

The International Journal of Robotics Research

<http://ijr.sagepub.com/>

Surveillance Strategies for a Pursuer with Finite Sensor Range

Rafael Murrieta-Cid, Teja Muppirala, Alejandro Sarmiento, Sourabh Bhattacharya and Seth Hutchinson

The International Journal of Robotics Research 2007 26: 233

DOI: 10.1177/0278364907077083

The online version of this article can be found at:

<http://ijr.sagepub.com/content/26/3/233>

Published by:



<http://www.sagepublications.com>

On behalf of:



Multimedia Archives

Additional services and information for *The International Journal of Robotics Research* can be found at:

Email Alerts: <http://ijr.sagepub.com/cgi/alerts>

Subscriptions: <http://ijr.sagepub.com/subscriptions>

Reprints: <http://www.sagepub.com/journalsReprints.nav>

Permissions: <http://www.sagepub.com/journalsPermissions.nav>

Citations: <http://ijr.sagepub.com/content/26/3/233.refs.html>

Rafael Murrieta-Cid

Centro de Investigación en Matemáticas CIMAT
Guanajuato México

Teja Muppirala

Alejandro Sarmiento

Sourabh Bhattacharya

Seth Hutchinson

University of Illinois at Urbana-Champaign Urbana,
IL 61801 USA

Surveillance Strategies for a Pursuer with Finite Sensor Range

Abstract

This paper addresses the pursuit–evasion problem of maintaining surveillance by a pursuer of an evader in a world populated by polygonal obstacles. This requires the pursuer to plan collision-free motions that honor distance constraints imposed by sensor capabilities, while avoiding occlusion of the evader by any obstacle. The paper extends the three-dimensional cellular decomposition of Schwartz and Sharir to represent the four-dimensional configuration space of the pursuer–evader system, and derive necessary conditions for surveillance (equivalently, sufficient conditions for escape) in terms of this new representation. A game theoretic formulation of the problem is then given, and this formulation is used to characterize optimal escape trajectories for the evader. A shooting algorithm is proposed that finds these trajectories using the minimum principle. Finally, noting the similarities between this surveillance problem and the problem of cooperative manipulation by two robots, several cooperation strategies are presented that maximize system performance for cooperative motions.

KEY WORDS—surveillance strategies, motion planning, pursuit–evasion

1. Introduction

In this paper, we consider the basic surveillance problem of planning motions for a pursuer such that it maintains visibility of a moving evader in a known workspace containing obstacles. This is a variant of the traditional pursuit–evasion problem (discussed in Section 2) in which the pursuer has the goal

of either catching or finding the evader. Our problem has application in several domains. In security applications, it may be desirable for a robot sentry to surveil a moving evader (either another robot or a human) as the evader moves through some sensitive area. In non-security applications, it may be desirable for a robot to monitor the performance of another robot or of a human worker (e.g. a mobile robot might follow a highway crew as it makes road repairs, monitoring the quality of the repair work). Furthermore, as we will discuss below, our methods for planning surveillance strategies are also applicable to situations in which robots cooperate, such as shared manipulation or maintaining robot formations.

For our problem, we assume that the evader is initially positioned within the pursuer's field of view, at a distance L_{init} from the pursuer. The pursuer's goal is to maintain visibility of the evader while maintaining a surveillance distance $L_{\text{min}} \leq L \leq L_{\text{max}}$, in which L_{min} and L_{max} are parameters determined by the capabilities of the sensor and by any pursuer safety concerns. This requires the pursuer to plan collision-free motions that prevent occlusion of the evader by obstacles in the workspace, while maintaining the surveillance distance in the interval $[L_{\text{min}}, L_{\text{max}}]$. We assume that the pursuer is provided with a map of the workspace, and that the workspace is populated with polygonal obstacles.

In the case of a visibility based pursuit–evasion problem, it is pertinent to analyze the case of bounded surveillance distance for the following reasons. First, commercially available sensors (laser and cameras) have upper and lower range limits. In particular, if the evader is farther from the pursuer than a

maximal sensor range then its location is unknown, and surveillance is broken. The lower bound on the surveillance distance may be due to sensor capabilities, but more likely it will derive either from safety concerns (e.g. if the evader has the ability to harm the pursuer) or from a desire for the pursuer to remain undetected throughout the surveillance. Hence, if the evader is within this minimal range, even if the pursuer can partially infer the location of the evader, the evader has already a great advantage: the pursuer may not be able to detect the evader, and furthermore, the evader has the ability to harm the pursuer.

A great deal of related research exists in the area of pursuit and evasion, much of it from the dynamics and control communities, and we review the most relevant of this in Section 2. This past work typically does not take into account constraints imposed on pursuer motion due to the existence of obstacles in the workspace, or visibility constraints that arise due to occlusion. In this paper, we focus on these often neglected geometric aspects of the problem.

In Section 3 we define the configuration space for our pursuer–evader system and derive an efficient combinatoric representation for this configuration space. If the pursuer and evader are constrained to remain separated by a fixed distance, the motion planning problem for the pursuer–evader system is analogous to the problem of moving a rod in a plane. The basic problem of moving a rod in a plane was solved by Schwartz and Sharir (1987) using an elegant combinatoric representation of the set of free configurations. This representation can be extended to the case when the surveillance distance is allowed to vary by noting that the qualitative structure of this representation changes only at a finite set of critical values of the surveillance distance. Thus, we can represent the full four-dimensional configuration space of our pursuer–evader system by using a finite set of Schwartz and Sharir decompositions, provided we connect them appropriately. The description of this process concludes Section 3.

The representation developed in Section 3 leads to a necessary condition for surveillance (symmetrically, a sufficient condition for the evader to escape), expressed in terms of the qualitative decomposition of the configuration space. In particular, in Section 4 we derive conditions for the existence of a surveillance strategy for the case of a pursuer with unbounded velocity. If no such strategy exists, then there is no strategy for a pursuer with bounded velocity.

In Section 5, we formulate the surveillance problem using the language of noncooperative, dynamic game theory. Here, we assume that the pursuer and evader have bounded velocities.

Conditions for escape are then formulated in terms of the time required by the evader to reach a particular configuration, and the corresponding time required by the pursuer to reach a position from which it can maintain surveillance.

For the case of bounded velocities, we solve the case corresponding to the following assumptions: The pursuer–evader

system's motion takes place in free space up to the moment that the system reaches a set of escapable configurations and the pursuer maintains a minimal surveillance distance.

The formulation in Section 5 relies on the ability of the evader to plan optimal escape trajectories. Therefore, in Section 6 we derive the optimal escape trajectories for the evader using Pontryagin's minimum principle. In this section we propose a shooting algorithm to compute optimal escape trajectories, and show several results. For any such trajectory, it is straightforward for the pursuer to determine the required motion to maintain surveillance. If the duration of this motion exceeds the duration of the optimal escape trajectory, the surveillance is broken.

The two-robot pursuer–evader system becomes a team of two cooperating robots if the evader acts in concert with the pursuer. Therefore, in Section 7 we describe several cooperative motion strategies. These strategies could be used, for example, by two robots manipulating a rigid object. In this section, we consider leader–follower motions, in which the evader assumes the role of leader and the pursuer assumes the role of follower.

Finally, Section 8 provides conclusions and discusses future work.

2. Previous Work

Our surveillance problem is related to pursuit–evasion games. A great deal of research exists in the area of pursuit and evasion, particularly in the area of dynamics and control in free space (without obstacles) (Hajek 1965; Isaacs 1975; Basar and Olsder 1982). This work typically does not take into account constraints imposed on pursuer motion due to the existence of obstacles in the workspace, nor visibility constraints that arise due to occlusion.

Within the robotic planning community, several versions of the pursuit–evasion problem have been considered. One such problem is that of *finding* an evader with one or more mobile pursuers that sweep the environment so that the evader does not eventually sneak into an area that has already been explored. Exact (Parsons 1976; Suzuki and Yamashita 1992; Guibas et al. 1997; LaValle and Hinrichsen 1999, Shas et al. 2004) and probabilistic algorithms (Vidal et al. 2002; Hespanha et al. 2000; Isler et al. 2005) have been proposed to solve this problem. Another problem is to actually “catch” the evader, that is, to move to a contact configuration or closer than a given distance (Isaacs 1975; Isler et al. 2004).

These problems are related to, but not the same as ours. We assume that initially the pursuer can establish visibility with the evader. Our problem consists of determining a pursuer motion strategy to always *maintain* that visibility. The problem of maintaining visibility of a moving evader has been traditionally addressed with a combination of vision and control techniques (Espiau 1992; Hutchinson et al. 1996). Pure control

approaches, however, are local by nature, and do not take into account the global structure of the environment. Our interest is in deriving pursuer strategies that guarantee successful surveillance, taking into account both constraints on motion due to obstacles, and constraints on visibility due to occlusion.

Some previous work has addressed the motion planning problem for maintaining visibility of a moving evader. Game theory is proposed in LaValle et al. (1997) as a framework to formulate the tracking problem and an online algorithm is presented. In Becker et al. (1997), an algorithm is presented that operates by maximizing the probability of future visibility of the evader. This algorithm is also studied with more formalism in LaValle et al. (1997). This technique was tested in a Nomad 200 mobile robot with good results.

The work in Fabiani and Latombe (1999) presents an approach that takes into account the positioning uncertainty of the robot pursuer. Game theory is again proposed as a framework to formulate the tracking problem, and an approach is proposed that periodically commands the pursuer to move into a region that has no localization uncertainty (a landmark region) in order to re-localize and better track the evader afterward.

In Gonzalez et al. (2002), a technique is proposed to track an evader without the need of a global map. Instead, a range sensor is used to construct a local map of the environment, and a combinatoric algorithm is used to compute a differential motion for the pursuer at each iteration.

The approach presented in Murrieta et al. (2002) computes a motion strategy by maximizing the shortest distance to escape – the shortest distance the evader needs to move in order to escape the pursuer's visibility region. In this work the evaders were assumed to move unpredictably, and the distribution of obstacles in the workspace is assumed to be known in advance. This planner has been integrated and tested in a robot system that includes perceptual and control capabilities. The approach has also been extended to maintain visibility of two evaders using two mobile pursuers.

Recently, some research has considered the problem of maintaining visibility of several evaders with multiple robots. In Parker (2002) a method is proposed to accomplish this task in uncluttered environments. The objective is to minimize the total time during which evaders escape observation by some robot team member. In Jung and Sukhatme (2002) an approach is proposed to maintain visibility of several evaders using mobile and static sensors. A metric for measuring the degree of occlusion, based on the average mean free path of a random line segment, is used.

3. Configuration Space for the Surveillance Problem

If we model the pursuer and evader as points in a plane, then our surveillance problem is very similar to the problem of plan-

ning the motion of a variable-length rod in a plane. The instantaneous distance between the pursuer and evader corresponds to the length of the rod, and occlusion of the evader by an obstacle corresponds to collision of the rod with an obstacle. To maintain surveillance, it is necessary and sufficient that the line segment connecting the pursuer and evader does not intersect any obstacle in the environment. In addition to this occlusion constraint, a particular surveillance problem may impose bounds on the surveillance distance. For example, we may require that $L \leq L_{\max}$ if the sensor range is bounded by L_{\max} . Further, there may be some minimum allowable surveillance distance L_{\min} . This minimum distance could be due to sensor capabilities, but it is more likely to derive from safety concerns (e.g. if the evader has the ability to harm the pursuer) or from the pursuer's desire to remain undetected by the evader. We express this constraint by bounding the surveillance distance, $L_{\min} \leq L \leq L_{\max}$.

The problem of planning the motion of a fixed-length rod in the plane has been addressed in Schwartz and Sharir (1987) and Banon (1990). While this solution is not directly applicable to our surveillance problem, the representation introduced there can be extended to the case of a rod with variable length, and this extended representation provides the basis for the sufficient conditions for escape given in Section 4. In the remainder of this section, we define the configuration space for our problem, briefly review the method of Schwartz and Sharir (1987) and Banon (1990), and show how this representation can be extended to our surveillance problem.

The position of the evader is given by (x, y) , its coordinates in the plane, and the instantaneous distance between the pursuer and evader is given by L . We specify the position of the pursuer implicitly by giving the orientation θ of the line of sight from the evader to pursuer. For a given θ , the position of the pursuer is given by

$$\begin{bmatrix} x_p \\ y_p \end{bmatrix} = \begin{bmatrix} x \\ y \end{bmatrix} + L \begin{bmatrix} \cos \theta \\ \sin \theta \end{bmatrix} \quad (1)$$

Figure 1 illustrates these definitions. Note that the figure also illustrates the evader velocity, parameterized by v_e and α ; these will be used in Section 6.

With these definitions, the configuration of the pursuer–evader system is given by (x, y, θ, L) , and the configuration space is $\mathcal{Q} = SE(2) \times [L_{\min}, L_{\max}]$, which corresponds exactly to the case of a variable-length rod moving in the plane. Under this formulation, the evader controls (x, y) and the pursuer controls θ and L .¹ Thus, the task of the pursuer is to choose a trajectory $\theta(t), L(t)$ such that there is never a collision of the rod with an obstacle, and such that $L_{\min} \leq L(t) \leq L_{\max}$ for all t .

1. When the velocity of the pursuer is bounded, control of the parameter L is shared by the pursuer and evader, since in this case, the pursuer cannot arbitrarily effect changes in the surveillance distance.

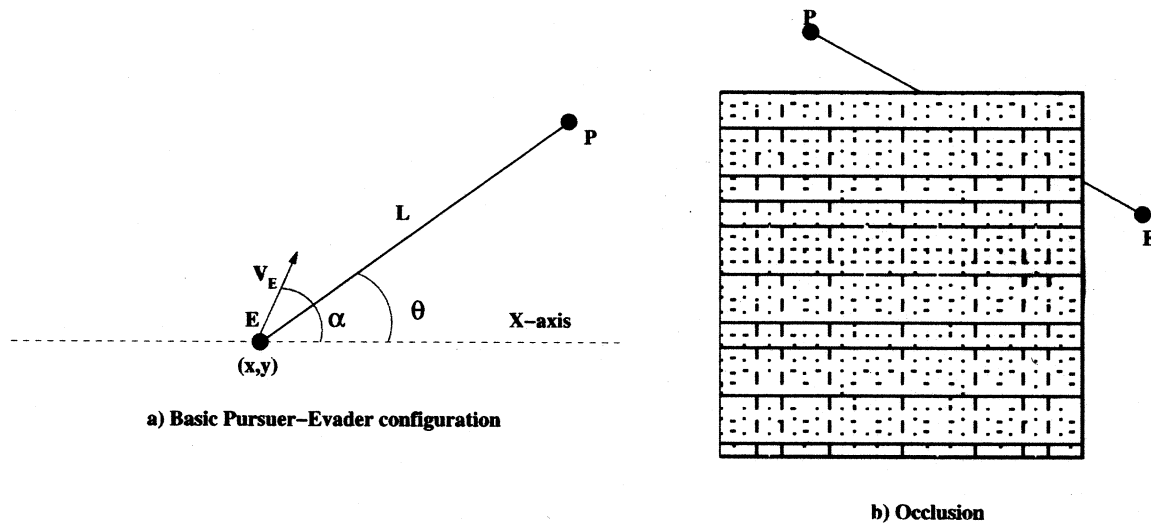


Fig. 1. Configuration variables for the pursuer–evader system

This problem lends itself to an elegant combinatoric representation. To explain this representation, we first describe the cylindrical decomposition given in Schwartz and Sharir (1987) for the case when L is fixed. Following this, we describe how the cylindrical decomposition can be extended to the case of variable L .

For a fixed surveillance distance L , we denote the corresponding three-dimensional slice of \mathcal{Q} by \mathcal{Q}_l . When building a representation of \mathcal{Q}_l , the key observation is that for a particular (x, y) , as θ varies, the rod will make contact with some fixed set of obstacle features (edges or vertices for the case of polygonal obstacles). If we vary (x, y) , this set of contact features will change only when (x, y) crosses a *critical curve*. Thus, the critical curves delineate *noncritical regions* in the plane such that, for a noncritical region R , the set of contact features remains constant for all $(x, y) \in R$. The set of critical curves is defined in terms of obstacle features (edges and vertices) and the length of the rod L . The catalog of critical curves is illustrated in Figure 2. More detailed explanations concerning their construction can be found in Schwartz and Sharir (1987) and Latombe (1991).

This partition of the plane into noncritical regions induces a cylindrical decomposition on \mathcal{Q}_l . For a noncritical region R we define the *cylinder above R* as the set of configurations $\{(x, y, \theta) \in \mathcal{Q}_l \mid (x, y) \in R\}$. In the cylinder above any noncritical region in the plane, there will be a set of simply connected cells, each of which lies either entirely in the free configuration space or entirely within the configuration space obstacle region.

For any collision free configuration (x, y, θ) in the cylinder above R , as the rod rotates it will make contact with two obsta-

cle features (called *stops*), one, say s_+ when θ increases, and the other, s_- , when θ decreases. The region R and these two features implicitly define a free cell κ in the cylinder above R : κ includes the set of configurations for which $(x, y) \in R$ such that increasing θ causes first contact with s_+ and decreasing θ causes first contact with s_- . For a fixed L , we denote by \mathcal{K}_l the decomposition $\{\kappa_i\}$.

Figure 3 shows the critical curves for a simple workspace consisting of a large free space and a long corridor. The noncritical regions are numbered 1–19. As an example, the cylinder above R_5 contains two free cells. We will adopt a simple naming scheme for the cells in this example: κ_{iL} denotes the free cell above region R_i for which the pursuer is in the corridor (generally to the left of the evader) and κ_{iR} denotes the free cell above region R_i for which the pursuer is not in the corridor (generally, but not always, to the right of the evader). In this simple example, there are at most two cells above any region. Thus, for region R_5 , the cell κ_{5L} includes all configurations (x, y, θ) such that $(x, y) \in R_5$ and increasing or decreasing θ causes collision with the walls in the corridor, while the cell κ_{5R} includes all configurations (x, y, θ) such that $(x, y) \in R_5$ and increasing or decreasing θ causes collision with the vertical walls bounding the free space. One configuration from each of these cells is illustrated in Figure 4. We can now represent the set of collision free configurations in \mathcal{Q}_l by a graph.

Definition: For a specific L , the connectivity graph $G_L = (V_L, E_L)$ has vertices that correspond to the cells $\kappa \in \mathcal{K}_l$ in the decomposition of \mathcal{Q}_l . An edge exists between vertices v and v' iff their corresponding cells κ and κ' are adjacent.




Fig. 2. The critical curves of Schwartz and Sharir (1987).




Fig. 3. A simple workspace and its critical curve decomposition.




Fig. 4. In the figure on the left, the rod's configuration is in cell κ_{5L} . In the figure on the right, the rod's configuration is in cell κ_{5R} .

Suppose two cells κ and κ' exist in the cylinders above regions R and R' , respectively, and that the regions R and R' are separated by the critical curve β . If κ and κ' are adjacent, any configuration in κ can be connected to any configuration in κ' by a free path whose projection onto the x - y plane crosses β transversally, with constant orientation in some neighborhood of the crossing point. The test for adjacency is a bit complicated, but it is formulated entirely in terms of the qualitative description of the cells (Schwartz and Sharir 1987; Latombe 1991). The complete graph for our example is illustrated in Figure 5. As an example, it is possible for the rod to move from κ_{5L} to κ_{4L} , but it is *not* possible to move from κ_{5L} to κ_{4L} .

To extend the representation of Schwartz and Sharir to the case of variable L , the following proposition is key.

Proposition 1: For $L \in [L_{\min}, L_{\max}]$, the structure of G_L changes only at a finite set of critical values $\{L_1, \dots, L_n\}$. If we order these values such that $L_1 < L_2 < \dots < L_n$, then G_L is constant for all $L \in (L_i, L_{i+1})$ for $i = 1, 2, \dots, n - 1$.

Proof: As L ranges from L_{\min} to L_{\max} , the critical curves move in the plane. The structure of G_L changes only when (a) a new critical curve appears, or (b) the qualitative structure of the intersection of two critical curves changes. In all other cases, the geometry of the noncritical regions and of the free cells may change, but the connectivity graph (which is derived from the qualitative structure of the decomposition) will not change.

For case (a), as L increases, new critical curves of type 3 and 5 will appear (critical curves of type 1, 2 and 4 exist for all values of L). By examining the catalog of critical curves shown in Figure 2, it can be seen that this will occur exactly when $L = d$, where d is the distance between two vertices for type 3 critical curves or the distance between a vertex and an edge for type 5 critical curves.

For case (b), as L increases, the critical curves move "away from" their corresponding obstacle features. In general, two critical curves will intersect for the first time when $L = d/2$, in which d is the distance between two obstacle features. For example, two critical curves of type 2 will intersect for the first time when $L = d/2$, where d is the distance between the two vertices. Likewise, a critical curve of type 1 will intersect a critical curve of type 2 for the first time when $L = d/2$, and d is the distance between the corresponding edge and vertex. The other combinations are similar, in that the critical value of L is uniquely determined for any two critical curves. ■

The main consequence of this proposition is that the entire four-dimensional configuration space of the pursuer-evader system can be represented using a finite collection of connectivity graphs $G_1 \dots G_n$. In particular, let $L_1 < L_2 < \dots < L_n$ be the sorted list of critical values of L , with $L_1 = L_{\min}$ and $L_n = L_{\max}$. Then $G_i = G_{L_i^q}$, where L_i^q is an arbitrarily chosen sample value such that $L_i < L_i^q < L_{i+1}$. All that remains is to add connections between the G_i to form the full connectivity graph for \mathcal{Q} .

Definition: The connectivity graph $G = (V, E)$ includes all vertices $V = \bigcup V_i$ and all edges $E = \bigcup E_i$. In addition, for vertices $v \in G_i$ and $v' \in G_{i+1}$ with corresponding noncritical regions R and R' and corresponding cells κ and κ' , an edge exists between v and v' iff $R' \cap R \neq \emptyset$ and the cell κ' is a free cell, i.e. if increasing the length of the rod without changing its position does not cause a collision.

In terms of this combinatoric representation of the surveillance problem, the evader controls the choice of x, y , and the pursuer controls the choice of θ and L (again, assuming for the moment that the pursuer velocity is unconstrained). With regard to the cylindrical decomposition of \mathcal{Q}_i and the

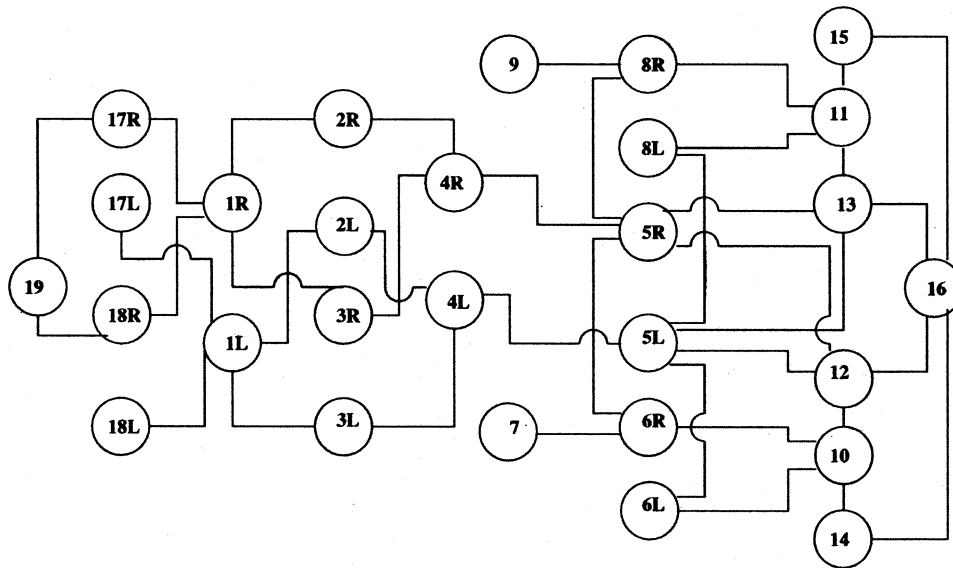


Fig. 5. Complete connectivity graph for the simple corridor example.

stratification induced by the G_i , the evader controls which cylinder in \mathcal{Q}_1 contains the system configuration, while the pursuer controls the choice of stratification G_i and the choice of which cell in the cylinder contains the system configuration.

Given this combinatoric representation for the surveillance problem, we will now establish sufficient conditions for escape.

4. Sufficient Conditions for Escape

In this section, we give a necessary and sufficient condition to escape based on the analysis of the four-dimensional configuration space of the system. Later, we impose a bounded speed constraint to both players. With the bounded speed constraint our necessary and sufficient condition to escape becomes only a sufficient condition, meaning that if this condition holds then the evader will escape. However, if this condition does not hold then the bounded speed constraint must be taken into consideration to know whether or not the pursuer is able to maintain surveillance of the evader.

It is possible to express a set of sufficient conditions for escape directly in terms of the connectivity graph G and the cylindrical decompositions computed for the L_i^σ . To derive these conditions, we first derive a sufficient condition for escape for a particular L_i^σ in terms of the graph G_i and the decomposition of $\mathcal{Q}_{L_i^\sigma}$. We then show that if an escape path exists for a particular value L_i^σ , then an escape path exists for all $L \geq L_i^\sigma$.

We begin by considering the case when L is fixed. Consider the example shown in Figures 3 and 5. Suppose the evader is in region R_6 . There are two cells in the cylinder above R_6 , cell κ_{6L} for which the pursuer is in the corridor, and cell κ_{6R} for which the pursuer is in a large empty space. If the current configuration is in cell κ_{6L} , then by moving to region R_7 the evader can break the surveillance (i.e. escape). There is no motion for the pursuer that prevents collision (the pursuer's line of sight to the evader is broken by a vertex). This can be seen from the connectivity graph in Figure 5. In the graph, there is no cell adjacent to κ_{6L} in the cylinder above R_7 . Thus, there is no free continuation into which the system can move if the evader chooses to move into region R_7 .

Suppose instead that the evader is in region R_5 and the system configuration is in cell κ_{5L} (i.e. the pursuer is in the corridor). It is not possible for the evader to break the surveillance by moving to an adjacent region.² For every region R' adjacent to R_5 , there is a cell in the cylinder above R' adjacent to κ_{5L} . The pursuer merely changes the orientation θ so that the system configuration moves into that free cell when the pursuer moves to region R' . However, the evader can move to region R_6 , forcing the system configuration into cell κ_{6L} , from which the evader can escape by moving to region R_7 . Thus, even though the evader cannot escape immediately, there is an escape path that the pursuer cannot prevent. This leads to the following recursive definition of an escapable cell.

2. The regions adjacent to R_5 are R_4, R_6, R_8, R_{12} and R_{13} . Note that regions R_7 and R_9 are not adjacent to R_5 , since the boundaries of R_5 and R_7 (resp. R_5 and R_9) share only a single point, an obstacle vertex v_2 (resp. v_1).

Definition: For the decomposition of \mathcal{Q}_1 corresponding to a particular sample value $L = L_i^q$, a cell $\kappa \in \mathcal{K}_1$ in the cylinder above a region R is an *escapable cell* if for some R' adjacent to R , in the cylinder above R'

- (a) there is no cell $\kappa' \in \mathcal{K}_L$ adjacent to κ , or
- (b) every cell $\kappa' \in \mathcal{K}_L$ that is adjacent to κ is an escapable cell.

If a cell satisfies condition (a) of the definition, the evader can escape immediately by moving to the adjacent region R' . This is the case for cells κ_{8L} and κ_{6L} of the example in Figures 3 and 5. If the current configuration lies in either of these cells, the evader can move into region R_9 or R_7 , respectively, to escape.

If condition (b) of the definition is satisfied, the evader can escape by moving through a sequence of adjacent regions, forcing the system configuration to move through a sequence of escapable cells, eventually reaching a cell that satisfies condition (a) of the definition. For the example of Figures 3 and 5, if the system is in configuration κ_{1L} , the evader can escape by following a path through regions R_3 , R_4 , R_5 , R_6 and R_7 . The pursuer has no choice but to follow, as the system configuration moves through cells κ_{3L} , κ_{4L} , κ_{5L} , κ_{6L} , and finally into collision when the evader moves into region R_7 . Thus, this sequence provides one guaranteed escape path for the evader.

By the definition above, the cell κ_{17L} is an escapable cell that satisfies both conditions (a) and (b) of the definition. If the evader is in region R_{17} and moves into region R_{19} , there is no motion possible for the pursuer that maintains the surveillance distance L . This illustrates a second way for the evader to “escape” – it can force the pursuer into a position from which it cannot maintain surveillance distance, or, the evader can “force the pursuer into a corner.” Both of these cases are taken into account by our definition of escapable cell.

As the above examples illustrate, for a fixed surveillance distance L_i^q , the evader can escape if at any time the system configuration enters an escapable cell. Thus, the pursuer must ensure that this does not occur. For the example of Figure 3, the pursuer strategy is simple: never enter the corridor, since doing so causes the system to enter some cell κ_{iL} , and no such cell exists in the reduced connectivity graph. This leads to the following sufficient condition for escape.

Proposition 2: For a fixed surveillance distance L_i^q , if there exists a region R such the cylinder above R contains only escapable cells, the evader can force the system into an escapable cell by moving to region R , from which there is a guaranteed escape path.

Proof: The proof is immediate from the recursive definition of escapable cells. ■

Comparing Figures 3 and 5, it can be seen that there is no such region for the simple corridor example. For every region R , there is some cell in the cylinder above R that is not an escapable cell. Thus, the pursuer can prevent escape. Suppose, however, that the simple corridor example is changed slightly, as shown in Figure 6. The connectivity graph is shown in Figure 7. Cells κ_{6L} , κ_{8L} , κ_{21R} and κ_{23R} satisfy condition (a) of the definition for escapable cells. By recursively applying the definition, it can be seen that every other cell in the connectivity graph satisfies condition (b) of the definition, i.e. the evader can escape regardless of which cell contains the initial configuration. The evader merely enters the corridor, and makes a sharp turn when exiting the other end of the corridor. For this example, there is no pursuer strategy that will maintain the surveillance.

Intuitively, planning collision-free paths for a variable length rod becomes more difficult as the length of the rod increases. Indeed, for a bounded workspace, there exists some length L' such that no free configuration exists for any $L > L'$. Likewise, for our surveillance problem the difficulty of maintaining surveillance increases as the surveillance distance increases. There are two reasons for this. First, as L increases the set of occlusion-free configurations becomes smaller (analogous to the set of free configurations for the rod becoming smaller as the length of the rod increases). Secondly, as L increases, the pursuer speed required to change θ increases. With respect to the sufficient conditions for escape, only the former is relevant, but as we will see later, for the task of planning optimal surveillance strategies for a pursuer with bounded velocity, the latter concern is pre-eminent.

Proposition 3: For a point (x, y) in the plane, if there exists an L such that in the decomposition \mathcal{K}_1 every cell in the cylinder above (x, y) is an escapable cell, then for every $L_i^q > L$, every cell in the cylinder above (x, y) in $\mathcal{K}_{L_i^q}$ is an escapable cell.

Proof: If the two cells κ_1 and κ_2 are contained in adjacent cylinders in \mathcal{K}_1 , then κ_1 and κ_2 are adjacent if and only if their ranges of free orientations overlap.³ The cell κ_1 is an escapable cell if no such κ_2 exists. For a particular (x, y) , the range of free orientations is nonincreasing with increasing L . Therefore, if the projection onto the plane of $\kappa_1 \in \mathcal{K}_1$ contains (x, y) , increasing the value of L will have the effect of deforming κ_1 , and the resulting cell will also be an escapable cell, since the range of free orientations for the deformed cell can be no greater than for the original cell, and thus no new adjacencies can arise. Note that for some values of L_i^q , the cell κ may actually split into multiple disconnected cells (each of which will have a range of free orientations that is no greater than the range for the original cell), and in this case, each of the resulting cells will be escapable cells. ■

3. Although the combinatoric representation described in Schwartz and Sharir (1987) is defined in terms of stops, the geometric condition for adjacency can be expressed in terms of overlapping ranges of free orientations.

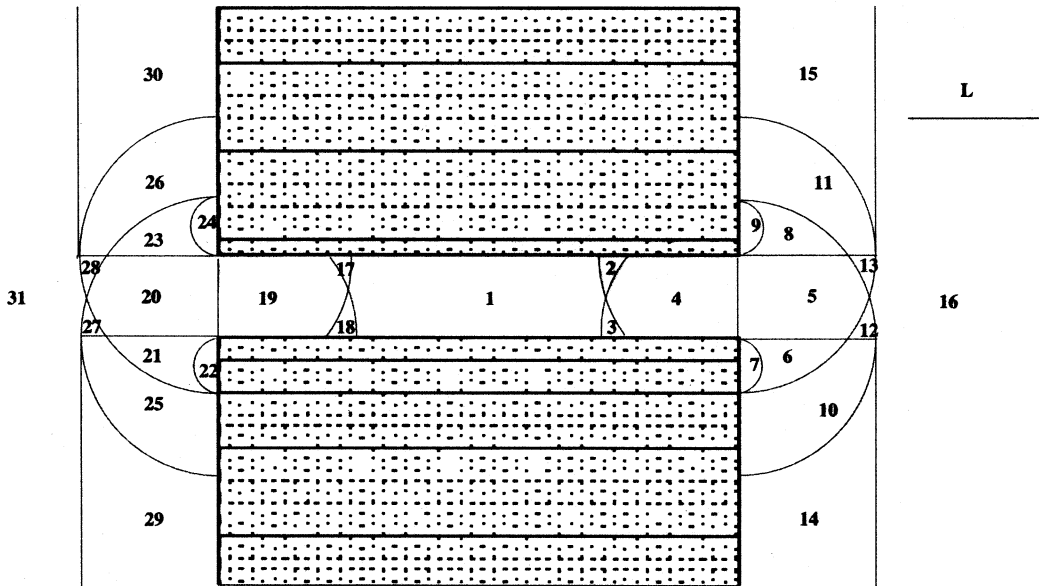


Fig. 6. Simple workspace in which a corridor connects two large empty spaces.

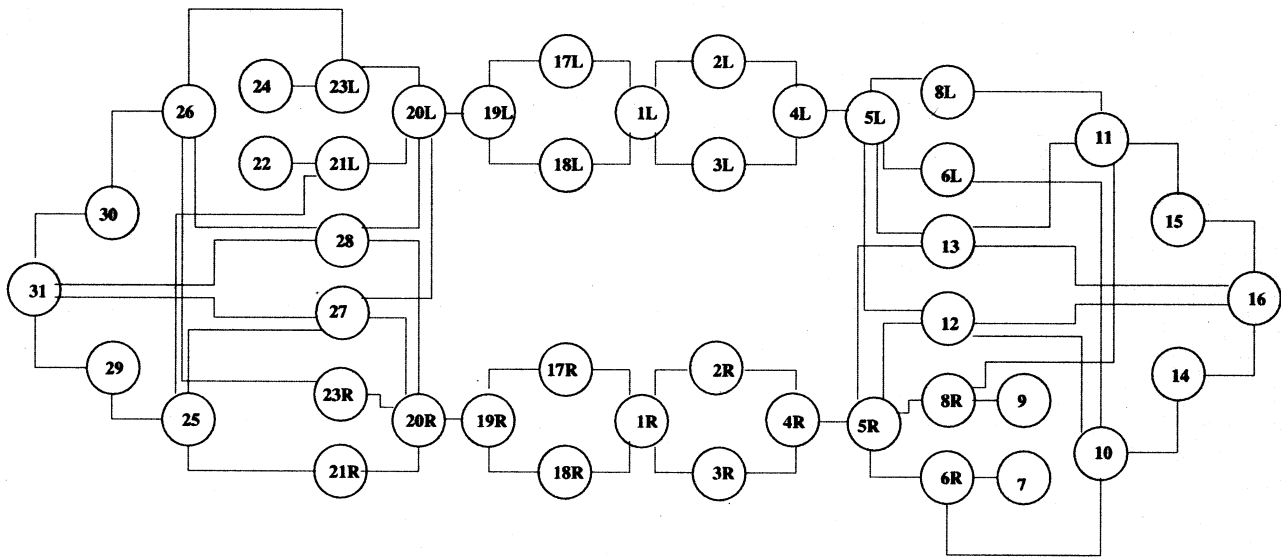
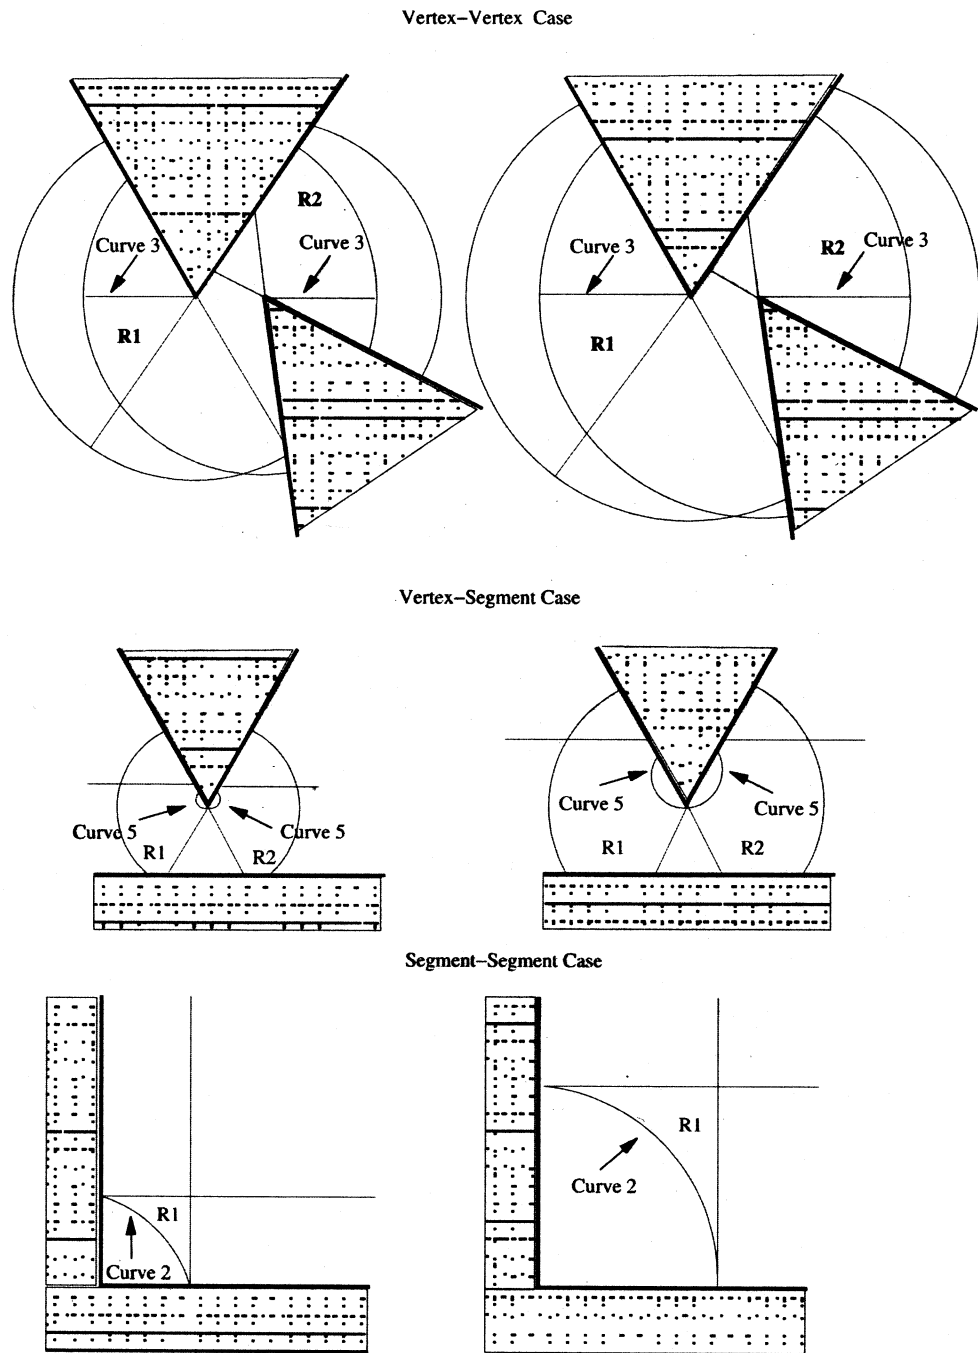


Fig. 7. Connectivity graph for the example of Figure 6.

Figure 8 illustrates three cases. In the top left example, an escapable cell exists in the cylinders over regions R_1 and R_2 . For example, if the evader is in region R_1 and the pursuer is in region R_2 , surveillance will be broken if the evader moves into the region directly above R_1 . This situation exists anytime there is a type 3 critical curve. It is easy to see that if a type

3 critical curve exists for some L' , then a type 3 critical curve will exist for all $L > L'$. In the example shown in the top right of the figure, it can be seen that when L is increased, escapable cells exist in the cylinders above regions R_1 and R_2 . The middle example in Figure 8 shows the similar case for a critical curve of type 5. The bottom example in Figure 8



Regions associated to escapable cells are labeled

Fig. 8. Bounded surveillance distance.

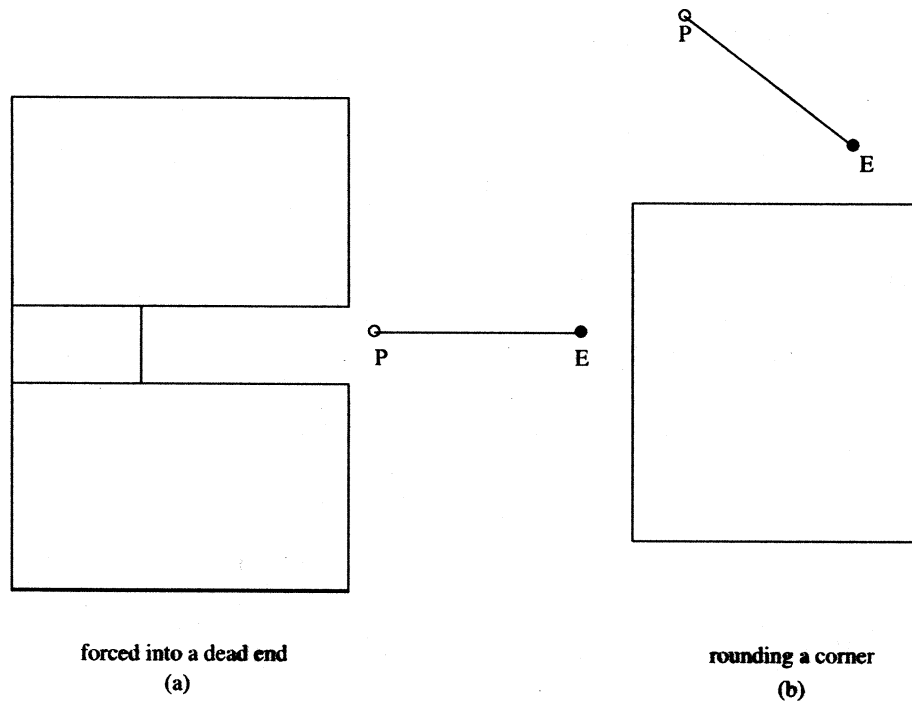


Fig. 9. If the pursuer is not able to change θ fast enough, motion of the evader directly toward the pursuer will force the system into an escapable cell (a). If the pursuer cannot move quickly enough, the evader can break the line of sight by forcing a collision of the rod with the corner (b)

illustrates the case for a type 2 critical curve. In this case, the escapable cell corresponds to forcing the pursuer into a corner.

The importance of Proposition 3 is that it allows us to express a sufficient condition for escape in terms of the combinatoric representation of \mathcal{Q} .

Proposition 4: For the decomposition $\mathcal{K}_{L_1^q}$ if there exists a region R such the cylinder above R contains only escapable cells, then there is no guaranteed surveillance strategy for the pursuer.

Proof: The proof follows immediately from the definition of escapable cells, Proposition 2 and Proposition 3, since $\mathcal{K}_{L_1^q}$ is the decomposition corresponding to the minimum surveillance distance. ■

If this sufficient condition is satisfied, then the evader can escape the surveillance regardless of the pursuer's capabilities. If the condition is not satisfied, it may still be possible for the evader to escape, depending on the constraints on the pursuer's velocity.

5. Escaping a Pursuer with Bounded Velocity

In this section, we consider the case of a pursuer with bounded velocity. The discussion above has assumed that the pursuer can move with unbounded velocity, but the sufficient condition for escape given above continues to hold when the pursuer velocity is bounded: if the evader can escape an infinitely fast pursuer, then of course it can escape a pursuer with finite speed. The concept of escapable cell remains key in our development, but when the pursuer's velocity is bounded, the pursuer must take more care in avoiding escapable cells; it must also take care to avoid collisions of the rod with obstacles in the workspace, even if the system configuration does not lie in an escapable cell.

Consider the example shown in Figure 9(a), for which $\theta = \pi$. If the evader moves directly toward the pursuer, the pursuer must change the orientation of the rod before the pursuer is forced into the corridor. If the pursuer cannot move fast enough to do so, the evader will have forced the system into an escapable cell, from which it has a guaranteed escape path.

In addition to avoiding escapable cells, when the velocity of the pursuer is bounded, it must also take care to avoid configurations from which the evader can force a collision of the rod with an obstacle. Consider the case illustrated in

Figure 9(b). If the pursuer cannot move quickly enough, the evader can break the line of sight by forcing a collision of the rod with the corner.

We can formulate this problem as a noncooperative dynamic game. Essentially, for a given configuration $q = (x, y, \theta, L)$, the evader's task is to find some location (x', y') in the plane, such that if the evader moves to (x', y') along an optimal trajectory (we define conditions for optimality in Section 6), the pursuer will not have sufficient time to change the orientation of the rod to avoid entering an escapable cell or to avoid a collision configuration.

In the previous section, it was seen that the difficulty of the surveillance problem increases as L increases. Above, we were concerned only with establishing necessary conditions for surveillance, but this same principle applies when planning surveillance strategies for a pursuer with bounded velocity. Assuming that the maximum evader speed is not greater than the maximum pursuer speed,⁴ the only escape strategy for the evader is to move in such a way as to occlude the pursuer's view. To prevent this occlusion the pursuer must change the orientation of the rod so that it does not contact an obstacle. Thus, the pursuer's effectiveness depends on the possible maximum value of $\dot{\theta}$. Using simple geometry, we have $L\dot{\theta} = \|V_p^\perp\|$, in which V_p^\perp is the component of the pursuer velocity that is perpendicular to the rod. Thus, we again see that as L increases, the surveillance task becomes more difficult, or alternatively, surveillance is the easiest when $L = L_{\min}$. For this reason, in the sequel we will specifically consider the case of $L = L_{\min}$, and we will derive pursuer strategies that attempt to maintain this minimum surveillance distance. In some cases, the pursuer may not be able to move fast enough to maintain this surveillance distance. In this case, if L exceeds L_{\max} , the evader escapes. Furthermore, since the pursuer's performance is maximized for $L = L_{\min}$, when we derive the optimal evader motion we will assume that $L = L_{\min}$ (a minimax type approach – the evader wishes to escape even when the pursuer's performance is maximal).

To formalize the problem, we define the following terminology. Let \mathcal{R} be the set of admissible positions in the plane for the evader (i.e. the collision-free points in the plane). Let $q = (x, y, \theta, L)$ be the current configuration, and $q' = (x', y', \theta', L')$ be a destination configuration. In this formulation, the evader will choose (x', y') and the pursuer will choose θ' and L' .

We define $t_e^*(q, q')$ to be the time required by the evader to follow its optimal path (as defined in Section 6) from (x, y) to (x', y') , and $t_p^*(q, q')$ to be the minimum time required by the pursuer to change the system orientation to θ' such that $q' \in \mathcal{Q}_{\text{free}}$, given that the evader follows its optimal path from

(x, y) to (x', y') . If we define the set of free orientations for a given $(x, y) \in \mathcal{R}$ to be

$$\Theta_f(x, y) = \{\theta \mid (x, y, \theta, L_{\min}) \in \mathcal{Q}_{\text{free}}\}$$

then the best choice for the pursuer is the free orientation that maximizes the difference between $t_e^*(q, q')$ and $t_p^*(q, q')$, and we define the corresponding evader loss (equivalently the pursuer gain) by

$$\mathcal{L}(q, (x', y')) = \sup_{\theta' \in \Theta_f(x', y')} t_e^*(q, q') - t_p^*(q, q') \quad (2)$$

The evader's task is to select the point (x', y') that minimizes the pursuer's gain,

$$\mathcal{L}^*(q) = \inf_{(x', y') \in \mathcal{R}} \mathcal{L}(q, (x', y')). \quad (3)$$

Note that in this formulation we have explicitly set $L = L_{\min}$. It is possible to derive analogous equations for the general case, but here we consider only the optimal pursuer strategy of maintaining the minimum surveillance distance.

With this formulation, the evader escapes when $\mathcal{L}^*(q) < 0$. This leads to the bounded-velocity analog of an escapable cell.

Definition: A configuration q is an *escapable configuration* if $\mathcal{L}^*(q) < 0$.

There are several limiting cases. If both (x, y) and (x', y') lie in the same region R , and R is a region with no stops (i.e. the rod can rotate freely), then $t_p^*(q, q') = 0$; the pursuer need make no effort to change the system orientation. If there is no collision free path for (x', y') , then we define $t_p^*(q, q') = \infty$. For the case of a pursuer with unbounded velocity, we have $t_p^*(q, q') = 0$ when q does not lie in an escapable cell, and $t_p^*(q, q') = \infty$ when q lies in an escapable cell and the evader follows a guaranteed escape path. In this case a configuration that lies in an escapable cell is trivially an escapable configuration.

For a given evader position (x, y) , $\mathcal{L}^*(q)$ depends not only on x and y , but also on the value of θ . This leads to the following definitions.

Definition: For a given evader position in the plane, the set of *admissible orientations* for the rod is given by

$$\Theta_{\text{adm}}(x, y) = \{\theta \mid q = (x, y, \theta, L_{\min}), \mathcal{L}^*(q) > 0\}$$

and a configuration $q = (x, y, \theta, L_{\min})$ is said to be an *admissible configuration* if $\theta \in \Theta_{\text{adm}}(x, y)$. We will denote by \mathcal{Q}_{adm} the set of admissible configurations. This leads to our sufficient condition for escape when the pursuer has bounded velocity.

Proposition 5: If there exists a point $(x, y) \in \mathcal{R}$ such that $\Theta_{\text{adm}}(x, y) = \emptyset$, then there is an escape path for the evader originating at (x, y) .

4. If this assumption does not hold, the evader has the trivial escape strategy of merely moving at its maximum speed away from the pursuer

If the condition of the proposition is not satisfied, then it is the task of the pursuer to keep the configuration of the rod in \mathcal{Q}_{adm} by choosing $\theta \in \Theta_{\text{adm}}(x, y)$ as the evader traces a path in the plane. In the next section, we describe the pursuer policy to keep $\theta \in \Theta_{\text{adm}}(x, y)$.

We solve the specific portion of the pursuit evasion game corresponding to the following assumptions: the rod's motion takes place in the free space up to the moment that the system reaches an escapable configuration, which corresponds to a contact with either an obstacle or an escapable cell. The pursuer maintains a minimal surveillance distance, and the set of escapable configurations is bounded by critical curves.

It is possible to formulate and solve this part of the whole game, even if the global trajectories of both evader and pursuer are unknown in advance. In other words, it is possible to deduce the local motion policies of both pursuer and evader and the required final rod orientation.

When the pursuer maintains the minimum surveillance distance, its trajectory can be decomposed into two components: one that is equal to the evader velocity and one that is perpendicular to the rod. This is true regardless of the evader's trajectory, and it does not require any cooperation on the part of the evader (i.e. it applies to the case of an antagonistic evader). In this case, the evader trajectory can influence the rate of change in θ , but not the initial and final orientations. Thus, the policy for the pursuer to change the rod from an initial orientation to a final one (avoiding an escapable configuration) will be the same regardless of the actual evader trajectory.

Since we know in advance both the map and the maximum velocities of the evader and the pursuer, then the escapable configurations are also known. If there is more than one set of escapable configurations then the evader can choose the optimal one with respect to a given criterion related to the capabilities of the pursuer. Thus, knowledge of the optimal final rod orientation can be deduced from the map.

The task of the evader is to find the optimal motion strategy to escape. In the section 6, we describe the proposed optimality criterion. We also explain the validity of this optimization criterion.

6. Optimal Escape Trajectories for the Evader

In the optimization of Equation (3), for a given initial configuration $q = (x, y, \theta, L)$ and for a specified point (x', y') in the plane, the evader should take the path from (x, y) to (x', y') that minimizes amount by which the pursuer can change the orientation of the rod in the time required for the evader to reach the escapable configuration.

If the evader's goal was to minimize the maximum angular speed of the rod (maximizing the time taken for the pursuer to change the rod orientation) then the evader should move directly toward or directly away from the pursuer. This evader policy is described in detail in the next section (6.1). However,

this evader motion may not reach an escapable configuration. On the other hand, a straight-line evader path minimizes the time to reach an escapable system configuration. This path is studied in section 7.2. But the straight-line evader path is not optimal with respect to the criterion of minimizing the change in θ that can be effected by the pursuer.

Because of the kinematic constraints (bounded speeds and surveillance distance), there is a trade-off between minimizing the time taken for the evader to reach the escapable configuration and maximizing the time taken for the pursuer to change the rod orientation. The solution of the game can be established as a function of the rod orientation. If the pursuer can change the rod angle to one corresponding to an admissible orientation then the pursuer wins. Thus, the key concern for the evader is to minimize this change of angle. We now formulate and solve this optimization problem.

Without loss of generality, suppose that the evader is located at the origin, and $(x', y') = (x_e, 0)$, i.e. the evader is moving to a point on the x -axis. In the development that follows, we first derive an expression for the rate at which the pursuer can change the orientation of the rod. Following this, we formulate the optimization problem using Pontryagin's minimum principle.

6.1. Maximum Rate of Orientation Change

We assume that both the pursuer and evader move at their maximum velocities. In this case, we can parameterize the evader's velocity by α , the angle from the x -axis to the evader velocity vector. This is illustrated in Figure 1. We denote by V_p and V_e the upper bounds on the norm of the pursuer and evader velocities, respectively. Using these conventions we have

$$\begin{bmatrix} \dot{x} \\ \dot{y} \end{bmatrix} = V_e \begin{bmatrix} \cos \alpha \\ \sin \alpha \end{bmatrix} \quad (4)$$

and the derivative of Equation (1) is therefore given by

$$\begin{bmatrix} \dot{x}_p \\ \dot{y}_p \end{bmatrix} = V_e \begin{bmatrix} \cos \alpha \\ \sin \alpha \end{bmatrix} + L\dot{\theta} \begin{bmatrix} -\sin \theta \\ \cos \theta \end{bmatrix} \quad (5)$$

Since the optimal strategy for the pursuer is to choose $L = L_{\min}$, we derive here the evader strategy corresponding to this choice. Thus, in the derivations that follow, we assume that $L = L_{\min}$ is a constant. If the escape strategy succeeds for $L = L_{\min}$, then it succeeds for any $L > L_{\min}$. Computing the squared norm of each side of Equation (5) gives

$$V_p^2 = L^2\dot{\theta}^2 + 2V_eL(\sin \alpha \cos \theta - \cos \alpha \sin \theta)\dot{\theta} + V_e^2$$

$$0 = L^2\dot{\theta}^2 + 2V_eL \sin(\alpha - \theta)\dot{\theta} + V_e^2 - V_p^2$$

and solving for $\dot{\theta}$ using the quadratic formula we obtain

$$\dot{\theta} = \frac{V_p}{L} \left(-\rho \sin(\alpha - \theta) \pm \sqrt{1 - \rho^2 \cos^2(\alpha - \theta)} \right) \quad (6)$$

in which $0 < \rho \leq 1$ is the ratio of V_e to V_p . Note, if $\rho > 1$, it is trivial for the evader to escape.

To simplify notation, define

$$\dot{\theta}^+ = \frac{V_p}{L} \left(-\rho \sin(\alpha - \theta) + \sqrt{1 - \rho^2 \cos^2(\alpha - \theta)} \right)$$

$$\dot{\theta}^- = \frac{V_p}{L} \left(-\rho \sin(\alpha - \theta) - \sqrt{1 - \rho^2 \cos^2(\alpha - \theta)} \right)$$

The decision of whether to use $\dot{\theta}^+$ or $\dot{\theta}^-$ is made by the pursuer, based on whether the pursuer wishes to maximize $|\dot{\theta}|$, to have $\dot{\theta} < 0$, or to have $\dot{\theta} > 0$. Since $0 < \rho \leq 1$,

$$|\rho \sin \phi| \leq \sqrt{1 - \rho^2 \cos^2 \phi}$$

and therefore $\dot{\theta}^+ > 0$ and $\dot{\theta}^- < 0$.

The magnitude of the angular change $|\dot{\theta}|$ is maximized by choosing $\dot{\theta}^+$ when $\sin(\alpha - \theta) < 0$, and choosing $\dot{\theta}^-$ when $\sin(\alpha - \theta) > 0$. When $\sin(\alpha - \theta) = 0$, either choice will yield

$$|\dot{\theta}| = + \frac{V_p}{L} \sqrt{1 - \rho^2 \cos^2(\alpha - \theta)}$$

The maximum angular speed is then given by

$$\max |\dot{\theta}| = \begin{cases} |\dot{\theta}^-| & 0 \leq \alpha - \theta \leq \pi \\ \dot{\theta}^+ & \pi \leq \alpha - \theta \leq 2\pi \end{cases}$$

In the interval $0 < \alpha - \theta < \pi$, the function $|\dot{\theta}^-|$ is concave. Likewise, in the interval $\pi < \alpha - \theta < 2\pi$, the function $\dot{\theta}^+$ is concave. Thus, the minimum is achieved at the transition point, where $|\dot{\theta}^-| = \dot{\theta}^+$, and where the derivative of $\max |\dot{\theta}|$ is discontinuous, i.e. for $(\alpha - \theta) = k\pi$. Therefore, if the evader wishes to choose a velocity that minimizes the maximum angular speed of the rod, the evader should move directly toward or directly away from the pursuer.

6.2. Formulating the Optimal Control Problem

In formulating the optimal control problem, it is convenient to let x be the independent variable (instead of time t) and to parameterize the control α by x , since the evader will move from the point $x = 0$ to the point $x' = x_e$. It is clear that the projection of the evader position onto the x -axis will be monotonically increasing, so this is a reasonable choice for a parameterization of the problem. This leads to a modified system description in which the state of the rod is given by $\zeta = (y, \theta)$, and the system equation is given by

$$\frac{d}{dx} \zeta = f(\zeta, \alpha)$$

The time derivatives for x and y are given by

$$\frac{dx}{dt} = V_e \cos \alpha$$

$$\frac{dy}{dt} = V_e \sin \alpha$$

With respect to the independent variable x , the state equations are then given by

$$\frac{dy}{dx} = \frac{dy}{dt} \left(\frac{dx}{dt} \right)^{-1} = \tan \alpha = f_1(\zeta, \alpha)$$

$$\begin{aligned} \frac{d\theta}{dx} &= \frac{d\theta}{dt} \left(\frac{dx}{dt} \right)^{-1} \\ &= \frac{-\rho \sin(\alpha - \theta) \pm \sqrt{1 - \rho^2 \cos^2(\alpha - \theta)}}{L \rho \cos \alpha} \\ &= f_2(\zeta, \alpha) \end{aligned}$$

In the remainder of this section, to simplify notation, we will use $\dot{\theta}$, $\dot{\lambda}$, and $\dot{\zeta}$, to denote the derivatives of θ , λ and ζ with respect to x rather than with respect to time. Further, in this section we will only solve the problem for the case $\dot{\theta} = \dot{\theta}^+$. The case for $\dot{\theta} = \dot{\theta}^-$ is analogous.

To minimize the angle by which the pursuer can change the orientation of the rod, we use the cost function

$$\begin{aligned} J &= \int_0^{x_e} \frac{d\theta}{dx} dx \\ &= \int_0^{x_e} f_2(\zeta, \alpha) dx \end{aligned}$$

Thus, we seek the control input α that minimizes J , subject to the system equation $\dot{\zeta} = f(\zeta, \alpha)$, and subject to the boundary conditions

$$\theta(0) = \theta, \quad y(0) = y(x_e) = 0$$

To solve this optimization problem, we use the minimum principle and solve the resulting equations numerically using a shooting method.

The system Hamiltonian is given by

$$\begin{aligned} \mathcal{H}(\zeta, \alpha, x) &= \frac{d\theta}{dx} + \lambda_1(x) f_1(\zeta, \alpha) + \lambda_2(x) f_2(\zeta, \alpha) \\ &= \lambda_1 f_1(\zeta, \alpha) + (1 + \lambda_2) f_2(\zeta, \alpha) \\ &= \lambda_1 \tan \alpha + (1 + \lambda_2) \\ &\quad \times \frac{-\rho \sin(\alpha - \theta) + \sqrt{1 - \rho^2 \cos^2(\alpha - \theta)}}{L \rho \cos \alpha} \end{aligned}$$

in which $\lambda_i(x)$ are the Lagrange multipliers, which are a function of the independent variable x .

Algorithm OptimalEvasion (θ, x_e)

1. Choose initial values for the Lagrange multipliers.
2. Let $j \leftarrow 0, \theta(0) \leftarrow \theta, y(0) \leftarrow 0, x(0) \leftarrow 0$.
3. Choose $\alpha(j)$ that minimizes the Hamiltonian,

$$\alpha(j) = \arg \min \mathcal{H}(\zeta(j), \alpha, x(j))$$
4. Integrate the state equations to determine $\zeta(j + 1)$

$$y(j + 1) \leftarrow y(j) + f_1(\zeta(j), \alpha(j))\Delta x$$

$$\theta(j + 1) \leftarrow \theta(j) + f_2(\zeta(j), \alpha(j))\Delta x$$
5. Integrate the adjoint equation for λ_2 to determine $\lambda_2(j + 1)$.
6. $x(j + 1) \leftarrow x(j) + \Delta x$.
7. If $y(j)$ and $y(j + 1)$ have different sign, then the system has crossed the x axis.
 If $|x(j + 1) - x_e| < \epsilon$, then the optimal trajectory is given by $\alpha(0), \dots, \alpha(j)$.
8. If $y(j)$ and $y(j + 1)$ have different sign but $|x(j + 1) - x_e| > \epsilon$, then we have missed the boundary condition. In this case, adjust the initial values for λ and go to step 2.
9. If $y(j)$ and $y(j + 1)$ have the same sign, then we have not crossed the x -axis, and we continue to iterate forward: $j \leftarrow j + 1$, go to step 3.

Fig. 10. Algorithm to find the optimal evader trajectory.

The adjoint equation for λ_1 is given by

$$\dot{\lambda}_1 = -\frac{\partial}{\partial y} \mathcal{H} = 0$$

which implies that λ_1 is a constant. This allows us to write the Hamiltonian as

$$\mathcal{H}(\zeta, \alpha, x) = K \tan \alpha + (1 + \lambda_2) \times \frac{-\rho \sin(\alpha - \theta) + \sqrt{1 - \rho^2 \cos^2(\alpha - \theta)}}{L \rho \cos \alpha}$$

The adjoint equation for λ_2 is given by

$$\dot{\lambda}_2 = -\frac{\partial}{\partial \theta} \mathcal{H} = -(1 + \lambda_2) \times \frac{\cos(\alpha - \theta)}{L \cos \alpha} \left(1 - \frac{\rho \sin(\alpha - \theta)}{\sqrt{1 - \rho^2 \cos^2(\alpha - \theta)}} \right)$$

We find the optimal evader trajectory α^* iteratively using the algorithm shown in Figure 10. Figures 11 and 12 show examples of optimal evader trajectories and the resulting pursuer trajectories. Values for the Lagrange multipliers and x_e are also shown in each figure. In each of these examples, the angle θ is increasing (the rod rotates in a counterclockwise direction), and thus the evader trajectory minimizes $\theta = \theta^+$.

Figure 11(a) shows an example for which the pursuer and the evader are moving in the same direction, and therefore the pursuer must “catch up” to the evader. Figure 11(b) shows an example for which the pursuer and the evader are moving in opposite directions. This could result, for example, from the evader trying to force the pursuer into a corner. For the example of Figure 12 there is no fixed escape point (reflex vertex, those with internal angle bigger than π). In this case, the final value of y is unconstrained and the evader merely chooses its instantaneous velocity to minimize the instantaneous value of θ .

The choice of the endpoint for the evader trajectory is a function of the nearest escape point. We consider two cases: the escape point is a reflex vertex of an obstacle in the workspace, or the escape point lies on a critical curve that bounds a region whose cylinder contains an escapable cell.

Reflex vertex: Two critical curves of type 4 (line segments) emerge from a reflex vertex. Let us call θ_i the angle of the line segment (critical curve) having the smallest angle of the two segments, and θ_f the angle of the other line segment. If the evader is at the reflex vertex and the rod is not at an orientation bounded by θ_i and θ_f then the pursuer would need unbounded speed to prevent the evader from escaping. Thus, the admissible rod orientations are bounded by these two orientations. The smallest change in the rod orientation in order to bring the rod to an admissible orientation corresponds to either θ_i or θ_f . The reflex vertex location defines the final boundary con-

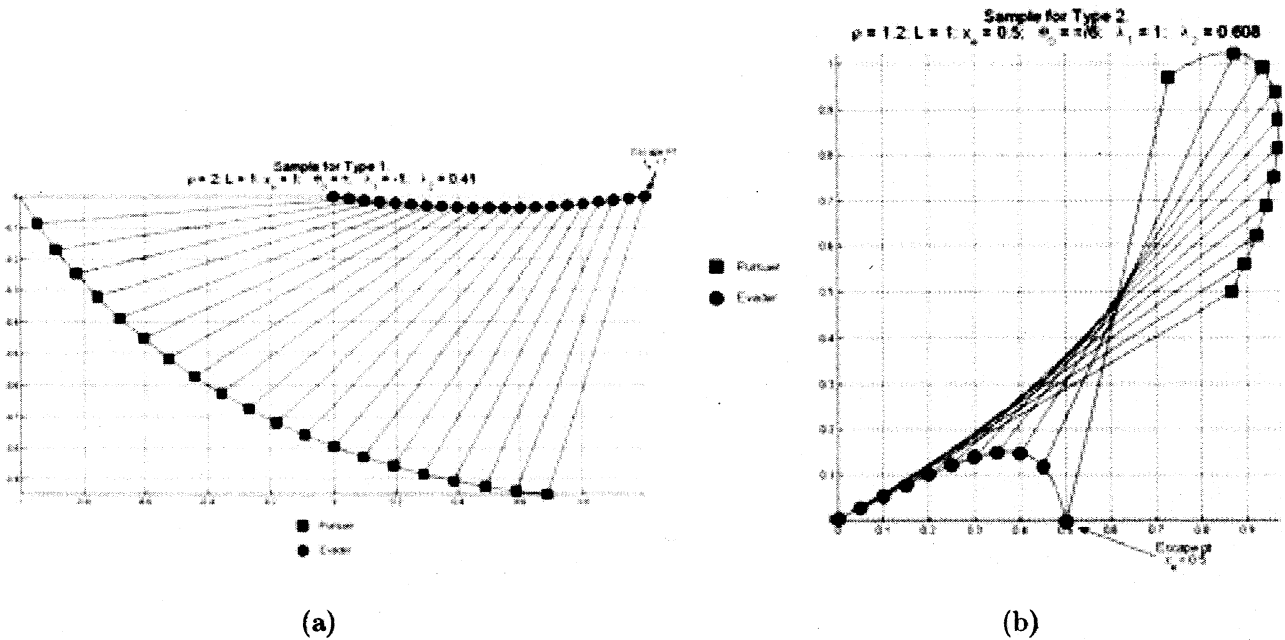


Fig. 11. (a) The goal is given by $x_e = 1$ and the initial values are given by $\theta_0 = \pi$, $\lambda_1 = -1$, and $\lambda_2 = 0.41$. (b) The goal is given by $x_e = 0.5$ and the initial values are given by $\theta_0 = \pi/6$, $\lambda_1 = 1$, and $\lambda_2 = 0.608$.




Fig. 12. In this case, the goal configuration is unspecified, and the evader chooses its instantaneous velocity to minimize the instantaneous rate of change in θ .

dition of the differential equations modeling the rod motion. This case is illustrated on the left in Figure 13. The evader is denoted with an E and the pursuer with a P. The critical curves are shown in dashed lines. Note that some critical curves are omitted for clarity.

A critical curve bounding an escapable cell: Analogous to the reflex vertex case, the orientations associated to an es-

capable cell are bounded by two angles θ_i and θ_f . If the evader location is in the region associated to the escapable cell and the rod is at an orientation within θ_i and θ_f then the evader escapes. In this case, the evader optimal trajectory is generated having the origin at the evader initial position and defining the x -axis perpendicular to the critical curve bounding the escapable cell. This x -axis orientation corresponds to the smallest distance trajectory between the evader and the critical curve, thus geometrically satisfying optimal time to reach the escapable cell. The initial conditions are the rod initial angle and the evader location. The point where the trajectory hits the critical curve is the endpoint of an optimal escape trajectory. This case is illustrated on the right in Figure 13.

Under the assumptions that we have presented in Section 5, the pursuer motion strategy is the same regardless of the evader trajectory. Note that, if the evader changes its mind and it takes another trajectory then it does not affect the pursuer's policy. If initially, the evader cannot escape by taking its optimal trajectory then it cannot escape with any other trajectory. The pursuer should still apply the same motion strategy. Note that given the invariance of this strategy it captures a large number of situations to avoid the evader escaping.

With the bounded speed constraint imposed over the evader and the pursuer, there are cases where the evader can escape pursuer surveillance if the rod moves in contact with the obstacles. We give in Sections 7.3 and 7.4 the general modeling




Fig. 13. The endpoints of optimal escape trajectories

of these motions. However, since we consider that in these cases, the evader and pursuer become a cooperative team, we have not analyzed the evader motion strategy needed to escape.

7. Leader-follower Cooperative Motion Plans

Although the development above has dealt with the case of an antagonistic evader that tries to escape surveillance, much of the analysis applies also to cooperative motion. Examples of such motions include cooperative manipulation of a rigid object or performing coverage tasks for which the distance between robots must be held at a constant value as the robots cover terrain, maintaining robot formations (our result applies to the special case of two-robot formations), and even planning motions of transportation systems such as trailers (although these systems include nonholonomic constraints in addition to the constraints that we consider; our work provides performance bounds on these systems). In such cases, it is more intuitive to refer to the evader as the *leader* and to the pursuer as the *follower*. In this section we describe four such cooperative motions: motions that minimize the follower's effort, motions in which the leader follows a straight-line path to the goal, motions for which the follower follows a straight line, and motions for which the follower and leader round a corner in a cooperative fashion.




Fig. 14. The Tractrix curve

7.1. Passive Pursuer Trajectories

If the leader is moving in a straight line, the follower can minimize its expended energy by following the path that it would take if it were "dragged by a string" held attached to the leader. This problem was first solved by Leibniz, and the resulting trajectory for the evader is known as the Tractrix curve (Mathworld; Lawrence 1971) (also called the hound curve). See Figure 14. The parametric equations that determine the follower position are

$$x_p(t) = t - \tanh t; \quad y_p(t) = \frac{1}{\cosh t}$$

7.2. Straight-line Path for the Leader

Suppose that the leader takes a straight line path to its goal position. What is the minimum time required by the follower to reach its goal position while maintaining the distance constraint? To formalize this question, we consider the case in which the leader moves in a straight line (i.e. α is a constant), and the follower must change the orientation of the system from θ_0 to θ_f . We then solve for the minimum time required by the follower to achieve this change.

The time derivative of θ is given by Equation (6). To compute the time T required to change the system orientation by $\Delta\theta$, we integrate dt over the required motion

$$\frac{dt}{d\theta} = \frac{L}{V_p - \rho \sin(\alpha - \theta) \pm \sqrt{1 - \rho^2 \cos^2(\alpha - \theta)}}$$

$$T = \int dt$$

$$= \frac{L}{V_p} \int_{\theta_0}^{\theta_1} \frac{d\theta}{-\rho \sin(\alpha - \theta) \pm \sqrt{1 - \rho^2 \cos^2(\alpha - \theta)}}$$

After some algebra, we obtain

$$T = \frac{L}{V_p} \frac{\rho}{1 - \rho^2} \int_{\theta_0}^{\theta_1} \sin(\alpha - \theta) d\theta$$

$$\pm \frac{L}{V_p} \frac{1}{1 - \rho^2} \int_{\theta_0}^{\theta_1} \sqrt{1 - \rho^2 \cos^2(\alpha - \theta)} d\theta$$

Under the change of variables $\phi = \frac{\pi}{2} - (\alpha - \theta)$ the second integral can be written as

$$\int_{\phi_0}^{\phi_1} \sqrt{1 - \rho^2 \sin^2(\phi)} d\phi \tag{7}$$

This integral is similar to the incomplete elliptic integral of the second kind (Abramowitz and Stegun 1964), which is defined as

$$E[\phi_1 | \rho^2] = \int_0^{\phi_1} \sqrt{1 - \rho^2 \sin^2(\phi)} d\phi$$

This integral corresponds to the arc length of a sector of an ellipse from 0 to ϕ_1 (the amplitude), and in which ρ^2 is called the elliptic module. Many modern mathematical programming languages (e.g. Mathematica (Mathworld)) provide a built-in function to solve this integral. The integral in Equation (7) can now be expressed as the difference

$$T = \int_{\phi_0}^{\phi_1} \sqrt{1 - \rho^2 \sin^2(\phi)} d\phi = E[\phi_1 | \rho^2] - E[\phi_0 | \rho^2]$$

As always, the choice of positive or negative sign for the second integral gives rise to two possible solutions, depending on whether the rod rotates in the counterclockwise or clockwise direction.

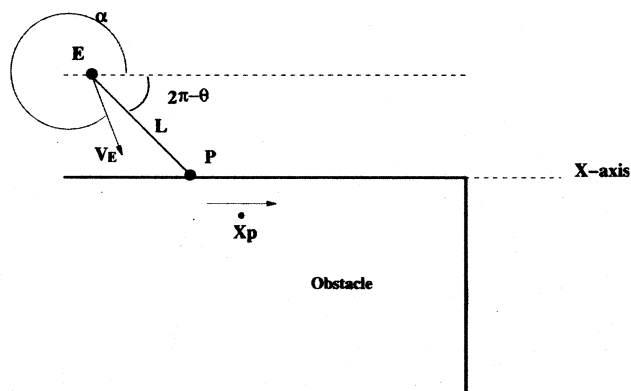


Fig. 15. In this figure, the pursuer attempts to follow the wall. Depending on the values of ρ and θ , it may or may not be possible for the follower to effect this strategy.

7.3. Wall Following

For this motion, the follower obeys a *wall following* strategy as the leader moves, possibly toward the wall, possibly parallel to the wall. Without loss of generality, assume that the wall lies along the x -axis so that wall following is achieved by the follower velocity $(\dot{x}_p, 0)$. As above, assume the leader moves with maximum velocity in direction α . In this case, Equation (5) can be written as the two equations

$$\dot{x}_p = V_e \cos \alpha - \dot{\theta} L \sin \theta$$

$$0 = V_e \sin \alpha + \dot{\theta} L \cos \theta$$

Solving the second equation for $\dot{\theta}$ yields

$$\dot{\theta} = -\frac{V_e \sin \alpha}{L \cos \theta}$$

and substituting this into the above equation for \dot{x}_p gives

$$\dot{x}_p = V_e \cos \alpha + V_e \sin \alpha \tan \theta$$

To find the leader trajectory that maximizes the required follower speed, we compute and set to zero the derivative of this equation with respect to α and obtain

$$0 = -V_e \sin \alpha + V_e \cos \alpha \tan \theta$$

$$\tan \alpha = \tan \theta$$

$$\alpha = \theta + k\pi, \quad k = 1, 2, \dots$$

i.e. if the leader moves directly along the follower's line of sight the result is to maximize the required follower velocity.

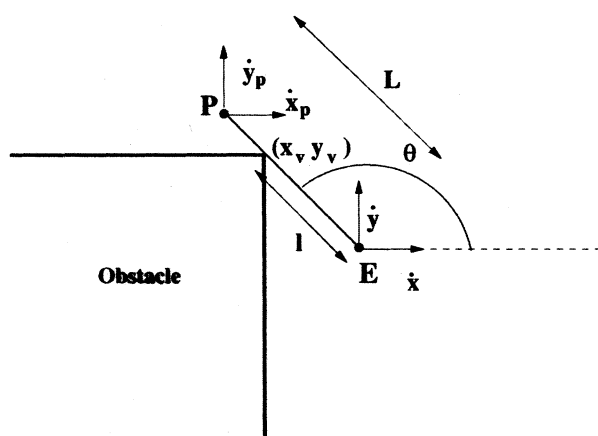


Fig. 16. Illustration of the required follower velocity when the leader rounds a corner

For the case $\alpha = \theta$, we obtain the maximum velocity by back substitution,

$$\begin{aligned} \dot{x}_p &= V_e \cos \theta + V_e \sin \theta \tan \theta \\ &= V_e \cos \theta + V_e \sin \theta \frac{\sin \theta}{\cos \theta} \\ &= \frac{V_e}{\cos \theta} (\cos^2 \theta + \sin^2 \theta) \\ &= \frac{V_e}{\cos \theta} \end{aligned}$$

The follower can effect this velocity provided

$$V_p \geq \frac{V_e}{\cos \theta} \rightarrow \cos \theta \geq \rho$$

Since $\sin \theta + \pi = -\sin \theta$ and $\cos \theta + \pi = -\cos \theta$, the same bound applies for motion directly away from the follower.

7.4. Rounding a Corner

Suppose the rod is in contact with a vertex whose coordinates are given by (x_v, y_v) as shown in Figure 16. If the leader moves with velocity (\dot{x}, \dot{y}) , what must be the pursuer's velocity to maintain the vertex contact?

The constraint that the rod is in contact with the vertex can be specified in terms of the angle θ

$$\cos \theta = \frac{x_v - x}{\sqrt{(x_v - x)^2 + (y_v - y)^2}}$$

$$\sin \theta = \frac{y_v - y}{\sqrt{(x_v - x)^2 + (y_v - y)^2}}$$

To simplify notation, let $l = \sqrt{(x_v - x)^2 + (y_v - y)^2}$ be the distance from the leader to the vertex. If we combine the above equations with Equation (1) we obtain

$$\begin{bmatrix} \dot{x}_p \\ \dot{y}_p \end{bmatrix} = \begin{bmatrix} \dot{x} \\ \dot{y} \end{bmatrix} + \frac{L}{l} \begin{bmatrix} x_v - x \\ y_v - y \end{bmatrix}$$

The pursuer velocity is obtained by differentiating the above

$$\begin{bmatrix} \dot{x}_p \\ \dot{y}_p \end{bmatrix} = \begin{bmatrix} \dot{x} \\ \dot{y} \end{bmatrix} + \frac{L}{l} \begin{bmatrix} -\sin^2 \theta \dot{x} + \cos \theta \sin \theta \dot{y} \\ \cos \theta \sin \theta \dot{x} - \cos^2 \theta \dot{y} \end{bmatrix}$$

8. Conclusions

This work proposes an approach to maintain visibility of a moving evader with a mobile robot in a polygonal environment. The evader moves continuously; its global trajectory is unknown but the distribution of obstacles in the workspace is known in advance. We give necessary conditions for the existence of a surveillance strategy.

The work presented in Efrat et al. (2003) gives an efficient algorithm to compute an optimal motion strategy to track an evader. The method assumes a predictable evader. Our work relaxes this constraint. Our approach is able to track the evader even if the global trajectories of both evader and pursuer are unknown in advance, since it is possible to deduce the optimal motion policies of both pursuer and evader. In Efrat et al. (2003), the authors proposed as an improvement to their work to include a range constraint. Our paper presents that extension.

In game theoretic terms, we have solved the game of degree (Isaacs 1975; Basar and Olsder 1982) of maintaining visibility at a fixed distance of a moving evader with a mobile robot in the presence of obstacles, for a single set of escapable configurations.

The decision problem, which is equivalent to answering the question *Can the evader escape?* depends on the optimal evader and pursuer motion policies. If the pursuer is able to track an evader that travels the optimal path then it will be able to track an evader that travels any other path.

Our analysis including bounded speeds holds for a single set of escapable configurations. If more than one set is considered then a combinatorial problem arises. This combinatorial problem corresponds to defining the optimal ordering to visit the escapable configurations. In this paper, we did not deal with this combinatorial problem, which we leave for future work. However, note that our sufficiency condition for escaping without the bounded speed constraint includes any ordering to visit escapable configuration sets.

In Murrieta et al. (2003b), we proposed a method able to determine the existence of a solution if there is a delay between evader and pursuer motion. In that work we assumed that the pursuer speed was unbounded. As future work, we would like to find a solution for the case of both delay and bounded pursuer speed. Finally, we want to consider dynamic and nonholonomic constraints.

Acknowledgment

The authors would like to thank Stephen Kloder and Jean Paul Laumond for helpful discussions.

References

- Abramowitz, M. and Stegun, I. A. (1964). *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. US National Bureau of Standards.
- Bañon, J. (1990). Implementation and extension of the ladder algorithm. *Proceedings of IEEE International Conference on Robotics and Automation*, Cincinnati Ohio, pp. 1548–1553.
- Başar, T. and Olsder, G. (1982). *Dynamic Noncooperative Game Theory*. Academic Press.
- Becker, C., González-Baños, H., Latombe, J.-C., and Tomasi, C. (1997). An intelligent observer. In *Lecture Notes in Control and Information Sciences 223* (eds O. Khatib and J.K. Salisbury), pp. 153–160, Springer, New York.
- Efrat, A., González-Baños, H. H., Kobourov, S. G., and Palaniappan, L. (2003). Optimal motion strategies to track and capture a predictable target. *Proceedings of IEEE International Conference on Robotics and Automation*, Taipei, Taiwan, pp. 411–423.
- Espiau, B., Chaumette, F., and Rives, P. (1992). A new approach to visual servoing in robotics. *IEEE Transactions on Robotics and Automation*, 8(3): 313–326.
- Fabiani, P. and Latombe, J.-C. (1999). Tracking a partially predictable object with uncertainty and visibility constraints: a game-theoretic approach. *Proceedings of International Joint Conference on Artificial Intelligence*, Stockholm, Sweden, pp. 942–947.
- González, H. H., Lee, C.-Y., and Latombe, J.-C. (2002). Real-time combinatorial tracking of a target moving unpredictably among obstacles. *Proceedings of IEEE International Conference on Robotics and Automation*, Arlington, VA, USA, pp. 1683–1690.
- Guibas, L., Latombe, J.-C., LaValle, S. M., Lin, D., and Motwani, R. (1997). Visibility-based pursuit–evasion in a polygonal environment. *Proceedings of 5th International Workshop on Algorithms and Data Structures (WADS'97)*, Halifax, Nova Scotia, Canada. Published as *Lecture Notes in Computer Science 1272* (eds F. Dehne, A. Rau-Chaplin, J.R. Sack and R. Tamassia), pp.17–30.
- Hájek, O. (1965). *Pursuit Games*. Academic Press, New York.
- Hespanha, J., Prandini, M., and Sastry, S. (2000). Probabilistic pursuit–evasion games: a one-step Nash approach. *Proceedings of IEEE Conference on Decision and Control*, Sydney, Australia, pp. 2272–2277.
- Hutchinson, S., Hager, G., and Corke, P. (1996). A tutorial on visual servo control. *IEEE Transactions on Robotics and Automation*, 12(5): 651–670.
- Isaacs, R. (1975). *Differential Games: A Mathematical Theory with Applications to Warfare and Pursuit, Control and Optimization*. Wiley, New York.
- Isler, V., Kannan, S., and Khanna, S. (2004). Locating and capturing an evader in a polygonal environment. *Sixth International Workshop on the Algorithmic Foundations of Robotics*, Utrecht/Zeist, The Netherlands, pp.351–367.
- Isler, V., Kannan, S., and Khanna, S. (2005). Randomized pursuit–evasion in a polygonal environment. *IEEE Transactions on Robotics*, 5(21): 875–884.
- Jung, B. and Sukhatme, G. (2002). Tracking targets using multiple robots: the effect of environment occlusion *Journal of Autonomous Robots*, 13(3): 191–205.
- LaValle, S. M., González-Baños, H. H., Becker, C., and Latombe, J.-C. (1997). Motion strategies for maintaining visibility of a moving target. *Proceedings of IEEE International Conference on Robotics and Automation*, Albuquerque, New Mexico, USA, pp.737–742.
- LaValle, S. M. and Hinrichsen, J. (1999) Visibility-based pursuit–evasion: An extension to curved environments. *Proceedings of IEEE International Conference on Robotics and Automation*, Detroit, Michigan, USA, pp.1677–1682.
- Latombe, J.-C. (1991). *Robot Motion Planning*. Kluwer Academic Publishers.
- Lawrence, J. D. (1971). *A Catalog of Special Plane Curves*. Dover, New York.
- Mathworld. <http://mathworld.woldfram.com/EllipticIntegraloftheSecondKind.html>
- Murrieta-Cid, R., González-Baños, H. H., and Tovar, B. (2002). A reactive motion planner to maintain visibility of unpredictable targets. *Proceedings of IEEE International Conference on Robotics and Automation*, Arlington VA, USA, pp.4242–4248.
- Murrieta-Cid, R., Sarmiento, A., and Hutchinson, S. (2003a). A motion planning strategy to maintain visibility of a moving target at a fixed distance in a polygon. *Proceedings of International Conference on Advanced Robotics*, Coimbra, Portugal, pp.228–233.
- Murrieta-Cid, R., Sarmiento, A., and Hutchinson, S. (2003b). On the existence of a strategy to maintain a moving target within the sensing range of an observer reacting with delay. *Proceedings of IEEE International Conference on Intelligent Robots and Systems*, Las Vegas, Nevada, USA, pp.1184–1191.

- Murrieta-Cid, R., Sarmiento, A., Bhattacharya, S., and Hutchinson, S. (2004) Maintaining visibility of a moving target at a fixed distance: the case of observer bounded speed. *Proceedings of IEEE International Conference on Robotics and Automation*, New Orleans, Louisiana, USA, pp.479–484.
- Muppirala, T., Murrieta-Cid, R., and Hutchinson, S. (2005). Optimal motion strategies based on critical events to maintain visibility of a moving target. *Proceedings of IEEE International Conference on Robotics and Automation*, Barcelona, Spain, pp.3837–3842.
- Parker, L. (2002). Algorithms for multi-robot observation of multiple targets. *Journal Autonomous Robots*, **12**(3): 231–255.
- Parsons, T. D. (1976). Pursuit-evasion in a graph. In *Theory and Application of Graphs* (eds Y. Alani and D.R. Lick), pp.426–441, Springer-Verlag, Berlin.
- Petitjean, S., Kriegman, D., and Ponce, J. (1992). Computing exact aspect graphs of curved objects: algebraic surfaces. *International Journal on Computer Vision*, **9**: 231–255.
- Schwartz, J. T. and Sharir, M. (1987). On the piano movers' problem: I. The case of a two-dimensional rigid polygon body moving amidst polygonal barriers. *Communications on Pure and Applied Mathematics*, **36**: 345–398.
- Shas, S., Rajko, S., and LaValle, S. M. (2004) Visibility-based pursuit–evasion in an unknown planar environment. *International Journal on Robotics Research*, **23**(1): 3–26.
- Suzuki, I. and Yamashita, M. (1992). Searching for a mobile intruder in a polygonal region. *SIAM Journal on Computing*, **21**(5): 863–888.
- Vidal, R., Shakernia, O., Jin, H., Hyunchul, D., and Sastry, S. (2002). Probabilistic pursuit–evasion games: theory, implementation, and experimental evaluation. *IEEE Transactions on Robotics and Automation*, **18**(5): 662–669.