.3 Wall-following

Given the paucity of sensors on the Create (and our coarse localization technique), we use a simple **wall-following** behavior using the IR sensors to traverse the environment. Clearly, **wall-following** may not generalize to other kinds of environments. It suffices for our purposes as a simple, even primitive, traversal behavior; more sophisticated behaviors can only improve the performance of our system. That said, our **wall-following** component includes several optimizations to ensure ro-

bustness: in our environment, at least, it always results in forward progress.

The robot continuously emits IR signals and reads the IR receiver continuously. As long as the IR values are between two thresholds and none of the bumpers are activated, the robot will move forward parallel to the wall using simple PD control.

If the IR reading falls below a threshold, the robot assumes it has “lost” the wall and attempts to **search** for it. It does this by executing a spiral (simultaneous rotation and translation) until the bumper senses contact (assumed to be with the wall). To avoid open office or elevator doors in the experiments, we use a “virtual wall” (essentially, an IR transmitter) available from iRobot. When the virtual wall is detected, the robot tries to avoid going through it by going back and turning to the left for a while without bumping, then going forward while turning to the right also without bumping. If the bump sensor is activated and the virtual wall is simultaneously detected, it invokes the **escape** behavior described below.

The **escape** behavior is invoked whenever the robot is stuck. In addition to the scenario discussed above, this behavior is invoked whenever the power consumption of the robot is above a predetermined threshold over a predefined time window (this might happen when the robot encounters a small obstacle that does not activate the bumper). To execute the escape behavior, the robot simply moves backwards for a short distance and then resumes **wall-following**.

Our robot’s wall following behavior uses the IR transmitter near the right side of the robot. As such, it keeps the wall to its right. To reverse direction, the robot implements a **detach** behavior which enables it to cross a corridor and find the opposite wall. The **detach** behavior is simple: the robot swivels slightly to the left and then moves in an arc until a bumper is activated. However, if the robot hits a wall in less than a preset time, it repeats this maneuver since it is unlikely to have crossed the corridor in that (short) time.

5.4 Navigation

The navigation component calculates the goal position, and invokes **wall-following** to move from one topological node to the adjacent one. Navigation is executed every time the robot changes its position (as well as when any evader changes its position), so the robot continuously updates its trajectory.

The navigation component calculates the goal position given its team’s position and that of its assigned evader. Using Algorithm 1, we pre-compute a state transition diagram for a given team configuration. This state

transition diagram is pre-loaded on all the robots, and each pursuer or evader makes a decentralized local state transition decision, given the current state. This decision tells the robot which topological node to move to next.

However, there exists an important subtlety in the navigation algorithm imposed by the minimality of our platform. Our robot has no inherent proprioception capability, it can only travel parallel to a wall, keeping the wall to its right (since the robot has only one IR wall sensor positioned on its right). As a result, the robot might actually move in a direction opposite to that intended by the navigation component. To rectify this, we add a simple *a posteriori* correction to the navigation component. If the **wall-following** moves the robot to a node that it does not expect to arrive at, it invokes the **detach** behavior described above to reverse direction.

6 Simulation

In this section, we evaluate the optimal and partition strategy by simulation. We implemented an idealized game simulator, which models time in discrete steps and implements the strategy discussed in Section 3.1. We exhaustively enumerate all configurations to compute a worst-case configuration.

6.1 Optimal Strategy

We choose four topologies for our case study (Table 1). The first topology is a two-dimensional grid graph, which is the graph Cartesian product of two path graphs [42]. The second topology is also a grid but embedded on the surface of a cylinder. Unlike the previous topology, the edge nodes of the grid are connected on the cylinder. The third topology is a grid embedded in a torus (a sphere with one handle) [26]. There are no edge nodes. Finally, the fourth topology is composed of two nodes connected by three disjoint paths (we call it a 3-flow graph). Table 1 shows the topologies graphical representations.

The Cylinder Grid is an interesting topology because a naive strategy to try to catch the closest evader by greedily following it would not work, as illustrated in Figure 6. The evader could always escape by going around the cylinder grid while pursuers try to catch it. At least one pursuer needs to travel in the opposite direction. The Toroidal Grid is an interesting topology because we need more than 2 pursuers to capture an evader. The 3-flow topology shows a topology where increasing the number of pursuers would decrease the

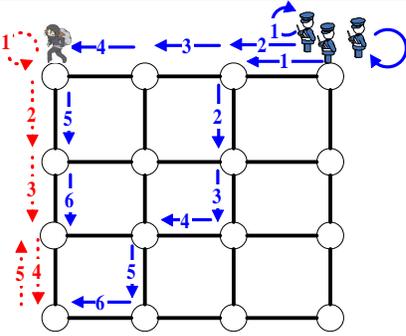
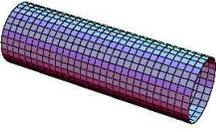
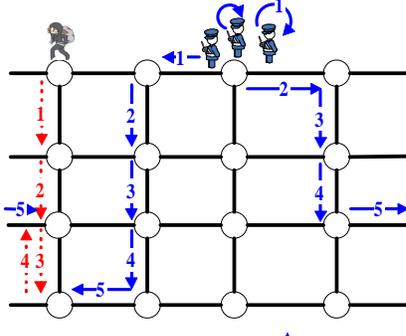
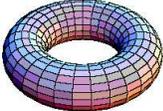
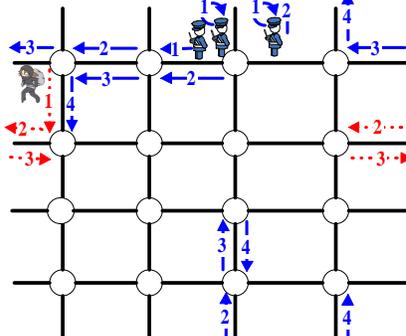
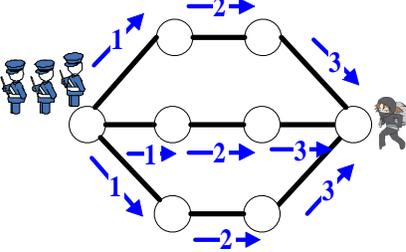
Topology Name	Graphical	Game	C(G)	2 pursuers	3 pursuers
Grid 2D 4x4			2	6	6
Cylinder Grid 4x4			2	5	5
Toroidal Grid 4x4			3	∞	4
3-flow	N/A		2	4	3

Table 1 Instances of Games and their properties

worst-case capture time. The 3-flow topology needs only 2 pursuers to guarantee the capture of an evader. But, with 3 pursuers, the worst-case capture time decreases from 4 to 3 steps.

Table 1 shows some instances of games and their properties. The third column illustrates a worst-case capture game on the given topology. The game shows the trajectory of each player. The nodes and solid undi-

rected edges connecting them represent the topological map. The solid directed lines (lines with arrows) show the pursuer path and the dashed lines illustrate the evader path. The edge labels represent the time sequence of the robots. The pursuer and evader's initial positions are indicated by the corresponding icons. The fourth column is the necessary number of pursuers to guarantee the termination of the game. The fifth and

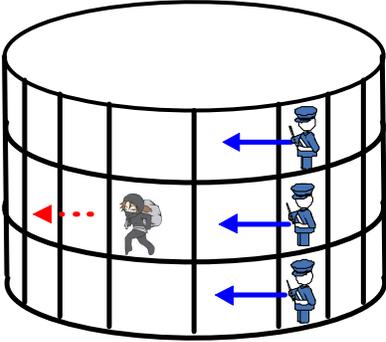


Fig. 6 A game where pursuer greedily follow the evader will not terminate in a cylinder Grid Topology.

sixth columns represent the maximum number of steps to terminate the game, given the number of pursuers. It is interesting to see that even though the torus needs at least 3 pursuers, if we play the game with 3 pursuers, the game will be shorter than in a 4x4 grid. Another interesting fact is increasing the number of edges from grid 2d 4x4 to cylinder 4x4 does not increase the necessary number of pursuers but it decreases the number of steps to terminate the game. Increasing the edges again (from cylinder to toroidal), increases the necessary number of pursuers.

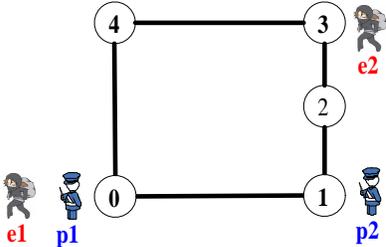


Fig. 7 Given our definition of capture, a 2-2 game will not terminate in a ring topology.

Finally, Figure 7 demonstrates a topology in which the game does not terminate for our definition of capture. In this ring topology with 5 nodes, a game with 2 pursuers and 2 evaders will not terminate since we defined the game to end only when all the evaders are captured. To illustrate this, suppose pursuer 1 stayed on the same location as evader 1. Pursuer 2 needs to capture evader 2, but evader 2 just needs to stay away from pursuer 2. If instead we had chosen to require an evader to exit the game when it is first captured, this game would have terminated.

Topology	$c(G)$	$t(c(G)-1 \text{ game})$	$t(2c(G)-2 \text{ game})$	
			Optimal	Partition
Grid 2D 3x3	2	4	4	4
Cylinder Grid 3x3	2	3	3	3
Toroidal Grid 4x4	3	4	4	4

Table 2 Games and their properties

6.2 Partition Strategy

We can also use the simulator to compare the partition strategy with the optimal (Table 2) for various topologies. The first column is the topology. The second column is the necessary number of pursuers to guarantee the termination of the game for that topology. The third column is the maximum number of steps to terminate the game (t) with 1 evader, which is the same in both strategies. The fourth and fifth columns give the maximum number of steps to terminate the game using optimal and partition strategy for two evaders. For these topologies, the partition strategy has the same completion time as optimal strategy.

The drawback of our partition strategy is we might use more pursuers than the minimum necessary. In a ring topology, 3 pursuers are sufficient to capture 2 evaders. In the partition strategy, we need 4 pursuers. We believe that this tradeoff is acceptable, since the partition strategy enables efficient capture.

It is not enough, however, to use the simulator alone. The simulator plays an idealized game and differs from the real environment in three ways. First, in the real testbed, errors in robot localization can shorten or prolong a game. The former can happen when two robots think they are at the same topological node, but are in reality far away from each other. The latter can happen when the converse is true, or when a pursuer takes a longer path because it is incorrectly localized.

Second, while our robots only follow walls in one direction (and must cross the corridor to reverse direction) and therefore must correct their orientation *afterwards* (Section 5), our simulator does not mimic this constraint. Effectively, the above differences help us understand how the games would have been different if our robots had expensive and precise positioning capability, and onboard orientation sensing capability.

The third difference arises from the synchronous execution in the simulator. In the real world, pursuer and evader decisions are not synchronized. Consider a pursuer and evader separated by one hop. In the simulator, because decisions are synchronized and localization is perfect, the evader can escape the pursuer. In the real world, however, it is quite possible for the pursuer to start moving towards the evader and (because the localization system is based on coarse-regions) “capture” the evader before it has had a chance to move outside

the region. In these infrequent cases, the game can terminate faster than it would in the idealized scenario.

For this reason, we played a few (specifically, 2 – 1, 4 – 2, and 6 – 3) games on our real world testbed, and we discuss the results below.

7 Experiments

In this section, we describe the results from several games played on the physical robot testbed described in Section 5.

We play the games on the floor plan shown in Figure 3, using the network whose nodes are shown in that figure. We run the experiments in the entire floorplan, which has about 50 m x 30m, providing us a realistic physical large environment. While our games are played only in the one environment shown (and future work needs to validate the generality of the claims we make later in this section), we emphasize two important points about the environment that lead us to believe that our results would hold generally in other similar environments. First, our building floor is representative of many office buildings (i.e., it does not have any unusual features that would materially affect our conclusions) and we do not alter the environment in any way other than to place virtual walls in some locations². Second, our network design is not optimized in any way for this particular application. Rather, the network was designed to mimic harsh wireless communication conditions to stress test wireless systems and applications.

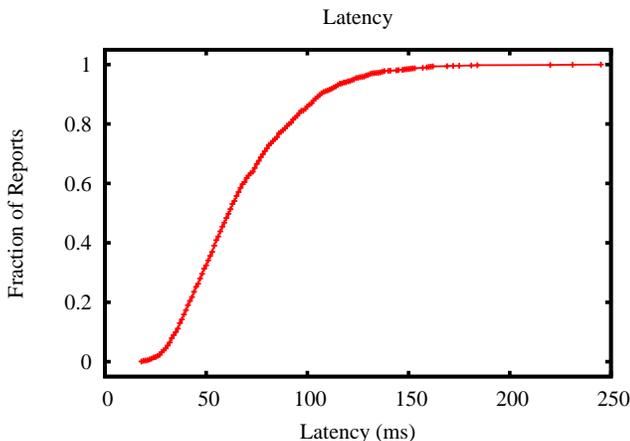


Fig. 8 The communication latency for detecting the robots.

² We place a virtual wall in front of the elevator doors to prevent the robot from entering an open elevator. We also place a virtual wall in front of the doors whose thresholds are lower than the height of the Create bumper.

Figure 8 shows the cumulative latency for beacon reports from robots. The latency measures the time from when the robot sends a beacon to when the robot location server receives it. This communication latency is at most 250 ms. We expect the latency of distributing the location estimate to the robots from the server to be comparable. Thus, our real-world implementation deviates from our simulations and theory in that robot poses may be slightly stale: however, given the robot speeds in our experiments, the location estimates will be off by about 0.1 m or so, a relatively small amount.

The convergence time of a game depends on the initial configuration. We play our games using a worst-case initial configuration (there can be many), which was calculated by the idealized game simulator.

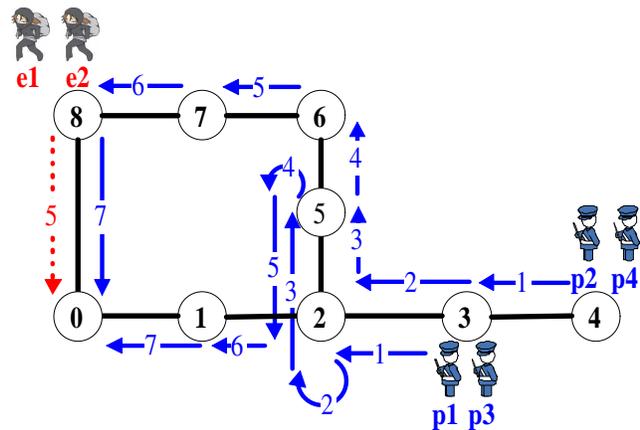


Fig. 9 An illustration of 4 pursuer, 2 evader game sequence. The nodes and solid undirected edges connecting them represent the topological map. The solid directed lines (lines with arrows) show the pursuer path as determined by the localization system. The dashed lines illustrate the evader path. The edge labels represent the time sequence of the robot. The pursuer and evader’s initial positions are indicated by the corresponding icons.

An Illustrative Game Instance. Figure 9 shows the trajectory of each pursuer during one instance of a 4 – 2 game. The game evolves as follows. Pursuers 1 to 4 start at position 3, 3, 4, and 4 respectively, and all evaders start at position 8 (this is a worst-case configuration, as determined by simulation). Pursuers execute the partition algorithm and independently decide on their teams. Pursuers 1 and 2 try to capture evader 1 and pursuers 3 and 4 will chase evader 2. Pursuers 2 and 4 try to corral the corresponding evader by approaching them through the alternate path in the topology, while pursuers 1 and 3 try to capture evaders by moving directly towards them. At time 2, pursuers 1 and 3 are at location 2. They do not have to move since either option will not affect the game convergence time. At the end of time 2, all pursuers are at near location 2. All

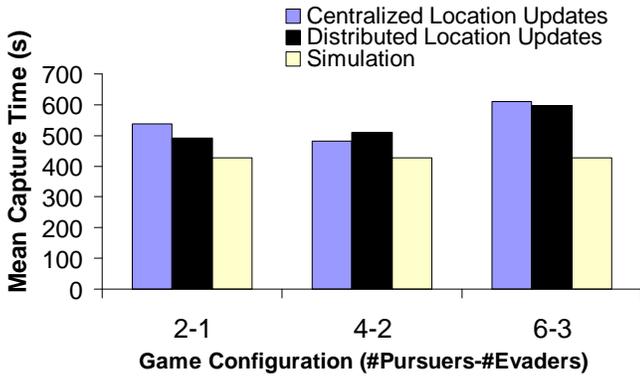


Fig. 10 The Mean Capture Time for three game configurations played with centralized and distributed location updates. The third bar is the average capture time from simulation.

pursuers keep moving and are at location 5. At time 4, pursuers 1 and 3 switch their orientation (turn around to another wall invoking *detach*). At time 5, evaders move to location 0 to be as far away from all pursuers as possible, otherwise pursuer 1 would have captured them. Pursuers are at 2, 2, 7, and 7. Finally, at steps 6 and 7, pursuers go toward the evaders and corral them since evaders do not have any option to escape.

In simulation, the capture time for the described game is 5 steps. Our results took 7 steps because our robots do not have a sense of direction and can only reverse direction by going in the wrong direction to determine that they have done so. This behavior is not captured in the simulator. Despite this (and other) non-idealities, our implementation works well in the real-world.

Results across Multiple Games. Finally, Figure 10 depicts the capture time across multiple games. The capture time is similar between the games, as one might expect (since each game is partitioned into a 2 – 1 game). However, there are small differences in capture times. For a 4 – 2 game, the game terminated even with a pursuer failure in localization, since only 2 robots are needed in our topology. Our 6 – 3 game had a longer convergence time because robots are not able to move at nominal speed continuously as a results of floor surface differences and because (as it happens in our environment) pursuers bump into each other and are slowed down during obstacle avoidance. The capture time is similar between games with centralized, distributed location updates.

The capture time is also similar between experiments and simulations (Figure 10). Although the simulations do not model some aspects of our real-world implementation (Section 6, such as not taking into account the extra-steps necessary to reverse direction),

the simulation results are in qualitative agreement with the implementation results.

A video demonstrating a 2-1 game is available as electronic supplementary material. The video `peg.mpg` was taken by attaching a video-capable cell phone to each robot, so it shows the progress of the game from the perspective of each robot. One can see the *wall-following* and *detach* behaviors at work. The video is composed of four screens. Two screens show the pursuer field of view. One screen shows the evader field of view. The evader wears a red hat to make is easily identifiable. The fourth screen shows the location of the players and explains what is happening during the game. The video playback is sped up three-fold to show the entire game. Even though the pursuers do not have the evader in their field of view, the wireless sensor network enhanced their sensing cover area, and the pursuers are able to capture the evader.

8 Related Work

Our broad interest is in the theme of using pervasive sensing and communication infrastructure with actuation. In the robotics and sensor networking communities this area has burgeoned recently. Examples include investigations on sensor-network guided robot localization [12], sensor-network guided robot navigation [21, 30, 4, 6] and exploration [24], robot-assisted sensor network localization [12], robot-assisted sensor network deployment [17, 11, 24] and sensor-network guided robot pursuit-evasion [29].

Like with body of work discussed above, we are interested specifically in the coordination and control of actuation using pervasive sensing and communication infrastructure. We distinguish our work (both in the choice of problem and the solution space we explore) from these approaches by 1. choosing a problem (PEG) which is only beginning to be addressed in the embedded network and robot context, yet has a rich history in the classical literature (using robots alone), 2. choosing an experimental testbed based on unmodified commodity hardware (the Creates from iRobot and the TelosB motes).

We note that the sensor networking community has championed (for a decade) several generations of the “mote” devices as minimal yet useful sensing and communication nodes. These have served as a baseline on which to implement and test ideas on sensor network routing, data-dissemination, localization, and many others. A recent trend in the robotics community parallels this; we cite as examples the following recent uses of the iRobot Create - the Microsoft Sumo Robot Project [1], the Roomba MADNet project [34], the SmURV project

	[38]	[3]	[37]	[16]	[18]	[39]	[29]	Our work
Pursuer to Evader Ratio	≥ 1	≥ 1	1	≥ 1	≥ 1	≥ 1	1	$\gg 1$
Pursuer Visibility	full	full	full	local	local	local	full	full
Evader Visibility	full	full	full	full	full	local	none	full
Information	full	full	full	full	full	no	pursuer full	full
Environment	Graph	Graph	Polygon	Polygon	Polygon	Polygon	Polygon	Graph
Capture	touch	touch	touch	see	touch	touch	touch	touch
Speed (faster entity)	any	same	same	same	evader	any	same	any
Time to Capture	yes	no	no	no	no	no	no	yes
Robot Implementation	None	None	None	None	None	Partial	Partial	Full

Table 3 Related Work in Pursuit-Evasion

[2], and the workbook associated with a recent text on robotics [23].

To situate our work in the existing literature, we classify the type of PEGs using seven criteria: the *ratio* of the number of pursuers to the number of evaders; whether pursuers and evaders have full and/or global *visibility*, or whether they can only see within a threshold distance or until occluded by an obstacle (usually modelled by the edge of a polygon in 2D); what additional *information* robots have with respect to the opponents’ strategy or planning algorithm; whether the *environment* is modelled as a graph (discrete) or a polygon (continuous half-space with lines in 2D as boundaries); how the evader is *captured*, whether by being surrounded, seen or sensed by the pursuer, or approached within a certain distance, or physically contacted; the relative *speed* between the pursuer and evader; and, if the *time to capture* is important.

Table 3 shows a classification of the related work in the literature along these dimensions. Our work is distinct from several pieces of prior work. Its novelty is clear in several dimensions: while other work has explored theoretical bounds on eventual capture [3, 37], or pursuit-evasion under constrained geometries [16, 25, 8, 10], or has examined sophisticated control strategies [39, 29], our proposed work attempts to minimize the time of capture of a multi-pursuer multi-evader under the pragmatic realization of physical multi-robot games.

Complementary to our work, [38] discusses a dynamic programming algorithm to maintain connectivity in a team of robotic routers in the case where there is one user moving in an adversarial trajectory. The problem is modeled as a pursuit-evasion game, with the goal of finding the shortest escape trajectory. There is no discussion on how to initially detect the evader. The algorithm has complexity $O(v^{3(p+1)})$. They presented simulation results and were seeking an improved running time.

In [7], an algorithm to determine if K pursuers are sufficient to capture an evader is presented. They also showed that every graph is topologically equivalent to a graph with pursuer number at most two. In the survey presented by Alspach [5], a number of references

on the necessary number of pursuers for a given graph class can be found. Aigner and Fromme [3] proved that in a planar graph G , 3 pursuers are sufficient for the pursuers to win the game. Quilliot [33] extended this result, giving an upper bound to the number of pursuers depending on the genus of the graph G . In [27], the necessary number of pursuers is studied under three graph product operations. In [15], given certain conditions, the complexity of pursuit on a graph is Exptime-complete.

For superior evaders, Seymour [36] showed that finding the necessary number of pursuers to capture a single evader with infinite velocity in a graph G when pursuers and evader move simultaneously is equivalent to finding the treewidth of G . Other works [41][40] focus on an Euclidean open plane and provide sub-optimal real-time approaches.

We are the first to present a scalable algorithm to minimize the time to capture of a multi-pursuer multi-evader game, and to validate the complete system in a real-world implementation.

9 Conclusions

In this paper, we have described a system that detects multiple evaders using wireless sensor network, disseminates the necessary information using a decentralized protocol, and executes an algorithm that captures multiple evaders in near-optimal capture time. We have presented a systematic derivation of this algorithm, by first presenting a provably optimal algorithm captures the evaders in the minimum capture time, when the evaders and pursuers are equipotent and play optimally. We have then described an optimal algorithm for the superior evader case. Our third practical algorithm, is an assignment algorithm that guarantees the game terminates, has bounded captured time, is robust, and is scalable in the number of robots. We have validated the feasibility of our algorithm by experimentally playing mobile robot-based pursuit evasion games on a physical testbed.

References

1. Sumo Robot (2008). URL <http://msdn2.microsoft.com/en-us/robotics/bb403184.aspx>
2. The SmURV Robotics Platform (2008)
3. Aigner, M., F.M.: A Game of Cops and Robber. Tech. rep. (1984)
4. Alankus, G., Atay, N., Lu, C., Bayazit, B.: Adaptive Embedded Roadmaps for Sensor Networks. In: IEEE International Conference on Robotics and Automation (2007)
5. Alspach, B.: Searching and Sweeping Graphs: a Brief Survey. *Le Matematiche (Catania)* **59**, 5–37 (2004)
6. Batalin, M.A., Sukhatme, G.S.: Coverage, Exploration and Deployment by a Mobile Robot and Communication Network. In: *Telecommunication Systems Journal, Special Issue on Wireless Sensor Networks*, pp. 376–391 (2003)
7. Berarducci, A., Intrigila, B.: On the Cop Number of a Graph. *Adv. Appl. Math.* **14**(4), 389–403 (1993). DOI <http://dx.doi.org/10.1006/aama.1993.1019>
8. Bhattacharya, S., Candido, S., Hutchinson, S.: Motion Strategies for Surveillance. In: *Robotics: Science and Systems* (2007)
9. Burkard, R.E., Çela, E.: Linear Assignment Problems and Extensions. *Handbook of Combinatorial Optimization* **4** (1999)
10. Cheung, W.: Constrained Pursuit-Evasion Problems in the Plane. Master Thesis, U.British Columbia (2005)
11. Corke, P.I., Hrabar, S.E., Peterson, R., Rus, D., Saripalli, S., Sukhatme, G.S.: Deployment and Connectivity Repair of a Sensor Net. In: 9th International Symposium on Experimental Robotics 2004 (2004)
12. Corker, P., Peterson, R., Rus, D.: Localization and Navigation Assisted by Networked Cooperating Sensors and Robots. *International Journal of Robotics Research* **24**(9), 771–786 (2005)
13. Gerkey, B.P., Vaughan, R.T., Støy, K., Howard, A., Sukhatme, G.S., Mataric, M.J.: Most Valuable Player: a Robot Device Server for Distributed Control. pp. 1226–1231. Maui, HI, USA (2001)
14. Gnawali, O., Greenstein, B., Jang, K.Y., Joki, A., Paek, J., Vieira, M., Estrin, D., Govindan, R., Kohler, E.: The TENET Architecture for Tiered Sensor Networks. In: *Proceedings of the ACM Conference on Embedded Networked Sensor Systems* (2006)
15. Goldstein, A.S., Reingold, E.M.: The Complexity of Pursuit on a Graph. *Theor. Comput. Sci.* **143**(1), 93–112 (1995). DOI [http://dx.doi.org/10.1016/0304-3975\(95\)80012-3](http://dx.doi.org/10.1016/0304-3975(95)80012-3)
16. Guibas, L.J., Latombe, J.C., LaValle, S.M., Lin, D., Motwani, R.: Visibility-Based Pursuit-Evasion in a Polygonal Environment. In: *WADS '97: Proceedings of the 5th International Workshop on Algorithms and Data Structures*, pp. 17–30. Springer-Verlag, London, UK (1997)
17. Howard, A., Mataric, M.J., Sukhatme, G.S.: An Incremental Self-Deployment Algorithm for Mobile Sensor Networks. *Autonomous Robots* **13**(2), 113–126 (2002)
18. Isler, V., Kannan, S., Khanna, S.: Randomized Pursuit-Evasion in a Polygonal Environment. *IEEE Transactions on Robotics* **5**(21), 864–875 (2005)
19. Jonker, R., Volgenant, A.: A Shortest Augmenting Path Algorithm for Dense and Sparse Linear Assignment Problems. *Computing* **38**(4), 325–340 (1987). DOI <http://dx.doi.org/10.1007/BF02278710>
20. Kuipers, B., Byun, Y.T.: A Robot Exploration and Mapping Strategy Based on a Semantic Hierarchy of Spatial Representations. Tech. Rep. AI90-120 (1990)
21. Li, Q., Rus, D.: Navigation Protocols in Sensor Networks. *ACM Transactions on Sensor Networks* **1**(1), 3–35 (2005)
22. Mataric, M.J.: Integration of Representation Into Goal-Driven Behavior-Based Robots. *IEEE Transactions on Robotics and Automation* (3) (1992)
23. Mataric, M.J.: *The Robotics Primer*. MIT Press (2007)
24. Maxim Batalin and Gaurav S. Sukhatme: The Design and Analysis of an Efficient Local Algorithm for Coverage and Exploration Based on Sensor Network Deployment. *IEEE Transactions on Robotics* **23**(4), 661–675 (2007)
25. Murrieta-Cid, R., Muppirla, T., Sarmiento, A., Bhattacharya, S., Hutchinson, S.: Surveillance Strategies for a Pursuer with Finite Sensor Range. *Int. J. Rob. Res.* **26**(3), 233–253 (2007). DOI <http://dx.doi.org/10.1177/0278364907077083>
26. Neufeld, E., Myrvold, W.: Practical Toroidality Testing. In: *SODA '97: Proceedings of the Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 574–580. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA (1997)
27. Neufeld, S., Nowakowski, R.: A Game of Cops and Robbers Played on Products of Graphs. *Discrete Math.* **186**(1-3) (1998). DOI [http://dx.doi.org/10.1016/S0012-365X\(97\)00165-9](http://dx.doi.org/10.1016/S0012-365X(97)00165-9)
28. Nowakowski, R.J., Winkler, P.: Vertex-to-Vertex Pursuit in a Graph. *Discrete Mathematics* **43**(2-3), 235–239 (1983)
29. Oh, S., Schenato, L., Chen, P., Sastry, S.: Tracking and Coordination of Multiple Agents using Sensor Networks: System Design, Algorithms and Experiments. *Proceedings of the IEEE* **95**(1) (2007)
30. O'Hara, K.J., Bigio, V., Dodson, E., Irani, A., Walker, D., Balch, T.: Physical Path Planning Using the Gnats. In: *IEEE International Conference on Robotics and Automation* (2005)
31. Parsons, T.D.: Pursuit-Evasion in a Graph. *Lecture Notes in Mathematics* **642**, 426–441 (1978). DOI [10.1007/BFb0070400](http://dx.doi.org/10.1007/BFb0070400)
32. Pferschy, U.: Solution Methods and Computational Investigations for the Linear Bottleneck Assignment Problem. *Computing* **59**(3), 237–258 (1997). DOI <http://dx.doi.org/10.1007/BF02684443>
33. Quilliot, A.: A Short Note About Pursuit Games Played on a Graph with a Given Genus. *J. Comb. Theory, Ser. B* **38**(1), 89–92 (1985)
34. Reich, J., Misra, V., Rubenstein, D.: Roomba MADNeT: a Mobile Ad-hoc Delay Tolerant Network Testbed. In: *MC2R: Mobile Computing and Communications Review*. ACM Sigmobile (2008)
35. Russell, S.J., Norvig, P.: *Artificial Intelligence: A Modern Approach*. Pearson Education (2003)
36. Seymour, P.D., Thomas, R.: Graph Searching and a Min-Max Theorem for Tree-Width. *J. Comb. Theory Ser. B* **58**(1), 22–33 (1993). DOI <http://dx.doi.org/10.1006/jctb.1993.1027>
37. Sgall, J.: Solution of David Gale's Lion and Man Problem. *Theor. Comput. Sci.* **259**(1-2), 663–670 (2001)
38. Tekdas, O., Isler, V.: Robotic Routers. In: *Proc. IEEE Int. Conf. on Robotics and Automation* (2008). URL <http://www.cs.rpi.edu/>. To appear
39. Vidal, R., Shakernia, O., Kim, H.J., Shim, D.H., Sastry, S.: Probabilistic Pursuit-Evasion Games: Theory, Implementation, and Experimental Evaluation. *Robotics and Automation, IEEE Transactions on* **18**(5), 662–669 (2002)
40. Wang, X., Jose B. Cruz, J., Chen, G., Pham, K., Blasch, E.: Formation Control in Multi-Player Pursuit Evasion Game with Superior Evaders. p. 657811. SPIE (2007). DOI [10.1117/12.723300](http://dx.doi.org/10.1117/12.723300). URL <http://link.aip.org/link/?PSI/6578/657811/1>

41. Wei, M., Chen, G., Cruz, J., Hayes, L., Chang, M.H.: A Decentralized Approach to Pursuer-Evader Games with Multiple Superior Evaders. Intelligent Transportation Systems Conference, 2006. ITSC '06. IEEE pp. 1586–1591 (2006). DOI 10.1109/ITSC.2006.1707450
42. Weisstein, E.W.: Grid graph. MathWorld – A Wolfram Web Resource. <http://mathworld.wolfram.com/GridGraph.html>. URL <http://mathworld.wolfram.com/GridGraph.html>