# Complete Information Pursuit Evasion in Polygonal Environments [*]

**Kyle Klein and Subhash Suri**

Department of Computer Science
University of California
Santa Barbara, CA 93106

## Abstract

Suppose an unpredictable evader is free to move around in a polygonal environment of arbitrary complexity that is under full camera surveillance. How many pursuers, each with the same maximum speed as the evader, are necessary and sufficient to guarantee a successful capture of the evader? The pursuers always know the evader's current position through the camera network, but need to physically reach the evader to capture it. We allow the evader the knowledge of the current positions of all the pursuers as well—this accords with the standard worst-case analysis model, but also models a practical situation where the evader has "hacked" into the surveillance system. Our main result is to prove that *three* pursuers are always *sufficient and sometimes necessary* to *capture* the evader. The bound is independent of the number of vertices or holes in the polygonal environment.

## Introduction

Pursuit-evasion games provide an elegant setting to study algorithmic and strategic questions of exploration or monitoring by autonomous agents. Their mathematical history can be traced back to at least 1930s when Rado posed the now-classical Lion-and-Man problem: a lion and a man in a closed arena have equal maximum speeds; what tactics should the lion employ to be sure of his meal? The problem was settled by Besicovitch who showed that the man can escape regardless of the lions strategy. An important aspect of this pursuit-evasion problem, and its solution, is the assumption of continuous time: each players motion is a continuous function of time, which allows the lion to get arbitrarily close to the man but never capture him. If, however, the players move in discrete time steps, taking alternating turns but still in continuous space, the outcome is different, as first conjectured by Gale (Guy 1991) and proved by Sgall (Sgall 2001).

A rich literature on pursuit-evasion problem has emerged since these initial investigations, and the problems tend to fall in two broad categories: discrete space, where the pursuit occurs on a graph, and continuous space, where the pursuit occurs in a geometric space. Our focus in this paper is on the latter: visibility-based pursuit in a polygonal environment in two dimensions. There exist simply-connected $n$-gons that may require $\Omega(\log n)$ pursuers in the worst-case to detect a single, arbitrarily fast moving evader, and $O(\log n)$ pursuers also always suffice for all $n$ vertex simple polygons (Guibas et al. 1999). When the polygon has $h$ holes, the number of necessary and sufficient pursuers turns out to be $O(\sqrt{h} + \log n)$ (Guibas et al. 1999). However, these results hold only for *detection* of the evader, not for the capture.

For capturing the evader, it is reasonable to assume that the pursuers and the evader all have the same maximum speed. Under this assumption, it is shown by Isler et al. (Isler, Kannan, and Khanna 2005) that two pursuers can capture the evader in a *simply-connected* polygon using a *randomized* strategy whose expected search time is polynomial in $n$ and the diameter of the polygon. When the polygon has holes, no non-trivial upper bound is known for capturing the evader. For instance, we do not even know if $O(h)$ pursuers are able to capture the evader. Because visibility-based pursuit allows *unbounded* line-of-sight visibility regardless of the distance, it is unclear how to map a detection strategy to a capture strategy.

In this paper, we attempt to *disentangle* these two orthogonal issues inherent in pursuit evasion: *localization*, which is purely an informational problem, and *capture*, which is a problem of planning physical moves. In particular, we ask how complex is the capture problem *if the evader localization is available for free*? In other words, suppose the pursuers have complete information about the evader's current position, how much does it help them to capture the evader? Besides being a theoretically interesting question, the problem is also a reasonable model for many practical settings. Given the rapidly dropping cost of electronic surveillance and camera networks, it is now both technologically and economically feasible to have such monitoring capabilities. These technologies enable cheap and ubiquitous detection and localization, but in case of intrusion, a physical capture of the evader is still necessary.

### Our Contribution

Our main result is that under such a complete information setting, *three pursuers* are always sufficient to capture an

---

[*] This research was supported in part by NSF grant IIS 0904501.

equally fast evader in a polygonal environment with holes, using a *deterministic* strategy. The bound is independent of the number of vertices $n$ or the holes of the polygon, although the capture time depends on both $n$ and the diameter of the polygon. Complementing this upper bound, we also show that there exists polygonal environments that require at least three pursuers to capture the evader even with full information.

### Related Work

There is an enormous literature on pursuit evasion and related problems (Parsons 1976; Aigner and Fromme 1984; Isler and Karnad 2008; Fomin, Golovach, and Kratochvíl 2008; Sachs, Rajko, and LaValle 2004; Halpern 1969; Bienstock and Seymour 1991; Isler, Kannan, and Khanna 2006; Park, Lee, and Chwa 2001; Suzuki and Yamashita 1992; LaPaugh 1993). This research tends to fall into two distinct categories: geometry-based and graph-based. The former assumes a continuous model of space, typically a polygon, while the latter assumes a discrete graph model where agents move along edges. The graphs provide a very general setting but can suffer from two shortcomings: one, the generality leads to weak upper bounds and, two, they fail to model many restrictions imposed by the geometry of two-dimensional world. Thus, for instance, determining the search number of cop-number or a general graph remains a difficult open problem despite decades of research.

In visibility-based pursuit, a seminal paper (Guibas et al. 1999) shows that $\Theta(\log n)$ pursuers are both necessary and sufficient in worst-case for a *simply-connected* $n$-vertex polygon. Most of the existing work in polygon searching, however, is on *detection* and not capture. The only relevant result on capture is by Isler et al. (Isler, Kannan, and Khanna 2005) showing that in *polygons without holes* two pursuers can achieve both detection and capture. When the environment has holes, it is not even known how many pursuers are sufficient to capture an evader, even though a tight bound of $\Theta(\sqrt{h} + \log n)$ for detection is known. In one important aspect, polygon searching is fundamentally different from graph searching: *re-contamination* is unavoidable in polygons, in general, while graphs can always be searched *optimally* without re-contamination (Guibas et al. 1999).

Our work bears some resemblance to, and is inspired by, the result of Aigner and Fromme (Aigner and Fromme 1984) on planar unweighted graphs, showing that graph searching on planar graph requires 3 cops. By contrast, our search occurs in continuous Euclidean plane, and players can move to any position within distance one. Thus, while our bounds are the similar, the proof techniques and technical details are quite different.

## The Problem Formulation

We assume that an evader $e$ is free to move in a two-dimensional closed polygonal environment $P$, which has $n$ vertices and $h$ holes. A set of pursuers, denoted $p_1, p_2, \ldots$, wish to capture the evader. All the players have the same maximum speed, which we assume is normalized to 1. The bounds in our algorithm depend on the number of vertices $n$ and the diameter of the polygon, $\mathrm{diam}(P)$, which is the maximum distance between any two vertices of $P$ under the shortest path metric.

For the sake of notational brevity, we also use $e$ to denote the current position of the evader, and $p_i$ to denote the position of the $i$th pursuer. We model the pursuit-evasion as a continuous space, discrete time game: the players can move anywhere inside the polygon $P$, but they take turns in making their moves, with the evader moving first. In each move, a player can move to any position whose shortest path distance from its current position is at most one; that is, within *geodesic disk* of radius one. On the pursuers' move, all the pursuers can move simultaneously and independently. We say that $e$ is successfully captured when some pursuer $p_i$ becomes collocated with $e$.

In order to focus on the complexity of the capture, we assume a complete information setup: each pursuer knows the location of the evader at all times. We also endow the evader the same information, so $e$ also knows the locations of all the pursuers. In addition, both sides know the environment $P$, but neither side knows anything about the *future* moves or strategies of the other side.

We begin with a high level description of the capture strategy, followed by its technical details and proof of correctness in the next section.

### The High Level Strategy for Capture

We show that three pursuers, denoted $p_1, p_2, p_3$, can always capture an evader using a deterministic strategy, regardless of the evader's strategy and the geometry of the environment. Our overall strategy is to progressively trap the evader in an ever-shrinking region of the polygon $P$. The pursuit begins by first choosing a path $\Pi_1$ that divides the polygon into sub-polygons (see Figure 1(a))—we will use the notation $P_e$ to denote the sub-polygon containing the evader. We show that, after an initialization period, the pursuer $p_1$ can successfully guard the path $\Pi_1$, meaning that $e$ cannot move across it without being captured.

In a general step, the sub-polygon $P_e$ containing the evader is bounded by two paths $\Pi_1$ and $\Pi_2$, satisfying a geometric property called *minimality*, each being guarded by a pursuer. We then choose a third path $\Pi_3$ splitting the region $P_e$ into two non-empty subsets. If both regions have holes, then we argue that the pursuer $p_3$ can guard $\Pi_3$, thereby trapping $e$ either between $\Pi_1$ and $\Pi_3$ (see Figure 1(b)), or between $\Pi_2$ and $\Pi_3$, in which case we pursuit iterates in a smaller region. If one of the regions is hole-free, then we show that the pursuer $p_3$ can *evict* the evader from this region, forcing it into a smaller region, where the search resumes.

### Visibility Graphs and Path Guarding

In order for this strategy to work, the paths $\Pi_i$ need to be carefully chosen, and must satisfy certain geometric conditions, which we briefly explain. First, although the pursuit occurs in continuous space, our paths will be computed from a *discrete* space, namely, the *visibility graph* of the polygon. The visibility graph $G(P)$ of a polygon $P$ is defined as follows: the nodes are the vertices of the polygonal environment (including the holes), and two nodes are joined by an edge if the line segment joining them lies entirely in the (closed) interior of the polygon. (In other words, the two vertices have line of sight visibility.) This *undirected* graph has $n$ vertices and at most $O(n^2)$ edges. We assign each edge a *weight* equal to the Euclidean distance between its two endpoints. In Figure 1(a), $y$ is visible from $x$ but not $z$; thus, $xy$ is an edge of $G(P)$ but not $xz$.

One can easily see that, given two vertices $u$ and $v$ of $P$, the *shortest path* from $u$ to $v$ in $G(P)$ is also the shortest Euclidean path constrained to lie inside $P$.[1] It is also easy to see that we cannot make such a claim for the *second*, or in general the $k$th, shortest path—one can create an infinitesimal "bend" in the shortest path $\Pi_1$ to create another path that is arbitrarily close to the first shortest path but does not belong to $G(P)$. Therefore, we will only consider paths that belong to $G(P)$ and are "combinatorially distinct" from $\Pi_1$—that is, they differ in at least one visibility edge. However, even then the $k$th shortest path between two nodes can exhibit counter-intuitive behavior. For instance, in graphs with non-negative weights, while the first shortest path is always loop-free, the *second*, or more generally $k$th, shortest path can have loops—this may happen if repeatedly looping around a small-weight cycle (to make the path distinct from others) is cheaper than taking an expensive edge (Hershberger and Suri 2002). Therefore, we will consider only shortest loop-free paths. One of our technical lemmas proves that these paths are also *geometrically* non-self-intersecting. (This is obvious for the shortest path $\Pi_1$ but not for subsequent paths.) In addition, we argue that these paths also satisfy a key geometric property, called *minimality*, which allows a pursuer to guard them against an evader.

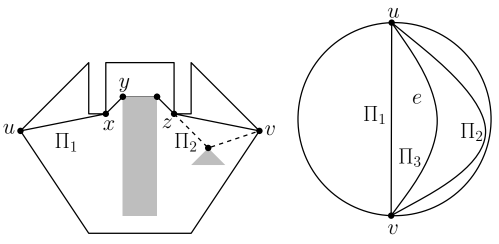

(a) (b)

Figure 1: The left figure shows a polygonal environment, with two holes (a rectangle and a triangle). $xy$ is a visibility edge of $G(P)$, while $xz$ is not. $\Pi_1$ and $\Pi_2$ are the first and the second shortest paths between anchors $u$ and $v$. The right figure illustrates the main strategy of trapping the evader through three paths.

## Proof of Sufficiency of 3 Pursuers

We begin with the discussion of how a single pursuer can guard a path in $P$, thus trapping the evader on one side. We then discuss the technically more challenging case of guarding the second and the third paths. In order to guarantee that a path in $P$ can be guarded, it must satisfy certain geometric properties. We begin by introducing two key ideas: a minimal path and the projection of evader on a path. In the following, we use the notation $d(x, y)$ to denote the shortest path distance between points $x$ and $y$. When we require that distance to be measured within a subset, such as restricted to a path $\Pi$, we write $d_\Pi(x, y)$. That is, $d_\Pi(x, y)$ is the length of path $\Pi$ between its points $x$ and $y$. Occasionally, we also use the notation $\Pi(x, y)$ to denote subpath of $\Pi$ between points $x, y$. We use the notation $x \prec y$ to emphasize that the point $x$ precedes $y$ on the path $\Pi$: that is, if $\Pi$ is the path from node $u$ to node $v$, then $x \prec y$ means that $d_\Pi(u, x) < d_\Pi(u, y)$. The following property is important for the patrolling of paths.

**Definition 1.** *(***Minimal Path:***) Suppose $\Pi$ is a path in $P$ dividing it into two sub-polygons, and $P_e$ is the sub-polygon containing the evader $e$. We say that $\Pi$ is* minimal *if*

$$d_\Pi(x, z) \;\leq\; d(x, y) + d(y, z)$$

*for all points $x, z \in \Pi$ and $y \in (P_e \setminus \Pi)$.*

Intuitively, a minimal path cannot be shortcut: that is, for any two points on the path, it is never shorter to take a detour through an interior point of $P_e$. (This is a weak form of triangle inequality, which excludes detours only through points contained in $P_e$.) The following definition introduces the projection of the evader on to a path, which turns out to be an important concept to ensure capture.

**Definition 2.** *(***Projection:***) Suppose $\Pi$ is a path in $P$ dividing it into two sub-polygons, and $P_e$ is the sub-polygon containing the evader $e$. Then, the* projection *of $e$ on $\Pi$, denoted $e_\pi$, is a point on $\Pi$ such that, for all $x \in \Pi$, $e$ is no closer to $x$ than is $e_\pi$.*

Thus, if a pursuer is able to position itself at the projection of $e$ at all times, then it guarantees that the evader cannot cross the path without being captured. With these definitions in place, we now discuss how to guard the first path $\Pi_1$.

### Guarding the First Path

We choose two non-neighbor vertices $u$ and $v$ on the outer boundary of $P$, and call them *anchors*. We let $\Pi_1$ be the shortest path from $u$ to $v$ in $G(P)$; this is also the shortest Euclidean path between $u$ and $v$ constrained to lie inside the environment. Our first observation is that the shortest path $\Pi_1$ is always minimal.

**Lemma 1.** *The path $\Pi_1$ between $u$ and $v$ is minimal.*

*Proof.* For the sake of contradiction, suppose there are two points $x, z \in \Pi_1$ that violate the minimality, and a point $y \notin \Pi_1$ is the witness. That is, $d(x, y) + d(y, z) < d_{\Pi_1}(x, z)$. But then $\Pi_1$ is not the shortest path from $u$ to $v$, because its subpath $\Pi_1(x, z)$ is sub-optimal. □

Next we show that for a minimal path, the projection of $e$ is always well defined.

---

[1] This follows because the shortest Euclidean path has corners only at vertices of $G(P)$.

**Lemma 2.** *Suppose $\Pi$ is a minimal path between the anchor nodes $u$ and $v$. Then, for every position of the evader $e$ in $P_e$, the projection $e_\pi$ exists.*

*Proof.* On the path $\Pi$, starting from $u$, we pick the furthest point $z$ such that for all $x \prec z$ we have $d_\Pi(x, z) \leq d(x, e)$. We note that necessarily for some $x \neq z$, this must be equality because otherwise there exists a point farther along $\Pi$. We claim that $z$ is a projection of $e$. Suppose not. Then there exists a point $y \succ z$ such that $d_\Pi(z, y) > d(e, y)$. This, however, violates the minimality of $\Pi$ because

$$\begin{aligned} d_\Pi(x, y) &= d_\Pi(x, z) + d_\Pi(z, y) = d(x, e) + d_\Pi(z, y) \\ &> d(x, e) + d(e, y) \end{aligned}$$

This completes the proof. ☐

The following lemma shows how a pursuer can successfully guard a minimal path.

**Lemma 3.** *Suppose $\Pi$ is a minimal path between the anchors $u, v$ in $P$, and a pursuer $p$ is located at the current projection of $e$. Suppose on its turn the evader moves from $e$ to $e'$. Then, the pursuer $p$ can either capture the evader or relocate to the new projection $e'_\pi$ in one move.*

*Proof.* First, suppose that the new position $e'$ is on different side of the path $\Pi$ than $e$; that is, the evader crosses the path. Because $e$ moves a distance of at most one, and suppose it crosses the path at a point $z$, we have $d(e, z) + d(z, e') \leq 1$. On the other hand, since $p$ is located at the projection of $e$ before the move, $d_\Pi(p, z) \leq d(e, z)$. Therefore, the new position of the evader $e'$ is within distance one of $p$, and the pursuer can capture the evader on its move.

Therefore, assume that the evader does not cross the path, and moves to a position $e'$ such that $e_\pi \neq e'_\pi$. Consider any two points $x \prec e_\pi \prec y$ (one on either side of the projection), and let $d(x, e') = d_\Pi(x, e_\pi) + c$ and $d(y, e') = d_\Pi(y, e_\pi) + c'$. Then we claim that $c + c' \geq 0$. We observe that by the minimality of $\Pi$, the following holds:

$$\begin{aligned} d_\Pi(x, y) &\leq d(x, e') + d(e', y) \\ &= d_\Pi(x, e_\pi) + d_\Pi(y, e_\pi) + c + c' \\ &= d_\Pi(x, y) + c + c' \end{aligned}$$

which implies that $c + c' \geq 0$.

This implies that $e$ can move closer to $x$ or $y$ but not both. Suppose it moves closer to $x$, meaning $c < 0$. Consider some $x \prec e_\pi$ with the smallest value of $c$. Then suppose for any $y \succ e_\pi$, $c' < |c|$. This would imply $c + c' < 0$, and so if $e$ moves a maximum of $c$ closer to any $x$, it moves at least $c$ farther from any $y$. We claim that $e'_\pi$ is at the point $c$ closer to $x$. To show this we argue that at the position $e'_\pi$, the pursuer is closer to all points on $\Pi$. But $e'$ cannot be closer to any $x \prec e'_\pi$ because we assumed $c$ was minimal. Similarly $e'$ must still be farther from any $y \succ p_e$ because $e$ moved $c$ or more away from those points. This leaves the points between $e'_\pi$ and $e_\pi$, but if $e'$ is closer to one of these points (say, $z'$), then $p$ can capture $e$ because $d(e, z') < d(e'_\pi, z')$. Therefore, the pursuer can move to the new projection because $|c| < 1$. This completes the proof. ☐

Finally, we show that a pursuer $p$ only requires $O(\mathrm{diam}(P))$ moves to either reach the current projection of $e$ or capture it.

**Lemma 4.** *Suppose $\Pi$ is a minimal path between anchors $u, v$ in $P$, and a pursuer $p$ is located at $u$. Then in $O(\mathrm{diam}(P))$ moves, $p$ can move to $e$'s projection.*

*Proof.* By Lemma 3, the projection of $e$ can only shift by distance at most one along the path $\Pi$. Thus, $p$'s strategy is simply to move along the path from one end to the other until it coincides with the current projection of $e$, or captures it. Since $p$ moves a distance of $1$ in each turn, and the path $\Pi$ is at most $\mathrm{diam}(P)$ long, the entire initialization phase takes $O(\mathrm{diam}(P))$ time. During this time, if the projection ever "crosses over" the current position of $p$, the pursuer immediately can move to the new projection point because at that moment it must be within distance one of the target location. This completes the proof. ☐

## Geometric Structure of Pursuer Paths

We now come to the main part of our pursuit strategy. The key idea is to progressively trap the evader in a region bounded by two minimal paths, which are guarded by two pursuers, and to use the third pursuer to further divide the current trap region. When the third pursuer subdivides the current region containing $e$, two possibilities emerge: either both regions of the subdivision have holes, in which case we show that the third path is necessarily minimal and thus guardable by the third pursuer, limiting the evader to a smaller region than before; or one of the regions is hole-free, in which case the third pursuer uses the capture strategy for a simply-connected polygon to evict the pursuer from this region (or capture it). In order to formalize our strategy, we first show a key geometric property of the second and third shortest paths between the anchors in the visibility graph, namely, that they are non-self-intersecting, and therefore lead to well-defined closed regions.

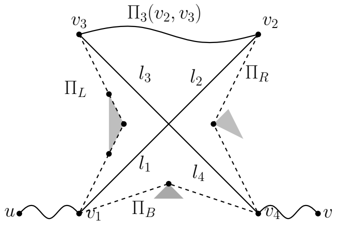

Figure 2: Non-self-crossing of shortest paths $\Pi_1, \Pi_2, \Pi_3$.

**Lemma 5.** *Let $\Pi_1$ be the shortest path between two anchor points $u$ and $v$ on $P$'s boundary, and focus on the subpolygon $P_e$ that lies on one side of $\Pi_1$. Let $\Pi_2$ and $\Pi_3$, respectively, be the second and the third simple (loop-free) shortest paths in the visibility graph $G(P_e)$ between $u$ and $v$. Then, $\Pi_2$ and $\Pi_3$ are non-self-crossing.*

*Proof.* Without loss of generality, suppose the path $\Pi_3$ violates the lemma, and that two of its edges $(v_1, v_2)$ and $(v_3, v_4)$ intersect. We first note that the intersection point cannot be a vertex of the visibility graph because otherwise the path has a cycle, and we assumed that $\Pi_3$ is loop-free. As shown in the figure, we break the segment $(v_1, v_2)$ into $l_1$ and $l_2$, and $(v_3, v_4)$ into $l_3$ and $l_4$. By the triangle inequality of the Euclidean metric, it is easy to see that $d(v_1, v_3) < l_1 + l_3$, and $d(v_2, v_4) < l_2 + l_4$. Similarly, $d(v_1, v_4) < l_1 + l_4$. Let $\Pi_L$, $\Pi_R$, $\Pi_B$, respectively, denote the shortest paths (in the graph $G(P_e)$) realizing these distances. Now consider the following three paths between $v_1$ and $v_4$, each contained in $G(P_e)$) and *non-self-intersecting*: $\Pi_L \cdot \Pi_3(v_2, v_3) \cdot \Pi_R$, $\Pi_B$, and the shorter of $\Pi_L \cdot (v_3, v_4)$ and $(v_1, v_2) \cdot \Pi_R$. They are all shorter than $\Pi_3$ and at least one of them must be distinct from both $\Pi_1$ and $\Pi_2$, thus contradicting the choice of $\Pi_3$. This completes the proof. ☐

## Shrinking, Guarding and Evicting

In a general step of the algorithm, assume that the evader lies in a region $P_e$ of the polygon bounded by two minimal paths $\Pi_1$ and $\Pi_2$ between two anchor vertices $u$ and $v$. (Strictly speaking, the region $P_e$ is initially bounded by $\Pi_1$, which is minimal, and a portion of $P$'s boundary, which is not technically a minimal path. However, the evader cannot cross the polygon boundary, and so we treat this as a special case of the minimal path to avoid duplicating our proof argument.) We also assume that $\Pi_1$ and $\Pi_2$ only share vertices $u$ and $v$; if they share a common prefix or suffix subpath, we can delete those and advance the anchor nodes to the last common prefix vertex and the first common suffix vertex. This ensures that the region $P_e$ is *non-degenerate*. Furthermore, we assume that the region $P_e$ contains at least one hole—otherwise, the evader is trapped in a simply-connected region, where a single (the third) pursuer can capture it.

The key idea of our proof is to show that, in the visibility graph $G(P_e)$, if we compute a *shortest path* from $u$ to $v$ that is distinct from both $\Pi_1$ and $\Pi_2$, then it divides $P_e$ into *only* two regions, and that the evader is trapped in one of those regions. We will call this new path the *third* shortest path $\Pi_3$. Specifically, $\Pi_3$ is the simple (loop-free) shortest path from $u$ to $v$ in $G(P_e)$ distinct from $\Pi_1$ and $\Pi_2$. (One can compute such a path using any of the algorithms for computing $k$ loop-free shortest paths in a weighted undirected graph (Hershberger and Suri 2002; Nardelli, Proietti, and Widmayer 2003; Yen 1971).)

**Lemma 6.** *The shortest path $\Pi_3$ between the anchor nodes $u$ and $v$ divides the current evader region $P_e$ into two connected regions.*

*Proof.* If the path is disjoint from $\Pi_1$ and $\Pi_2$ except at endpoints, then $P_e$ is clearly subdivided into two regions. If $\Pi_3$ shares vertices only with $\Pi_1$ or only with $\Pi_2$, *but in multiple disjoint subpaths creating multiple regions*, then minimality of those paths means that we can contract all but one region and shorten $\Pi_3$, contradicting the choice of this third path. Therefore, let us suppose that $\Pi_3$ shares vertices with both the paths, and so "hops" between $\Pi_1$ and $\Pi_2$, sharing common subpaths with them, and creates three or more regions. In that case, $\Pi_3$ must leave and rejoin $\Pi_1$ and $\Pi_2$ at least once, as shown by points $x, y, z$ in Figure 3(a). We observe that $d_{\Pi_2}(y, v)$ is no longer than $d(y, z) + d_{\Pi_1}(z, v)$, otherwise we can contradict the choice of $\Pi_2$. Thus the third region can be removed by altering $\Pi_3$ to use the subpath $\Pi_2(y, v)$. (A symmetric case arises when the roles of $\Pi_1$ and $\Pi_2$ are swapped.) Thus, we conclude that $\Pi_3$ can create only two subregions. ☐

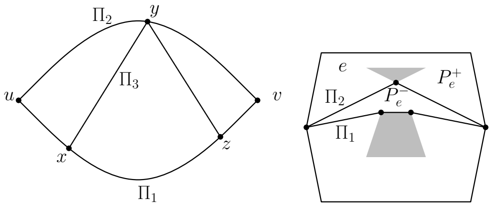

(a) (b)

Figure 3: The left figure illustrates the proof of Lemma 6; the right figure illustrates the two subregions created by a path, $\Pi_2$ in this case.

Because $P_e$ contains at least one hole, one of the regions created by the third shortest path $\Pi_3$ also must contain a hole. The following lemma argues that $\Pi_3$ is minimal with respect to that region. (The other region may be hole-free, and we will argue that the evader can be evicted from such a hole-free region or captured.) Due to space limitation, we omit the proofs of the following lemmas from this abstract.

**Lemma 7.** *Suppose $\Pi_3$ divides the region $P_e$ into two subregions $P_e^+$ and $P_e^-$, and assume that $P_e^+$ contains at least one hole. Then, $\Pi_3$ is a minimal path within the region $P_e^+$.*

When one of the regions created by $\Pi_3$ is hole-free, then $\Pi_3$ has a very simple structure, consisting of only two distinct edges as seen in Figure 3(b), allowing it to be cleared using the search strategy of a simply-connected polygon.

**Lemma 8.** *Suppose the evader lies in a hole-free region bounded by $\Pi_3$ and some minimal path. Then, in $O(n \cdot \mathrm{diam}(P)^2)$ moves, a single pursuer $p$ can either capture the evader or force it out of the region and place itself on $e$'s projection on the path $\Pi_3$.*

## Completing the Capture

We can now summarize our main result.

**Theorem 1.** *Three pursuers are always sufficient to capture an evader in $O(n \cdot \mathrm{diam}(P)^2)$ moves in a polygon with $n$ vertices and any number of holes.*

*Proof.* Whenever a new path is introduced, the size (number of vertices) of the region $P_e$ containing $e$ shrinks by at least one. Thus, the number of different paths guarded during the course of the pursuit before $e$ is trapped in a hole-free region is at most $n$. Guarding each path requires $O(\mathrm{diam}(P))$ moves for a minimal path, and $O(k \cdot \mathrm{diam}(P)^2)$ moves when in a hole-free region with $k$ vertices. Since the evader cannot reenter hole-free regions once they have been guarded, the total cost of guarding all the hole-free regions during the course of the algorithm sums to $O(n \cdot \mathrm{diam}(P)^2)$. ☐

## Necessity of 3 Pursuers

**Theorem 2.** *There exists an infinite family of polygons with holes that require at least three pursuers to capture an evader even with complete information about the evader's location.*

*Proof.* The proof is based on a reduction from searching in *planar graphs*. Consider a planar graph $G$, with minimum degree 3, and without any cycles of length three or four (see Figure 4(c)). Using Fary's Theorem, we can embed such a graph so that each edge maps to a straight line segment. By suitable scaling, assume that the longest edge in the embedding has length 1. (See Figure 4(a) for an example.)

We now transform this straight-line embedding into a polygon with holes, by converting each edge into a "corridor." Each corridor is constructed to ensure that the shortest path through it has length 1. In particular, the edges of length 1 map to straight corridors, while shorter edges correspond to corridors with multiple turns, as shown in Figure 4 (b). It is easy to see that such a construction can ensure that all the corridors are non-overlapping. With this transformation, the outer face of the graph becomes the boundary of the polygon $P$, while each face of the plane graph becomes a hole.

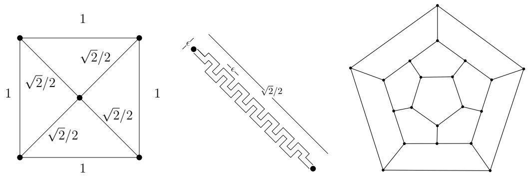

(a) (b) (c)

Figure 4: Embedding of a planar graph (a), corridor construction (b), and a planar graph with min-degree 3 (c).

It is known that, in any graph with minimum degree $k$ and no cycles of length three or four, the evader has a winning strategy against $k-1$ pursuers (Aigner and Fromme 1984), as follows: the evader only moves when at least one of its neighbors is occupied by a pursuer, and in that case it moves to a vertex that is not a neighbor of the pursuer-occupied vertices. We can mimic the same strategy in our polygon setting, and show that the evader has a winning strategy against two pursuers. The precise details of the strategy, which are mostly technical, are omitted from this conference version of the paper due to lack of space. ☐

## Closing Remarks

In this paper, we proved that three pursuers are always sufficient to capture an evader in a polygonal environment of arbitrary complexity, under the assumption that pursuers have access to evader's location at all times. We also proved a matching lower bound, showing that three pursuers are also necessary in the worst-case. Traditionally, the papers on continuous space, visibility-based pursuit problem have focussed on simply detecting the evader, and not on capturing it. One of our contributions is to isolate the *intrinsic* complexity of the capture from the associated complexity of detection or localization. In particular, while $\Theta(\sqrt{h} + \log n)$ pursuers are necessary (and also sufficient) for detection or localization of an evader in a $n$-vertex polygon with $h$ holes (Guibas et al. 1999), our result shows that full localization information allows capture with only *3* pursuers. On the other hand, it still remains an intriguing open problem whether $\Theta(\sqrt{h} + \log n)$ pursuers can *simultaneously* perform localization and capture. We leave that as a topic for future research.

## References

Aigner, M., and Fromme, M. 1984. A game of cops and robbers. *Discrete Applied Mathematics* 8(1):1–12.

Bienstock, D., and Seymour, P. 1991. Monotonicity in graph searching. *J. Algorithms* 12(2):239–245.

Fomin, F. V.; Golovach, P. A.; and Kratochvíl, J. 2008. On tractability of cops and robbers game. In *TCS*, 171–185.

Guibas, L. J.; Latombe, J.-C.; LaValle, S. M.; Lin, D.; and Motwani, R. 1999. Visibility-based pursuit-evasion in a polygonal environment. *IJCGA* 9(5):471–494.

Guy, R. K. 1991. Unsolved problems in combinatorial games. In *Games of No Chance*, 475–491.

Halpern, B. 1969. The robot and the rabbit–a pursuit problem. *The American Mathematical Monthly* 76(2):140–145.

Hershberger, J., and Suri, S. 2002. "Vickrey pricing and shortest paths: What is an edge worth?". In *43th FOCS*.

Isler, V., and Karnad, N. 2008. The role of information in the cop-robber game. *TCS* 399(3):179 – 190.

Isler, V.; Kannan, S.; and Khanna, S. 2005. Randomized pursuit-evasion in a polygonal environment. *Robotics, IEEE Transactions on* 21(5):875 – 884.

Isler, V.; Kannan, S.; and Khanna, S. 2006. Randomized pursuit-evasion with local visibility. *SIAM Journal on Discrete Mathematics* 1:26–41.

LaPaugh, A. 1993. Recontamination does not help to search a graph. *J. ACM* 40(2):224–245.

Nardelli, E.; Proietti, G.; and Widmayer, P. 2003. Finding the most vital node of a shortest path. *TCS.* 296(1):167–177.

Park, S.-M.; Lee, J.-H.; and Chwa, K.-Y. 2001. Visibility-based pursuit-evasion in a polygonal region by a searcher. In *ICALP*, 281–290. Springer-Verlag.

Parsons, T. D. 1976. Pursuit-evasion in a graph. In Alavi, Y., and Lick, D. R., eds., *Theory and Application of Graphs*. Berlin: Springer-Verlag. 426–441.

Sachs, S.; Rajko, S.; and LaValle, S. M. 2004. Visibility-based pursuit-evasion in an unknown planar environment. *International Journal of Robotics Research* 23(1):3–26.

Sgall, J. 2001. Solution of david gale's lion and man problem. *Theor. Comput. Sci.* 259(1-2):663–670.

Suzuki, I., and Yamashita, M. 1992. Searching for a mobile intruder in a polygonal region. *SIAM J. Comp.* 21:863–888.

Yen, J. 1971. Finding the k shortest loopless paths in a network. *Management Science* 17(11):pp. 712–716.