The International Journal of Robotics Research
30(14) 1709–1727

DOI: 10.1177/0278364911415885
ijr.sagepub.com
SAGE

A cell decomposition approach to visibility-based pursuit evasion among obstacles

Sourabh Bhattacharya and Seth Hutchinson

Abstract

In this paper, we address the problem of surveillance in an environment with obstacles. We consider the problem in which a mobile observer attempts to maintain visual contact with a target as it moves through an environment containing obstacles. This surveillance problem is a variation of traditional pursuit–evasion games, with the additional condition that the pursuer immediately loses the game if at any time it loses sight of the evader. We analyze this tracking problem as a game of kind. We use the method of explicit policy to compute guaranteed strategies for surveillance for the observer in an environment containing a single corner. These strategies depend on the initial positions of the observer and the target in the workspace. Based on these strategies a partition of the visibility polygon of the players is constructed. The partitions have been constructed for varying speeds of the observer and the target. Using these partitions we provide a sufficient condition for escape of a target in a general environment containing polygonal obstacles. Moreover, for a given initial target position, we provide a polynomial-time algorithm that constructs a convex polygonal region that provides an upper-bound for the set of initial observer positions from which it does not lose the game. We extend our results to the case of arbitrary convex obstacles with differentiable boundaries. We also present a sufficient condition for tracking and provide a lower-bound on the region around the initial position of the target from which the observer can track the target. Finally, we provide an upper bound on the area of the region in which the outcome of the game is unknown.

Keywords

Surveillance, visibility based pursuit-evasion, game theory

1. Introduction

Surveillance is monitoring of behavior, activities, or other information in an environment, often carried out in a surreptitious manner. An important problem in surveillance is target tracking. It involves maintaining knowledge of the current location of a target. In the case of visibility-based target tracking, an observer must constantly maintain a line of sight with a target. A challenging problem in this scenario is to plan motion strategies for the observer in the presence of environmental occlusions. In this work, we address the problem of a mobile observer trying to maintain a line of sight with a mobile target in the presence of obstacles in the environment. Both the observer and the target are holonomic and have bounded speeds. The observer has no knowledge about the future actions of the target. In this scenario, we address the following problem: given an initial position of the observer and the target, is it possible for the observer to track the target forever and in case it can do so, what should be its strategy?

Apart from surveillance applications, a mobile robot might be required to continuously follow and monitor at a distance a target performing a task not necessarily related to the target tracking game, such as relaying signals to and from the target (Tekdas et al. 2009). The observer may also be monitoring the target for quality control, verifying that the target does not perform some undesired behavior, or ensuring that the target is not in distress. In applications that involve automated processes that need to be monitored, such as in an assembly work cell, parts or sub-assemblies might need to be verified for accuracy or determined to be in correct configurations. Visual monitoring tasks are also suitable for mobile robot applications (Briggs and Donald 1996). In home-care settings, a tracking robot can follow elderly people and alert caregivers of emergencies (Hsu et al. 2008). Target-tracking techniques in the presence of obstacles have been proposed for the graphic animation of digital actors, in order to select the successive viewpoints under which an actor is to be displayed as it moves

Department of Electrical and Computer Engineering, University of Illinois at Urbana Champaign, Urbana, Illinois, USA

Corresponding author:
Sourabh Bhattacharya, Department of Electrical and Computer Engineering, University of Illinois at Urbana Champaign, Everitt Laboratory, MC-702, Urbana, IL 61801-2918, USA.
Email: sbhattac@illinois.edu

in its environment (Li et al. 1997). In surgical applications, controllable cameras could keep a patient's organ or tissue under continuous observation, despite unpredictable motions of potentially obstructing people and instruments. In wildlife monitoring applications, autonomous underwater vehicles use target-tracking algorithms to navigate in cluttered environments while tracking marine species.

Target-tracking using sonar and infrared sensors has been studied traditionally in the field of automatic control for naval and missile applications (Sung and Um 1996). With the emergence of computer vision, a combination of vision and control techniques were used to design control laws to track a target using vision sensors (Espiau et al. 1992; Hutchinson et al. 1996; Malis et al. 1999; Marchand et al. 1999). A major drawback of pure control approaches is that they are local by nature and it is difficult to take into account the global structure of the environment such as the configuration of workspace obstacles.

In the case of a completely predictable target, the problem can be addressed using techniques from optimization. Such techniques have been used by Efrat et al. (2003) and LaValle et al. (1997) to provide algorithms for an observer to track a predictable target among obstacles. In the case of an unpredictable target the hardness of the problem increases due to the lack of information about the current as well as the future strategies of the target. A plausible way to reduce the hardness of the problem is to solve the problem for specific environments. For instance, Cheung (2005) solved the problem of target-tracking around a regular polygonal obstacle for a specific initial position of the observer and the target. In a similar vein, in this work we have shown that for an environment having a single corner, the problem is completely decidable (Sipser 1997). Although many computationally intensive approximation techniques (Hsu et al. 2008) have been used to address the target-tracking problem, the decidability in a general environment still remains an open problem.

In the past, various techniques have been proposed to devise strategies for an observer that optimizes a local cost function based on the current configuration of the target and observer in the environment. Gonzalez-Banos et al. (2002) and Bandyopadhyay et al. (2004, 2006, 2009) formulate a risk function that takes into account the position of the target and the observer with respect to the occluding vertices of the environment. The strategy for the observer is to move in a direction that minimizes the risk function at every instant. Fabiani and Latombe (1999) designed a planner for target tracking that takes into account the positioning uncertainty of an observer that has a map of the environment. The observer tries to minimize a utility function that maximizes the probability of future visibility of the target and minimizes the uncertainty in its own position. Murrieta-Cid et al. (2002) obtained a motion strategy for the observer by maximizing the target's *shortest distance to escape* from the observer's field of view. Owing to the greedy nature of the above techniques, the completeness or optimality of the resulting strategies cannot be guaranteed for the observer.

Maintaining visibility of a moving target can also be cast as a connectivity problem on a graph that encodes a pertinent cell decomposition of the workspace. Murrieta-Cid et al. (2007) drew the similarity between the target-tracking problem and piano-mover's problem. They extended the three-dimensional cellular decomposition of Schwartz and Sharir (1987) to represent the four-dimensional configuration space of an observer trying to maintain a fixed distance from a target. They reduced the problem to a recursive update and reachability problem on a graph that is constructed using the cellular decompositions. Murrieta-Cid et al. (2008) introduced the notion of *strong mutual visibility* and *accessibility*. Using these two notions, they modeled the problem of maintaining visibility of a moving evader by means of a pair of graphs. They showed that the decision problem of whether a pursuer is able to maintain *strong mutual visibility* of the evader is NP complete. In this work, we present a complete cell decomposition of the free workspace around a single corner and extend these decompositions to general environments. Hence, we feel that the underlying theme of our work belongs to this category.

There have been some efforts in the past to address the target-tracking problem in the scenario where multiple observers try to track multiple targets. Parker (2002) presented a method of tracking several targets with multiple observers. Unlike our work, they did not view the problem from the perspective of computing geometric visibility. Instead they investigated the power of a weighted force vector approach distributed across robot teams in simple, uncluttered environments that are either obstacle free or have a random distribution of simple convex obstacles. Jung and Sukhatme (2002) addressed the problem of tracking multiple targets using a network of communicating robots and stationary sensors. A region-based approach is introduced which controls robot deployment at two levels, namely, a coarse deployment controller and a target-following controller. Kolling and Carpin (2006, 2007) presented a behavior-based solution to the problem of observing multiple targets using multiple robots. They proposed a distributed behavior-based control system where robots share workload by assuming responsibilities concerning the observation of certain targets. Tang and Ozguner (2005) investigated the scenario in which the number of trackers is strictly less than the number of targets. A gradient-approximation algorithm is proposed to generate paths for mobile agents to traverse a sequence of target points. Luke et al. (2005) proposed centralized algorithms for many mobile agents to stay within an 'observation range' of as many targets as possible in the absence of sensing constraints. The algorithms are based on K-means clustering and hill-climbing algorithms. None of the previous authors (except Jung and Sukhatme (2002)) considered the effect of occlusion in visibility due to the presence of obstacles.

Target tracking is related to the game of pursuit–evasion. The goal of the pursuer is to maintain a line of sight to the evader that is not occluded by any obstacle. The goal of the evader is to escape the visibility polygon of the pursuer (and break this line of sight) at any instant of time. Pursuit–evasion games are mainly classified into two categories based on the nature of the cost function. In a *game of kind*, there are only two possible outcomes at the end of the game. The pursuer favors one of the possible outcomes and the evader favors the other possible outcome. The set of initial positions of the players that leads to a favorable outcome for the pursuer is called the *capture set*. The set of initial positions of the players that leads to a favorable outcome for the evader is called the *escape set*. On the other hand, in a *game of degree*, there is a continuum of possible outcomes. The objective of one player is to maximize the outcome and the other player wants to minimize it. In our previous work (Bhattacharya et al. 2009; Bhattacharya and Hutchinson 2010), we modeled target tracking as a *game of degree* with the terminal time as the payoff function of the game. Using techniques from differential game theory (Başar and Olsder 1995), the necessary and sufficient conditions for escape were presented in terms of the saddle point strategies of the players. From these strategies, trajectories were obtained that are optimal in the vicinity of termination situations. In order to predict the outcome of the game from any given initial position of the players, a complete construction of the optimal trajectories is required. Despite some recent progress towards characterizing the optimal trajectories (Bhattacharya et al. 2011), a complete solution to the problem of constructing the optimal trajectories is still unknown for general environments. In contradistinction, this paper poses the target-tracking problem as a *game of kind*. We provide a complete spatial decomposition of the workspace for a simple environment based on the *method of explicit policy* (Isaacs 1965). In the method of explicit policy, strategies are provided for one player to ensure a favorable outcome for it irrespective of the other player's strategies. Extending these strategies to a general environment provides us with a lower bound on the size of the *escape set* and the *capture set* (Lewin 1994).

This work is an extension of our previous work (Bhattacharya et al. 2007; Bhattacharya and Hutchinson 2008). The contributions of this work are as follows. First, we show in Section 2 that in an environment with one corner, the target-tracking problem is completely decidable. Second, we prove in Section 3 that in an environment containing obstacles, the initial positions of the pursuer from which it can track the evader is bounded. Although this result is trivially true for a bounded workspace, for an unbounded workspace it is intriguing. Third, while the general problem of deciding whether the evader can escape or the pursuer can track the evader forever in any arbitrary polygonal environment is still, as far as we know, an open problem, we offer partial solutions to it. In Section 3, we provide polynomial-time approximation schemes to bound the set of initial positions of the pursuer from which it might be able to track successfully. If the initial position of the pursuer lies outside this region, the evader escapes. The size of the region depends on the geometry of the environment and the ratio of the maximum evader speed to the maximum pursuer speed. Fourth, in Section 4, we address the problem of target tracking in an environment containing non-polygonal obstacles. In the past, researchers (LaValle and Hinrichsen 2001) have addressed the problem of searching an evader in non-polygonal environments. However, we do not know of any prior work that addresses the problem of tracking an evader in non-polygonal environments. Fifth, in Section 5, we present a sufficient condition for tracking. Based on this sufficient condition we provide a region around the initial position of the evader from which the pursuer can track the evader. Finally, in Section 6, we provide an upper bound on the size of the region from which the outcome of the game cannot be decided from the aforementioned approximation schemes in environments containing multiple obstacles. In addition to the ratio of the maximum speeds of the players, the bound depends on the geometry of the environment and the initial position of the evader.

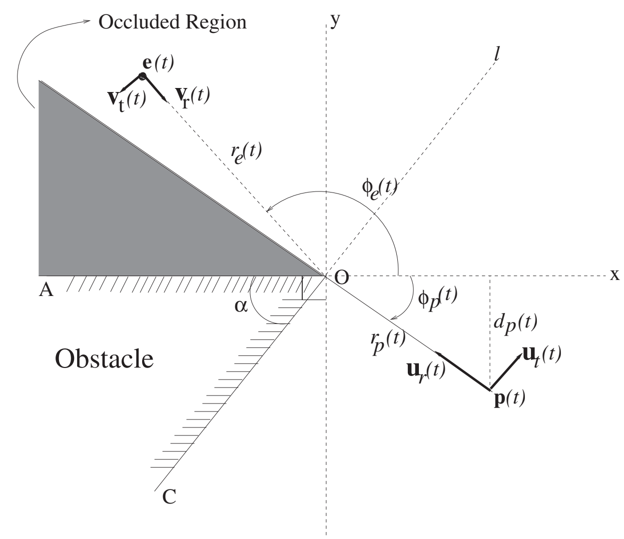

Fig. 1. The pursuer and the evader in an environment having a single corner. The region shaded in a brick-like pattern is the semi-infinite obstacle. The region shaded in gray is the part of the free space that lies out of the visibility polygon of the pursuer.

2. Analysis of a corner

In this section, we address the problem of target tracking in a simple environment containing one corner. The workspace contains a semi-infinite obstacle with one corner that restricts pursuer and evader motions and may occlude the pursuer's line of sight to the evader. Without loss of generality, this corner is placed at the origin and one of the sides lies along the $-x$-axis as shown in Figure 1. A mobile pursuer and evader exist on a plane and move with velocities $\mathbf{v_p}(t)$ and $\mathbf{v_e}(t)$, respectively. Their speeds are bounded by $\bar{v}_p$ and $\bar{v}_e$, respectively. The positions of the pursuer and the evader are expressed in polar coordinates as $\mathbf{p}(t) = (r_p(t), \phi_p(t))$ and $\mathbf{e}(t) = (r_e(t), \phi_e(t))$, respectively. They can also be expressed in Cartesian coordinates as $\mathbf{p}(t) = (x_p(t), y_p(t))$ and $\mathbf{e}(t) = (x_e(t), y_e(t))$, respectively. Let the initial position of the pursuer and the evader be denoted by $\mathbf{p}_0$ and $\mathbf{e}_0$. The tangential velocities of the pursuer and the evader are denoted as $\mathbf{u}_t(t)$ and $\mathbf{v}_t(t)$, respectively. The tangential velocities are considered to be positive in the direction shown in the figure. Here $\mathbf{u}_r(t)$ and $\mathbf{v}_r(t)$ describe the radial velocities of the pursuer and the evader, respectively. The radial velocities are considered to be positive if they point away from the origin. In Figure 1, the radial velocities of the pursuer and the evader are in the negative direction. The pursuer and the evader know each other's current position as long as they can see each other. Moreover, the pursuer knows the evader's current velocity. The initial position of the pursuer and the evader is such that they are visible to each other. Both of the players have a complete map of the environment.

The unshaded region is the visibility region of the pursuer. Visibility extends uniformly in all directions and is only terminated by workspace obstacles (omnidirectional, unbounded visibility). To prevent the evader from escaping, the pursuer must keep the evader in its visibility polygon, $V(\mathbf{p}(t))$. The visibility polygon of the pursuer is the set of points from which a line segment from the pursuer to that point does not intersect the obstacle region. The evader escapes if at *any* instance of time it can break the line of sight to the pursuer.

The two obstacle edges meeting at this corner are considered to extend for an infinite length, so that there is no other geometry that the evader can hide behind in the workspace. The two sides of the obstacle form an angle α. If $\alpha \geq \pi$ then every point in the free workspace is visible to every other point and the pursuer will trivially be able to track the evader indefinitely. Thus, we only consider obstacles where $\pi > \alpha \geq 0$.

Analogous to a star domain (de Berg et al. 1997) in computational geometry, we define the *star region* associated with a vertex as the region in the free workspace bounded by the lines supporting the vertex of the obstacle. The shaded region in Figure 2 shows the star region associated with the vertex v.

The concept of a star region is only applicable for a convex vertex (a vertex of angle less than π). As can be seen in Figure 1, in the case of a semi-infinite obstacle having a single corner, the star region extends outward from the corner of the obstacle. It is semi-infinite and bounded by the ray l and the x-axis. In the case of a single corner, the entire free space is visible from any point in the star region. If the pursuer can enter the star region before losing sight of the evader, it will trivially be able to track the evader at all future times.

In this setting, we address the following problem. Given $\mathbf{p}_0$, $\mathbf{e}_0$, $\bar{v}_e$ and $\bar{v}_p$, does there exist a policy for the evader to

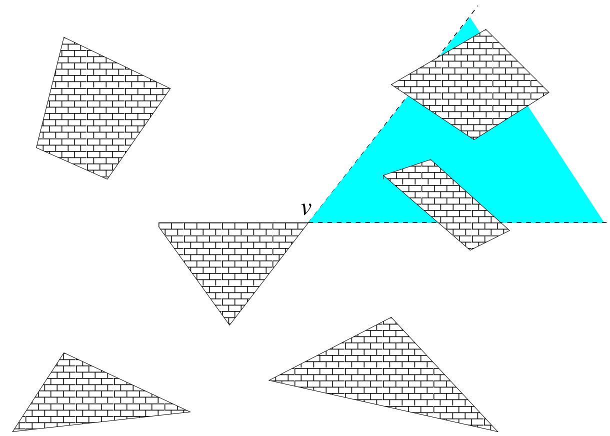

Fig. 2. *Star Region* associated with the vertex.

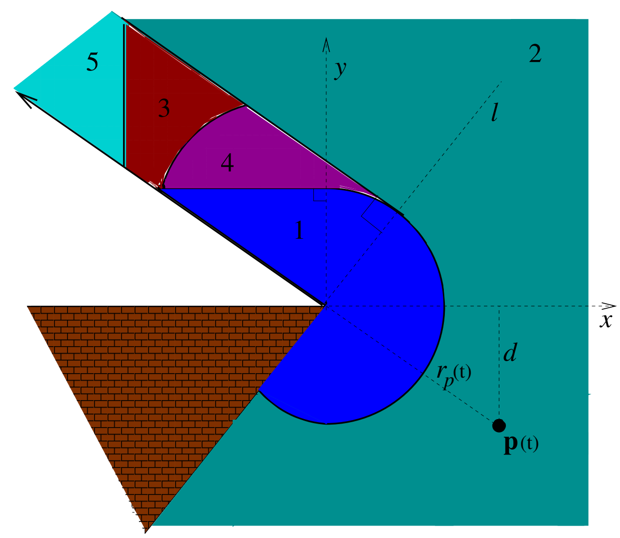

Fig. 3. The partition of the visibility polygon of the pursuer based on the strategies of the players shown in Table 1.

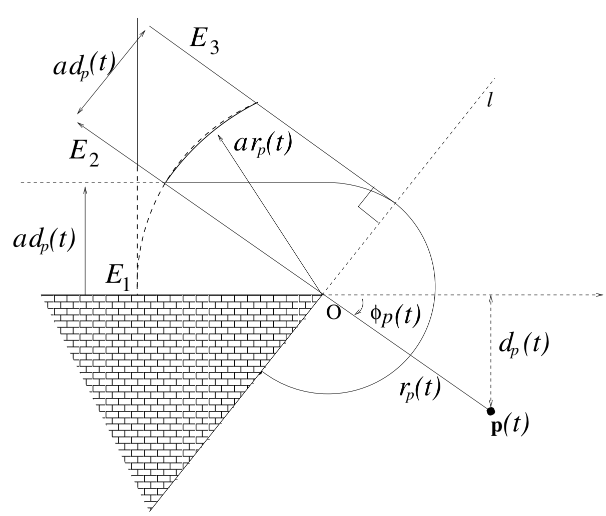

Fig. 4. The geometrical parameters associated with the partitions shown in Figure 3.

escape the visibility region of the pursuer in finite time or does there exist a policy for the pursuer to track the evader for all time. In the following sections, we present a partition of the workspace for an environment having a single corner so that we can answer the above question depending on the ratio $\overline{v}_e/\overline{v}_p$, $\mathbf{p}_0$ and $\mathbf{e}_0$.

2.1. Pursuer-based partition

We now present a decomposition of $V(\mathbf{p}_0)$, the visibility region of the pursuer at an initial position, into regions in which the evader may lie based on the outcome of the game. These partitions can be constructed at any time during the game with the current knowledge of the pursuer's position. Depending on the partition in which the evader lies currently, we present instantaneous strategies for the winner of the game.

The number of partitions and their geometry depend on the initial position of the pursuer. If the initial position of the pursuer is in the star region of the corner, the pursuer can see the entire workspace at all times. Hence, for any initial position of the evader, the pursuer wins the game. In the remaining section, we consider the initial positions of the pursuer in which it does not lie inside the star region. Owing to symmetry of the environment, the analysis is the same if the initial position of the pursuer lies below the x-axis or if it lies in the left half-space of l. Without loss of generality, we analyze the former situation.

Let us first consider the case of a corner for which $\alpha < \pi/2$ and $\mathbf{p}_0 = (r_p(0), \phi_p(0))$ is such that $\phi_p(0) \in [-\pi/2, 0)$. Define $a = \overline{v}_e/\overline{v}_p$ and let $d_p(t)$ denote the minimum distance of the pursuer from x-axis. Let $d = d_p(t)|_{t=0}$. For $\mathbf{x} = (x, y) \in \mathbb{R}^2$, we define the minimum distance from $\mathbf{x}$ to a segment, ray or line as $d(\mathbf{x}, E) = \min_{\mathbf{y} \in E} \|\mathbf{x} - \mathbf{y}\|_2$, where E denotes an edge, ray or line. Here E_1 is the edge of the obstacle that lies along the x-axis. The ray forming the free boundary of the visibility region of the pursuer is denoted by E_2.

Figure 3 shows the partition of $V(\mathbf{p}(t))$ and Figure 4 shows the geometry of the partitions. Here $V(\mathbf{p}(t))$ is decomposed into the following regions.

1. Region 1 = $\{\mathbf{x} \mid d(\mathbf{x}, E_1) < a d_p(t)\}$.
2. Region 2 = $\{\mathbf{x} \mid d(\mathbf{x}, E_2) \geq a d_p(t)\}$.
3. Region 3 = $\{\mathbf{x} \mid d(\mathbf{x}, E_2) \leq a d_p(t), \|\mathbf{x}\|_2 \geq a r_p(t), x \leq -a r_p(t)\}$.
4. Region 4 = $\{\mathbf{x} \mid d(\mathbf{x}, E_2) \leq a d_p(t), \|\mathbf{x}\|_2 \leq a r_p(t), d(\mathbf{x}, E_1) \geq a d_p(t)\}$.
5. Region 5 = $\{\mathbf{x} \mid d(\mathbf{x}, E_2) \leq a d_p(t), x \leq -a r_p(t)\}$.

Further, we define Region 6 as the set of points in the free workspace not belonging to $V(\mathbf{p}(t))$. Before we give a set of theorems that define the winning strategy for each region in the partition, these strategies are summarized in Table 1.

Theorem 1. *If the evader lies in Region 1 of $V(\mathbf{p}_0)$ and follows Policy A, no pursuer policy exists that can prevent the escape of the evader.*

Table 1. Strategies that define the winning strategy for each region in the partition.

Evader Policies	Evader Region	Control Law
A	1 and $\phi_e \in [\alpha - \pi, \pi/2]$	$\dot{r}_e(t) = -\overline{v}_e$
	1 and $\phi_e \in [\pi/2, \pi + \phi_p]$	$\dot{y}_e(t) = -\overline{v}_e$
Pursuer Policies	**Evader Region**	**Control Law**
B	2, 4	$\dot{y}_p(t) = \overline{v}_p$
C	3	$\dot{\phi}_p(t) = \lvert\dot{\phi}_e(t)\rvert$
		$\dot{r}_p(t) = -\frac{r_p(t)}{r_e(t)} \lvert\dot{r}_e(t)\rvert$
D	5	$\dot{\phi}_p(t) = \frac{\overline{v}_p}{r_p(0)}$

Proof. If the evader lies in Region 1, the maximum time required by the evader to reach E_1 by following Policy A is $t_e < ad/\overline{v}_e = d/\overline{v}_p$. The minimum time required by the pursuer to reach x-axis with any policy is at least $t_p > d/\overline{v}_p$. Since $t_p > t_e$ the evader reaches E_1 before the pursuer can reach the x-axis. If the evader lies on E_1 and the pursuer has not yet reached the x-axis the evader will be outside the visibility region of the pursuer. Hence, the evader escapes. □

Theorem 2. *If the evader lies in Region 2 of $V(\mathbf{p}_0)$ and the pursuer follows Policy B, no evader policy exists that can escape the visibility region of the pursuer.*

Proof. The time required by the pursuer to reach the x-axis by following Policy B is $t_p = d/\overline{v}_p$. If the evader lies in Region 2, the minimum time required by the evader to reach E_2 is $t_e > ad/\overline{v}_e = d/\overline{v}_p$. Thus, $t_e \geq t_p$. If the pursuer follows Policy B, $V(\mathbf{p}_0) \subset V(\mathbf{p}(t))|_{t>0}$, i.e. the visibility region for the pursuer is monotonically increasing during the execution of this policy. Since the evader cannot reach E_2, the only free boundary of $V(\mathbf{p}_0)$, before the pursuer reaches the boundary of the star region, $\mathbf{e}(t) \in V(\mathbf{p}(t))$, $\forall t \in [0, t_p]$. Once the pursuer reaches the x-axis, the entire free workspace belongs to $V(\mathbf{p}(t_p))$ and the evader remains in sight of the pursuer for all future times. □

Theorem 3. *For all initial positions of the evader in Regions 3 and 4 of $V(\mathbf{p}_0)$, the pursuer can track the evader by following a reactive motion and switching between Policies B, C and D appropriately.*

Proof. In order to prove the theorem, we need Lemmas 1, 2 and 3.

Lemma 1. *If the evader lies in Region 3 of $V(\mathbf{p}(t))$ and the pursuer follows Policy C, for every evader policy the evader can either stay in Region 3 or move to Region 2 or Region 5 of $V(\mathbf{p}(t))$.*

Proof. If the pursuer follows Policy C, then it follows both the radial and angular movements of the evader. According to the control law of the pursuer in Region 3, $|\mathbf{v_p}(t)| = |\mathbf{v_e}(t)|\frac{r_p(t)}{r_e(t)}$. The maximum speed of the evader is $\overline{v}_e$ and the geometry of Region 3 is such that $r_p(t)/r_e(t) \leq 1/a$. Hence, $|\mathbf{v_p}(t)| \leq \overline{v}_e/a = \overline{v}_p$. Thus, the pursuer velocities of Policy C are always attainable in Region 3.

If order to keep the evader in the visibility polygon of the pursuer and prevent it from entering Region 6, the following inequality must hold at all times before the pursuer can enter the *star region*.

$$\phi_e(t) - \phi_p(t) \leq \pi.$$

If the evader lies in Region 3, from the geometry of Region 3 we can see that $\phi_e(t) > \phi_p(t)$. The angular speeds of the players obey the following control law:

$$|\dot{\phi}_e(t)| \quad = \quad \dot{\phi}_p(t).$$

Integrating both sides of the equation gives us the following equations and further using the fact that $\left|\int_0^t \dot{\phi}_e(t)\, dt\right| \leq \int_0^t |\dot{\phi}_e(t)|\, dt$, we obtain the following equation:

$$\left|\int_0^t \dot{\phi}_e(t)\, dt\right| \leq \int_0^t \dot{\phi}_p(t)\, dt$$
$$\Rightarrow \quad |\phi_e(t) - \phi_e(0)| \leq \phi_p(t) - \phi_p(0).$$

Since $\phi_e(t) - \phi_e(0) \leq |\phi_e(t) - \phi_e(0)|$

$$\Rightarrow \phi_e(t) - \phi_e(0) \leq \phi_p(t) - \phi_p(0),$$

$$\Rightarrow \phi_e(t) - \phi_p(t) \leq \phi_e(0) - \phi_p(0).$$

From the assumption that the pursuer and the evader are visible to each other at the beginning of the game, we obtain the following:

$$\phi_e(0) - \phi_p(0) \leq \pi.$$

This leads to the following inequality:

$$\phi_e(t) - \phi_p(t) \leq \pi.$$

Hence, the evader cannot escape the visibility region of the pursuer if the pursuer follows Policy C. The radial component of the velocities obeys the following equation:

$$\frac{|\dot{r}_e(t)|}{r_e(t)} = -\frac{\dot{r}_p(t)}{r_p(t)}$$
$$\Rightarrow \frac{\dot{r}_e(t)}{r_e(t)} \geq \frac{\dot{r}_p(t)}{r_p(t)}$$
$$\Rightarrow \frac{r_e(t)}{r_p(t)} \geq \frac{r_e(0)}{r_p(0)} \geq a.$$

Thus, the evader cannot enter Region 4. Hence, for any policy the evader can either stay in Region 3 or it can enter Region 2 or Region 5 of $V(\mathbf{p}(t))$. □

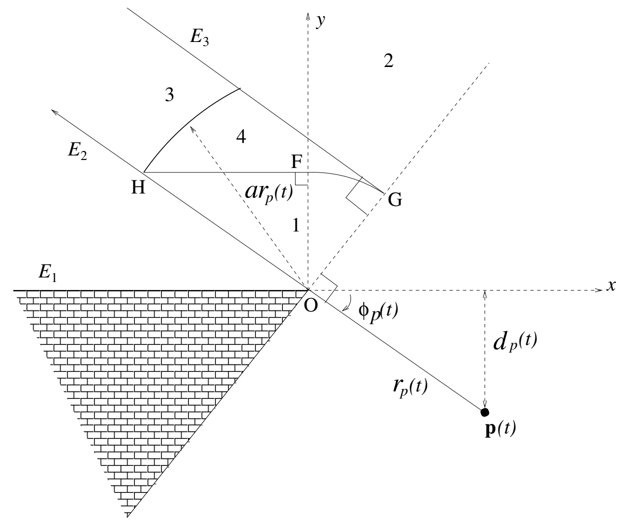

Fig. 5. The geometry of Region 4 in the partition for a given pursuer position.

Lemma 2. *If the evader lies in Region 4 of $V(\mathbf{p}(t))$ and the pursuer follows Policy B, for every evader policy the evader can either stay in Region 4 or move to Regions 2 or 3 of $V(\mathbf{p}(t))$.*

Proof. Refer to Figure 5. If the pursuer follows Policy B, all points on segment HF move with velocity $a\overline{v}_p = \overline{v}_e$ towards the edge E_1. Similarly, all points on the arc FG move with radial velocity $\overline{v}_e$ toward O. In order to enter Region 1 from Region 4, the evader must move toward the boundary of Region 1 with a velocity greater than the velocity at which the boundary is receding away from the evader. That is not possible since the boundary of Region 1 moves with velocity $\overline{v}_e$, the maximum possible evader velocity, away from the evader. Hence, the evader cannot enter Region 1 from Region 4. Hence, for all evader policies, the evader can only reach Region 3 or Region 2 from Region 4. □

Lemma 3. *For all initial positions of the evader in Region 5 of $V(\mathbf{p}_0)$, the pursuer can track the evader by following Policy D.*

Proof. Refer to Figure 6. After time t, the evader lies in the closure of a circle of radius $\overline{v}_e t$ centered at $\mathbf{e}_0$. Let OL denote the tangent from the origin to the circle. A sufficient condition for the pursuer to keep the evader in sight for all future times is to keep the magnitude of the angular velocity of the line of the sight, OP, to be greater than the magnitude of the angular velocity of the line tangent to the growing circle, OL, for all future time until the pursuer reaches the *star region*. The pursuer moves in a circle of radius $r_p(0)$ with tangential velocity of $\overline{v}_p$ while it follows Policy D. Hence, the magnitude of the angular velocity of the line OP is given by $\omega_p = \overline{v}_p/r_p(0)$. The magnitude of the angular velocity of OL is given by $\omega_{OL} = -\overline{v}_e/OL$. Here ω_{OL} is maximum when the radial distance of L is minimum. This happens when the circle touches the edge OA. This length

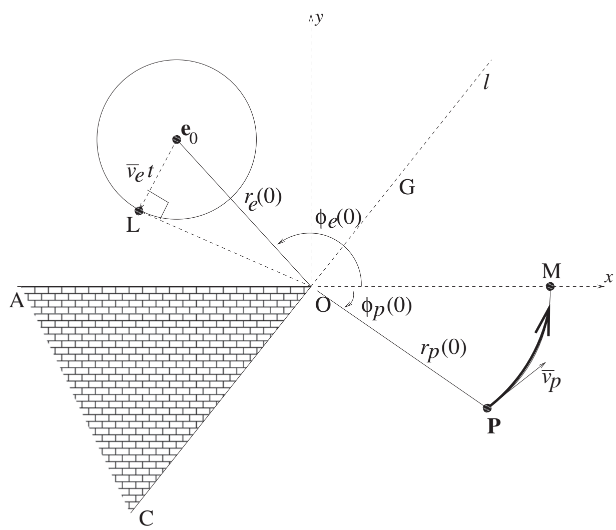

Fig. 6. The pursuer's policy and the resulting path when the evader lies in Region 5 of the partition.

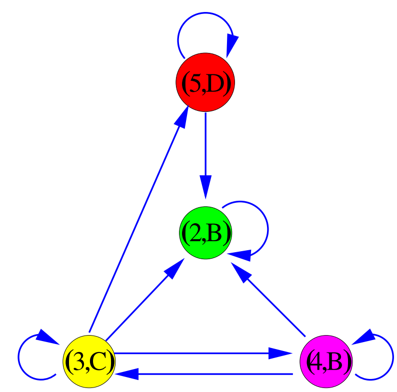

Fig. 7. A finite automaton that models the transitions of the evader based on partition in which it lies and the policies followed by the pursuer.

is given by $r_e(0)\cos(\phi_e(0))$. Hence, the maximum value of ω_{OL} is given by $\omega^*_{OL} = -\bar{v}_e/r_e(0)\cos(\phi_e(0))$. Solving for $\omega_p \geq \omega^*_{OL}$ leads to the following condition:

$$r_e(0) \geq -\frac{ar_p(0)}{\cos(\phi_e(0))}.$$

Since $\cos(\phi_e(0)) \leq 0$, we obtain the following condition

$$x_e(0) \leq -ar_p(0)$$

which is satisfied for all points in Region 5. □

Now we return to the proof of Theorem 3. If the evader starts in Region 3, it can either stay in the same region or leave it while the pursuer follows Policy C. If the evader remains in Region 3 forever, then the pursuer can keep the evader in sight for all future times since Region 3 lies in the visibility polygon of the pursuer. Based on the connectivity of the cells, we can conclude that the evader lies in one of the Regions among 2, 4, 5 and 6 once it leaves Region 3. The following statements provide an argument that the pursuer eventually keeps the evader in its sight for the scenario in which the evader leaves Region 3.

- Lemma 1 proves that the evader cannot enter Regions 6 and 4 if the pursuer follows Policy C.
- If the evader enters Region 2, Proposition 2 proves that the pursuer can track the evader by switching to Policy B.
- If the evader enters Region 5, Lemma 3 proves that the pursuer can keep track of the evader by switching to policy D.

Finally, we consider the case in which evader starts in Region 4. There are two following possibilities. Either the evader stays in Region 4 forever in which case the pursuer can keep the evader in sight for all future times since Region 4 lies in the visibility polygon of the pursuer. Otherwise, the evader can lie in one of the Regions among 1, 2, 3 and 6 since they share a boundary with Region 4. The following statements provide an argument that the pursuer eventually keeps the evader in its sight for the scenario in which the evader leaves Region 4.

- Lemma 2 proves that the evader cannot enter Regions 1 and 6 if the pursuer follows Policy B.
- If the evader enters Region 2, Proposition 2 proves that the pursuer can track the evader by switching to Policy B.
- If the evader enters Region 3, then the previous argument proves that the pursuer can track the evader by switching to Policy C.

This concludes the proof of the theorem. □

Figure 7 summarizes Theorems 2 and 3. Each state is a 2-tuple. The first element is the region in the partition of $V(\mathbf{p}(t))$ in which the evader lies. The second element is the policy followed by the pursuer when the evader occupies the corresponding region. The arrows show the allowable transitions of the evader under the respective policy of the pursuer. Hence, given the initial position of the pursuer and the evader, we can construct the partition of $V(\mathbf{p}_0)$ and use Figure 7 to obtain the instantaneous strategy of the pursuer if it can track the evader.

From Table 1, we can conclude that the control laws are discontinuous along the boundary of the regions. Let us first consider the case in which the evader lies on the boundary shared by two regions in which the pursuer has a winning strategy. In this case, the pursuer has the freedom to choose between the two strategies that are valid in the regions forming the boundary. This follows from the fact that the proofs of Theorems 2 and 3, as well as Lemmas 1, 2 and 3 provide control laws that lead to winning strategies for the pursuer on the boundary of the respective regions. Now, let us consider the case in which evader lies on the boundary of Region 1 except on E_2. The winning strategy for the evader in Region 1 requires it to travel for a time t_e strictly less

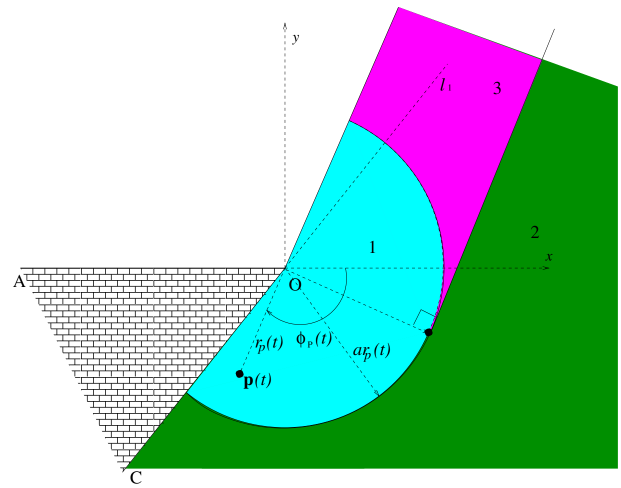

Fig. 8. The partition of $V(\mathbf{p}(t))$ when $\phi_p(t) < -\pi/2$.

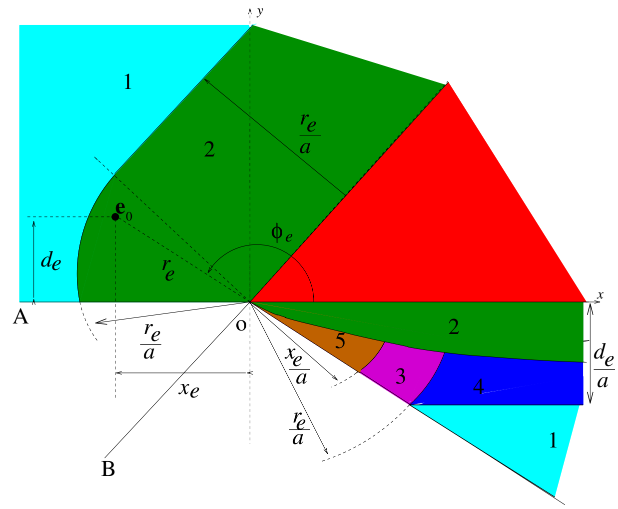

Fig. 9. The evader is nearer to the side of the obstacle than the corner.

than $d/\bar{v}_p$. Therefore, the evader cannot win if it lies on the boundary by following its winning strategy in Region 1. On the other hand, Theorem 2 as well as Lemma 2 hold true even if the evader is on the boundary of Region 1. Therefore, the pursuer has a winning strategy if the evader lies on the boundary of Region 1.

The above analysis was for the case when $\phi_p(0) \in [-\pi/2, 0)$. For the case when $\phi_p(0) < -\pi/2$, the analysis still holds. The only changes are that Region 1 expands, the area of Region 4 is reduced to zero and Region 5 ceases to exist. Figure 8 shows the partition of the visibility region of the pursuer in this case.

The analysis we have presented so far assumed that $\alpha \in [0, \pi/2]$. Refer to Figure 1. If $\alpha \in [\pi/2, \pi]$, then $\phi_p(0)$ must lie in the fourth quadrant and hence $\phi_p(0)$ must be greater than $-\pi/2$. Hence, it reduces to the problem we analyzed in this section.

We now provide a decomposition of $V(\mathbf{e}_0)$ into regions in which the pursuer may lie based on the outcome of the game.

2.2. Evader-based partition

In the previous section, a partition of $V(\mathbf{p}_0)$ has been given based on the policies used by the players to win the game. In this section, we use the same policies as used by the players in Table 1. We fix the position of the evader and compute the boundaries across which the policies of the winner changes. These curves partition $V(\mathbf{e}_0)$ into regions in which the pursuer may lie depending on the policy of the winner. The geometry of the partitions is a function of the velocity ratio between the pursuer and the evader.

To determine the partition of $V(\mathbf{e}_0)$, we must consider three cases depending on whether (a) the closest point to the evader on the obstacle lies on the corner (b) the closest point belongs uniquely to one of the sides (c) the evader lies inside the star region. Figure 9 shows the partition of $V(\mathbf{e}_0)$ for the case when the closest point to the evader on the obstacle belongs to the side AO. In the rest of this section, we analyze this case.

Since we are computing the partitions by fixing the initial position of the evader while retaining the policies of the players, the geometry of the regions in this case is different from that given in Table 1. Moreover, in the previous section, we saw that the result of the game depends on the initial position of the pursuer and the evader. Hence, the configuration variables in this section denote their values at the beginning of the game.

First, let us consider the case in which the pursuer lies in the star region. In this case, the entire free workspace is visible to the pursuer and it can track the evader by remaining stationary. Hence, if the pursuer lies in the star region, it wins the game and its policy is to remain stationary. Now we present the derivation of each region of the partition in the remaining part of $V(\mathbf{e}_0)$.

Region 1 From the previous section, Region 1 consists of all of those points in $V(\mathbf{p}_0)$ from which the evader wins the game irrespective of the pursuer's policy.

First, let us consider the case in which the pursuer lies below the x-axis. The strategy of the evader is to move directly towards the obstacle so that it can reach AO before the pursuer can reach the boundary of the star region which is the x-axis in this case. Since we are considering the case where the closest point to the evader on the obstacle belongs to side AO, the evader lies in Region 1 of $V(\mathbf{p}_0)$ if $d_e \leq ad_p \Rightarrow d_p \geq d_e/a$.

Now let us consider the case in which the pursuer lies above the x-axis and outside the star region. In this case, the evader wins the game if the time taken by the evader to reach the corner is less than the time taken by the pursuer to reach the star region. Let d_e denote the perpendicular distance of the evader from the edge AO. Hence, Region 1 consists of points such that $r_e < ad_p \Longrightarrow d_p > r_e/a$.

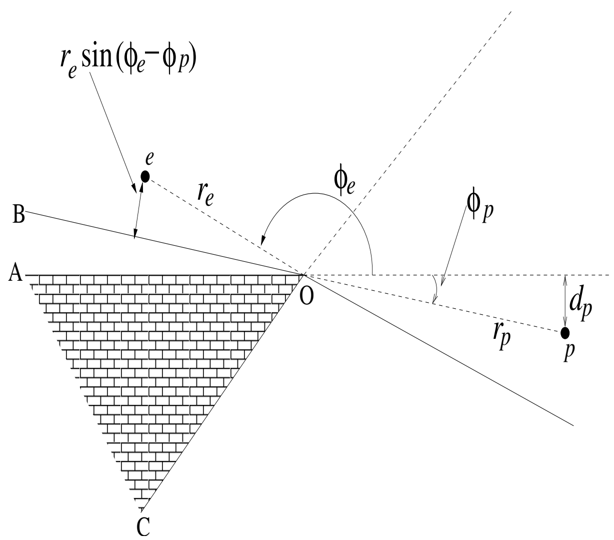

Fig. 10. Distance of the evader from the line of sight of the pursuer

Region 2 Let us first consider the case in which the pursuer lies below the x-axis. From Figure 10, we can see that the shortest distance of the evader from line OB is $r_e \sin(\phi_e - \phi_p)$. Refer to Figure 4. We can see that the evader lies in Region 2 of $V(\mathbf{p}_0)$ if the shortest distance of the evader from line OB is greater than ad_p. This leads to the following inequality:

$$-ar_p \sin\phi_p \leq r_e \sin(\phi_e - \phi_p) .$$

Since $\phi_p < 0$, the above inequality can be written as

$$r_p \leq -\frac{r_e \sin(\phi_e - \phi_p)}{a \sin\phi_p}$$
$$\Rightarrow r_p \leq \frac{r_e \sin\phi_e(\cot\phi_e - \cot\phi_p)}{a}.$$

Now let us consider the case when the pursuer lies above the x-axis and outside the star region. From Figures 4 and 8, we can conclude that the evader lies in Region 2 of $V(\mathbf{p}_0)$ if $r_e \geq a\min\{r_p, d_p\} \Rightarrow r_e/a \geq \min\{r_p, d_p\}$.

Region 3 Refer to Figure 4. The evader lies in Region 3 of $V(\mathbf{p}_0)$ if $r_e \geq ar_p$, $x_e \geq -ar_p$ and least distance of the evader from line OB is less than ad_p. This implies that $r_p \leq r_e/a$, $r_p \geq -x_e/a$ and $r_p \geq r_e/a \sin\phi_e(\cot\phi_e - \cot\phi_p)$. Hence, $\max\{-x_e/a, r_e/a \sin\phi_e(\cot\phi_e - \cot\phi_p)\} \leq r_p \leq r_e/a$.

Region 4 From Figure 4, we see that the evader lies in Region 4 of $V(\mathbf{p}_0)$ if $r_e \leq ar_p$, $\min\{d_e, r_e\} \geq ad_p \Rightarrow \min\{d_e, r_e\} \geq -ar_p \sin\phi_p$ and the shortest distance of the evader from line OB is less than ad_p. This leads to the following condition:

$$-\frac{\min\{d_e, r_e\}}{a \sin\phi_p} \geq r_p \geq \max\left\{\frac{r_e}{a}, r_e \sin\phi_e(\cot\phi_e - \cot\phi_p)\right\}.$$

Region 5 From Figure 4, we see that the evader lies in Region 5 of $V(\mathbf{p}_0)$ if $x_e \leq -ar_p \Rightarrow r_p \leq -x_e/a$.

All of the above partitions are shown in Figure 9. Figure 11(a) shows the partition of $V(\mathbf{e}_0)$ when the nearest point of the obstacle to the evader is corner O but the evader is outside the star region and Figure 11(b) shows the partition of $V(\mathbf{e}_0)$ when the evader is in the star region.

Based on the partition of $V(\mathbf{e}_0)$, we present a sufficient condition of escape for the evader in the next section that is used to bound the set of initial position of the pursuer from which it might win the game.

3. Approximation schemes for a polygonal environment

In the previous section, we provided a partition of $V(\mathbf{e}_0)$ to decide the outcome of the target tracking game. From the previous section, we can conclude that if the pursuer lies in Region 1 of $V(\mathbf{e}_0)$, then the evader has a strategy to win irrespective of the pursuer's strategy. The presence of other obstacles does not affect this result. This leads to the following sufficient condition for escape of the evader in any general environment.

Sufficient condition If the time required by the pursuer to reach the star region associated with a vertex is greater than the time required by the evader to reach the vertex, the evader has a strategy to escape the pursuer's visibility region.

The relation between the time taken by the pursuer and evader can be expressed in terms of the distances traveled by the pursuer and the evader and their speeds. In a general environment, if d_e is the length of the shortest path of the evader from a corner, d_p is the length of the shortest path of the pursuer from the star region associated with the corner and a is the ratio of the maximum speed of the evader to that of the pursuer, the sufficient condition can also be expressed in the following way:

SC: If $d_e < ad_p$, the evader wins the game.

For convenience, we refer to the sufficient condition as SC in the rest of the paper. Using the SC, we show that in any environment containing polygonal obstacles, the set of initial positions from which a pursuer can track the evader is bounded. First, we prove the statement for an environment containing a single convex polygonal obstacle. Then we extend the results to a general polygonal environment containing multiple obstacles. This leads to our first approximation scheme.

Consider an evader in an environment with a single convex polygonal obstacle having n edges $e_1, e_2, \ldots, e_n$. Every edge e_i is a line segment that lies on a line l_{e_i} in the plane. Let $\{h_i\}_1^n$ denote a family of lines, each given by the equation $h_i(x, y, \mathbf{e}_0, a) = 0$. The presence of the terms $\mathbf{e}_0$ and a in the equation imply that the equation of the line depends on the initial position of the evader and the speed ratio respectively. Each line h_i divides the plane into two

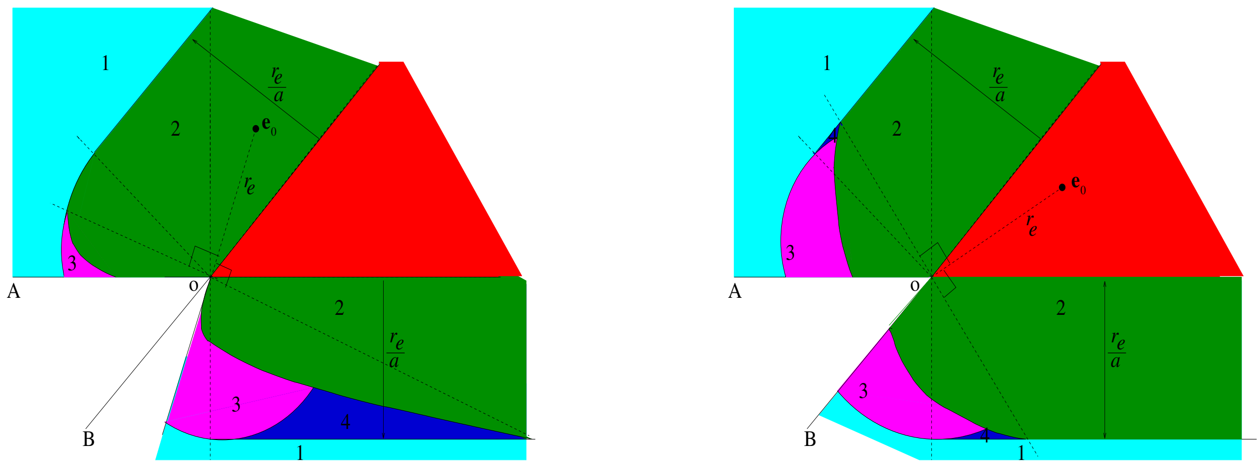

Fig. 11. (a) Evader outside the star region and (b) evader inside the star region.

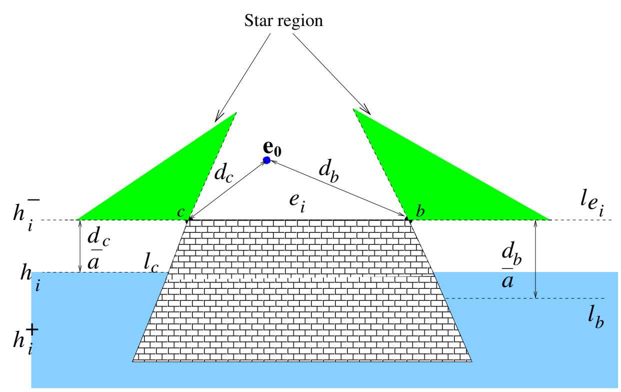

Fig. 12. Proof of Lemma 4.

half-spaces, namely, $h_i^+ = \{(x, y) \mid h_i(x, y, \mathbf{e}_0, a) > 0\}$ and $h_i^- = \{(x, y) \mid h_i(x, y, \mathbf{e}_0, a) < 0\}$. Now we use the SC to prove a property related to the edges of the obstacle.

Lemma 4. *For every edge e_i, there exists a line h_i parallel to e_i and a corresponding half-space h_i^+ such that the pursuer loses the game if $\mathbf{p}_0 \in h_i^+$.*

Proof. Consider an edge e_i of a convex obstacle as shown in Figure 12. Since the obstacle is convex, it lies in one of the half-spaces generated by the line l_{e_i}. Without loss of generality, let the obstacle lie in the half-space below the line l_{e_i}. Let d_c and d_b be the length of the shortest path of the evader from vertices c and b of the edge e_i respectively. Since the obstacle lies in the lower half-space of l_{e_i}, the star region associated with vertices c and b are in the upper half-space of l_{e_i} as shown by the green shaded region. Let l_c and l_b be the lines at a distance of d_c/a and d_b/a, respectively, from the line l_{e_i}. If the pursuer lies at a distance greater than $\min(d_c/a, d_b/a)$ below the line l_{e_i}, then the time taken by the pursuer to reach the line l_{e_i} is $t_p \geq 1/\overline{v}_p \min(d_c/a, d_b/a)$. The minimum time required by the evader to reach corner c or b, whichever is nearer, is given by $t_e = 1/\overline{v}_p \min(d_c, d_b)$. From the expressions of t_e and t_p we can see that $t_p > t_e$. Hence, the pursuer will reach the nearer of the two corners before the evader reaches line l_{e_i}. Hence, from SC, we conclude that if the pursuer lies below the line h_i parallel to e_i at a distance of $\min(d_c/a, d_b/a)$, then the evader wins the game by following the shortest path to the nearer of the two corners. In Figure 12, since $d_b > d_c$ the line h_i coincides with line l_c. □

Given an edge e_i and the initial position of the evader, the proof of Theorem 4 provides an algorithm to find the line h_i and the corresponding half-plane h_i^+ as long as the length of the shortest path of the evader to the corners of an edge is computable. For example, in the presence of other obstacles, the length of the shortest path of the evader to the corners can be obtained by Dijkstra's algorithm.

Now we present some geometrical constructions required to prove the next theorem. Refer to Figure 13. Consider a convex obstacle. Consider a point c strictly inside the obstacle. For each vertex v_i, extend the line segment v_ic to infinity in the direction $\vec{v_ic}$ to form the ray cv_i'. Define the region bounded by rays cv_i' and cv_{i+1}' as *sector* $v_i'cv_{i+1}'$. The *sectors* possess the following properties.

1. Any two sectors are mutually disjoint.
2. The union of all the sectors is the entire plane.

We use this construction to prove the following theorem.

Theorem 4. *In an environment containing a single convex polygonal obstacle, given the initial position of the evader, the set of initial positions of the pursuer from which it can win the game is a bounded subset of the free workspace.*

Proof. Refer to Figure 14. Consider an edge e_i of the convex obstacle with end points v_i and v_{i+1}. Without loss of generality, the obstacle lies below l_{e_i}. Let c be a point strictly inside the convex polygon. Extend the line segments v_ic and $v_{i+1}c$ to form sector $v_i'cv_{i+1}'$. Using Lemma 4, given the initial position of the evader, we can construct a line h_i parallel

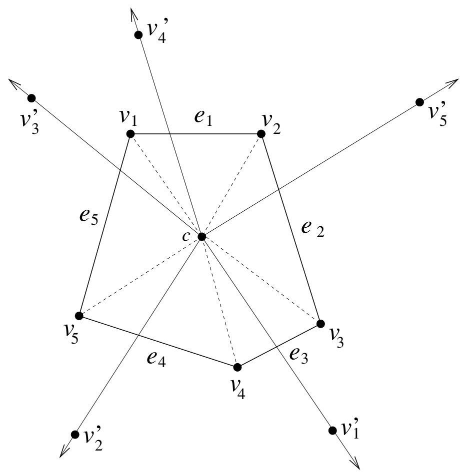

Fig. 13. A polygon and its sectors.

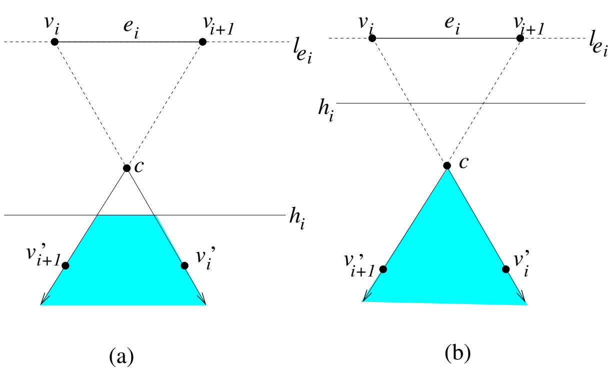

Fig. 14. Proof of Proposition 4.

to e_i such that if the initial pursuer position lies below h_i, the evader wins the game. In case the line h_i intersects sector $v_i'cv_{i+1}'$, as shown in Figure 14(a), the evader wins the game if the initial pursuer position lies in the shaded region. In case the line h_i does not intersect sector $v_i'cv_{i+1}'$, as shown in Figure 14(b), the evader wins the game if the initial pursuer position lies anywhere in the sector. Hence, for every sector, there is a region of finite area such that if the initial pursuer position lies in that region then it might win the game. Every edge of the polygon has a corresponding sector associated with it. Since each sector has a region of finite area such that if the initial pursuer position lies in it, the pursuer might win the game, the union of all of these regions is finite. Hence, the theorem follows. □

In the proof of Theorem 4, we generate a bounded set for each convex polygonal obstacle such that the evader wins the game if the initial position of the pursuer lies outside this set. Figure 15 shows the evader in an environment containing a single hexagonal obstacle. The polygon in the center bounded by thick lines shows the region of possible pursuer win. In a similar way, we can generate a bounded set for a non-convex obstacle. Given a non-convex obstacle, we construct its convex hull. We can prove that Lemma 4 holds for

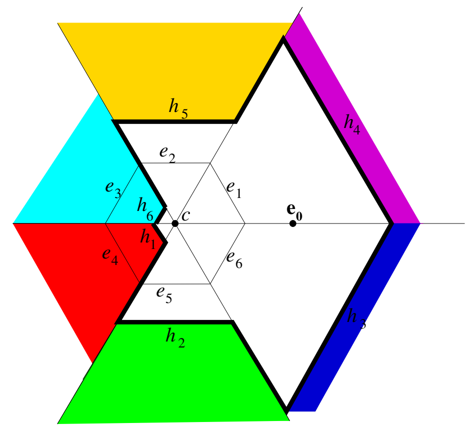

Fig. 15. *B set* for an environment consisting of a regular hexagonal obstacle and $a = 0.5$.

the convex hull. Finally, we can use Theorem 4 to prove the existence of a bounded set.

From the previous discussions, we conclude that any polygonal obstacle, convex or non-convex, restricts the set of initial positions from which the pursuer might win the game to a bounded set. Moreover, given the initial position of the evader and the ratio of the maximum speed of the evader to the pursuer, the bounded set can be obtained from the geometry of the obstacle by the construction used in the proof of Theorem 4. For any polygon in the environment, let us call the bounded set generated by it the *B set*. If the initial position of the pursuer lies outside the *B set*, the evader wins the game. For an environment containing multiple polygonal obstacles, we can compute the intersection of all *B sets* generated by individual obstacles. Since each *B set* is bounded, the intersection is a bounded set. Moreover, the intersection has the property that if the initial position of the pursuer lies outside the intersection, the evader wins the game. This leads to the following theorem.

Theorem 5. *Given the initial position of the evader, the set of initial positions from which the pursuer might win the game is bounded for an environment consisting of polygonal obstacles.*

Proof. The bounded set referred to in this theorem is the intersection of the *B sets* generated by the obstacles. If the initial pursuer position does not lie in the intersection it implies that it is not contained in all of the *B sets*. Hence, there exists at least one polygon in the environment for which the initial pursuer position does not lie in its *B set*. By Proposition 4, the evader has a winning strategy. Hence, the theorem follows. □

However, we still do not know the result of the game for all initial positions of the pursuer inside the intersection. However, we can find better approximation schemes and reduce the size of the region in which the result of the game is unknown. In the next section, we present one such approximation scheme.

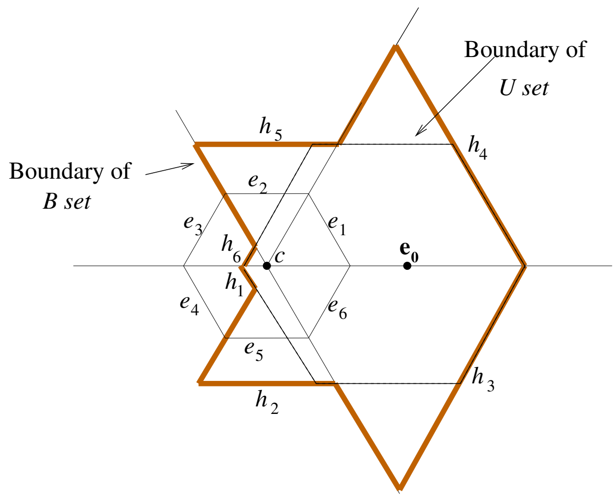

Fig. 16. *B set* and *U set* for an environment containing a regular hexagonal obstacle and $a = 0.5$. The polygon bounded by thick lines is the *B set* and the polygon bounded by thin lines is the *U set*

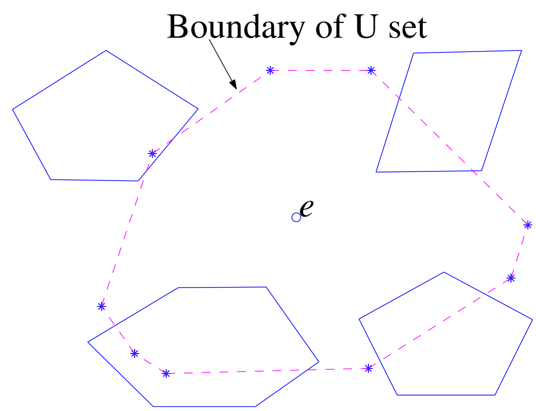

Fig. 17. *U set* for a general environment.

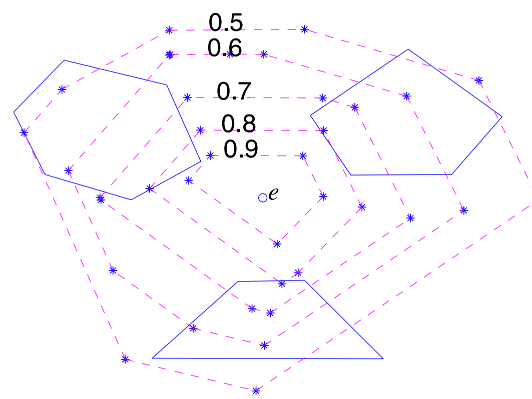

Fig. 18. *U set* for a various speed ratios of the evader to that of the pursuer.

3.1. *U set*

Now we present an approximation scheme that gives a tighter bound on the initial positions of the pursuer from which it might win the game. From Lemma 4, the evader wins the game if $\mathbf{p}_0 \in h_i^+$ for any edge. We can conclude that if $\mathbf{p}_0 \in \bigcup_{i=1}^n h_i^+$, the evader wins the game. Since $(\bigcup_{i=1}^n h_i^+)^c = \bigcap_{i=1}^n (h_i^+)^c = \bigcap_{i=1}^n h_i^-$, where S^c denotes the complement of set S, if $\mathbf{p}_0$ lies outside $\bigcap_{i=1}^n h_i^-$, the evader wins the game. Hence, the set of initial positions from where the pursuer might win the game is contained in $\bigcap_{i=1}^n h_i^-$. We call $\bigcap_{i=1}^n h_i^-$ the *U set*. An important point to note is that the intersection can be taken among any number of half-spaces. If the intersection is among the half-spaces generated by the edges of an obstacle, we call it the *U set* generated by the obstacle. If the intersection is among the half-spaces generated by all of the edges in an environment, we call it the *U set* generated by the environment.

The next theorem proves that the *U set* generated by a single obstacle is a subset of the *B set* and hence a better approximation.

Theorem 6. *For a given convex obstacle, the U set is a subset of the B set and hence bounded.*

Proof. Consider a point q that does not lie in the *B set*. From the construction of the *B set*, q must belong to some half-plane h_j^+. If $q \in h_j^+$, then $q \notin h_j^- \implies q \notin \cap_{i=1}^n h_i^-$. This implies that the complement of the *B set* is a subset of the complement of the *U set*. This implies that the *U set* is a subset of the *B set*. □

Figure 16 shows the *B set* and *U set* for an environment containing a regular hexagonal obstacle. In the appendix, we present a polynomial-time algorithm to compute the *U set* for an environment with polygonal obstacles. The overall time complexity of this algorithm is $O(n^2 \log n)$ where n is the number of edges in the environment. Figure 17 shows the evader in a polygonal environment. The region enclosed by the dashed lines is the *U set* generated by the environment for the initial position of the evader. The *U set* for any environment having polygonal obstacles is a convex polygon with at most n sides. Figure 18 shows the *U set* for an environment for various ratio of the maximum speed of the evader to that of the pursuer. In Figure 18, it can be seen that as the speed ratio between the evader and the pursuer increases, the size of the *U set* decreases. The size of the *U set* diminishes to zero at a critical speed ratio. At speed ratios higher than the critical ratio, the evader has a winning strategy for any initial position of the pursuer.

Before we proceed to the next theorem, we prove the following lemma.

Lemma 5. *For $a \leq 1$, the evader lies inside the U set.*

Proof. For $a \leq 1$, $\bar{v}_p \geq \bar{v}_e$. If the pursuer lies at the same position as the evader, its strategy to win is to maintain the same velocity as that of the evader. Hence, if the pursuer and the evader have the same initial position, the pursuer can track the evader successfully. Since all of the initial positions from which the pursuer can win the game must

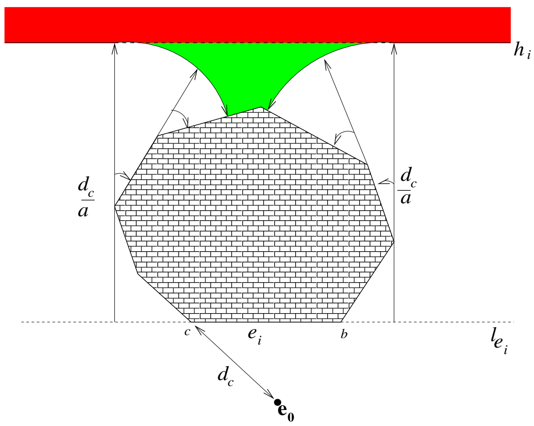

Fig. 19. A polygon in free space. The shaded Region 1 is obtained by using Lemma 4. The shaded Region 2 is added by using a better approximation scheme.

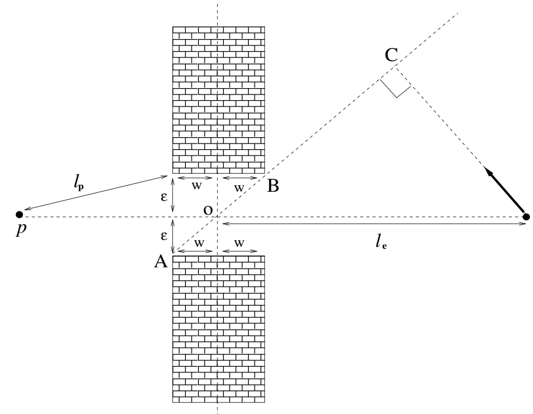

Fig. 20. An example to show that the SC is not a necessary condition for escape.

be contained inside the *U set*, the evader position must also be inside the *U set*. □

The following theorem provides a sufficient condition for escape of the evader in an environment containing obstacles using the *U set*.

Theorem 7. *If the U set does not contain the initial position of either the pursuer or the evader, the evader wins the game.*

Proof. From the definition of the *U set*, if the pursuer lies outside the *U set*, it loses. If the evader lies outside the *U set*, Lemma 5 implies $a > 1$. If $a > 1$, $\bar{v}_e > \bar{v}_p$. If $\bar{v}_e > \bar{v}_p$, the evader wins the game in any environment containing obstacles. Its winning strategy is to move on the convex hull of any obstacle. □

3.2. Discussion

In the previous sections, we have provided a simple approximation scheme for computing the set of initial pursuer positions from which the evader can escape based on the intersection of a family of half-spaces. A slight modification to the proposed scheme leads to a better approximation. In the proof of Lemma 4, we presented an algorithm to find a half-space for every edge of the polygon such that if the initial position of the pursuer lies in the half-space, the evader wins the game. All of the points in the half-space are at a distance greater than d_c/a from l_{e_i}. By imposing the condition that the minimum distance of the desired set of points from l_{e_i} in the free workspace should be greater than d_c/a, we can include more points in the decidable regions as shown in Figure 19. The figure shows an obstacle in free space. From the proof of Lemma 4, we obtain the half-space above h_i. By adding the new condition, Region 2 is included. When we repeat this for every edge, the set of initial positions from which the pursuer might win the game is reduced and leads to a better approximation. The boundary of the shaded region consists of straight lines and arcs of circles. The boundary of the desired set is obtained by computing the intersections among a collection of rays and arcs of circles generated by each edge. In this case a better approximation comes at the cost of expensive computation.

None of the approximation schemes we have suggested so far restrict the initial position of the pursuer to be in the evader's visibility region. This condition can be imposed by taking an intersection of the output of the approximation algorithm with the visibility polygon at the evader's initial position. Efficient algorithms exist for computing the visibility polygon of a static point in an environment (Goodman and Rourke 1997).

If $\mathbf{p_0}$ lies outside the *U set*, then the evader has at least one strategy to win the game. There might exist additional strategies for the evader to win the game. Since we are only concerned about the final outcome of the game, the evader is free to choose any of the strategies that leads to a favorable outcome.

The SC is not a necessary condition for escape. Figure 20 illustrates a scenario in which the evader breaks the line of sight with the pursuer even when SC is not satisfied. Consider the evader strategy in which it moves along the line perpendicular to line AB with speed $\bar{v}_p$. For the pursuer to keep the evader in sight , the pursuer must reach the vertex A before the evader can reach the point C. The minimum time required by the pursuer to reach A is $l_p/\bar{v}_p$. The time required for the evader to reach C is $\epsilon l_e/(\bar{v}_e\sqrt{\epsilon^2+w^2})$. Therefore, in order to break the line of sight, the evader must satisfy the following condition

$$\frac{\epsilon l_e}{\bar{v}_e\sqrt{\epsilon^2+w^2}} < \frac{l_p}{\bar{v}_p}$$
$$\Rightarrow \epsilon < \frac{w}{\sqrt{\frac{l_e^2}{a^2 l_p^2}-1}}.$$

From the above inequalities, we can conclude that we can choose an ϵ small enough for the evader to break the line of

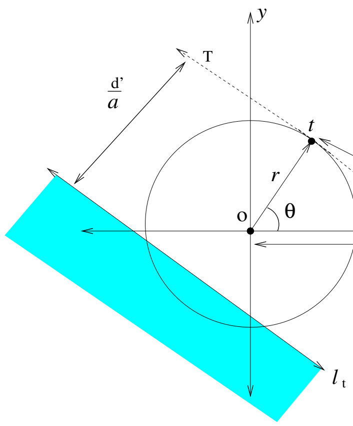

Fig. 21. A circular obstacle in the plane. The shaded region denotes the initial positions of the pursuer from which the evader can win the game by breaking the line of sight at the point t.

sight with the pursuer irrespective of the value of $\frac{l_e}{l_p}$. Therefore, the evader can win the game even if $\frac{l_e}{l_p}$ is large enough to violate the SC. Since the SC is not a necessary condition for escape, the U set provides an upper bound on the size of the capture set. In Section 5, we derive a lower bound on the size of the capture set from a sufficient condition for tracking.

In the next section, we present the construction of the U set for non-polygonal environments.

4. *U set* for specific environments

In the real world we encounter many non-polygonal obstacles in the environment. One of the common obstacles in an environment are circular columns and pillars that project to a disk in a plane. In this section we compute the *U set* for a disc in a plane and then extend the procedure to compute the *U set* for obstacles whose boundary have a well defined tangent at each point.

4.1. Disk in a plane

Consider an environment consisting of an obstacle in the shape of a disk of radius r in free space. Refer to Figure 21. Let C denote the boundary of the obstacle. Let $\mathbf{e}_0$ denote the initial position of the evader. Let O be the center of the circular obstacle. The distance between O and $\mathbf{e}_0$ is d_0. O is also the origin of the world reference frame. The x-axis of the world reference frame passes through $\mathbf{e}_0$ and O. Let t be a point on the boundary of the obstacle such that Ot makes an angle θ with the x-axis. Let d' denote the distance between t and $\mathbf{e}_0$. Let T denote the tangent to the circle at the point t. Let l_t be a line at a distance of d'/a from T in the same half-space of T as the obstacle. By SC, the evader will win the game if the pursuer lies in the half-space shown by the shaded region. The equation of line l_t is $y + x\cot\theta - (r - d'/a)\csc\theta = 0$. For each point t on the circle, we can find such a line l_t and the corresponding half-space l_t^+. The *U set* is defined as $\bigcap_{t\in C} l_t^-$. If the initial position of the pursuer lies outside the *U set*, the evader wins the game. Let $l(x, y, \theta)$ denote the family of lines l_t generated by all points t lying on C. Owing to symmetry of the environment about the x-axis, the *U set* is symmetric about the x-axis. We construct the part of the *U set* generated as θ increases from 0 to π.

Let B denote the boundary of the *U set*.

Theorem 8. *The boundary B is the envelope of the family of lines $l(x, y, \theta)$.*

Proof. Consider any point q on B. The point q belongs to some line in the family of lines $l(x, y, \theta)$ since it belongs to the boundary. Let that line be l_q. Here l_q has to be tangent to the boundary B or otherwise there is a neighborhood around q in which B lies in both the half-spaces generated by l_q. Since q is any point on B, it is true for all points q on B that the tangent to B at q belongs to the family of lines $l(x, y, \theta)$. A curve satisfying this property is the envelope to the family of lines $l(x, y, \theta)$. □

We can find the envelope of a family of lines $l(x, y, \theta)$ by solving the following equations simultaneously

$$l(x, y, \theta) = y + x\cot\theta - (r - \frac{d'}{a})\csc\theta = 0, \quad (1)$$

$$\frac{\partial l}{\partial \theta} = 0. \quad (2)$$

The distance d' as a function of θ is given by

$$d'(\theta) = \begin{cases} \sqrt{r^2 + d_0^2 - 2rd_0\cos\theta} & \text{if } \theta \le \theta_0 \\ \sqrt{d_0^2 - r^2} + r(\theta - \theta_0) & \text{if } \theta \ge \theta_0 \end{cases}$$

where $\theta_0 = \cos^{-1}(r/d)$

4.1.1. Case 1 ($\theta \le \theta_0$) Substituting Equation (1) into Equation (2) gives us

$$x = \left(r - \frac{\sqrt{r^2 + d_0^2 - 2rd_0\cos\theta}}{a} \right)\cos\theta + \frac{rd_0\sin^2\theta}{a\sqrt{r^2 + d_0^2 - 2rd_0\cos\theta}},$$

$$y = \left(r - \frac{\sqrt{r^2 + d_0^2 - 2rd_0\cos\theta}}{a} \right)\sin\theta - \frac{rd_0\sin\theta\cos\theta}{a\sqrt{r^2 + d_0^2 - 2rd_0\cos\theta}}.$$

4.1.2. Case 2 ($\theta \ge \theta_0$) Substituting Equation (1) into Equation (2) gives us

$$x = \left(r - \frac{\sqrt{d_0^2 - r^2} + r(\theta - \theta_0)}{a} \right)\cos\theta + \frac{\sin\theta}{a},$$

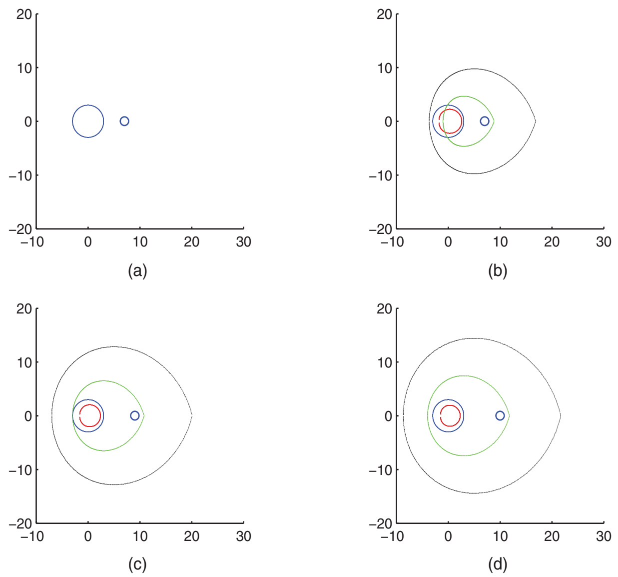

Fig. 22. (a) A disk-like obstacle with the initial position of the evader. The smaller circle is the evader. (b)–(d) The boundary of the *U sets* for the obstacle with increasing distance between the evader and the center of the disk. In (b), (c) and (d), the black boundary is for the case when $a = 0.5$, the cyan boundary is for the case when $a = 1$ and the red boundary is for the case when $a = 10$.

$$y = \left(r - \frac{\sqrt{d_0^2 - r^2} + r(\theta - \theta_0)}{a}\right)\sin\theta - \frac{\cos\theta}{a}.$$

Since B is symmetrical about the x-axis, the other half of B is obtained by reflecting the above curves about the x-axis. Figure 21 shows the boundary of the *U set* for a disk of radius 3 units. Figures 22(b)–(d) show the boundary of the *U set* for varying distance between the evader and the obstacle. In each of these figures, the boundary of the *U set* is shown for three different values of a. We can see that for $a \leq 1$, the evader lies inside the *U set*.

The above procedure can be used to construct the *U set* for any convex obstacle whose boundary has a well-defined tangent at every point. If the boundary is given by the equation $f(x, y) = 0$ where $f(x, y)$ is such that $\partial f/\partial x$ and $\partial f/\partial y$ exist for all points, the procedure to generate the boundary of the *U set* is as follows.

1. Given any point t on the boundary, find the equation of the line l_t as defined above.
2. Find the family $l(x, y, \theta)$ of lines generated by l_t as t moves on the boundary of the obstacle. θ is a parameter that defines t.
3. The envelope of the family $l(x, y, \theta)$ is the boundary of the *U set*. This is true since the proof of Proposition 8 does not depend on the shape of the obstacle and hence Proposition 8 is true for any obstacle.

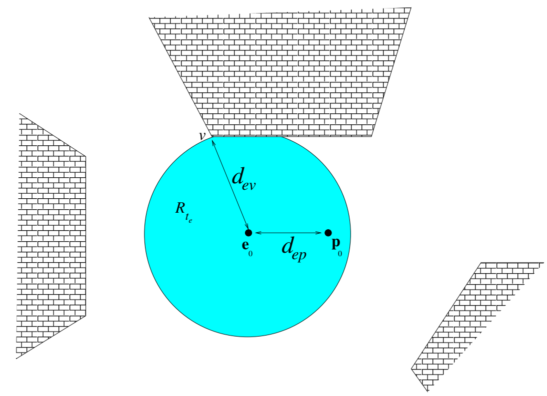

Fig. 23. The shaded portion of the disk centered at the initial position of the evader shows the region in which the pursuer can capture the evader and therefore maintain line of sight forever. The radius of the disk depends on the ratio of the maximum speeds between the pursuer and the evader and the distance between the evader and the nearest convex vertex.

In the next section we present an approximate bound on the initial positions of the pursuer from which it can track the evader.

5. Sufficient condition for surveillance

In this section we present a sufficient condition for a pursuer to track the evader. If $\overline{v}_e > \overline{v}_p$, the evader wins the game for any initial position of the pursuer. So a necessary condition for successful tracking is $\overline{v}_e \leq \overline{v}_p$. A plausible strategy for the pursuer to track the evader would be to catch the evader in a finite time and then move with the same velocity as that of the evader. The latter is possible since we assumed that the pursuer knows the instantaneous velocity of the evader at all times. Using the above ideas, we obtain the following sufficient condition of tracking.

Sufficient condition for tracking. Let d_{ev} denote the distance to the nearest convex vertex from $\mathbf{e}_0$ and $d_{ep} = \|\mathbf{e}_0 - \mathbf{p}_0\|$ (Figure 23 shows an example). A sufficient condition for tracking is

$$\min\left\{\frac{1-a}{a}, 1\right\} > \frac{d_{ep}}{d_{ev}}.$$

Proof. The minimum time required by the evader to reach the nearest convex vertex is $t_e = d_{ev}/\overline{v}_e$. Let R_{t_e} denote the set of points in the free workspace reachable by the evader, starting at $\mathbf{e}_0$, in time t_e, i.e. R_{t_e} consists of points $\mathbf{x} \in \mathbb{R}^2$ such that $\|\mathbf{x} - \mathbf{e}_0\| \leq d_{ev}$. Here R_{t_e} is convex since it cannot contain any convex vertex of the environment in its interior as t_e is the time required by the evader to reach the nearest convex vertex. Hence, any motion of the pursuer that keeps it inside R_{t_e} until time t_e ensures that it can see the evader until time t_e due to convexity of R_{t_e}. This leads to the following condition at the beginning of the game: $d_{ep}/d_{ev} < 1$.

Therefore, if the pursuer has a strategy to remain inside R_{t_e} until it catches the evader, it can win the game. Consider a strategy for the pursuer in which it moves directly towards the evader with speed $\bar{v}_p$. At any time $t \le t_e$, the following holds:

$$\begin{aligned}\|\mathbf{p}(t)-\mathbf{e}_0\| &\le \|\mathbf{p}(t)-\mathbf{e}(t)\| + \|\mathbf{e}(t)-\mathbf{e}_0\| \\ &\le \bar{v}_e t + d_{ep} - (\bar{v}_p - \bar{v}_e)t \\ &\le d_{ep} + (2\bar{v}_e - \bar{v}_p)t_e \\ &\le d_{ep} + (2\bar{v}_e - \bar{v}_p)\frac{d_{ev}}{\bar{v}_e} \\ &= d_{ep} + \frac{(2a-1)}{a} d_{ev} \\ &\le \frac{(1-a)}{a} d_{ev} + \frac{(2a-1)}{a} d_{ev} \\ &= d_{ev}. \end{aligned} \tag{3}$$

Hence, at all times $t \le t_e$, the pursuer remains inside R_{t_e} which ensures that the pursuer can see the evader until time t_e. If the pursuer follows the strategy to move directly towards the evader with speed $\bar{v}_p$, the time required by the pursuer to catch the evader is $t_p \le d_{ep}/(\bar{v}_p - \bar{v}_e)$. If $t_p \le t_e(\Rightarrow (1-a)/a > d_{ep}/d_{ev})$, the pursuer remains inside R_{t_e} until it catches the evader. This ensures that the pursuer can see the evader for all times before catching it while following the strategy. Therefore, the two conditions needed to be satisfied for the above strategy to work are $d_{ep}/d_{ev} < 1$ and $(1-a)/a > d_{ep}/d_{ev}$. Hence, the sufficient condition follows. □

From the sufficient condition, we can conclude that if the pursuer lies within a disk of radius $\min\{(1-a)d_{ev}/a, d_{ev}\}$ centered on the initial position of the evader, the pursuer can track the evader for all times. We can see that as $a \to 1$, the radius of the disk tends to zero. Moreover, for any $a < 0.5$, the radius of the disk is d_{ev}.

6. Bounds on the capture region

Given the initial position of the evader, we presented the construction of the *U set* in Section 3 that provides an upper bound on the set of initial positions of the pursuer from which it can track the evader. In the previous section, we presented the construction of a region that provides a lower bound on the size of the initial positions of the pursuer from which it can capture the evader and track it forever. In this section, we propose a measure to quantify the gap between the upper and the lower bound and, further, compute it.

Let A_c^u denote the area of the *U set* generated by the obstacles in the environment. Let A_c^l denote the area of the capture region presented in the previous section. Let ρ be defined as the following quantity

$$\rho = \frac{A_c^u}{A_c^l}. \tag{4}$$

From the above definition, the value of ρ cannot be less than one. If $\rho = 1$, the outcome of the game is completely known for all initial positions of the pursuer. As ρ increases, the gap between the upper and lower bound grows larger. Therefore, the quantity $\rho - 1$ is a measure of the effectiveness of the approximation scheme. In the rest of the section, we present an analysis to compute an upper bound for ρ based on the initial position of the evader and the geometry of the obstacles present in the environment.

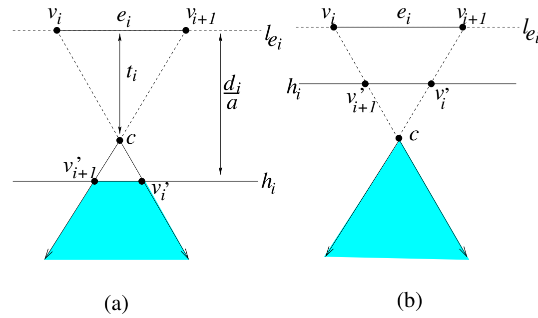

Fig. 24. The sector associated with an edge e_i. In (a), the edge e_i contributes to the area of the *B set*. In (b), the edge does not contribute any area to the *B set*.

From Section 3, we know that the *U set* generated by all of the obstacles is a subset of the *U set* generated by a single obstacle. From Theorem 6, we know that for a single obstacle, the area of the *B set* is an upper bound on the area of the *U set*. Therefore, computing the area of the *B set* generated by a single obstacle provides an upper bound on the area of the *U set* generated by all of the obstacles in the environment, A_c^u.

In order to compute the area of the *B set* for a single obstacle, we consider a convex obstacle in the environment. Let the vertices of the obstacle be denoted as $v_1, \ldots, v_n$ and the edges be denoted as $e_1, \ldots, e_n$. Let (x_i, y_i) be the coordinates of the vertex v_i and $y = m_i x + c_i$ be the equation of the edge e_i with respect to a coordinate axes in the plane. Let h_i be the line parallel to e_i at a distance d_i/a as defined in Section 3. The results in Section 3 hold for any arbitrary point c chosen inside the obstacle. In the subsequent analysis, we choose c to be the centroid of the polygon. The coordinates of c are given by the following expression

$$x_c = \frac{1}{n}\sum_{i=1}^{n} x_i, \quad y_c = \frac{1}{n}\sum_{i=1}^{n} y_i.$$

Refer to Figure 24. Let t_i be the perpendicular distance from c to e_i. Therefore, the perpendicular distance of h_i from c is $|\frac{d_i}{a} - t_i|$. Let A_i and A'_i denote the area of $\Delta c v_i v_{i+1}$ and $\Delta c v'_i v'_{i+1}$, respectively. Since $A_i \sim A'_i$, $A'_i = A_i[\frac{d_i}{a t_i} - 1]^2$. If c lies between h_i and e_i then A'_i contributes to the area of the *B set*. This case is shown in Figure 24(a). On the other hand, if h_i lies between c and e_i then e_i does not contribute any area to the *B set*. This case is shown in Figure 24(b). In either case, the maximum area contributed by e_i to the *B set* is equal to A'_i. Therefore, by construction, the area of the *B* *set* is bounded from above by the following expression

$$\sum_{i=1}^{n} A_i \left[\frac{d_i}{at_i} - 1\right]^2.$$

The area of the entire polygon A is related to A_i by the relation, $\sum_{i=1}^{n} A_i = A$. Therefore, an upper bound on the area of the *B set* is given by the following expression

$$\begin{aligned}\sum_{j=1}^{n}\left[\frac{d_j}{at_j} - 1\right]^2 A_j &\le \sum_{j=1}^{n}\left(\max_i\left[\frac{d_i}{at_i} - 1\right]^2\right) A_j \\ &\le \left(\max_i\left[\frac{d_i}{at_i} - 1\right]^2\right)\sum_{i=1}^{n} A_j \\ &\le \left(\max_i\left[\frac{d_i}{at_i} - 1\right]^2\right) A \\ &\le \left(\max_i\left[\frac{d_i^2}{a^2 t_i^2}\right] + 1\right) A \\ &\le \left(1 + \frac{1}{a^2}\frac{\max_i d_i^2}{\min_i t_i^2}\right) A. \qquad (5)\end{aligned}$$

The perpendicular distance of a point (x', y') from the line $y = m_i x + c_i$ is given by the following expression Beyer (1987)

$$\left|\frac{y' - m_i x' - c_i}{\sqrt{1 + m_i^2}}\right|.$$

Since c is chosen as the centroid of the polygon and all of the vertices of a convex polygon lie on the same side of l_{e_i}, t_i is given by the following expression

$$t_i = \frac{1}{n}\left|\sum_{k=1}^{n}\frac{y_k - m_i x_k - c_i}{\sqrt{1 + m_i^2}}\right|.$$

The above expression is the average distance of all of the vertices of the polygon from e_i and, therefore, is positive for all edges. This implies that $\min_i t_i > 0$ because the minimum is computed over a finite set of positive real numbers. From the above discussion, we can conclude that A_c^u is bounded from above by the following expression:

$$A_c^u \le A\left(1 + \frac{1}{a^2}\frac{\max_i d_i^2}{\min_i t_i^2}\right). \qquad (6)$$

Since $\min_i t_i > 0$, the right-hand side is finite for all values of $a > 0$. From the previous section, we can also conclude the following:

$$A_c^l \ge k\pi d_{ep}^2 \qquad (7)$$

where $k = [\min\{(1 - a)/a, 1\}]^2$. From Equations (4), (6) and (7), an upper bound for the quantity ρ is given by the following expression

$$\rho \le \frac{1}{2k\pi d_{ep}^2}\left(1 + \frac{1}{a^2}\frac{\max_i d_i^2}{\min_i t_i^2}\right) A.$$

A similar bound can be computed for a polygon that is non-convex by repeating the above analysis for the convex hull of the polygon. All of the quantities in the bound are parameters of the environment and the ratio of the maximum speed of the evader to that of the pursuer. For values of a near to zero, the term $(1+(1/a^2)(\max_i d_i^2/\min_i t_i^2))$ dominates the right-hand side. For values of a near to one, the term $1/2k\pi d_{ep}^2$ dominates the right-hand side. The above bound has been computed for a single obstacle. Therefore, it can be made tighter by computing the expression on the right-hand side for all of the obstacles in the environment and choosing the minimum value.

7. Conclusion and future research

In this paper we have addressed the problem of visibility-based target tracking in an environment with obstacles. Prior work in this area has mainly focused on obtaining tracking strategies by using different techniques to optimize a metric that partially encodes the effect of environmental occlusion due to obstacles. In this work, we have viewed this problem as a pursuit–evasion problem and used purely algorithmic techniques to analyze it. For an environment having a single corner, we have shown that the target-tracking problem is completely decidable. We have presented a partition of the visibility region of the observer. Depending on the cell of the partition in which the target lies, we have provided strategies for the target to escape the visibility region of the pursuer or for the observer to track the target for all future time. Given the initial position of the target, we also presented the solution to the problem of computing the positions of the observer for which the target can escape the visibility region of the observer. These results have been provided for varying speeds of the observer and the target. Based on these partitions we presented a sufficient condition for escape of the target and a polynomial time algorithm to approximate the *escape set*. In addition, we have presented an extension of the approximation schemes to obstacles with smooth boundaries. We have also provided a sufficient condition for tracking that leads to a region around the target from which the observer can catch the evader and track it. Finally, we provided an upper bound on the area of the region from which the outcome of the game is unknown using the proposed approximation schemes.

Given the complete map of the environment, our results depend only on the initial position and the maximum speeds of the pursuer and evader. Hence, our results hold for various settings of the problem such as an unpredictable or predictable evader (LaValle et al. 1997) or localization uncertainties in the future positions of the players (Fabiani and Latombe 1999). Although the bounds presented in this paper are conservative for all of the above cases.

This work unfolds many research directions for our future work and we would like to highlight some of them. One direction of future research is to take into account the constraints in the motion and visibility of the observer. In this work we assume a kinematic system for the observer. A more appropriate approach is to include dynamics in the motion model of the observer. Another direction of research deals with tracking a target in environments for which the observer has local or limited information. In the past, researchers have addressed the problem of minimizing the time to enter the star region for a robot without any knowledge about the geometry of the corner (Icking et al. 1993). As a future work, it might be possible to extend the results for the case in which a mobile target is present in the environment. Finally, in real-time systems, actuation takes place at discrete moments in time but the dynamics of the system are continuous in time. Such systems lie in the category of hybrid systems or discrete-event dynamic systems. These areas provide an entirely new set of techniques for analysis.

Funding

This material is based in part upon work supported by the National Science Foundation under award CNS 0931871.

Acknowledgment

Parts of this work have appeared in *Robotics: Science and Systems III* and *Robotics: Science and Systems IV*. The first author is thankful to Salvatore Candido for his contribution in Section 2 of the paper, and Professor Steven LaValle for insightful comments regarding the initial sections of the paper.

References

Başar T and Olsder GJ (1995) *Dynamic Noncooperative Game Theory*, 2nd ed. London: Academic Press.

Bandyopadhyay T, Li Y, Ang M Jr and Hsu D (2004) Stealth Tracking of an Unpredictable Target among Obstacles. *Proceedings of the International Workshop on the Algorithmic Foundations of Robotics*.

Bandyopadhyay T, Li Y, Ang M Jr and Hsu D (2006) A greedy strategy for tracking a locally predicatable target among obstacles. In *Proceedings IEEE International Conference on Robotics and Automation (ICRA'02)*, pp. 2342–2347.

Bandyopadhyay T, Rong N, Ang M Jr, Hsu D and Lee W (2009) Motion planning for people tracking in uncertain and dynamic environments. In *IEEE International Conference on Robotics and Automation, Workshop on People Detection and Tracking*.

Beyer WH (1987) *CRC Standard Mathematical Tables*. Boca Raton, FL: CRC Press.

Bhattacharya S, Başar T and Hovakimyan N (2011) Non-cooperative strategies for connectivity maintenance among autonomous mobile agents. In *IEEE Conference on Decision and Control* (under review).

Bhattacharya S, Candido S and Hutchinson S (2007) Motion strategies for surviellance. In *Robotics: Science and Systems III*.

Bhattacharya S and Hutchinson S (2008) Approximation schemes for two-player pursuit evasion games with visibility constraints. In *Proceedings of Robotics: Science and Systems IV*, Zurich, Switzerland.

Bhattacharya S and Hutchinson S (2010) On the existence of nash equilibrium for a two player pursuit–evasion game with visibility constraints. *Int J Robotics Res* 29: 831–839.

Bhattacharya S, Hutchinson S and Başar T (2009) Game theoretic analysis of pursuit–evasion game with visibility constraints. In *American Control Conference*, pp. 373–378.

Briggs AJ and Donald BR (1996) Robust geometric algorithms for sensor planning. In *Proceedings of the Second Workshop on Algorithmic Foundations of Robotics*.

Cheung W (2005) *Constrained Pursuit-Evasion in the Plane*. Master's thesis, University of British Columbia.

de Berg M, van Kreveld M, Overmars M and Schwarzkopf O (1997) *Computational Geometry—Algorithms and Applications*. Heidelberg: Springer-Verlag.

Efrat A, Gonzalez-Banos H, Kobourov S and Palaniappan L (2003) Optimal strategies to track and capture a predictable target. *Proceedings IEEE International Conference on Robotics and Automation, 2003 (ICRA'03)*, Vol. 3.

Espiau B, Chaumette F and Rives P (1992) A new approach to visual servoing in robotics. *IEEE Transactions on Robotics and Automation* 8: 313–326.

Fabiani P and Latombe J (1999) Dealing with geometric constrints in game-theoretic planning. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence*, Vol. 2, Tolouse, France, pp. 731–736.

Gonzalez-Banos H, Lee C and Latombe J (2002) Real-time combinatorial tracking of a target moving unpredictably among obstacles. *Proceedings of the IEEE International Conference on Robotics and Automation, 2002 (ICRA'02)*, Vol. 2.

Goodman JE and Rourke JO (1997) *Handbook of Discrete and Computational Geometry*. New York: CRC Press.

Hsu D, Lee W and Rong N (2008) A point-based POMDP planner for target tracking. In *Proceedings IEEE International Conference on Robotics and Automation*, pp. 2644–2650.

Hutchinson SA, Hager GD and Corke PI (1996) A tutorial on visual servo control. *IEEE Trans Robotics Automat* 12: 651–670.

Icking C, Klien R and Ma L (1993) The optimal way for looking around a corner. In *IEEE–IEE Vehicle Navigation and Information Systems Conference*, Ottawa, Canada, pp. 547–550.

Isaacs R (1965) *Differential Games: A Mathematical Theory with Applications to Warfare and Pursuit, Control and Optimization*. New York: Wiley.

Jung B and Sukhatme G (2002) Tracking targets using multiple robots: the effect of environment occlusion. *Autonomous Robots* 13: 191–205.

Kolling A and Carpin S (2006) Multirobot cooperation for surveillance of multiple moving targets - a new behavioral approach. In *IEEE Conference on Robotics and Automation*, pp. 1311–1316.

Kolling A and Carpin S (2007) Cooperative observation of multiple moving targets: an algorithm and its formalization. *Int J Robotics Res* 26: 935–953.

LaValle SM, Gonzalez-Banos HH, Becker C and Latombe JC (1997) Motion strategies for maintaining visibility of a moving target. In *Proceedings 1997 IEEE International Conference on Robotics and Automation*, Vol. 1, Albuquerque, NM, pp. 731–736.

LaValle SM and Hinrichsen J (2001) Visibility-based pursuit–evasion: The case of curved environments. *IEEE Trans Robotics Automat* 17: 196–201.

Lewin J (1994) *Differential Games: Theory and Methods for Solving Game Problems with Singular Surfaces*. London: Springer-Verlag.

Li T, Lien J, Chiu S and Yu T (1997) Automatically generating virtual guided tours. In *Computer Animation Conference*, pp. 99–106.

Luke S, Sullivan K, Panait L and Balan G (2005) Tunably decentralized algorithms for cooperative target observation. In *International Joint Conference on Autonomous Agents and Multiagent Systems*, pp. 911–917.

Malis E, Chaumette F and Boudet S (1999) 2D 1/2 Visual Servoing. *IEEE Trans Robotics Automat* 15: 238–250.

Marchand E, Bouthemy P, Chaumette F and Moreau V (1999) Robust real-time visual tracking using a 2D–3D model-based approach. In *IEEE International Conference on Computer Vision (ICCV'99)* 1: 262–268.

Murrieta-Cid R, Gonzalez-Banos HH and Tovar B (2002) A reactive motion planner to maintain visibility of unpredictable targets. In *Proceedings IEEE International Conference on Robotics and Automation, 2002 (ICRA '02)*, Vol. 4, pp. 4242–4248.

Murrieta-Cid R, Monroy R, Hutchinson S and Laumond JP (2008) A complexity result for the pursuit–evasion game of maintaining visibility of a moving evader. In *IEEE International Conference on Robotics and Automation*.

Murrieta-Cid, R., Muppirala, T., Sarmiento, A., Bhattacharya, S., and Hutchinson, S. (2007) Surveillance strategies for a pursuer with finite sensor range. *Int J Robotics Res* 26: 233–253.

Parker L (2002) Algorithms for Multi-Robot Observation of Multiple Targets. *J Autonomous Robots* 12: 231–255.

Schwartz JT and Sharir M (1987) On the piano movers' problem: 1. the case of a two-dimensional rigid polygon body moving amidst polygonal barriers. *Commun Pure Appl Math* 36: 345–398.

Sipser M (1997) *Introduction to Theory of Computation*. Boston, MA: PWS.

Sung TL and Um TY (1996) Practical guidance for homing missiles with bearings-only measurements. *IEEE Trans Aerospace Electron Syst* 32: 434–443.

Tang Z and Ozguner U (2005) Motion planning for multitarget surveillance with mobile sensor agents. *IEEE Trans Robotics Automat* 21: 898–908.

Tekdas O, Yang W and Isler V (2009) Robotic routers: Algorithms and implementation. *Int J Robotics Res*, accepted.

Appendix: Algorithm for generating the *U set*

Algorithm CONSTRUCTUSET($S, a, (x_e, y_e)$)
Input: A set S of disjoint polygonal obstacles, the evader position $r_e = (x_e, y_e)$, ratio of maximum evader speed to maximum pursuer speed a
Output: The coordinates of the vertices of the *U set*

1. For every edge e_i in the environment with end points a_i, b_i
2. $l_1 = \text{DIJKSTRA}(\text{VG(S)}, r_e, a_i)$
3. $l_2 = \text{DIJKSTRA}(\text{VG(S)}, r_e, b_i)$
4. $d_{e_i} = \frac{\min(l_1, l_2)}{a}$
5. Find the equation of h_i using Lemma 4.
6. INTERSECTHALFPLANES($h_1^-, \ldots, h_n^-$).

The subroutine VG(S), computes the visibility graph of the environment S. The subroutine DIJKSTRA(G,I,F) computes the shortest distance between nodes I and F in graph G. The subroutine INTERSECTHALFPLANES($h_1^-, \ldots, h_n^-$) computes the intersection of the half planes $h_1^-, \ldots, h_n^-$ (de Berg et al. 1997). The time complexity of the above algorithm is $O(n^2 \log n)$, where n is the number of edges in the environment.