

A Graph Search Algorithm for Indoor Pursuit / Evasion

Athanasios Kehagias

Geoffrey Hollinger

Sanjiv Singh

CMU-RI-TR-08-38

July 2008

Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213

© Carnegie Mellon University

Abstract

Using concepts from both robotics and graph theory, we formulate the problem of indoor pursuit / evasion in terms of searching a graph for a mobile evader. We present an *offline, greedy, iterative* algorithm which performs *guaranteed search*, i.e. no matter how the evader moves, it will eventually be captured; in fact our algorithm can be applied to either an adversarial (actively trying to avoid capture) or randomly moving evader. Furthermore the algorithm produces an *internal* search (the searchers move only along the edges of the graph, “teleporting” is not used) and can accommodate “extended” (across nodes) visibility and finite or infinite evader speed. We present search experiments for several indoor environments, some of them quite complicated, in all of which the algorithm succeeds in *clearing the graph* (i.e. capturing the evader).

1 Introduction

In this paper we present a formulation of *pursuit / evasion* as a *graph search* problem and introduce an algorithm to solve this problem.

In broad terms *pursuit / evasion (PE)* involves a team of K *pursuers* trying to locate / capture a mobile *evader*. Several more specific variations of the problem can be defined, depending on the capabilities of the pursuers and the evader (and the information available to each of them), the environment in which PE takes place, the type of solution sought and so on. In this paper we are concerned with the following version of the PE problem.

1. The pursuit takes place in an indoor environment (e.g. a house, an office building) and the evader cannot leave the environment (i.e. no exits are available).
2. The evader is *adversarial* (i.e. wants to avoid capture), *arbitrarily fast* and has complete knowledge of the environment map and the pursuer locations.
3. The pursuers know the map of the environment and the locations of their teammates. They do *not* know the evader’s position unless it is within their *visibility region*.

What is required is a *search schedule* (i.e. a plan for the movement of each pursuer) which will be computed *offline* (i.e. before the PE process starts) and will *guarantee* the eventual capture of the evader (i.e. every possible sequence of evader moves leads to capture). To produce such a schedule we examine a *discretized* version of the problem by introducing the following assumptions.

1. The PE process evolves in discrete time steps $t = 1, 2, 3, \dots$ and at every time step movement is *sequential*, i.e. first the evader moves and then one pursuer at a time.
2. The map of the environment is decomposed into convex *cells* and each cell is represented as a *node* of a *finite graph*; the *edges* of the graph correspond to the connectivity of cells in the original environment; it is assumed that at every time step pursuers and evader reside in the nodes of the graph.¹
3. *Capture* takes place (with certainty) when the evader and at least one pursuer reside in the same node.

We present an algorithm which guarantees capture by ensuring that the “*dirty set*” (the set of nodes which *may* contain the evader) decreases progressively until it reduces to the empty set, i.e. the evader is left without any escape moves² (when this happens we say that the graph has been *cleared*). Our algorithm combines concepts and methods from two distinct research themes: *graph search* and *robotic pursuit / evasion*.

PE on a graph was first studied by Parsons [24] in connection to *cave searching*. Papadimitriou and his collaborators [16, 22] placed the problem in the context of graph theory and initiated a vigorous line of research which continues to this day; for recent reviews see [1, 8]. In this context the problem is usually described as “graph search” but it should be distinguished from both (a) searching a graph for an *immobile* target and (b) pursuing a *visible* mobile evader (“cops-and-robbers” problems [12, 23]).

The graph theoretic work mentioned in the previous paragraph is rather abstract. For example, much of this work is concerned with studying the *search number* of a graph (i.e. the minimum number

¹This is a naive way to convert the map to a graph; more sophisticated ways are available but our focus in this paper is on the search algorithm, not on the discretization method.

²Note that the search schedule is not about the evader per se, but about the dirty set; in fact the search can be executed even in the absence of an evader.

of searchers required to clear a graph) and establishing its relationship to other graph parameters (e.g. pathwidth [5, 6], vertex separation [16] etc.). Another topic of interest is to establish the NP-completeness of the various versions of graph search. Several search algorithms have been presented but they are complex and hard to implement. Simpler algorithms exist for special types of graphs (e.g. for trees [2]) but the relevance of such graphs to robotic PE is rather limited.

In fact a major difference between robotic pursuit / evasion and graph search concerns *what* is being searched. Because a cave is essentially a collection of tunnels, in his original formulation Parsons assumed the evader to be hiding in the *edges* of the graph; this version of the problem is the one still studied in most of the graph theoretic literature, under the name *edge search*³. However, in robotic indoor PE we are usually interested in searching *rooms* and it is natural to model these as *nodes* of a graph.

Furthermore, the search schedules studied in the graph theoretic literature allow a pursuer to move between non-adjacent nodes (i.e. not necessarily following the edges of the graph). We call this “*teleporting*”; obviously, teleporting is not appropriate for robotics problems. Only recently has some research been published [2, 3, 4] that concentrates on *non-teleporting* search where the pursuers move only along the graph edges (such a search is called *internal* in the graph theoretic literature). In addition, most graph theoretic search algorithms produce *monotonic* searches, i.e. the dirty set decreases at *every* step of the search. In other words, *recontamination* (of edges and/or nodes) is not allowed. In fact it has been proved [17] that in teleporting edge search recontamination does not “help” (does not decrease the search number). However we will give examples which show that recontamination *may* help in non-teleporting node search.

Hence we see that the graph theoretic literature on graph search does not fit all aspects of robotic PE. However, we feel that the graph theoretic point of view has much to offer to robotic PE and that this connection has not been sufficiently explored by robotic researchers. We will elaborate this point in Section 5.

Let us now turn to robotic PE. The indoor version of the problem was first studied by Lavelle and his collaborators [18, 19, 20]; this work was continued by several other researchers [9, 11, 13, 25, 26]⁴. A very extensive exposition of robotic PE appears in the book [21]. While robotics researchers make use of the graph representation, apparently they have not made much use of the graph theoretic research mentioned in the previous paragraphs. Rather, the approach most often used consists of the following steps.

1. The indoor environment is discretized into cells, each of which corresponds to a node of an (undirected) *navigation graph*. The discretization is based on *critical visibility events*.
2. A (directed) *information space graph* is obtained from the navigation graph. The information graph essentially is a *state transition diagram* of the search process.
3. Search consists in finding a path *in the information graph* from the starting state to a *clear* state (using, for example, Dijkstra’s algorithm).

The basic disadvantage of the above approach is that it does not scale well, either with the size of the environment or the number of searchers. Because discretization is based on visibility changes, the navigation graph can involve a large number of nodes, even for simple environments. Furthermore, the size of the information graph grows exponentially with the number of nodes in the navigation graph *and* with the number of searchers. Hence the shortest path computation is viable only for simple environments and one searcher; for more complicated cases various approximations are used.

³Strangely, in the graph theoretic literature a variant appears which is named “node search” [16] but again assumes the evader to be hiding in the edges.

⁴All of these papers deal with guaranteed search; “probabilistic” PE, where the evader is assumed to move in a random manner, has also been studied [13, 15, 26] but will not concern us here.

Our goal in the current paper is to reduce the robotic PE problem to a graph *node* search problem and to obtain an algorithm which does not use teleporting, is not restricted to monotonic searches and can solve problems involving large indoor environments. The paper is organized as follows: in Section 2 we formulate the problem; in Section 3 we present our algorithm; in Section 4 we present experiments to evaluate the algorithm; in Section 5 we present our conclusions and directions for future research.

2 Problem Formulation

As already mentioned we will study PE in indoor environments. Consider the map of one such environment, illustrated in Fig.1.a. This map can be reduced to a *graph* by discretization. For example the map of Fig.1a can be reduced to the graph of Fig.1b.

Recall that a graph $G = (V, E)$ is fully described by (a) the set of nodes V and (b) the set of edges $E \subseteq V \times V$. We will label nodes by integers, i.e. for a graph of N nodes we will take $V = \{1, 2, \dots, N\}$. We also have (using the *set cardinality* notation $|\cdot|$) that $|V| = N$. The *map-to-graph reduction* is performed by

1. decomposing the map of the environment into convex cells (typically, but not necessarily, rectangles),
2. associating a node to each cell,
3. inserting edges between nodes which correspond to adjacent cells (and also an edge from each node to itself).




Figure 1. (a) An indoor environment partitioned into rectangular cells. (b) A graph obtained by corresponding the convex cells of 1.a. to nodes. The graph also includes edges from each node to itself (loops) which are not displayed for simplicity of presentation.

Note that the graph reduction is *not* unique, since the cell decomposition is not either. In other words the same map can be reduced to more than one graph. For example the map of Fig.1a is

reproduced in Fig.2a with a different cell decomposition and the corresponding graph appears in Fig.2b.

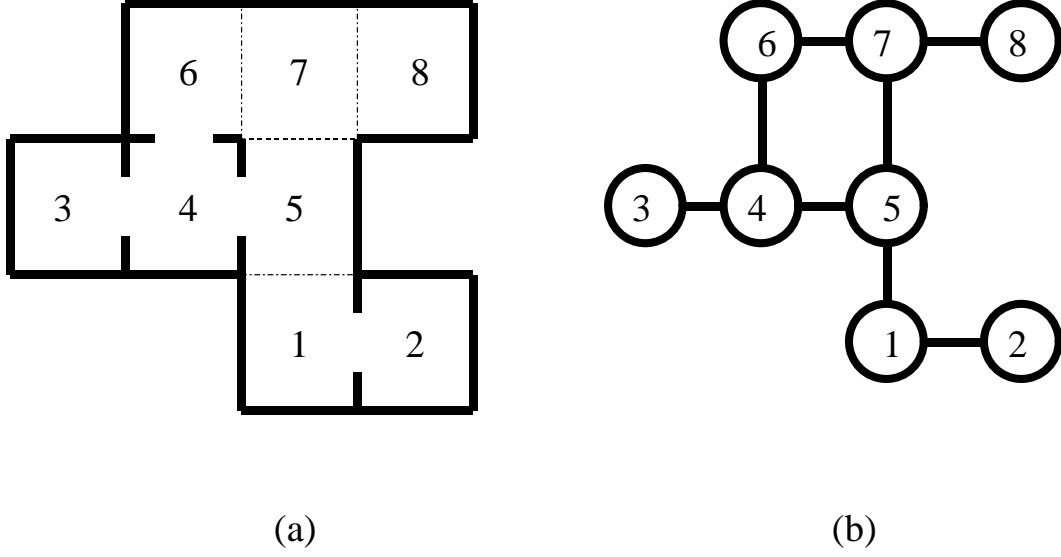


Figure 2. (a) The same environment as in Fig.1.a., partitioned into different cells. (b) The corresponding graph. Compare with 1.b. The graph again includes non-displayed looping edges.

Given a graph $G = (V, E)$ with N nodes, the *neighborhood* of node $n \in V$ is denoted by $N(n)$ and defined by

$$N(n) = \{m \in V : (n, m) \in E\}.$$

The $N \times N$ *adjacency matrix* A of the graph is defined by

$$A_{mn} = \begin{cases} 1 & \text{iff } n \in N(m); \\ 0 & \text{else.} \end{cases}$$

For example, the adjacency matrices of the graphs in Fig.1b and Fig.2b are, respectively,

$$A = \begin{bmatrix} 1 & 0 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 & 0 \\ 0 & 1 & 1 & 1 & 1 \\ 1 & 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 \end{bmatrix}, \quad C = \begin{bmatrix} 1 & 1 & 0 & 0 & 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \end{bmatrix}. \quad (1)$$

Given a graph obtained from a map, the *visibility region* of node $n \in V$ is denoted by $V(n)$ and defined by

$$V(n) = \{m \in V : m \text{ is visible from } n\}.$$

The condition “ m is visible from n ” is evaluated under the assumption of “straight-line-visibility” and, obviously, depends on the map and cell decomposition from which the graph was obtained. More specifically, “ m is visible from n ” holds iff *every* point of cell m is visible from every point of cell n (under straight line visibility). (In other words, if part of cell m is visible from n then $m \notin V(n)$; similarly if all of cell m is visible only from some points of n then $m \notin V(n)$). The $N \times N$ *visibility matrix* B is defined by

$$B_{mn} = \begin{cases} 1 & \text{iff } n \in V(m); \\ 0 & \text{else.} \end{cases}$$

Again, the visibility matrix of a graph depends on the map and cell decomposition from which the graph was obtained. For example, for the graphs of Figs.1b and 2b the visibility matrices are, respectively,

$$B = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}, \quad D = \begin{bmatrix} 1 & 0 & 0 & 0 & 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 \\ 1 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 \end{bmatrix}. \quad (2)$$

Note that the visibility matrix B (corresponding to Fig.1) is simply the identity matrix. Also note that the visibility matrix is always symmetric.

In the rest of this paper we will assume that PE takes place in a graph $G = (V, E)$ with $V = \{1, 2, \dots, N\}$, by a team of K pursuers. We will denote the position of the k -th pursuer at time t by $x_k(t)$. The K -long vector of pursuer positions will be denoted by $\mathbf{x}(t) = [x_1(t), \dots, x_K(t)]$; it takes values in the space of (pursuer) *configurations*, denoted by $\mathbf{X} = \{1, 2, \dots, N\}^K$.

If a node *might* contain the evader, the node is called *dirty*, otherwise *clear* (for example, if a pursuer resides in node n , then n is a clear node). A node n is *cleared* when visited by a pursuer or falls inside the pursuer's visibility region; it remains clear as long as there is no free (from pursuer inspection) path from it to a dirty node. (Hence the clear set is more than simply the union of the pursuers' visibility regions.) If at some time a free path from the clear node n to a dirty node m becomes available, then n is *recontaminated*. The set of all dirty nodes (resp. clear nodes) is the *dirty set*, denoted by D (resp. *clear set*, denoted by C). Obviously $C \cup D = V$, $C \cap D = \emptyset$. We will often describe sets by their *indicator vectors*. For example the indicator vector of D is $\mathbf{d} = [d_1, d_2, \dots, d_N]$, where $d_n = 1$ iff $n \in D$, 0 else; similarly \mathbf{c} is the indicator vector of C . The *length* of \mathbf{d} is $|\mathbf{d}| = \sum_{n=1}^N d_n$ (note that $|\mathbf{d}| = |D|$); similarly for \mathbf{c} .

Since the evader is assumed to be arbitrarily fast and to have complete knowledge of the pursuers' positions, it will stay outside the pursuers' visibility regions for as long as possible. The evader will only be captured when it has no escape moves left or, in other words, when the *dirty set becomes empty*. Hence the pursuers must move in such a manner that the dirty set is progressively (but not necessarily *monotonically*) reduced, until it becomes the empty set. Note that this schedule is really concerned with the dirty set, not with the evader; in fact it can be executed even in the absence of an evader.

We assume the PE process evolves in discrete time $t = 1, 2, \dots$. At time t , the essential information available to each pursuer is contained in the vector

$$\mathbf{z}(t) = [x_1(t), \dots, x_K(t), d_1(t), \dots, d_N(t)],$$

i.e. every pursuer knows the location of itself and its teammates and the *possible* locations of the evader. We call $\mathbf{z}(t)$ the *state vector* of the PE process (from the pursuers' point of view). For every t , $\mathbf{z}(t)$ takes values in the *state space* $\mathbf{Z} = \mathbf{X} \times \mathbf{D}$, where $\mathbf{X} = \{1, \dots, N\}^K$ (the set of all possible *pursuer configurations*) and $\mathbf{D} = \{0, 1\}^N$ (the set of all possible dirty sets). Hence every element $\mathbf{z} \in \mathbf{Z}$ can be written as $\mathbf{z} = (\mathbf{x}, \mathbf{d})$, where $\mathbf{x} \in \mathbf{X}$ is a configuration vector and $\mathbf{d} \in \mathbf{D}$ is a dirty set indicator vector. The \mathbf{Z} elements of the form $\mathbf{z} = (\mathbf{x}, \mathbf{0})$ are the *clear states*, i.e. the ones for which the clear set $C = V$. Since the state space \mathbf{Z} is finite we can also visualize the state evolution equation in terms of a *directed* state transition graph S (this is different from the original graph G) where the nodes are states $\mathbf{z} = (\mathbf{x}, \mathbf{d})$ and the directed edges are state transitions.

We will also use the *control vector* $\mathbf{u}(t) = [u_1(t), \dots, u_K(t)]$, with $u_k(t)$ denoting the next move of the k -th pursuer, i.e. $x_k(t+1) = u_k(t)$.

We will now obtain the *PE state evolution equation*. This will have the form

$$\mathbf{z}(t+1) = F(\mathbf{z}(t), \mathbf{u}(t)), \quad (3)$$

i.e. it will yield the next state $\mathbf{z}(t+1)$ in terms of the previous state $\mathbf{z}(t)$ and the control $\mathbf{u}(t)$.

The movement of the evader (or, more precisely, the evolution of the dirty set) is modeled by *max-min matrix multiplication*, denoted by $*$ and defined as follows: for arbitrary matrices P (with dimensions $L \times M$) and Q (with dimensions $M \times N$) the matrix $R = P * Q$ is defined by

$$R_{ln} = \max_{m=1,2,\dots,M} [\min(P_{lm}, Q_{mn})] \text{ for } l \in \{1, 2, \dots, L\} \text{ and } n \in \{1, 2, \dots, N\}.$$

The i -th *max-min power* of a square matrix P is denoted by $P^{[i]}$ and defined by

$$P^{[i]} = \underset{i \text{ times}}{P * P * \dots * P}.$$

Suppose for a moment that the pursuers are removed from the graph and the evader moves with unit speed (i.e. it traverses exactly one edge per time step). Then, the possible locations of the evader at time $t+1$ are the ones adjacent to its possible locations at time t . Mathematically this is expressed by

$$\mathbf{d}(t+1) = \mathbf{d}(t) * A, \quad (4)$$

If the evader has speed M (it crosses M edges in a single time step) then instead of (4) we have

$$\mathbf{d}(t+1) = \mathbf{d}(t) * A^{[M]} \quad (5)$$

For the case of an arbitrarily fast evader it suffices to set $M = N$, the number of nodes in the graph.

Now suppose a single pursuer is placed at node x_1 . The pursuer blocks certain paths which, if taken by the evader, would result in capture. In particular, the evader will not move into any node which falls in the visibility region of the pursuer, because it would be captured immediately after its move. Similarly, the evader will not move out of any node in the visibility region, because it would have been captured immediately *before* the move. Similar rules hold in case more than one pursuer is located in the graph.

These facts can be modeled as follows. Recall that $\mathbf{x} = [x_1, x_2, \dots, x_K]$ is the vector containing the locations of the K pursuers. Now we introduce the *modified adjacency matrices* $\overline{A}(\mathbf{x})$, defined as follows.

$$\overline{A}_{ij}(\mathbf{x}) = \begin{cases} A_{ij} & \text{if } i \notin V(x_k) \text{ and } j \notin V(x_k) \text{ for any } k \in \{1, 2, \dots, K\}; \\ 0 & \text{if } i \in V(x_k) \text{ or } j \in V(x_k) \text{ for some } k \in \{1, 2, \dots, K\}. \end{cases}$$

In other words, if a pursuer is located at node n , then we obtain $\overline{A}(\mathbf{x})$ by setting to zero the rows and columns of A which correspond to the nodes in the visibility region of x_1 (i.e. the nodes visible by the pursuer). Another way to think of this is that we remove from the graph the nodes in the visibility region of the pursuer (and all their incident edges).

Putting all of the above pieces together, the PE state evolution equation for K pursuers and evader speed equal to M is

$$\begin{aligned} \mathbf{x}(t+1) &= \mathbf{u}(t) \\ \mathbf{d}(t+1) &= \mathbf{d}(t) * [\overline{A}(\mathbf{u}(t))]^{[M]} \\ \text{subject to : } &u_k(t) \in N(x_k(t)) \text{ for } k = 1, 2, \dots, K. \end{aligned} \quad (6)$$

Example. Let us illustrate the above ideas using the graph of Fig.1b, with adjacency matrix A of eq.(1) and the visibility matrix B of eq.(2). Typical modified adjacency matrices are, for example,

$$\bar{A}(3) = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}, \quad \bar{A}(4) = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

Assume a PE with an evader of *unit speed* and two pursuers, starting at nodes 2 and 3 and moving, respectively, to nodes $u_1(1)=3$ and $u_2(1)=4$. Then the initial state vector is

$$\mathbf{x}(1) = [2, 3] \\ \mathbf{d}(1) = [1 \ 0 \ 0 \ 1 \ 1].$$

The new pursuer locations are $\mathbf{x}(2) = [3, 4]$, the modified adjacency matrix is

$$\bar{A}([3, 4]) = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

and the new dirty set is

$$\mathbf{d}(2) = \mathbf{d}(1) * \bar{A}([3, 4]) = [1 \ 0 \ 0 \ 1 \ 1] * \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} = [1 \ 0 \ 0 \ 0 \ 1].$$

Hence, the clear nodes are 3 and 4 (because occupied by pursuers) and 2 (because it was previously clear and there is no free path from 2 to the dirty nodes 1 and 5). Now suppose the next pursuer moves are $u_1(2)=4$ and $u_2(2)=1$. Then we have $\mathbf{x}(3) = [4, 1]$, and the new dirty set is

$$\mathbf{d}(3) = \mathbf{d}(2) * \bar{A}([4, 1]) = [1 \ 0 \ 0 \ 0 \ 1] * \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 \\ 0 & 1 & 1 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 1 \end{bmatrix} = [0 \ 0 \ 1 \ 0 \ 1].$$

Nodes 1 and 4 are clear, because they are occupied by pursuers. Node 3 is recontaminated, because there is a free path from 3 to the dirty node 5. Node 2 is clear because, even though there is a free path from 5 to 2, it has length two and in this example we have assumed that the evader has unit speed. The reader can continue in this manner to compute dirty sets after application of various controls $\mathbf{u}(t)$, $t = 3, 4, \dots$

We complete the problem formulation by introducing a *cost function* which is nonnegative and takes the value zero exactly when the evader is captured (or, more precisely, when its location is known with certainty). The cost function is simply the cardinality of the dirty set:

$$J(t) = |\mathbf{d}(t)|. \tag{7}$$

Indeed, if $J(t) = 0$ then the dirty set has size 0, i.e. it is empty or, in other words, the evader has been discovered.

In (7) we have written J as a function of time. Actually it is more relevant to write J as a function of the controls $\mathbf{u}(1), \dots, \mathbf{u}(t)$: as the pursuers apply a search schedule (described by $\mathbf{u}(1), \dots, \mathbf{u}(t)$) the size of the dirty set changes, so at time t (after the entire control sequence has been applied) we have

$$J(\mathbf{u}(1), \dots, \mathbf{u}(t)) = |\mathbf{d}(t)|$$

and this is the form we will use in what follows.

We conclude this section by introducing a concept that we will use in our algorithm: *distance of pursuer configuration from dirty set*. This is denoted by $dist(\mathbf{x}, \mathbf{d})$, where \mathbf{x} is a pursuer configuration and \mathbf{d} the indicator vector of a dirty set. To define $dist(\mathbf{x}, \mathbf{d})$, first recall that the distance of nodes m, n in a graph, is the length of the shortest path joining m and n ; denote this by $dist(m, n)$. Second, we denote by $dist(U, W)$ the distance between *sets* of nodes U and W and define it to be the *sum* of distances between nodes of the two sets:

$$dist(U, W) = \sum_{m \in U} \sum_{n \in W} dist(m, n).$$

Finally, given \mathbf{x} , a pursuer configuration and \mathbf{d} , the indicator vector of a dirty set, denote by X the set of nodes occupied by the pursuers (corresponding to \mathbf{x}) and by D the dirty set (corresponding to \mathbf{d}); then define

$$dist(\mathbf{x}, \mathbf{d}) = dist(X, D). \tag{8}$$

3 The Search Algorithm

3.1 Search Algorithm: a Preliminary Version

In the state transition graph S of the PE process, suppose that the pursuers start at an initial node $\mathbf{z}_{in} = (\mathbf{x}_{in}, \mathbf{d}_{in})$. We want to move the pursuers through nodes $\mathbf{z} = (\mathbf{x}, \mathbf{d})$ (and along edges of S) in a manner such that the 1's in the \mathbf{d} part progressively decrease (the dirty set shrinks) until eventually we reach one of the clear states, say $\mathbf{z}_{fin} = (\mathbf{x}_{fin}, \mathbf{d}_{fin})$, where $\mathbf{d}_{fin} = \mathbf{0}$. To this end we could, theoretically, find a path from \mathbf{z}_{in} to some \mathbf{z}_{fin} ; indeed we could find a *shortest* path using, for example, Dijkstra's algorithm. Note that this approach does not take into account any specific moves of the evader, it is a *worst case* approach. While this approach solves the graph search problem *in principle*, it is computationally intractable: even for moderate values of N and K the state space is too big. For example, with $N = 10$ and $K = 2$ we have $|\mathbf{Z}| = 10^2 \times 2^{10} = 102\,400$. To even build the state transition graph is computationally taxing; to find shortest paths in it is computationally prohibitive (recall that Dijkstra's algorithm is quadratic in the number of nodes of S [7]).

To reduce computational burden we use *two* approximations.

1. First, we project \mathbf{Z} to a smaller set, namely \mathbf{X} or \mathbf{D} . Several \mathbf{Z} elements are mapped to the same \mathbf{X} element – because the same pursuer configuration, reached through different paths, can result in many different dirty sets.
2. \mathbf{X} is still not sufficiently small to be searched by an exact shortest path algorithm. So we use a second approximation: to each element $\mathbf{x} \in \mathbf{X}$ we assign a cost, namely the size of the *least known* dirty set associated with \mathbf{x} . By “*least known*” we mean that we will gradually explore more and more paths terminating at a given pursuer configuration \mathbf{x} and will always store the dirty set of least cost (along with the corresponding search path). This computation will be iterative: least cost of a configuration is updated in terms of previously computed paths and dirty sets of other configurations.

We illustrate the above ideas in pseudocode as Algorithm 1 (on p.10). The algorithm is presented for the case of $K = 2$ (two pursuers); extension to $K > 2$ is obvious. The following notations are used in the pseudocode representation:

1. the *path vector* $\mathbf{p}^{(n_1, n_2)}$ has the form $(i_1 i_2 \dots i_T, j_1 j_2 \dots j_T)$ where $i_1 i_2 \dots i_T$ is the path of the first pursuer, and $j_1 j_2 \dots j_T$ is the path of the second pursuer;
2. the components of $\mathbf{p}^{(n_1, n_2)}$ are denoted by $\mathbf{p}_1^{(n_1, n_2)} = i_1 i_2 \dots i_T$ and $\mathbf{p}_2^{(n_1, n_2)} = j_1 j_2 \dots j_T$;
3. the *cost* associated with configuration (n_1, n_2) is denoted by $v^{(n_1, n_2)}$;
4. for every set of nodes A , $\text{Ind}(A)$ denotes the indicator vector of A .

In summary, Algorithm 1 operates as follows. The algorithm computes, for every searcher configuration (n_1, n_2) , the vector of best paths $\mathbf{p}^{(n_1, n_2)}$, the vector of best dirty set $\mathbf{d}^{(n_1, n_2)}$ and the corresponding best cost $v^{(n_1, n_2)}$. This computation is *iterative* and *greedy*: in every iteration the cost of one or more configurations may be reduced. Note that, while the cost of a specific configuration is a monotonic function of algorithm iteration, *it does not have to be a monotonic function of path length*, i.e. *recontamination* is possible. The algorithm terminates when some configuration (n_1, n_2) achieves zero cost ($v^{(n_1, n_2)} = 0$). Because of greediness, termination is not guaranteed; a condition can be added for maximum number of iterations, to break out of infinite repetition.

Algorithm 1 Graph Search, Preliminary Version

Input: Graph $G = (V, E)$; Visibility matrix B ; Starting nodes i, j

Initialization

For $n = 1$ to N

 For $m = 1$ to N

$\mathbf{p}^{(n,m)} = ()$

$\mathbf{d}^{(n,m)} = \text{Ind}(V)$

$v^{(n,m)} = N$

 Next m

Next n

$\mathbf{p}^{(i,j)} = (i, j)$

$\mathbf{d}^{(i,j)} = \text{Ind}(V - V(i) - V(j))$

$v^{(i,j)} = |\mathbf{d}^{(i,j)}|$

$v_{Best} = v^{(i,j)}$

$\mathbf{p}_{Best} = (i, j)$

Main

While $v_{Best} > 0$

 For $n_1 = 1$ to N

 For $n_2 = 1$ to N

$\hat{\mathbf{x}} = [n_1, n_2]$

 For $m_1 \in N(n_1)$

 For $m_2 \in N(n_2)$

$\hat{\mathbf{u}} = [m_1, m_2]$

$\hat{\mathbf{p}} = (\mathbf{p}_1^{(n_1, n_2)} m_1, \mathbf{p}_2^{(n_1, n_2)} m_2)$

$\hat{\mathbf{d}} = \mathbf{d}^{(n_1, n_2)} * [\overline{A}(\hat{\mathbf{u}})]^{[N]}$

$\hat{v} = |\hat{\mathbf{d}}|$

 If $\hat{v} < v^{(m_1, m_2)}$

$\mathbf{p}^{(m_1, m_2)} = \hat{\mathbf{p}}$

$\mathbf{d}^{(m_1, m_2)} = \hat{\mathbf{d}}$

$v^{(m_1, m_2)} = \hat{v}$

 EndIf

 If $v^{(m_1, m_2)} < v_{Best}$

$v_{Best} = v^{(m_1, m_2)}$

$\mathbf{p}_{Best} = \hat{\mathbf{p}}$

 EndIf

 Next m_2

 Next m_1

 Next n_2

 Next n_1

End While

Output: \mathbf{p}_{Best}

In more detail, Algorithm 1 operates as follows.

1. *Input* is the graph to be searched (and its visibility matrix B) and nodes i, j , the *initial positions* of the pursuers.
2. *Initialization* consists of three parts.
 - (a) In the two nested loops every pursuer configuration (n, m) is associated with the empty path, the “full” dirty set ($D = V$) and the “full” cost N .
 - (b) Then the starting configuration (i, j) is associated to (one-step long) path (i, j) , dirty set $V - V(i) - V(j)$ (i.e. all nodes outside the visibility regions of nodes i and j) and the corresponding cost $v^{(i,j)}$.
 - (c) Also the overall best cost v_{Best} is set to $v^{(i,j)}$ and the overall best path is set to $\mathbf{p}_{Best} = (i, j)$.
3. In the *main* part of the program, the while-loop is executed for as long as best cost v_{Best} is greater than zero (or until the maximum number of iterations is exceeded).
 - (a) Inside the while-loop we have two nested for-loops which examine each possible pursuer configuration (n_1, n_2) .
 - (b) Furthermore, we have two inner for-loops which find every configuration (m_1, m_2) which is a possible successor of (n_1, n_2) . For each (m_1, m_2) the *expanded* path $\hat{\mathbf{p}} = (\mathbf{p}_1^{(n_1, n_2)} m_1, \mathbf{p}_2^{(n_1, n_2)} m_2)$ is formed and the new dirty set $\hat{\mathbf{d}}$ and associated cost \hat{v} are computed
 - (c) The first if statement tests whether $\hat{v} < v^{(m_1, m_2)}$; if this is the case, then the path, dirty set and cost of (m_1, m_2) are updated.
 - (d) The second if statement tests whether $v^{(m_1, m_2)}$ is less than the overall best cost v_{Best} ; if this is the case, then v_{Best} and also the overall best path \mathbf{p}_{Best} are updated.
 - (e) If v_{Best} becomes equal to zero, the graph has been cleared; the algorithm terminates and returns \mathbf{p}_{Best} .

Here is a final remark, which will be important for the next section. Algorithm 1 operates on pursuer configurations $\mathbf{x} = (n_1, \dots, n_K)$. For every such configuration we store the corresponding paths and dirty sets. Another way to say this is that stored triples $(\mathbf{x}, \mathbf{d}, \mathbf{p})$ (configurations, dirty sets and paths) are *indexed* by their \mathbf{x} 's; at most one triple with a particular \mathbf{x} is stored.

3.2 Search Algorithm: an Improved Version

We now present an improved version of our algorithm, which can search a graph more efficiently than Algorithm 1. The improved Algorithm 2 involves two modifications.

1. *Indexing by dirty sets, not by pursuer configurations.* As already mentoned, Algorithm 1 indexes triples $(\mathbf{x}, \mathbf{d}, \mathbf{p})$ by their \mathbf{x} 's. We have found, however, that it is more efficient to index triples $(\mathbf{x}, \mathbf{d}, \mathbf{p})$ by \mathbf{d} , the dirty set⁵; as a consequence, at any given time we store at most one (the best) triple with a given dirty set. Suppose one such $(\mathbf{x}, \mathbf{d}, \mathbf{p})$ is already stored and now we reach

⁵Indexing by dirty sets (rather than configurations) requires less storage and less computation. To see this consider the two configurations (n_1, n_2) and (n_2, n_1) ; these two pairs correspond to the same PE situation (one pursuer located at n_1 and one at n_2) but they appear as two different elements of \mathbf{X} ; each of (n_1, n_2) and (n_2, n_1) will be stored separately, and a corresponding best path will also be stored for each and expanded. This results in considerable duplication of computation, especially when we deal with $K > 2$ pursuers (for example, for $K = 3$ the six configurations $(3,4,6)$, $(3,6,4)$, $(4,3,6)$, $(4,6,3)$, $(6,3,4)$, $(6,4,3)$ correspond to the same physical placement of pursuers).

some $(\hat{\mathbf{x}}, \hat{\mathbf{d}}, \hat{\mathbf{p}})$ with $\hat{\mathbf{d}} = \mathbf{d}$. Should we retain $(\mathbf{x}, \mathbf{d}, \mathbf{p})$ or replace it by $(\hat{\mathbf{x}}, \hat{\mathbf{d}}, \hat{\mathbf{p}})$ (we cannot keep both, since $\mathbf{d} = \hat{\mathbf{d}}$)? In other words, which of the two triples has better chances to lead to graph clearing? We answer this question by a *heuristic* criterion, which involves the distance function of eq.(8).

- (a) If $dist(\mathbf{x}, \mathbf{d}) > dist(\hat{\mathbf{x}}, \hat{\mathbf{d}})$ we replace $(\mathbf{x}, \mathbf{d}, \mathbf{p})$ by $(\hat{\mathbf{x}}, \hat{\mathbf{d}}, \hat{\mathbf{p}})$;
- (b) else we retain $(\mathbf{x}, \mathbf{d}, \mathbf{p})$ and drop $(\hat{\mathbf{x}}, \hat{\mathbf{d}}, \hat{\mathbf{p}})$.

The rationale of the criterion is that, between two states which have same dirty set, the one in which the pursuers are closer to the dirty nodes is preferable; this can be seen as a way to force the pursuers to move closer to the dirty nodes.

2. *Searching serially, not in parallel.* Algorithm 1 involves two *nested* loops “For $n_1 = 1$ to N ”, “For $n_2 = 1$ to N ”; this corresponds to *parallel* search of the configuration space (and the graph G), i.e. the two pursuers are assumed to move simultaneously. An alternative implementation would be *serial*: the first loop is executed and *then* the second. This corresponds to *serial* search: first pursuer no.1 moves while no.2 remains stationary and then their roles are reversed. The same distinction can be applied for the general case with $K \geq 2$. Computationally, parallel implementation / nested loops is more intensive than serial. For K searchers a rough estimate of the computational complexity is $O(N^{2K})$ for parallel and $O(KN^2)$ for serial implementation. To speed up execution time, our improved algorithm uses serial search⁶. There is a price to pay however: given a particular configuration $\mathbf{x} = [x_1, x_2, \dots, x_K]$, parallel search examines all possible successor configurations, while serial search examines only successor configurations which differ from \mathbf{x} by at most one coordinate. This can be seen either as a disadvantage of serial search (because some configurations are not examined) or as an advantage (exactly because some configurations are not examined, the algorithm operates faster).

Introducing the above two modifications of Algorithm 1, we obtain Algorithm 2, which is listed on p.13 for the case of $K = 2$ pursuers (the extension to $K > 2$ is obvious). We only describe the operation of Algorithm 2 in summary (the reader may consult the listing of p.13 for details). A set Q is used to store all dirty sets discovered. For every $\mathbf{d} \in Q$ we store the corresponding cost $v^{(\mathbf{d})} = |\mathbf{d}|$ and the best known path and configuration ($\mathbf{p}^{(\mathbf{d})}$ and $\mathbf{x}^{(\mathbf{d})}$, respectively) which reach this \mathbf{d} . At initialization Q contains only the dirty set corresponding to the initial positions (i, j) of the pursuers. In the main part of Algorithm 2, for every $\mathbf{d} \in Q$ the corresponding configuration $\mathbf{x}^{(\mathbf{d})} = (n_1, n_2)$ is expanded, i.e. all neighbors *first* of n_1 and *then* of n_2 (recall that search is *serial*) are considered as possible next searcher moves. For each such move a new dirty set $\hat{\mathbf{d}}$ is computed and the algorithm checks whether $\hat{\mathbf{d}} \in Q$. If $\hat{\mathbf{d}} \notin Q$, then $\hat{\mathbf{d}}$ is inserted in Q and the corresponding cost, path and pursuer configuration are stored; if $\hat{\mathbf{d}} \in Q$, then the distance criterion (eq.(8) of Section 2) is used to determine if the old or new path and pursuer configuration should be stored as $\mathbf{p}^{(\hat{\mathbf{d}})}$ and $\mathbf{x}^{(\hat{\mathbf{d}})}$, respectively. The algorithm terminates when an empty dirty set is generated (the graph has been cleared) and returns the corresponding path.

Algorithm 2 (similarly to Algorithm 1) is greedy, iterative and not guaranteed to terminate.

⁶For completeness let us mention a third mode of search, which we can call *serial / parallel*. This involves breaking down the K searchers into J groups of K_1, K_2, \dots, K_J searchers; the j -th group is implemented by K_j nested loops, and the J groups are run serially. Presumably this gives execution time between serial and parallel.

Algorithm 2 Graph Search, Improved Version

Input: Graph $G = (V, E)$; Visibility matrix B ; Starting nodes i, j

Initialization

$\mathbf{d} = \text{Ind}(V - V(i) - V(j))$

$Q = \{\mathbf{d}\}$

$\mathbf{x}^{(\mathbf{d})} = (i, j)$

$\mathbf{p}^{(\mathbf{d})} = (i, j)$

$v^{(\mathbf{d})} = |\mathbf{d}|$

$\mathbf{p}_{Best} = \mathbf{p}^{(\mathbf{d})}$

$v_{Best} = v^{(\mathbf{d})}$

Main

While $v_{Best} > 0$

 For $\mathbf{d} \in Q$

 For $i \in \{1, 2\}$

$n_i = \mathbf{x}_i^{(\mathbf{d})}$

 Next i

 For $i \in \{1, 2\}$

$\hat{\mathbf{x}} = \mathbf{x}^{(\mathbf{d})}$

$\hat{\mathbf{p}} = (\mathbf{p}_1^{(\mathbf{d})} n_1, \mathbf{p}_2^{(\mathbf{d})} n_2)$

$\hat{\mathbf{u}} = \mathbf{x}^{(\mathbf{d})}$

 For $m_i \in N(n_i)$

$\hat{\mathbf{u}}_i = m_i$

$\hat{\mathbf{x}}_i = m_i$

$\hat{\mathbf{p}}_i = \mathbf{p}_i^{(\mathbf{d})} m_i$

$\hat{\mathbf{d}} = \mathbf{d} * [\bar{A}(\hat{\mathbf{u}})]^{[N]}$

$\hat{v} = |\hat{\mathbf{d}}|$

 If $\hat{\mathbf{d}} \notin Q$

$Q = Q \cup \{\hat{\mathbf{d}}\}$

$\mathbf{x}^{(\hat{\mathbf{d}})} = \hat{\mathbf{x}}$

$\mathbf{p}^{(\hat{\mathbf{d}})} = \hat{\mathbf{p}}$

$v^{(\hat{\mathbf{d}})} = \hat{v}$

 ElseIf $\text{dist}(\hat{\mathbf{x}}, \hat{\mathbf{d}}) < \text{dist}(\mathbf{x}^{(\mathbf{d})}, \hat{\mathbf{d}})$

$\mathbf{x}^{(\hat{\mathbf{d}})} = \hat{\mathbf{x}}$

$\mathbf{p}^{(\hat{\mathbf{d}})} = \hat{\mathbf{p}}$

$v^{(\hat{\mathbf{d}})} = \hat{v}$

 EndIf

 If $v^{(\hat{\mathbf{d}})} < v_{Best}$

$\mathbf{p}_{Best} = \mathbf{p}^{(\hat{\mathbf{d}})}$

$v_{Best} = v^{(\hat{\mathbf{d}})}$

 EndIf

 Next m_i

 Next i

 Next \mathbf{d}

End While

Output: \mathbf{p}_{Best}

3.3 Search Algorithm: the Final Version

In this section we present some heuristics which can be incorporated in Algorithm 2. The first two heuristics are intended to reduce computation time by pruning the list of stored triples $(\mathbf{x}, \mathbf{d}, \mathbf{p})$; the third heuristic reduces the probability of the algorithm getting trapped into nonclearing paths.

1. We introduce an algorithm parameter L_0 (*maximum list size*) and, if the number of elements of Q grows to more than L_0 elements, we retain only the L_0 best elements (here “best” is understood in terms of lowest cost).
2. We introduce an algorithm parameter v_0 (*cost tolerance*) and store only triples with cost v no greater than $v_{Best} + v_0$ (keep in mind that v_{Best} changes during the execution of the algorithm, hence previously stored triples may be discarded).
3. To avoid getting stuck, we also introduce randomness in our algorithm, in a manner similar to stochastic steepest descent. Namely, given triples $(\mathbf{x}, \mathbf{d}, \mathbf{p})$ and $(\hat{\mathbf{x}}, \hat{\mathbf{d}}, \hat{\mathbf{p}})$ with $\mathbf{d} = \hat{\mathbf{d}}$ and $dist(\mathbf{x}, \mathbf{d}) = dist(\hat{\mathbf{x}}, \hat{\mathbf{d}})$ we will replace $(\mathbf{x}, \mathbf{d}, \mathbf{p})$ with $(\hat{\mathbf{x}}, \hat{\mathbf{d}}, \hat{\mathbf{p}})$ if $|\hat{\mathbf{d}}| < |\mathbf{d}| + \varepsilon$, where ε is a zero mean, normally distributed random variable with standard deviation σ (and σ is a parameter of the algorithm)⁷.

The “full” search algorithm, which we have used in the experiments of the next section, is Algorithm 2 plus the above described heuristics. A pseudocode version is not given here, for economy of space.

4 Experiments

In this section we present several indoor PE experiments. Each of the following subsections has the same general structure: a particular floorplan is selected, it is discretized (sometimes in more than one ways) and Algorithm 2, enhanced with the heuristics of Section 3.3, is used to search the resulting graph. The search algorithm is implemented in Matlab 7.0 and all the experiments reported here were run on a PC with Intel double core Pentium 4 processor running at 2.13 GHz; the operating system was Windows XP.

⁷This idea is similar to the stochastic selection of search schedules in Gerkey’s Parish algorithm [10].

4.1 Floorplan no.1

The first floorplan we use appears in Fig. 3.a and has been previously used by Sarmiento and Murrieta [26]. The discretization we use also appears in Fig. 3.a and the resulting graph in Fig. 3.b. We assume local visibility (in the discretization we have used, no cell affords full visual coverage of its neighbors).




Figure 3. (a) Floorplan no.1. The discretization cells are indicated with dashed lines. (b) The graph resulting from the floorplan and discretization of Figure 3.a. Dots indicate the final positions of the two pursuers.

We apply the graph search algorithm with two searchers, both starting at node 1, and the following parameter values: *maximum list size* $L_0 = 12$, *cost tolerance* $v_0 = 0$ and standard deviation (of the noise term ϵ) $\sigma = 1$. We obtain the following search schedule which clears the graph in 22 moves (computation time was 0.90 s). The k -th row of \mathbf{p}_{Best} shows the path of the k -th searcher ($k \in \{1, 2\}$).

$$\mathbf{p}_{Best} = \begin{bmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 4 & 3 & 7 & 8 & 11 & 11 & 12 & 13 & 10 & 9 & 9 & 9 & 9 & 5 & 6 & 2 \\ 1 & 4 & 3 & 7 & 8 & 11 & 12 & 12 & 12 & 12 & 12 & 12 & 15 & 15 & 15 & 15 & 15 & 14 & 16 & 17 & 17 & 17 & 17 \end{bmatrix}$$

We have experimented with other starting positions which also result in graph clearing. We have also experimented with various parameter values; for some of these the algorithm terminates successfully; for some combinations of values the algorithm fails to clear the graph. We find this acceptable: when given a particular graph to clear, it is only required to find a clearing schedule for *some* combination of parameter values; since such a combination is relatively easy to find by trial-and-error (especially after the user has become familiar with the algorithm) we believe appropriate selection of parameter values is not a problem.

4.2 Floorplan no.2

The next floorplan we use has been previously used by Lavalley et al. [18]. It is illustrated in Fig.4.a along with a possible discretization which yields *local visibility*; the resulting graph appears in Fig. 4.b.

The search algorithm with two searchers starting at node 1 (and $L_0 = 12$, $v_0 = 0$, $\sigma = 1$) yields the following search schedule which clears the graph in 13 moves (computed in 0.15 s).

$$\mathbf{p}_{Best} = \begin{bmatrix} 1 & 1 & 3 & 1 & 4 & 5 & 6 & 6 & 6 & 6 & 2 & 2 & 2 & 2 \\ 1 & 4 & 4 & 4 & 4 & 4 & 4 & 8 & 4 & 5 & 5 & 6 & 9 & 7 \end{bmatrix}.$$

The graph of Fig. 4.b has search number 2, but the floorplan of Fig. 4.a can be cleared by one searcher. To see this, we plot the same floorplan with a different discretization in Fig. 5.a. The new discretization yields *extended* (straight-line visibility). The resulting graph appears in Fig. 5.b.




Figure 4. (a) Floorplan no.2. The discretization cells are indicated with dashed lines. (b) The graph resulting from the floorplan and discretization of Figure 4.a. Dots indicate the final positions of the two pursuers.




Figure 5. Floorplan no.2 with an alternative, finer discretization. Note that this discretization affords *extended*, straight-line visibility. (b) The graph resulting from the discretization of Figure 5.a.

It is not immediately obvious but, under extended visibility, the graph of Figure 5.b can be cleared by a single searcher. Assuming the searcher starts at node 7 (and using $L_0 = 12$, $v_0 = 5$, $\sigma = 1$) our algorithm finds the following clearing schedule of length 19 (in 0.25 s). Note the use of a large-size list Q ($L_0 = 40$) and the storage of many suboptimal values ($v_0 = 5$); these choices are necessary because the algorithm must explore a large of state sequences until it finds a clearing path.

$$\mathbf{p}_{Best} = [7 \ 3 \ 2 \ 1 \ 6 \ 9 \ 12 \ 15 \ 12 \ 9 \ 10 \ 11 \ 8 \ 5 \ 8 \ 11 \ 14 \ 19 \ 18 \ 17].$$

Hence we see that appropriate discretization of a floorplan (and use of extended visibility) may reduce the numbers of searchers used to clear the floorplan.

4.3 Floorplan no.3

The next floorplan (also used by Lavalley et al. [18]) appears in Fig.6.a along with a discretization yielding *local visibility*; the resulting graph appears in Fig. 6.b.




Figure 6. (a) Floorplan no.3 and a discretization, the cells of which are denoted by dashed lines. (b) The graph obtained from the discretization of Figure 6.b.

While the graph of Fig. 6.a is 3-searchable (as can be seen by inspection), our algorithm fails to find a clearing path for three searchers even for large values of L_0 (the largest value we have tried is $L_0 = 250$). It can find clearing paths with *four* searchers and, for example, $L_0 = 12$, $v_0 = 0$ and $\sigma = 1$.

We believe that the algorithm fails to clear the graph of Fig. 6.b with three searchers because of the graph symmetry which results in storage of many essentially similar paths. To test this hypothesis, we used a different discretization (Fig. 7.a) which breaks some of the graph symmetries.




Figure 7.(a) The same floorplan as in Fig. 6.a, but with a different discretization. (b) The graph corresponding to the discretization of Fig. 7.a.

We ran our algorithm on the graph of Fig.7.b, with three searchers starting at node 1 and with $L_0 = 50$, $v_0 = 2$, $\sigma = 1$. The algorithm finds a clearing path of length 67, in 6 s. This is one more case where the search number of the floorplan may be reduced by appropriate discretization. Note that again use of a large-size list Q and storage of many suboptimal values ($v_0 = 5$) are necessary for a more complete exploration of the relatively complex state space of this example.

4.4 Floorplan no.4

In this experiment we use the floorplan of Fig. 8.a. First we use a relatively coarse discretization which yields local visibility (Fig.8.a and Fig.8.b).




Figure 8. (a) Floorplan no.4 and a discretization (cells denoted by dashed lines) which yields local visibility. (b) The graph obtained from the discretization of Figure 8.a.

It is easy to see that the graph of Fig.14 can be cleared with three searchers. Using $K = 3$, $L_0 = 12$, $v_0 = 0$, $\sigma = 1$ we obtain a clearing path of length 24 in 0.34 s.

We now use a finer discretization (Fig.9) which yields extended visibility.




Figure 9. (a) Floorplan no.4 and a finer discretization (cells denoted by dashed lines) which yields extended visibility. (b) The graph obtained from the discretization of Figure 9.a.

Using the finer discretization of Fig. 9.a and the corresponding graph of Fig.9.b (with extended visibility) and $L_0 = 30$, $v_0 = 2$, $\sigma = 1$ we are able to find a clearing schedule (of length 60) with *two* searchers (the schedule was computed in 1.84 s). Note that, in this example too, finding a solution with a smaller number of searchers requires a more extensive exploration of the state space, which in turns requires higher values of L_0 and v_0 .

4.5 Floorplan no.5

In this experiment we use the floorplan of Fig. 10.a. First we use the discretization of Fig.10.a which yields the graph of Fig.10.b. This graph can be cleared with 3 searchers (*using recontamination*; the search number under monotonic search is 4) but our algorithm could not find this solution




Figure 10. (a) Floorplan no.5 and a discretization (cells denoted by dashed lines) which yields local visibility. (b) The graph obtained from the discretization of Figure 10.a.




Figure 11. (a) Floorplan no.5 with a different discretization (cells denoted by dashed lines) used to break the symmetries of the resulting graph. (b) The graph obtained from the discretization of Figure 11.a.

We conjecture the reason the algorithm failed is again the high degree of symmetry of the graph. Hence we used the discretization of Figure 11.a (i.e. splitting the top northeast cell in two) which yields the asymmetric graph of Fig.11.b; for this graph (in 6.38 s of computation) we find a three searcher clearing schedule (of length 51). Parameter values are $L_0 = 30$, $v_0 = 2$, $\sigma = 1$.

4.6 Floorplan no.6

In this experiment we use a real-world floor plan, illustrated in Fig. 12 (it is the third floor of Newell-Simon Hall at Carnegie-Mellon University – this floorplan has been previously used in [13, 14]). The discretization used is also illustrated in Fig. 12 and the resulting graph in Fig.13. It can be seen that the graph has more nodes (58) than the previously used graphs. It is not obvious what the search number of this graph is (a clearing search schedule with 3 searchers has been found in [14]).




Figure 12. Floorplan no.6, a real world office building. The discretization (which yields local visibility) is denoted by dashed lines and cell numbers.




Figure 13. The graph obtained from the discretization of Figure 12.

We have been able to clear the graph with 6 searchers. The parameters used are $L_0 = 3$, $v_0 = 0$, $\sigma = 1$; the clearing schedule has length 137 and was computed in 14 s.

4.7 Floorplan no.7

In this experiment we use a real-world floor plan, illustrated in Fig. 14 (it is the floorplan of a museum, hence the many cycles – this floorplan has been previously used in [13, 14]). The discretization used is also illustrated in Fig. 14 and the resulting graph in Fig.15. This is a highly complex graph with 69 nodes and many cycles. It is not obvious what the search number of this graph is (a clearing search schedule with 5 searchers has been found in [14]).




Figure 14. Floorplan no.7, a real world museum. The discretization (which yields local visibility) is denoted by cell numbers.




Figure 15. The graph obtained from the discretization of Figure 14.

We have been able to clear the graph with 6 searchers. The parameters used are $L_0 = 5$, $v_0 = 0$, $\sigma = 1$; the clearing schedule has length 120 and was computed in 15 s.

4.8 Discussion

From the above experiments we see that that our graph search algorithm has cleared every floorplan studied, in most cases using the minimum number of searchers (for floorplans no.6 and 7 we do not know the actual minimum number of searchers but we know they can be cleared with less than the six searchers we have used). The computation of the clearing schedules usually takes a few seconds or less.

While we do not have an exact method to determine optimal parameter values, the general approach we have used is to start with small values of L_0, v_0, σ and increase these only in case a clearing schedule cannot be found with the smaller values. This approach usually yields a clearing schedule with a few trials. Note that increased values of L_0, v_0 require longer computation times.

When our algorithm gets stuck (i.e. when repeated iterations yield the same, nonzero minimum cost and the list of dirty sets does not change any longer) we have used the following heuristic: change the discretization so as to obtain slightly asymmetric graphs. This is done in an ad hoc manner. In general, the experiments indicate that the role of discretization can be quite significant: recall the examples of Sections 4.2 and 4.4, where an appropriate discretization reduces the search number; also recall that in Sections 4.3 and 4.5 a solution was obtained only after changing the discretization.

In general, the choice of discretization and algorithm parameters must be further investigated to establish principled approaches to these issues; the matter will be further discussed in Section 5.

The experiments also indicate the importance of recontamination. Most of the search schedules our algorithm computes are recontaminating and in at least some cases recontamination *does* help. For example, the graph (and floorplan) of Section 4.5 can only be cleared by three searchers if recontamination is used; without recontamination four searchers are necessary.

Finally, let us note that our algorithm does not address the issue of *time optimality*; in other words no effort is made to find the shortest possible clearing schedule. This is an issue we plan to address in the future.

5 Conclusion

We have presented an approach to robotic indoor PE which is based on (a) discretization of the indoor environment and conversion to a graph, (b) search of the graph. Let us emphasize that the two components (discretization and graph search) are quite independent. We have used a naive discretization into rectangular nonoverlapping cells and proposed a greedy iterative graph search algorithm.

The search algorithm searches the graph nodes in a non-teleporting and non-monotonic fashion; it can deal with non-local visibility and with an adversarial evader of finite or infinite speed. While the algorithm is not guaranteed to always find a solution, it has done so in all the example problems we have presented, including some quite complex ones. We have used the algorithm in conjunction with our naive discretization, but it can be used with any other discretization, for example the sophisticated one used by Lavalley et al. [18, 19, 20].

We believe our algorithm is superior to the ones previously proposed in the robotic literature [9, 18, 19, 20, 26] for guaranteed search; our algorithm has better scaling properties than the rest and hence can clear more complex floorplans (such as the office building of Section 4.6 and the museum of Section 4.7). We are aware of only one algorithm which is competitive to the one of the current paper: the unpublished algorithm appearing in [14], which can clear complex environments (such as the office building and the museum) with fewer searchers than ours. However this algorithm does not allow for recontamination, cannot handle variable evader speed and does not allow for extended visibility.

Our algorithm is also preferable to the ones appearing in the graph theoretic literature, at least as far as robotics problems are concerned. Indeed, as already remarked in the Introduction, graph theoretic search algorithms suffer from several limitations from the from the robotics point of view. They search edges (rather than nodes), they produce monotonic and teleporting search chedules, they

cannot handle variable evader speed (they always assume the evader is arbitrarily fast) and assume “local” visibility.

However we believe that the graph theoretic literature has much to offer to robotic PE: the current graph theoretic search algorithms can be used as starting points for the development of robotic PE algorithms and, perhaps most important, the theoretical issues addressed in graph theoretic search (monotonicity, computation of search number etc.) provide useful guidelines for the theoretical analysis of robotic PE, an area which (in our opinion) has not been sufficiently studied from the theoretical point of view (see item 7 below).

We conclude this paper by listing some research directions which we intend to pursue in the future.

1. Several practical modifications of the search algorithm are possible, for example the online modification of the parameters σ, v_0, L_0 depending on the evolution of the search. For example, to control the size of the list Q , we could start with small values of v_0, L_0 and increase them in case the algorithm gets stuck. A way to control the random component of the search by dynamically adjusting the σ value is also of interest.
2. Perhaps more important is the dynamic adjustment of the number of searchers K . A simple way to do this would be to start with a single searcher and use him until no further reduction of the dirty set is possible, at which time an additional searcher could be introduced. In this manner searchers could be introduced one by one until the graph is cleared. More sophisticated schemes for controlling the number of searchers may also be devised.
3. Discovering *shortest length* clearing schedules is of obvious importance.
4. We have treated discretization in an *ad hoc* manner, using hand-crafted discretizations. Optimal or, at any rate, principled discretization of the floorplan is of great interest, since appropriate discretization can reduce the number of pursuers necessary to clear the graph. A useful alternative would be an application which allows the user to *interactively* discretize a floorplan.
5. It is easy to modify the problem formulation to handle edge search (this is done by introducing “auxiliary” nodes, corresponding to edges). It will be interesting to apply our algorithm to edge search.
6. We can modify the state evolution equation to model *random evader movement*. A simple way to do this is to assume the evader is not adversarial, i.e. does not try to avoid the pursuers but instead performs a random walk on the graph. A cost function appropriate to this case is the *entropy* of the probability distribution of the evader’s location: $\hat{J} = \sum_{n=1}^N p_n \log p_n$. Our algorithm can use \hat{J} in place of J with minimal modifications. It will be useful to apply our algorithm to this probabilistics PE problem.
7. On a more theoretical note, practically all the graph theoretic work done so far considers the problem of the evader hiding in the *edges* of the graph. We have essentially introduced a *new* graph search problem, where the evader is assumed to be hiding in the *nodes* of the graph. Very little is known about the theoretical properties of this problem. For example, is edge search equivalent to node search? Is node search NP-complete (edge search is)? Many more theoretical questions arise and we are interested in investigating them.
8. Finally, we want to modify our algorithm to handle unknown or partially known environments.

References

- [1] B. Alspach. “Searching and sweeping graphs: a brief survey”. *Le Matematiche*.

- [2] L. Barriere, P. Flocchini, P. Fraigniaud and N. Santoro. “Capture of an intruder by mobile agents”. *Proceedings of the fourteenth annual ACM symposium on ...*, 2002.
- [3] L. Barriere, P. Fraigniaud, N. Santoro and D. Thilikos. “Searching is not jumping”. *Graph-theoretic concepts in computer science, Lecture Notes in Comput. Sci.*, vol. 2880, pp. 34–45, 2003.
- [4] L. Barriere. “Connected and Internal Graph Searching”. Tech Report TR-03-69, 2003.
- [5] H.L. Bodlaender. “A Tourist Guide through Treewidth”. *Acta Cybernetica*, 1993
- [6] H.L. Bodlaender. “Efficient and Constructive Algorithms for the Pathwidth and Treewidth of Graphs”. *J. of ALGORITHMS*, vol. 21, pp. 358-402, 1996.
- [7] E. W. Dijkstra. “A note on two problems in connexion with graphs”. *Numerische Mathematik*, vol.1, pp. 269–271, 1959.
- [8] F.V. Fomin and D. Thilikos. “An annotated bibliography on guaranteed graph searching”. <http://www.ii.uib.no/~fomin/fedor/articles1/abgs.pdf>, 2007.
- [9] B.P. Gerkey, S. Thrun and G. Gordon. “Visibility-based pursuit-evasion with limited field of view”. *Proc. of the Natl. Conf. on Artificial Intelligence (AAAI 2004)*, pp. 20-27, 2004.
- [10] B.P. Gerkey, S. Thrun and G. Gordon. “Parallel stochastic hill-climbing with small teams”. In *Multi-Robot Systems: From Swarms to Intelligent Automata III*, pp. 65-77, 2005.
- [11] B.P. Gerkey, S. Thrun and G. Gordon. “Visibility-based pursuit-evasion with limited field of view”. *Intl. Journal of Robotics Research*, vol. 25, pp. 299-316, 2006.
- [12] G. Hahn and G. MacGillivray. “A note on k -cop, l -robber games on graphs”. *Discrete Math.* vol. 306, pp. 2492–2497, 2006.
- [13] G. Hollinger, Ath. Kehagias and S. Singh. “Probabilistic strategies for pursuit in cluttered environments with multiple robots”. In *Proc. ICRA 2007*, 2007.
- [14] G. Hollinger. “Guaranteed search on spanning trees”, in preparation, 2008.
- [15] V. Isler, S. Kannan and S. Khanna. “Randomized Pursuit-Evasion in a Polygonal Environment”. *IEEE Transactions on Robotics*, vol. 5, pp. 864–875, 2005.
- [16] L. Kirousis and C. Papadimitriou. “Searching and pebbling”. *Theoret. Comput. Sci.*, vol. 47, pp. 205–218, 1986.
- [17] A.S. LaPaugh. “Recontamination does not help to search a graph”. *Journal of the ACM*, vol. 40, pp. 224-245, 1993.
- [18] S. M. LaValle, D. Lin, L. J. Guibas, J.-C. Latombe and R. Motwani. “Finding an unpredictable target in a workspace with obstacles”. *Proceedings IEEE International Conference on Robotics and Automation*, pp. 737–742, 1997.
- [19] L. J. Guibas, J.-C. Latombe, S. M. LaValle, D. Lin and R. Motwani. “Visibility-based pursuit-evasion in a polygonal environment”. *Lecture Notes in Computer Science*, vol. 1272, pp.17–30, 1997.
- [20] L. J. Guibas, J.-C. Latombe, S. M. LaValle, D. Lin and R. Motwani. “Visibility-based pursuit-evasion in a polygonal environment”. *International Journal of Computational Geometry and Applications*, vol. 9, pp. 471–494, 1999.

- [21] S.M. LaValle. *Planning Algorithms*. Cambridge University Press, 2008.
- [22] N. Megiddo, S.L. Hakimi, M.R. Garey, DS Johnson and C.H. Papadimitriou. “The complexity of searching a graph”. *Journal of the ACM*, vol. 35, pp. 18-44, 1988.
- [23] R. Nowakowski and P. Winkler. “Vertex-to-vertex pursuit in a graph”. *Discrete Math.*, vol. 43, pp. 235–239, 1983.
- [24] T.D. Parsons. “Pursuit-evasion in a graph”. *Theor. Appl. Graphs, Proc.* 1976.
- [25] D. Pellier and H. Fiorino. “Coordinated exploration of unknown labyrinthine environments applied to the Pursuit-Evasion problem”. AAMAS, pp.895-902, 2005.
- [26] A. Sarmiento, R. Murrieta, S.A. Hutchinson. “An Efficient Strategy for Rapidly Finding an Object in a Polygonal World”.