Chasing Puppies:
Mobile Beacon Routing on Closed Curves

Mikkel Abrahamsen [*] *Jeff Erickson* [†] *Irina Kostitsyna* [‡] *Maarten Löffler* [§] *Tillmann Miltzow* [§] *Jérôme Urhausen* [§] *Jordi Vermeulen* [§] *Giovanni Viglietta* [¶]

Abstract. We solve an open problem posed by Michael Biro at CCCG 2013 that was inspired by his and others' work on beacon-based routing. Consider a human and a puppy on a simple closed curve in the plane. The human can walk along the curve at bounded speed and change direction as desired. The puppy runs along the curve (faster than the human) always reducing the Euclidean straight-line distance to the human, and stopping only when the distance is locally minimal. Assuming that the curve is smooth (with some mild genericity constraints) or a simple polygon, we prove that the human can always catch the puppy in finite time. Our results hold regardless of the relative speeds of puppy and human, and even if the puppy's speed is unbounded.

[*] *Basic Algorithms Research Copenhagen (BARC), University of Copenhagen, Denmark*, `miab@di.ku.dk`

[†] *University of Illinois at Urbana-Champaign, USA*jeffe@illinois.edu, `jeffe@illinois.edu`

[‡] *Eindhoven University of Technology, Netherlands*, `i.kostitsyna@tue.nl`

[§] *Department of Information and Computing Sciences, Utrecht University, Netherlands*, `[m.loffler|t.miltzow|j.e.urhausen|j.l.vermeulen]@uu.nl`

[¶] *Japan Advanced Institute of Science and Technology, Nomi City, Ishikawa, Japan*, `johnny@jaist.ac.jp`

1 Introduction

You have lost your puppy somewhere on a simple closed curve. Both of you are forced to stay on the curve. You can see each other and both want to reunite. The problem is that the puppy runs faster than you, and it believes naively that it is always a good idea to minimize its straight-line distance to you. What do you do?

To be more precise, let $\gamma\colon S^1 \hookrightarrow \mathbb{R}^2$ be a simple closed curve in the plane, which we informally call the *track*. Two special points move around the track, called the *puppy* p and the *human* h. The human can walk along the track at bounded speed and change direction as desired. The puppy runs with unbounded speed along the track as long as its Euclidean straight-line distance to the human is decreasing, until it reaches a point on the curve where the distance is locally minimized. As the human moves along the track, the puppy moves to stay at a local distance minimum. The human's goal is to move in such a way that the puppy and the human meet. See Figure 1 for a simple example.

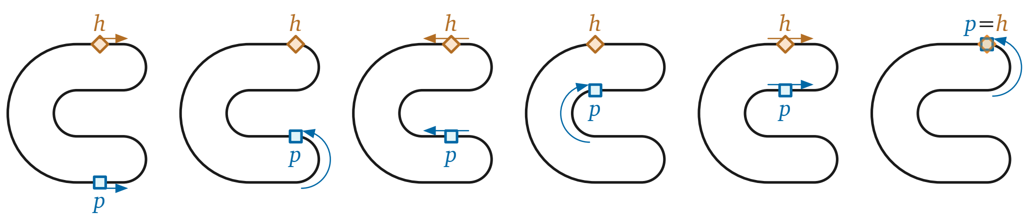

Figure 1: Catching the puppy.

In this paper we show that it is always possible to reunite with the puppy under the assumption that the curve is well-behaved (in a sense to be defined), or if the curve is a polygon. From this result it easily follows that catching a puppy that moves at any bounded speed is also possible: the strategy is essentially the same as for the unbounded-speed case, except that the human may have to move at a lower speed or occasionally stop, in order to let the puppy reach a point of minimal distance before continuing.

The problem was posed in a different guise at the open problem session of the 25th Canadian Conference on Computational Geometry (CCCG 2013) by Michael Biro. In Biro's formulation, the track was a railway, the human a locomotive, and the puppy a train carriage that was attracted to an infinitely strong magnet installed in the locomotive.

Returning to our formulation of catching a puppy, it was also asked if the human will always catch the puppy by choosing an arbitrary direction and walking only in that direction. This turns out not to be the case; consider the star-shaped track in Figure 2. Suppose the human and puppy start at points h_1 and p_1, respectively, and the human walks counterclockwise around the track. When the human reaches h_2, the puppy runs from p_2 to p_2'. When the human reaches h_3, the puppy runs from p_3 to p_3'. Then the pattern repeats indefinitely. Examples of this type, where the human walking in the wrong direction will never catch the puppy, were independently discovered during the conference by some of the authors and by David Eppstein.

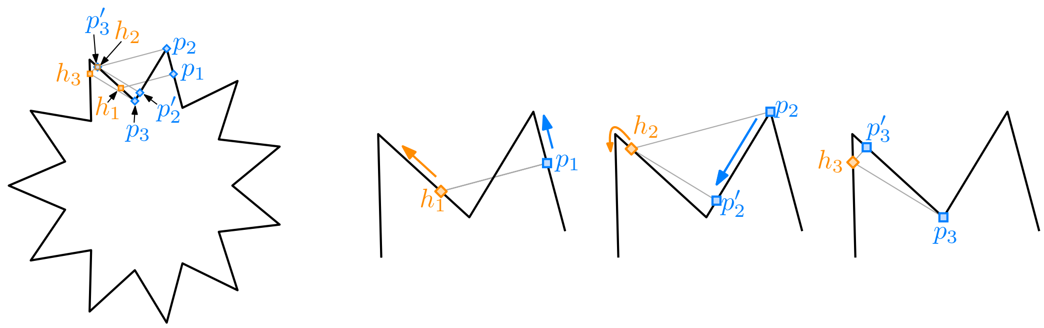

Figure 2: If the human keeps walking counterclockwise from h_1, the human and the puppy will never meet. To the right are closeups of two of the spikes of the star.

1.1 Related work

Biro's problem was inspired by his and others' work on *beacon-based geometric routing*, a generalization of both greedy geometric routing and the art gallery problem introduced at the 2011 Fall Workshop on Computational Geometry [7] and the 2012 Young Researchers Forum [8], and further developed in Biro's PhD thesis [6] and papers [9, 10]. A *beacon* is a stationary point object that can be activated to create a "magnetic pull" towards itself everywhere in a given polygonal domain P. When a beacon at point b is *activated*, a point object p moves moves greedily to decrease its Euclidean distance to b, alternately moving through the interior of P and sliding along its boundary, until it either reaches b or gets stuck at a "dead point" where Euclidean distance is minimized. By activating different beacons one at a time, one can route a moving point object through the domain. Initial results for this model by Biro and his colleagues [6–10] sparked significant interest and subsequent work in the community [2, 3, 5, 13, 18, 20–22, 26]. More recent works have also studied how to utilize objects that repel points instead of attracting them [11, 24].

Biro's problem can also be viewed as a novel variant of classical *pursuit* problems, which have been an object of intense study for centuries [25]. The oldest pursuit problems ask for a description of the *pursuit curve* traced by a *pursuer* moving at constant speed directly toward a *target* moving along some other curve. Pursuit curves were first systematically studied by Bouguer [12] and de Maupertuis [14] in 1732, who used the metaphor of a pirate overtaking a merchant ship; another notable example is Hathaway's problem [16], which asks for the pursuit curve of a dog swimming at unit speed in a circular lake directly toward a duck swimming at unit speed around its circumference. In more modern *pursuit-evasion* problems, starting with Rado's famous "lion and man" problem [23, pp.114–117], the pursuer and target both move strategically within some geometric domain; the pursuer attempts to *capture* the target by making their positions coincide while the target attempts to evade capture. Countless variants of pursuit-evasion problems have been studied, with multiple pursuers and/or targets, different classes of domains, various constraints on motion or visibility, different capture conditions, and so on. Biro's problem can be naturally described as a *cooperative pursuit* or *pursuit-attraction* problem, in which a strategic target (the human) *wants* to be captured by a greedy pursuer (the puppy).

Kouhestani and Rappaport [19] studied a natural variant of Biro's problem, which we can recast as follows. A *guppy* is restricted to a closed and simply-connected *lake*, while the human is restricted to the boundary of the lake. The guppy swims with unbounded speed to decrease its Euclidean distance to the human. Kouhestani and Rappaport described a polynomial-time algorithm that finds a strategy for the human to catch the guppy, if such a strategy exists, given a simple polygon as input; they also conjectured that a capturing strategy always exists. Abel, Akitaya, Demaine, Demaine, Hesterberg, Korman, Ku, and Lynch [1] recently proved that for some polygons and starting configurations, the human cannot catch the guppy, even if the human is allowed to walk in the exterior of the polygon, thereby disproving Kouhestani and Rappaport's conjecture. Their simplest counterexample is an orthogonal polygon with about 50 vertices.

1.2 Our results

Before describing our results in detail, we need to carefully define the terms of the problem. The *track* is a simple closed curve $\gamma\colon S^1 \hookrightarrow \mathbb{R}^2$. We consider the motion of two points on this curve, called the *human* (or *beacon* or *target*) and the *puppy* (or *pursuer*). A *configuration* is a pair $(x, y) \in S^1 \times S^1$ that specifies the locations $h = \gamma(x)$ and $p = \gamma(y)$ for the human and puppy, respectively. Let $D(x, y)$ denote the straight-line Euclidean distance between these two points. When the human is located at $h = \gamma(x)$, the puppy moves from $p = \gamma(y)$ to greedily decrease its distance to the human, as follows.

- If $D(x, y + \varepsilon) < D(x, y)$ for all sufficiently small $\varepsilon > 0$, the puppy runs forward along the track, by increasing the parameter y.
- If $D(x, y - \varepsilon) < D(x, y)$ for all sufficiently small $\varepsilon > 0$, the puppy runs backward along the track, by decreasing the parameter y.

If both of these conditions hold, the puppy runs in an arbitrary direction. While the puppy is running, the human remains stationary. If neither condition holds, the configuration is *stable*; the puppy does not move until the human does. When the configuration is stable, the human can walk in either direction along the track; the puppy walks along the track in response to keep the configuration stable, until it is forced to run again. The human's goal is to *catch* the puppy; that is, to reach a configuration in which the two points coincide.

Our main result is that the human can always catch the puppy in finite time, starting from any initial configuration, provided the track is either a generic simple smooth curve or an arbitrary simple polygon.

The remainder of the paper is structured as follows. We begin in Section 2 by considering some variants and special cases of the problem. In particular, we give a simple self-contained proof of our main result for the special case of orthogonal polygons.

We consider generic smooth tracks in Sections 3 and 4. Specifically, in Section 3 we define two important diagrams, which we call the *attraction diagram* and the *dual attraction diagram*, and prove some useful structural results. At a high level, the attraction diagram is a decomposition of the configuration space $S^1 \times S^1$ according to the puppy's behavior, similar to the *free space diagrams* introduced by Alt and Godau to compute Fréchet distance [4]. We show that for a sufficiently generic smooth track, the attraction diagram consists of a finite number of disjoint simple closed *critical* curves, exactly two of which are topologically nontrivial. Then in Section 4, we argue that the human can catch the puppy on any track whose attraction diagram has this structure.

In Section 5, we describe an extension of our analysis from smooth curves to simple polygonal tracks. Because polygons do not have well-defined tangent directions at their vertices, this extension requires explicitly modeling the puppy's direction of motion in addition to its location. We first prove that the human can catch the puppy on a polygon that has no acute vertex angles and where no three vertices form a right angle; under these conditions, the attraction diagram has exactly the same structure as for generic smooth curves. We then reduce the problem for arbitrary simple polygons to this special case by *chamfering*—cutting off a small triangle at each vertex—and arguing that any strategy for catching the puppy on the chamfered track can be pulled back to the original polygon.

Finally, we close the paper by suggesting several directions for further research.

Open-source software demonstrating several of the tools developed in this paper is available at https://github.com/viglietta/Chasing-Puppies or https://archive.softwareheritage.org/swh:1:dir:58dd270b0896aa11024666b5cbd2481068e8eab9 .

2 Warmup: other settings and a special case

In this section, we discuss two variants of Biro's problem and the special case of orthogonal polygons.

In the first variant, both the human h and the puppy p are allowed to move anywhere in the interior and on the boundary of a simple polygon P. Here, as in beacon routing and Kouhestani and Rappaport's variant [1, 19], the puppy moves greedily to decrease its Euclidean distance to the human, alternately moving through the interior of P and sliding along its boundary.

As we will show in Theorem 1, h has a simple strategy to catch p in this setting, essentially by walking along the dual graph of any triangulation. This is an interesting contrast to the proof by Abel et al. [1] that h and p cannot always meet when h is restricted to the *exterior* of P and p to the interior. Our main result that h and p can meet when both are restricted to the *boundary* of P (even for a much wider class of simple closed curves), somehow sits in between these other two variants.

When both h and p are restricted to the interior of P, we propose the following strategy for h; see Figure 3. Let $\mathcal{T}$ be a triangulation of P and let $t_1, \ldots, t_k$ be the path of pairwise adjacent triangles in $\mathcal{T}$ such that $h \in t_1$ and $p \in t_k$. Let e_i be the common edge of t_i and t_{i+1} and let d_i be the midpoint of e_i. Let $\pi = hd_1d_2 \ldots d_{k-1}$ be a path from h to d_{k-1}, which is contained in the triangles $t_1, \ldots, t_{k-1}$. The human starts walking along π. As soon as the puppy enters a new triangle, the human recomputes π as described and follows the new path.

Theorem 1. *The proposed strategy will make h and p meet.*

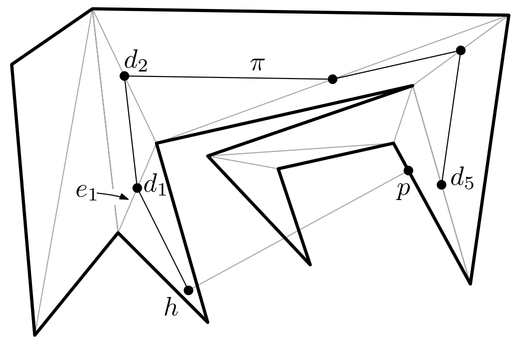

Figure 3: The proposed strategy when h and p are restricted to the interior of a simple polygon P. The human h will follow the path π. Note that the triangle containing p will change before h reaches d_1, and π will be updated accordingly.

Proof. First, we observe that if the puppy ever enters the triangle t_1 that is occupied by the human, then the puppy and the human will meet immediately. Assume that the human does not meet the puppy right from the beginning. The region $P \setminus t_1$ consists of one, two, or three polygons, one of which P_p contains p. Thus, whenever the human moves from one triangle to another, the set of triangles that can possibly contain p shrinks. We conclude that the human and the puppy must meet eventually. □

In our second variant, the human and the puppy are both restricted to a simple, closed curve γ in $\mathbb{R}^3$. Here it is easy to construct curves on which h and p will never meet; the simplest example is a "double loop" that approximately winds twice around a planar circle, as shown in Figure 4.

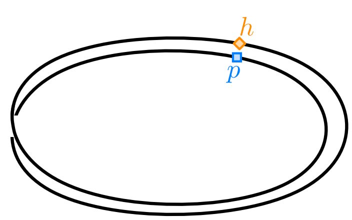

Figure 4: A double loop in $\mathbb{R}^3$; the human and puppy will never meet.

Finally, we consider the special case of Biro's original problem where the track γ is the boundary of an orthogonal polygon in the plane. This special case of our main results admits a much simpler self-contained proof.

Theorem 2. *The human can catch the puppy on any simple orthogonal polygon, by walking counterclockwise around the polygon at most twice.*

Proof. Let P be an arbitrary simple orthogonal polygon. Let u_1 be its leftmost point with the maximum y-coordinate, and u_2 be the next boundary vertex of P in the clockwise order (see Figure 5). Finally, let ℓ be the horizontal line supporting the segment $u_1 u_2$.

We break the motion of the human into two phases. In the first phase, the human moves counterclockwise around P from the starting location to u_1. If the human catches the puppy during this phase, we are done, so assume otherwise. In the second phase, the human walks counterclockwise around P starting from u_1 to u_2.

We claim that the puppy p is never in the interior of the segment u_1u_2 during the second phase; thus, p always lies on the closed counterclockwise subpath of P from h to u_2 (or less formally, "between h and u_2"). This claim implies that the human and the puppy are united during the second phase.

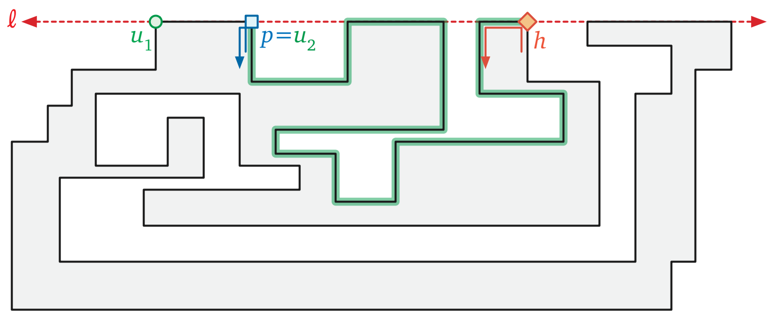

Figure 5: Proof of Theorem 2. During the human's second trip around P, the puppy lies between u_2 and the human.

The puppy must first cross the point u_2 if it ever enters the interior of u_1u_2. So consider any moment during the second phase when p moves upward to the vertex u_2. At that moment, h must be on the line ℓ to the right of p. (For any point a below ℓ, there is a point b on the segment below u_2 that is closer to a than u_2.) Thus, the puppy will stay on u_2 as long as h is on ℓ. As soon as h leaves ℓ the puppy will leave u_2 downward. Thus the puppy can never go to the interior of the edge u_1u_2. □

The star-shaped track in Figure 2 shows that this simple argument does not extend to arbitrary polygons, even with a constant number of edge directions.

3 Diagrams of smooth tracks

We first formalize both the problem and our solution under the assumption that the track is a generic smooth simple closed curve $\gamma\colon S^1 \hookrightarrow \mathbb{R}^2$. In particular, for ease of exposition, we assume that γ is regular and C^3, meaning it has well-defined continuous first, second, and third derivatives, and its first derivative is nowhere zero. We also assume γ satisfies some additional genericity constraints, to be specified later. We consider polygonal tracks in Section 5.

3.1 Configurations and genericity assumptions

We analyze the behavior of the puppy in terms of the *configuration space* $S^1 \times S^1$, which is the standard torus. Each configuration point $(x, y) \in S^1 \times S^1$ corresponds to the human being located at $h = \gamma(x)$ and the puppy being located at $p = \gamma(y)$.

For any configuration (x, y), recall that $D(x, y)$ denotes the straight-line Euclidean distance between the points $\gamma(x)$ and $\gamma(y)$. We classify all configurations $(x, y) \in S^1 \times S^1$ into three types, according to the sign of the partial derivative of distance with respect to the puppy's position.

- (x, y) is a *forward* configuration if $\frac{\partial}{\partial y} D(x, y) < 0$.
- (x, y) is a *backward* configuration if $\frac{\partial}{\partial y} D(x, y) > 0$.
- (x, y) is a *critical* configuration if $\frac{\partial}{\partial y} D(x, y) = 0$.

Starting in any forward (resp. backward) configuration, the puppy automatically runs forward (resp. backward) along the track γ. Genericity implies that there are a finite number of critical configurations (x, y) with any fixed value of x, or with any fixed value of y. We further classify the critical configurations as follows:

- (x, y) is a *stable* critical configuration if $\frac{\partial^2}{\partial y^2} D(x, y) > 0$.
- (x, y) is an *unstable* critical configuration if $\frac{\partial^2}{\partial y^2} D(x, y) < 0$.
- (x, y) is a *forward pivot* configuration if $\frac{\partial^2}{\partial y^2} D(x, y) = 0$ and $\frac{\partial^3}{\partial y^3} D(x, y) < 0$.
- (x, y) is a *backward pivot* configuration if $\frac{\partial^2}{\partial y^2} D(x, y) = 0$ and $\frac{\partial^3}{\partial y^3} D(x, y) > 0$.

In any stable configuration, the puppy's distance to the human is locally minimized, so the puppy does not move unless the human moves. In any unstable configuration, the puppy can decrease its distance by running in either direction. Finally, in any forward (resp. backward) pivot configuration, the puppy can decrease its distance by moving in one direction but not the other, and thus automatically runs forward (resp. backward) along the track.

Critical points can also be characterized geometrically as follows. Refer to Figure 6. A configuration (x, y) is critical if the human $\gamma(x)$ lies on the line $N(y)$ normal to γ at the puppy's location $\gamma(y)$. Let $C(y)$ denote the center of curvature of the track at $\gamma(y)$. Then (x, y) is a pivot configuration if $\gamma(x) = C(y)$, a stable critical configuration if the open ray from $C(y)$ through the human point $\gamma(x)$ contains the puppy point $\gamma(y)$, and an unstable critical configuration otherwise.

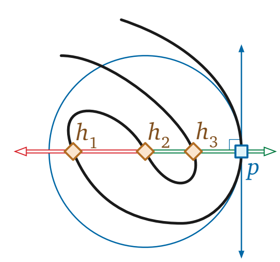

Figure 6: Three critical configurations: (h_1, p) is unstable; (h_2, p) is a pivot configuration, and (h_3, p) is stable.

Genericity of the track γ implies that this classification of critical configurations is exhaustive, and moreover, that the set of pivot configurations is finite. In particular, our analysis requires that in any pivot configuration (x, y), the puppy point $\gamma(y)$ is not a local curvature minimum or maximum.[1] Otherwise, we would need higher derivatives to disambiguate the puppy's behavior. In the extreme case where γ contains both an open circular arc α and its center c, all configurations where $h = c$ and $p \in \alpha$ are stable.

3.2 Attraction diagrams

The ***attraction diagram*** of the track γ is a decomposition of the configuration space $S^1 \times S^1$ by critical configurations. Our genericity assumptions imply that the set of critical points—the common boundary of the forward and backward configurations—is the union of a finite number of disjoint simple closed curves, which we call *critical cycles*. At least one of these critical cycles, the main diagonal $x = y$, consists entirely of stable configurations; critical cycles can also consist entirely of unstable configurations. If a critical cycle is neither entirely stable nor entirely unstable, then its points of vertical tangency are pivot configurations, and these points subdivide the curve into x-monotone paths, which alternately consist of stable and unstable configurations.

Figure 7 shows a sketch of the attraction diagram of a simple closed curve. We visualize the configuration torus $S^1 \times S^1$ as a square with opposite sides identified. Green and red paths indicate stable and unstable configurations, respectively; blue dots indicate pivot configurations; and backward configurations are shaded light gray. Figure 8 shows the attraction diagram for a more complex polygonal track, with slightly different coloring conventions. (Again, we will discuss polygonal tracks in more detail in Section 5.)

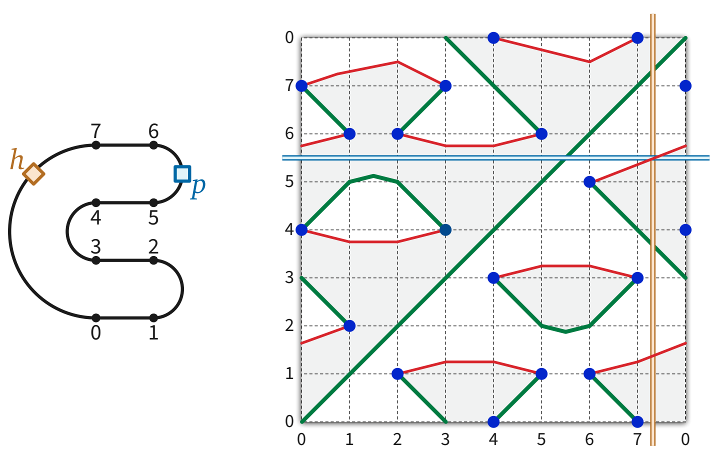

Figure 7: The attraction diagram of a simple closed curve, with one unstable critical configuration emphasized.

The cycles in any attraction diagram have a simple but important topological structure. A critical cycle in the attraction diagram is *contractible* if it is the boundary of a simply

[1]More concretely, we assume the track γ intersects its evolute (the locus of centers of curvature) transversely, away from its cusps.

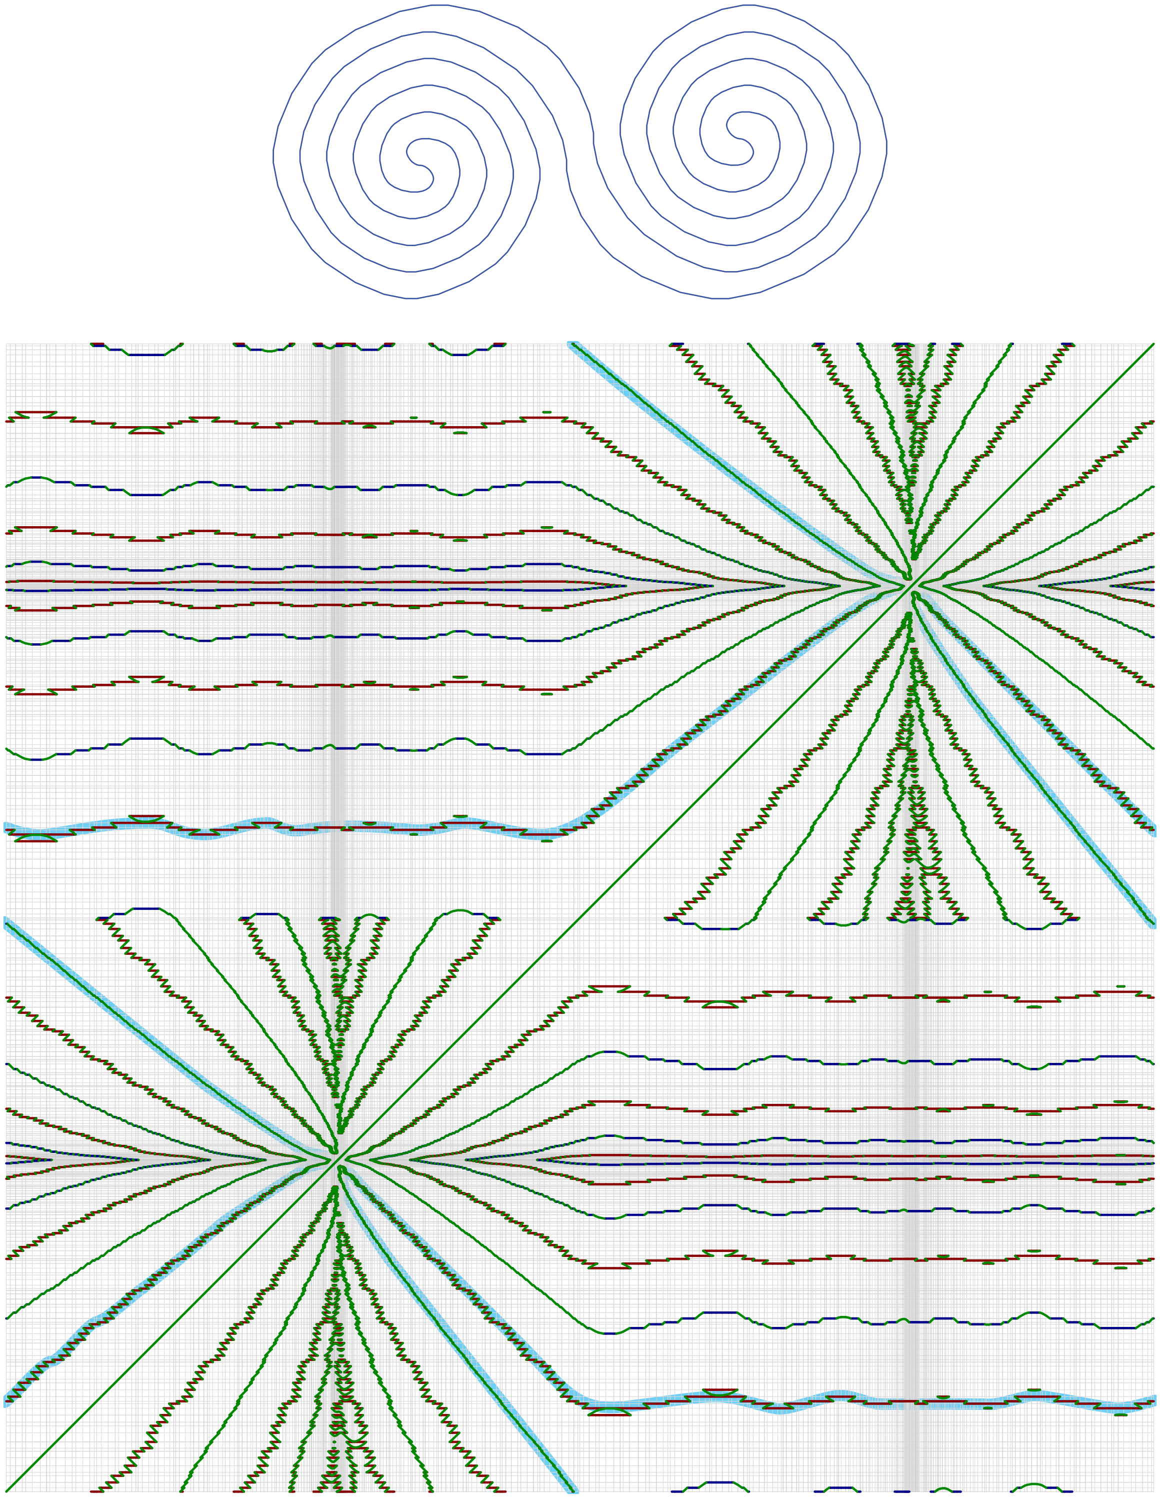

Figure 8: The attraction diagram of a complex simple polygon. Serrations in the diagram are artifacts of the curve being polygonal instead of smooth. The river is highlighted in blue.

connected subset of the torus $S^1 \times S^1$ and *essential* otherwise. For example, the main diagonal is essential, and the attraction diagram in Figure 7 contains two contractible critical cycles and two essential critical cycles.

Lemma 3. *The attraction diagram of any generic closed curve contains an even number of essential critical cycles.*

Proof. This lemma follows immediately from standard homological arguments, but for the sake of completeness we sketch a self-contained proof.

Fix a generic closed curve γ. Let α and β denote the horizontal and vertical cycles $S^1 \times \{0\}$ and $\{0\} \times S^1$, respectively. Without loss of generality, assume α and β intersect every critical cycle in the attraction diagram of γ transversely.

A critical cycle C in the attraction diagram is contractible if and only if α and β each cross C an even number of times. (Indeed, this parity condition characterizes all simple contractible closed curves in the torus.) On the other hand, α and β each cross the main diagonal once. It follows that α and β each cross *every* essential critical cycle an odd number of times; otherwise, some pair of essential critical cycles would intersect.

Because the critical cycles are the boundary between the forward and backward configurations, α and β each contain an even number of critical points. The lemma now follows immediately. $\square$

We emphasize that this lemma does *not* actually require the track γ to be simple; the argument relies only on properties of generic functions over the torus that are minimized along the main diagonal.

3.3 Dual attraction diagrams

Our analysis also relies on a second diagram, which we call the ***dual attraction diagram*** of the track. We hope the following intuition is helpful. While the attraction diagram tells us the possible positions of the puppy depending on the position of the human, the dual attraction diagram gives us the possible positions of the human depending on the position of the puppy. For each puppy configuration $y \in S^1$, we consider the normal line $N(y)$. We are interested in the intersection points of γ with $N(y)$, as those are the possible positions of the human. The idea of the dual attraction diagram is to trace the positions of the human as a function of the position of the puppy, see Figure 10.

Let $T(y)$ denote the directed line tangent to γ at the point $\gamma(y)$. For any configuration (x, y), let $\boldsymbol{\ell(x, y)}$ denote the distance from $\gamma(x)$ to the tangent line $T(y)$, signed so that $\ell(x, y) > 0$ if the human point $\gamma(x)$ lies to the left of $T(y)$ and $\ell(x, y) < 0$ if $\gamma(y)$ lies to the right of $T(y)$. More concisely, assuming without loss of generality that the track γ is parameterized by arc length, $\ell(x, y)$ is twice the signed area of the triangle with vertices $\gamma(x)$, $\gamma(y)$, and $\gamma(y) + \gamma'(y)$.

Let $L\colon S^1 \times S^1 \to S^1 \times \mathbb{R}$ denote the function $L(x, y) = (y, \ell(x, y))$. The dual attraction diagram is the decomposition of the infinite cylinder $S^1 \times \mathbb{R}$ by the points

$\{L(x,y) \mid (x,y) \text{ is critical}\}$. At the risk of confusing the reader, we refer to the image $L(x,y) \in S^1 \times \mathbb{R}$ of any critical configuration (x,y) as a critical point of the dual attraction diagram.

The dual attraction diagram can also be described as follows. For any $y \in S^1$ and $d \in \mathbb{R}$, let $\Gamma(y,d)$ denote the point on the normal line $N(y)$ at distance d to the left of the tangent vector $\gamma'(y)$. More formally, assuming without loss of generality that γ is parameterized by arc length, we have $\Gamma(y,d) = \gamma(y) + d \left[\begin{smallmatrix} 0 & -1 \\ 1 & 0 \end{smallmatrix}\right] \gamma'(y)$. We emphasize that $\Gamma(y,d)$ does not necessarily lie on the curve γ. The dual attraction diagram is the decomposition of the cylinder $S^1 \times \mathbb{R}$ by the preimage $\Gamma^{-1}(\gamma)$ of γ.

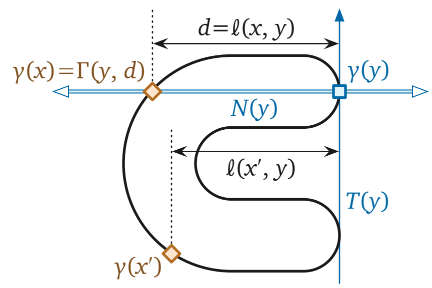

Figure 9: Examples of the functions ℓ and Γ used to define the dual attraction diagram.

Because γ is simple and regular, the dual attraction diagram is the union of simple disjoint closed curves. The function L continuously maps each critical cycle in the attraction diagram to a closed curve in the cylinder $S^1 \times \mathbb{R}$; we also call this image curve a *critical cycle*. Thus, the restriction of L to the set of critical configurations is a homeomorphism onto its image in the dual attraction diagram. In particular, L maps the main diagonal $x = y$ to the horizontal axis $\ell(x,y) = 0$ of the dual attraction diagram. We emphasize, however, that the two diagrams are not topologically equivalent. Figure 10 shows the dual attraction diagram of the same track whose attraction diagram is shown in Figure 7; here preimages of points inside the track are shaded.

Just as in the attraction diagram, a critical cycle in the dual attraction diagram is *contractible* if it is the boundary of a simply connected subset of the cylinder $S^1 \times \mathbb{R}$ and *essential* otherwise.

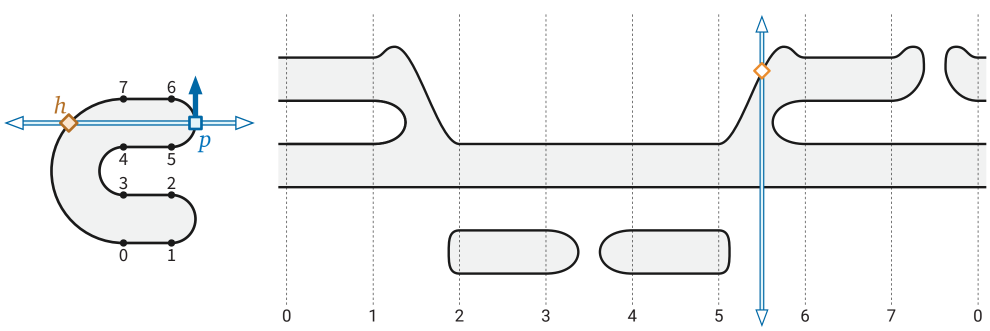

Figure 10: The dual attraction diagram of a simple closed curve, with one critical configuration emphasized. Compare with Figure 7.

Lemma 4. *The function L bijectively maps essential critical cycles in the attraction diagram to essential critical cycles in the dual attraction diagram. In particular, the two diagrams have the same number of essential critical cycles.*

Proof. Let $\alpha = S^1 \times \{0\}$ and $\alpha' = S^1 \times \{0\}$ denote the horizontal cycles in the torus $S^1 \times S^1$ and in the infinite cylinder $S^1 \times \mathbb{R}$, respectively. Let C be any critical cycle on the attraction diagram, and let $C' = L(C)$ be the corresponding critical cycle in the dual attraction diagram.

Recall from the proof of Lemma 3 that C is contractible on the torus if and only if $|C \cap \alpha|$ is even. Similarly, C' is contractible in the cylinder if and only if $|C' \cap \alpha'|$ is even. The map $L\colon S^1 \times S^1 \to S^1 \times \mathbb{R}$ maps $C \cap \alpha$ bijectively to $C' \cap \alpha'$. We conclude that C is essential if and only if C' is essential. □

With this correspondence in hand, we can now more carefully describe the topological structure of the *attraction* diagram when the track is simple.

Lemma 5. *The attraction diagram of a* ***simple*** *generic closed curve contains* ***exactly two*** *essential critical cycles.*

Proof. Fix a generic closed curve γ. Lemma 3 implies that the attraction diagram of γ contains at least two essential critical cycles, one of which is the main diagonal. Thus, to prove the lemma, it remains to show that there are *at most* two essential critical cycles, in either the attraction diagram or the dual attraction diagram.

Let $\Sigma \subset S^1 \times \mathbb{R}$ denote the set of essential critical cycles in the *dual attraction* diagram. Any two cycles in Σ are homotopic—meaning one can be continuously deformed into the other—because there is only one nontrivial homotopy class of simple cycles on the infinite cylinder $S^1 \times \mathbb{R}$. It follows that the cycles in Σ have a well-defined vertical total order. In particular, the highest and lowest intersection points between any vertical line and Σ always lie on the *same* two essential cycles in Σ.

Without loss of generality, suppose $\gamma(0)$ is a point on the boundary of the convex hull of γ. Let C be any essential critical cycle in the attraction diagram of γ, and let $C' = L(C)$ denote the corresponding essential cycle in the dual attraction diagram. The cycle C must pass through all possible puppy positions *and* all possible human positions; thus, C contains a configuration $(0, y)$ for some parameter $y \in S^1$. Recall that $N(y)$ denotes the line normal to γ at $\gamma(y)$. Then $\gamma(0)$ must be an endpoint of the convex hull of $\gamma \cap N(y)$, which is a line segment. We conclude that C' must be either the highest or lowest essential critical cycle in the dual attraction diagram. Therefore, there are at most two critical cycles, completing the proof. □

In the rest of the paper, we mnemonically refer to the two essential critical cycles in the attraction diagram of a simple track as the ***main diagonal*** and the ***river***.

We emphasize that the converse of Lemma 5 is false; there are non-simple tracks whose attraction diagrams have exactly two essential critical cycles. (Consider the figure-eight curve ∞.) Moreover, we conjecture that Lemma 5 can be generalized to all (smooth) tracks with turning number ± 1.

4 Dexter and sinister strategies

We can visualize any strategy for the human to catch the puppy as a path through the attraction diagram, consisting entirely of segments of stable critical paths and vertical segments, that ends on the main diagonal, as shown in Figure 11. We refer to the vertical segments as *pivots*. Every pivot (except possibly the first) starts at a pivot configuration, and every pivot ends at a stable configuration.

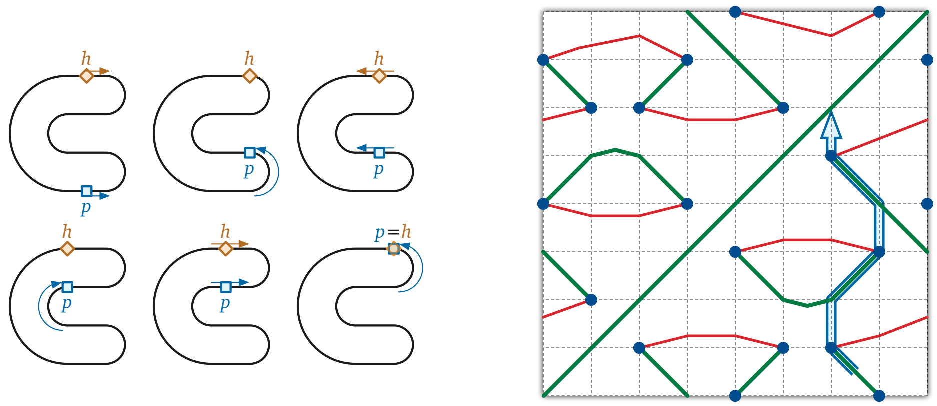

Figure 11: A sinister strategy for catching the puppy; compare with Figures 1 and 7.

We call a strategy ***dexter*** if it ends with a backward pivot—a *downward* segment, approaching the main diagonal to the *right*—and we call a configuration (x, y) *dexter* if there is a dexter strategy for catching the puppy starting at (x, y). Similarly, a strategy is ***sinister*** if it ends with a forward pivot—a *skyward* segment, approaching the main diagonal to the *left*—and a configuration is sinister if it is the start of a sinister strategy.[2] A single configuration can be both dexter and sinister; see Figure 12.

Theorem 6. *Let γ be a generic track whose attraction diagram has exactly two essential critical cycles. Every configuration on γ is dexter or sinister, or possibly both; thus, the human can catch the puppy on γ from any starting configuration.*

Before giving the proof, we emphasize that Theorem 6 does not require the track γ to be simple. Also, it is an open question whether having exactly two essential critical cycle curves is a *necessary* condition for the human to always be able to catch the puppy. (We conjecture that it is not.)

Proof. Fix a generic track γ whose attraction diagram has exactly two essential critical cycles, which we call the *main diagonal* and the *river*. Assume γ has at least one pivot configuration, since otherwise, from any starting configuration, the puppy runs directly to the human.

Let D be the set of all dexter configurations, and let S be the set of all sinister configurations. We claim that D and S are both annuli that contain both the main diagonal and the river. Because S and D meet on opposite sides of the main diagonal, this claim

[2] *Dexter* and *sinister* are Latin for right (or skillful, or fortunate, or proper, from a Proto-Indo-European root meaning "south") and left (or unlucky, or unfavorable, or malicious), respectively.

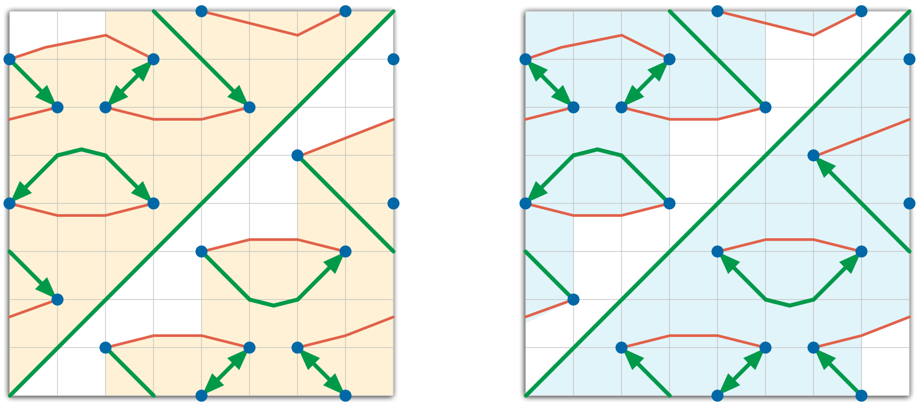

Figure 12: Dexter (orange) and sinister (cyan) configurations in the example attraction diagram. Arrows on the stable critical paths describe dexter and sinister strategies for catching the puppy.

implies that $D \cup S$ is the entire torus, completing the proof of the lemma. We prove our claim explicitly for D; a symmetric argument establishes the claim for S.

For purposes of argument, we partition the attraction diagram of γ by extending vertical segments from each pivot configuration to the next critical cycles directly above and below. We call the cells in this decomposition *trapezoids*, even though their top and bottom boundaries may not be straight line segments. At each forward pivot configuration p, we color the vertical segment above (x, y) *green* and the vertical segment below p *red*; the colors are reversed for backward vertical segments, see Figure 13.

The first step of any strategy is a (possibly trivial) pivot onto a stable critical path. Because the human and puppy can move freely within any stable critical path σ, either every point in σ is dexter, or no point in σ is dexter. Similarly, for any green pivot segment π, either every point in π is dexter or no point in π is dexter.

Consider any trapezoid τ, and let σ be the stable critical path on its boundary. Starting in any configuration in τ, the puppy immediately moves to a configuration on σ. Thus, if any point in τ is dexter, then σ is dexter, which implies that *every* point in τ is dexter. Thus, we can describe entire trapezoids as dexter or not dexter. It follows that D is the union of trapezoids.

If two trapezoids share a stable critical path *other than the main diagonal*, then either both trapezoids are dexter or neither is dexter. Similarly, if the green pivot segment leaving a pivot configuration p is dexter, then all four trapezoids incident to p are dexter; otherwise, either two or none of these four trapezoids are dexter.

We conclude that aside from the main diagonal, the boundary of D consists entirely of unstable critical paths, pivot configurations, and red vertical segments. Moreover, for every pivot configuration p on the boundary of D, the green pivot segment leaving p is *not* dexter.

By definition, every point in D is connected by a (dexter) path to the main diagonal,

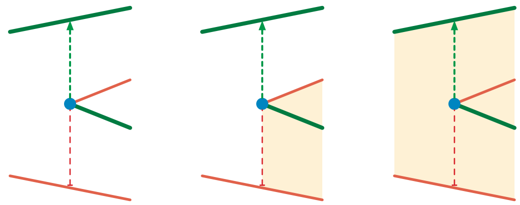

Figure 13: Possible arrangements of dexter trapezoids near a forward pivot configuration.

so D is non-empty and connected. On the other hand, D excludes a complete cycle of forward configurations just below the main diagonal. For any $x \in S^1$, let $D(x)$ denote the set of dexter configurations (x, y); this set consists of one or more vertical line segments in the attraction diagram.

Suppose for the sake of argument that some set $D(x)$ is disconnected. Because D is connected, the boundary of D must contain a *concave vertical bracket*: A vertical boundary segment π whose adjacent critical boundary segments both lie (without loss of generality) to the right of π, but D lies locally to the left of π. See Figure 14. Let p be the pivot configuration at one end of π. The green vertical segment on the other side of p is dexter, which implies that *all* trapezoids incident to p are dexter, contradicting the assumption that π lies on the boundary of D. We conclude that for all x, the set $D(x)$ is a single vertical line segment; in other words, D is a *monotone* annulus.

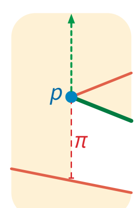

Figure 14: A hypothetical concave vertical bracket on the boundary of D.

The bottom boundary of D is the main diagonal. The monotonicity of D implies that the top boundary of D is a monotone "staircase" alternating between upward red vertical segments and rightward unstable critical paths. Every trapezoid immediately above the top boundary of D contains only forward configurations. Thus, there is a complete essential cycle ϕ of forward configurations just above the upper boundary of D. Because ϕ contains only forward configurations, ϕ must lie entirely above the river. It follows that D contains the entire river.

Symmetrically, S is an annulus bounded above by the main diagonal and bounded below by a non-contractible cycle of backward configurations; in particular, the entire river lies inside S. We conclude that $D \cup S$ is the entire configuration torus. □

If the attraction diagram of γ has more than two essential critical cycles curves, then D and S are still monotone annuli, each bounded by the main diagonal and an essential cycle of red vertical segments and unstable paths, and thus S and D each contain at least one essential critical cycle other than the main diagonal. However, $D \cup S$ need not cover the entire torus.

Corollary 7. *The human can catch the puppy on any generic simple closed track, from any starting configuration.*

5 Polygonal tracks

Our previous arguments require, at a minimum, that the track has a continuous derivative that is never equal to zero. We now extend our results to polygonal tracks, which do not have well-defined tangent directions at their vertices.

5.1 Polygonal attraction diagrams

Throughout this section, we fix a simple polygonal track P with n vertices. We regard P as a continuous piecewise linear function $P\colon S^1 \hookrightarrow \mathbb{R}^2$, parameterized by arc length. Without loss of generality $P(0)$ is a vertex of the track. We index the vertices and edges of P in order, starting with $v_0 = P(0)$, where edge e_i connects v_i to v_{i+1}; all index arithmetic is implicitly performed modulo n.

To properly describe the puppy's behavior, we must also account for the direction that the puppy is facing, even when the puppy lies at a vertex. To that end, we represent the track using both a continuous *position* function $\pi\colon S^1 \to \mathbb{R}^2$ and a continuous *direction* function $\theta\colon S^1 \to S^1$. Intuitively, the two functions describe the position and orientation of the puppy as it makes a complete circuit along P: it advances at constant speed along each edge, and it stops at each vertex to modify its direction vector, again at constant speed.

To be precise, both $\pi(y)$ and $\theta(y)$ are piecewise linear functions of the puppy's parameter $y \in S^1$. The curve $\pi(y)$ is a re-parameterization of P such that, when $\pi(y)$ is in the interior of an edge e_i of P, its derivative $\pi'(y)$ is a constant positive multiple of $\theta(y) = (v_{i+1} - v_i)/\|v_{i+1} - v_i\|$. Moreover, for each vertex v_i of P, the preimage $\pi^{-1}(v_i)$ is a non-degenerate interval $[a_i, b_i] \subset S^1$ such that $\pi'(y) = 0$ whenever $a_i < y < b_i$; also, $\theta(a_i) = (v_i - v_{i-1})/\|v_i - v_{i-1}\|$, $\theta(b_i) = (v_{i+1} - v_i)/\|v_{i+1} - v_i\|$, and $\theta(y)$ is linear and injective on $[a_i, b_i]$, turning clockwise if the edges e_{i-1} and e_i define a clockwise turn, and vice versa. (The ratio of the speeds at which the puppy moves along edges and turns around at vertices is not relevant.)

We classify any human-puppy configuration $(x, y) \in S^1 \times S^1$ as *forward*, *backward*, or *critical*, if the dot product $(P(x) - \pi(y)) \cdot \theta(y)$ is negative, positive, or zero, respectively. In any forward configuration (x, y), the puppy moves to increase the parameter y; in any backward configuration, the puppy moves to decrease the parameter y. (The human's direction is irrelevant.) The *attraction diagram* is the set of all critical configurations $(x, y) \in S^1 \times S^1$. We further classify critical configurations (x, y) as follows:

- *final* if $P(x) = \pi(y)$,
- *stable* if $(x, y - \varepsilon)$ is forward and $(x, y + \varepsilon)$ is backward for all suffic. small $\varepsilon > 0$,

- *unstable* if $(x, y - \varepsilon)$ is backward and $(x, y + \varepsilon)$ is forward for all suffic. small $\varepsilon > 0$,
- *forward pivot* if $(x, y - \varepsilon)$ and $(x, y + \varepsilon)$ are both forward for all suffic. small $\varepsilon > 0$, or
- *backward pivot* if $(x, y - \varepsilon)$ and $(x, y + \varepsilon)$ are both backward for all suffic. small $\varepsilon > 0$.

A straightforward case analysis implies that this classification is exhaustive.

To define the attraction diagram of P, we decompose the torus $S^1 \times S^1$ into a $2n \times n$ grid of rectangular cells, where each column corresponds to an edge e_j containing the human, and each row corresponds to either a vertex v_i or an edge e_i containing the puppy. The *main diagonal* of the attraction diagram is the set of all final configurations. Strictly speaking, in this case the "main diagonal" is not just a straight line, but consists of alternating diagonal and vertical segments. We can characterize the critical points inside each cell as follows:

Each edge-edge cell $e_i \times e_j$ contains at most one boundary-to-boundary path of stable critical configurations (x, y). Refer to Figure 15.

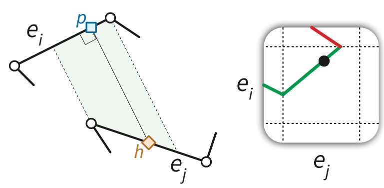

Figure 15: All edge-edge critical configurations are stable.

Each vertex-edge cell $v_i \times e_j$ contains at most one boundary-to-boundary path of stable critical configurations and at most one boundary-to-boundary path of unstable critical configurations. If the cell contains both paths, they are disjoint. A configuration (x, y) with $\pi(y) = v_i$ is stable if and only if $P(x)$ lies in the outer normal cone at v_i, and unstable if and only if $P(x)$ lies in the inner normal cone at v_i; see Figure 16.

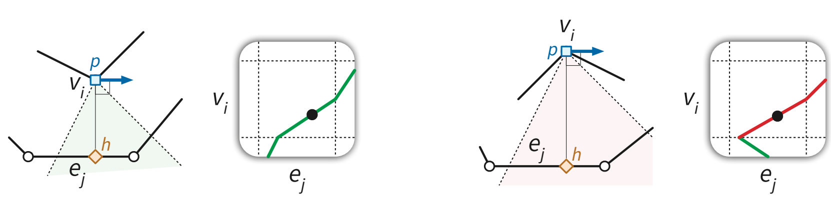

Figure 16: Stable and unstable vertex-edge critical configurations.

5.2 Polygonal pivot configurations

Unlike the attraction diagrams of generic smooth curves defined in Section 3.2, the attraction diagrams of polygons are not always well-behaved. In particular, a pivot configuration may be incident to more (or fewer) than two critical curves, and in extreme cases, pivot configurations need not even be discrete. We call such a configuration a *degenerate* pivot configuration.

In any pivot configuration (x, y), the puppy $\pi(y)$ lies at some vertex v_i, the puppy's direction $\theta(y)$ is parallel to either e_i (or e_{i+1}). Generically, each pivot configuration is a shared endpoint of an unstable critical path in cell $v_i \times e_j$ and a stable critical path in cell $e_i \times e_j$ (or $e_{i-1} \times e_j$); see Figure 17.

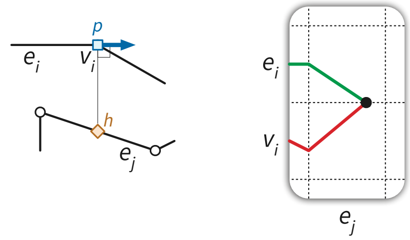

Figure 17: Near a non-degenerate pivot configuration.

There are three distinct ways in which degenerate pivot configurations can appear.

A ***type-1 degeneracy*** is caused by an acute angle on P. Specifically, let v_i be a vertex of P. The configuration (x, y) with $P(x) = \pi(y) = v_i$ is degenerate if the angle between e_{i-1} and e_i is strictly acute. In the attraction diagram of a type-1 degeneracy, two stable critical curves and two unstable critical curves end on a single vertical section of the main diagonal (corresponding to the human and the puppy being both at v_i, but the puppy facing in different directions). Refer to Figure 18.

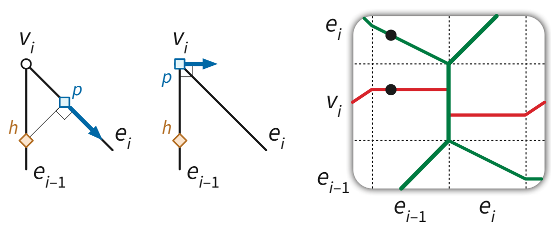

Figure 18: Stable and unstable configurations near an acute vertex angle.

A ***type-2 degeneracy*** is caused by a more specific configuration. Let e_i be an edge of P, and let ℓ be the line perpendicular to e_i through v_i (or, symmetrically, through v_{i+1}). Let v_j be another vertex of P which lies on ℓ. The configuration (x, y) with $P(x) = v_j$ and $\pi(y) = v_i$ is degenerate if:

- v_{i-1} and v_j lie in the same open halfspace of the supporting line of e_i; **and**
- v_{j-1} and v_{j+1} lie in the same open halfspace of ℓ.

A type-2 degeneracy corresponds to a vertex (pivot configuration) of degree 4 or 0 in the attraction diagram. We further distinguish these as *type-2a* and *type-2b*. Refer to Figure 19.

Finally, a ***type-3 degeneracy*** is essentially a limit of both of the previous types of degeneracies. Let e_i be an edge of P, let ℓ be the line perpendicular to e_i through v_i, and let e_j be another edge of P which lies on ℓ. The configuration (x, y) with $P(x) \in e_j$ and

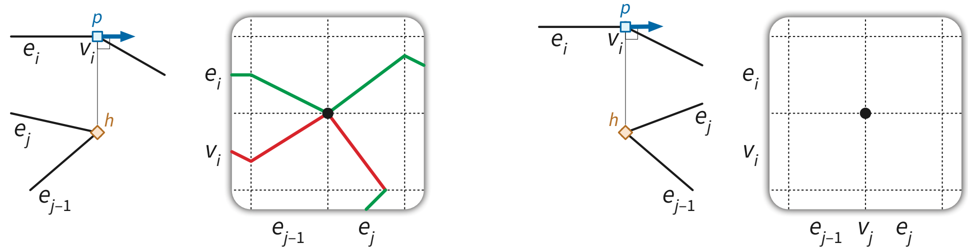

Figure 19: Type-2a and type-2b degenerate pivot configurations.

$\pi(y) = v_i$ is degenerate if vertices v_{i-1} and v_j lie in the same open halfspace of the supporting line of e_i. When this degeneracy occurs, pivot configurations are not discrete, because the point $P(x) \in e_j$ can be chosen arbitrarily. Moreover, the vertex-vertex configurations (v_j, v_i) and (v_{j-1}, v_i) have odd degree in the attraction diagram. A type-3 degeneracy can be connected to (two or more) other critical curves, or be isolated. We further distinguish these as *type-3a* and *type-3b*, depending on whether v_i is an endpoint of e_j. See Figure 20.

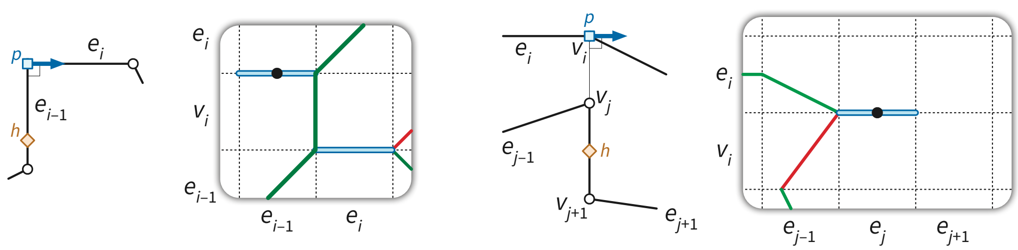

Figure 20: Type-3a and type-3b degenerate pivot configurations.

In Section 5.3 we first consider polygonal tracks which do not have any degeneracies of these three types. To simplify exposition, we forbid degeneracies by assuming that no vertex angle in P is acute and that no three vertices of P define a right angle. In Section 5.5 we lift these assumptions by *chamfering* the polygon, cutting off a small triangle at each vertex.

5.3 Catching puppies on generic obtuse polygons

Generic obtuse polygonal tracks behave almost identically to smooth tracks, once we properly define the attraction diagram and dual attraction diagram.

Lemma 8. *Let P be a simple polygon with no acute vertex angles, in which no three vertices define a right angle. The attraction diagram of P is the union of disjoint simple critical cycles.*

Proof. Each edge-edge cell $e_i \times e_j$ contains at most one section of stable critical configurations (x, y) (Figure 15). For each such configuration, the points $\pi(y) \in e_i$ and $P(x) \in e_j$ are connected by a line perpendicular to e_i. Because no three vertices of P define a right angle, these points cannot both be vertices of P; thus, any critical path inside the cell $e_i \times e_j$ avoids the corners of that cell.

Each vertex-edge cell $v_i \times e_j$ contains at most one section of a stable and one section of an unstable path (Figure 16). Again, because no three vertices of P define a right angle, these paths avoid the corners of the cell $v_i \times e_j$.

It follows from the definition of pivot that, in any pivot configuration (x, y), the puppy lies at a vertex $\pi(y) = v_i$, and the puppy's direction $\theta(y)$ is parallel to either e_i (or e_{i+1}). Also, by the above, the human lies in the interior of some edge: $P(x) \in e_j$. Moreover, our assumptions on P imply that there are no degenerate pivot configurations; thus, each pivot configuration is a shared endpoint of exactly one unstable critical path in cell $v_i \times e_j$ and exactly one stable critical path in cell $e_i \times e_j$ (or $e_{i-1} \times e_j$).

Thus, the set of unstable critical configurations is the union of x-monotone paths whose endpoints are pivot configurations. Similarly, the set of stable critical configurations is also the union of x-monotone paths whose endpoints are pivot configurations. Moreover, each unstable critical path lies in a single vertex strip.

Because every vertex angle in P is obtuse, every configuration (x, y) where the human $P(x)$ lies on an edge e_i and the puppy $\pi(y)$ lies on the previous edge e_{i-1} is either forward of final. Similarly, if $P(x) \in e_{i-1}$ and $\pi(y) \in e_i$, then the configuration (x, y) is either backward or final. Thus, the main diagonal is disjoint from all other critical cycles; in fact, no other critical cycle intersects any grid cell that touches the main diagonal.

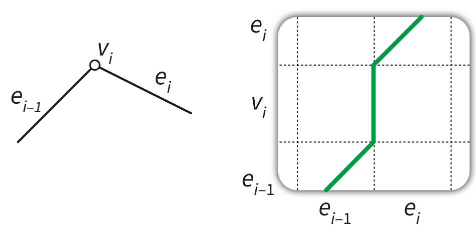

Figure 21: Near the main diagonal.

This completes the classification of all critical configurations. We conclude that the attraction diagram consists of the (simple, closed) main diagonal and possibly other simple closed curves composed of stable and unstable critical paths meeting at pivot configurations. All these critical cycles are disjoint. □

Lemma 9. *Let P be a simple polygon with no acute vertex angles, in which no three vertices define a right angle. If the attraction diagram of P has exactly two essential critical cycles, then the human can catch the puppy on P, starting from any initial configuration.*

The remainder of the proof is essentially unchanged from the smooth case. For any configuration (x, y), let $T(y)$ denote the directed "tangent" line through $\pi(y)$ in direction $\theta(y)$, and let $L(x, y)$ denote the signed distance from $P(x)$ to $T(y)$, signed positively if $P(x)$ lies to the left of $T(y)$ and negatively if $P(x)$ lies to the right of $T(y)$. The *dual attraction diagram* of P consists of all points $(y, L(x, y)) \in S^1 \times \mathbb{R}$ where (x, y) is a critical configuration.

As in the smooth case, the map $(x, y) \mapsto (y, L(x, y))$ is a homeomorphism from the critical cycles in the attraction diagram to the curves in the dual attraction diagram; moreover, this map preserves the contractibility of each critical cycle.

Lemma 10. *Let P be a simple polygon with no acute vertex angles, in which no three vertices define a right angle. The attraction diagram of P contains exactly two essential critical cycles.*

Theorem 11. *Let P be a simple polygon with no acute vertex angles, in which no three vertices define a right angle. The human can catch the puppy on P, starting from any initial configuration.*

We can easily extend this theorem to polygons with degenerate pivot configurations of type 2b and type 3b. Since these correspond to vertically isolated forward or backward pivot configurations in the attraction diagram, they do not impact the existence of a strategy to catch the puppy. The puppy will just move over them as if they were normal forward or backward configurations.

Corollary 12. *Let P be a simple polygon with no degeneracies of type 1, type 2a, or type 3a. The human can catch the puppy on P, starting from any initial configuration.*

5.4 Chamfering

We now extend our analysis to arbitrary simple polygons. We define a *chamfering* operation, which transforms a polygon P into a new polygon $\bar{P}$. First we show that $\bar{P}$ has no degenerate pivot configurations of type 1, 2a, or 3a (although it may still have degeneracies of type 2b and type 3b). Hence there is a strategy to catch the puppy on $\bar{P}$. Finally, we show that such a strategy can be correctly translated back to a strategy on P.

Let P be an arbitrary simple polygon, and let $\varepsilon > 0$ be smaller than half of any distance between two non-incident features of P. Then the ε-*chamfered* polygon $\bar{P}$ is another polygon with twice as many vertices as P, defined as follows. Refer to Figure 22. For each vertex v_i of P, we create two new vertices v_i' and v_i'', where v_i' is placed on e_{i-1} at distance ε from v_i, and v_i'' is placed on e_i at distance ε from v_i. Edge e_i' in $\bar{P}$ connects v_i'' to v_{i+1}', and a new *short edge* s_i connects v_i' to v_i''. Note that the condition on ε implies that $\bar{P}$ is itself a simple (i.e., not self-intersecting) polygon.

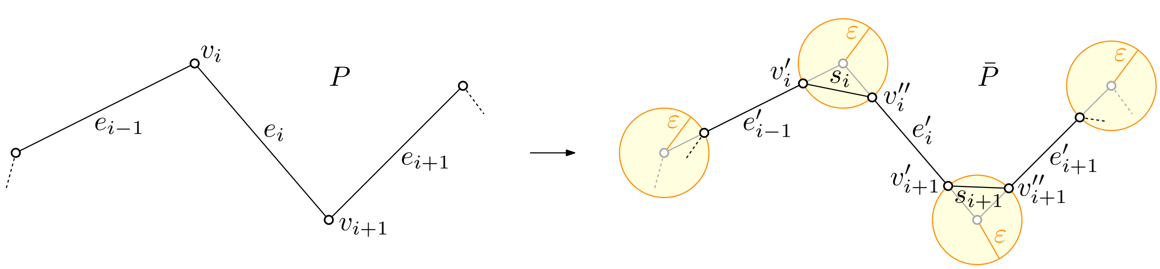

Figure 22: The chamfering operation.

The chamfering operation alters the local structure of the attraction diagram near every vertex. The idea is that at non-degenerate configurations, the change will not influence the behavior of the puppy, and as such will not influence the existence of any catching strategies. However, at degenerate configurations, the change in the structure is significant. We will argue in Section 5.5 that the changes are such that every strategy in the chamfered polygon translates to a strategy in the original polygon.

Here we review again the different types of degenerate pivot configurations, and how the ε-chamfering operation, for a small-enough ε, affects the local structure of the attraction diagram in each case. Refer to Figure 23.

- Near type-1 degeneracies, the higher-degree vertices on the main diagonal disappear. Instead, two separate critical curves almost touch the main diagonal: one from above and one from below.
- Near type-2a degeneracies, the degree-4 vertex disappears. Instead, the two incident critical curves coming from the left are connected, and the two incident curves coming from the right are connected.
- Near type-2b degeneracies, the isolated pivot vertex simply disappears.
- Near type-3 degeneracies, the degenerate pivot "vertex" disappears. Any connected critical curve is locally rerouted away from the degenerate location.

5.5 Catching puppies on arbitrary simple polygons

Even when the chamfering radius ε is arbitrarily small, the attraction diagram of the chamfered polygon $\bar{P}$ may have type-2b and type-3b degeneracies, and even new non-degenerate critical curves, that are not present in the original attraction diagram. See Figures 24 and 25 for examples. We argue in the next lemma that these are the only degeneracies that can appear in $\bar{P}$.

Lemma 13. *Let P be an arbitrary simple polygon. For all sufficiently small ε, the ε-chamfered polygon $\bar{P}$ has no degenerate pivot configurations of type 1, type 2a, or type 3a.*

Proof. First, note that $\bar{P}$ has no type-1 or type-3a degeneracies: we replace each vertex v_i with angle α_i by two new vertices v_i' and v_i'' with angles $\alpha_i' = \alpha_i'' = \pi - \frac{1}{2}(\pi - \alpha_i) = \frac{1}{2}\pi + \frac{1}{2}\alpha_i > \frac{1}{2}\pi$.

Next, we consider the type-2 degeneracies, which may occur for some values of ε. We argue that each potential type-2a degeneracy only occurs for at most one value of ε; since there are finitely many potential degeneracies, the lemma then follows.

Note that, as we vary ε, all vertices of $\bar{P}$ move linearly and with equal speed. Thus, if more than one value of ε gives rise to a type-2a degeneracy, then all of them do. There are two configurations in $\bar{P}$ that could potentially give rise to infinitely many type-2a degeneracies. We argue that, in fact, such configurations cannot satisfy all requirements of a type-2a degeneracy.

- An edge e_i' has endpoint v_i' (or symmetrically, v_{i-1}'') such that the line ℓ through v_i' and perpendicular to e_i' contains another vertex v_j' (or v_{j-1}''). Refer to Figure 26. Then, as

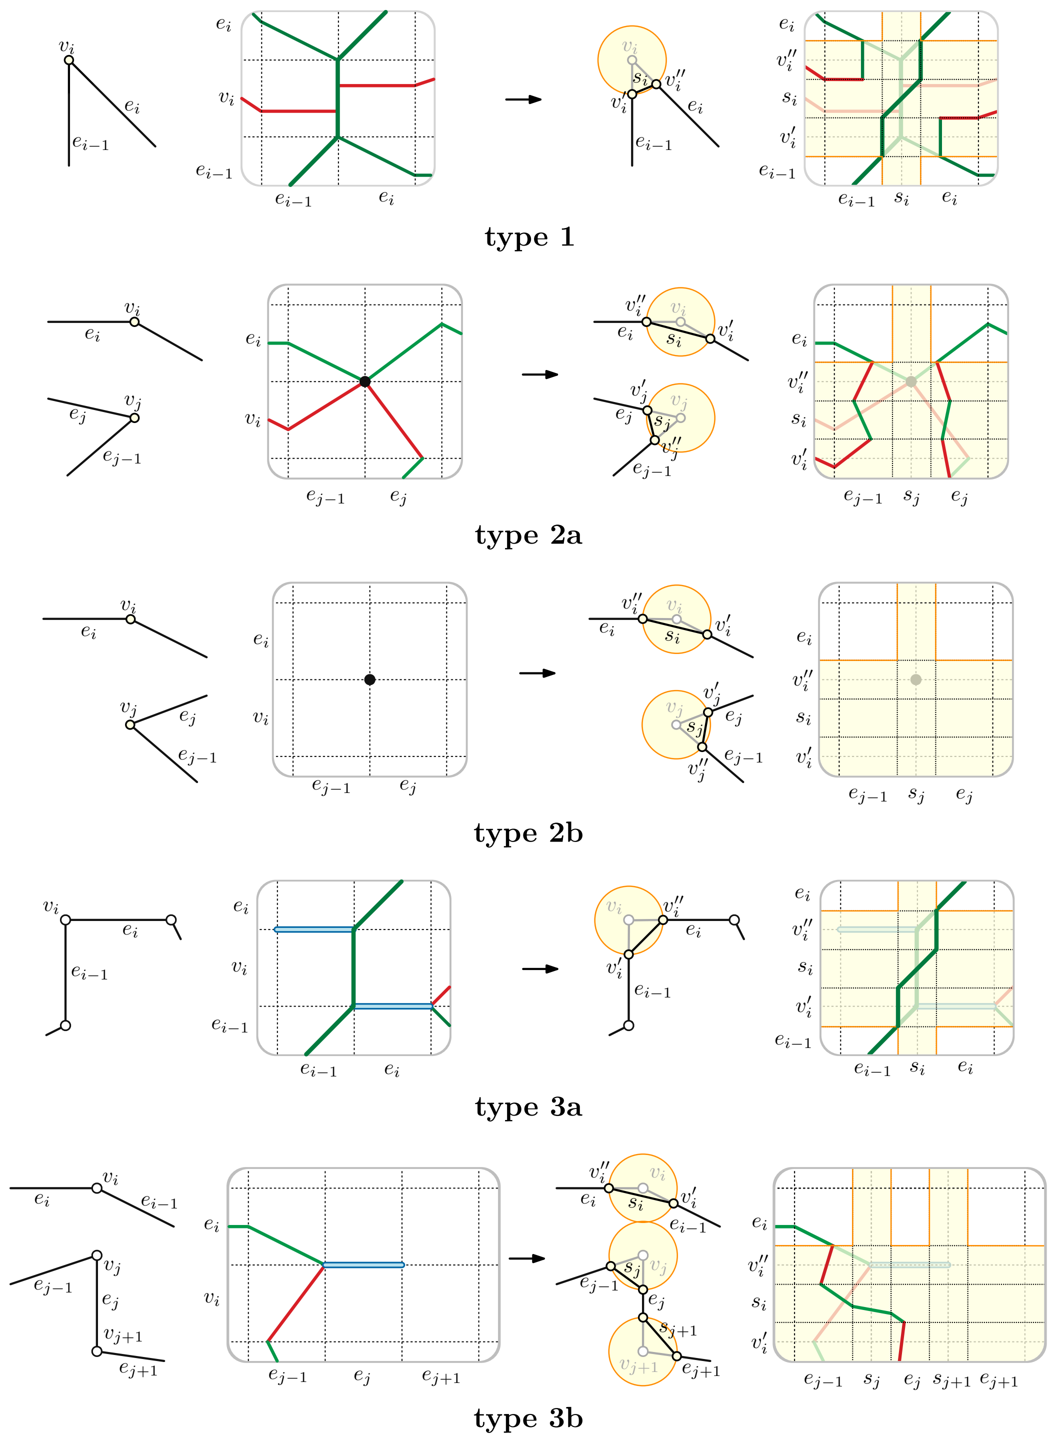

Figure 23: Effect of the chamfering operation on the attraction diagram near degenerate pivot configurations. The size of ε is exaggerated; the figures show the combinatorial structure of the chamfered diagram for a much smaller value of ε. Only the effect of chamfering vertices relevant for the degeneracy is shown.

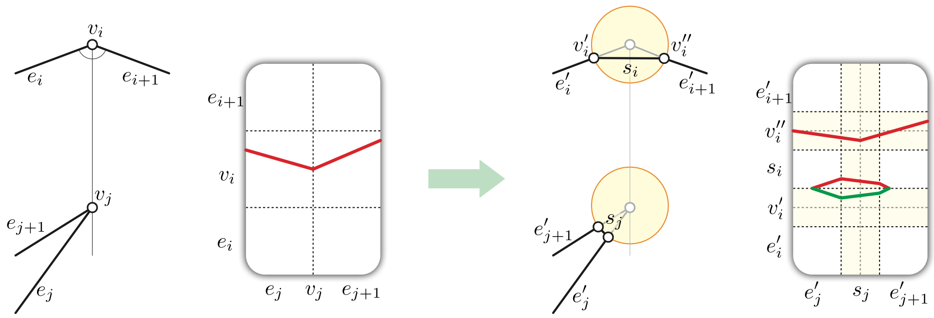

Figure 24: Chamfering P can create a new non-degenerate critical curve when one vertex of P lies on the angle bisector of another.

v_i' moves along e_i', ℓ moves at the same speed as v_i', and v_j' moves in the same direction at the same speed along e_j'. So e_j' is parallel to e_i'. But since the angles of $\bar{P}$ are obtuse, we conclude that v_{j-1}'' and v_j'' lie on the opposite sides of ℓ; thus, this cannot be a type-2 degeneracy.

- A short edge s_i of $\bar{P}$ has an endpoint v_i' (or symmetrically, v_i'') such that the line ℓ through v_i' and perpendicular to s_i contains another vertex v_j' (or v_{j-1}''). Refer to Figure 27. In this case, vertex v_j must lie on the angle bisector of edges e_i and e_{i+1}, and edges e_i and e_j must be parallel. Because the angles of $\bar{P}$ are obtuse, s_i and e_i' lie on opposite sides of ℓ. Now, as ε varies, v_i' moves along e_i', the slope of s_i does not change, and thus ℓ remains parallel to itself. Since v_j' moves in a direction concordant with ℓ's direction, e_j' lies on the same sides of ℓ as e_i'. Thus, this cannot be a type-2a degeneracy. Note that it is possible that v_j'' lies on the same side of ℓ as e_j', in which case we have a degeneracy of type 2b (Figure 27 (left)), or that v_j'' lies on ℓ, in which case we have a degeneracy of type 3b (Figure 27 (middle)). If v_j'' lies on the opposite side of ℓ, there is no degeneracy (Figure 27 (right)). □

Note that it may be tempting to define a different chamfering parameter ε for each vertex of P, in order to eliminate also the type-2b and type-3b degeneracies from $\bar{P}$. The reason why we insist on having the same ε for all vertices will become apparent shortly, when proving Lemma 14.

Let P be an arbitrary simple polygon and $\bar{P}$ an ε-chamfered copy without degeneracies of type 1, type 2a, or type 3a. We say a parameter value x is *verty* whenever $P(x)$ is at distance at most ε from a vertex of P. We say a parameter value x is *edgy* if it is not verty. We reparameterize $\bar{P}$ such that $P(x) = \bar{P}(x)$ whenever x is edgy; the parameterization of $\bar{P}$ is uniformly scaled for verty parameters. We say a configuration (x, y) is edgy when x and y are both edgy.

We say a path in the attraction diagram is *valid* if it describes a human and puppy behavior that obeys the rules imposed on the puppy and the human, as explained in Section 1. For polygonal tracks, it is not restrictive to assume that a valid path is piecewise linear, and that the derivative of the human's parameter value x only changes sign at pivot configurations (that is, the human may invert direction along the curve only when the configuration is a pivot one).

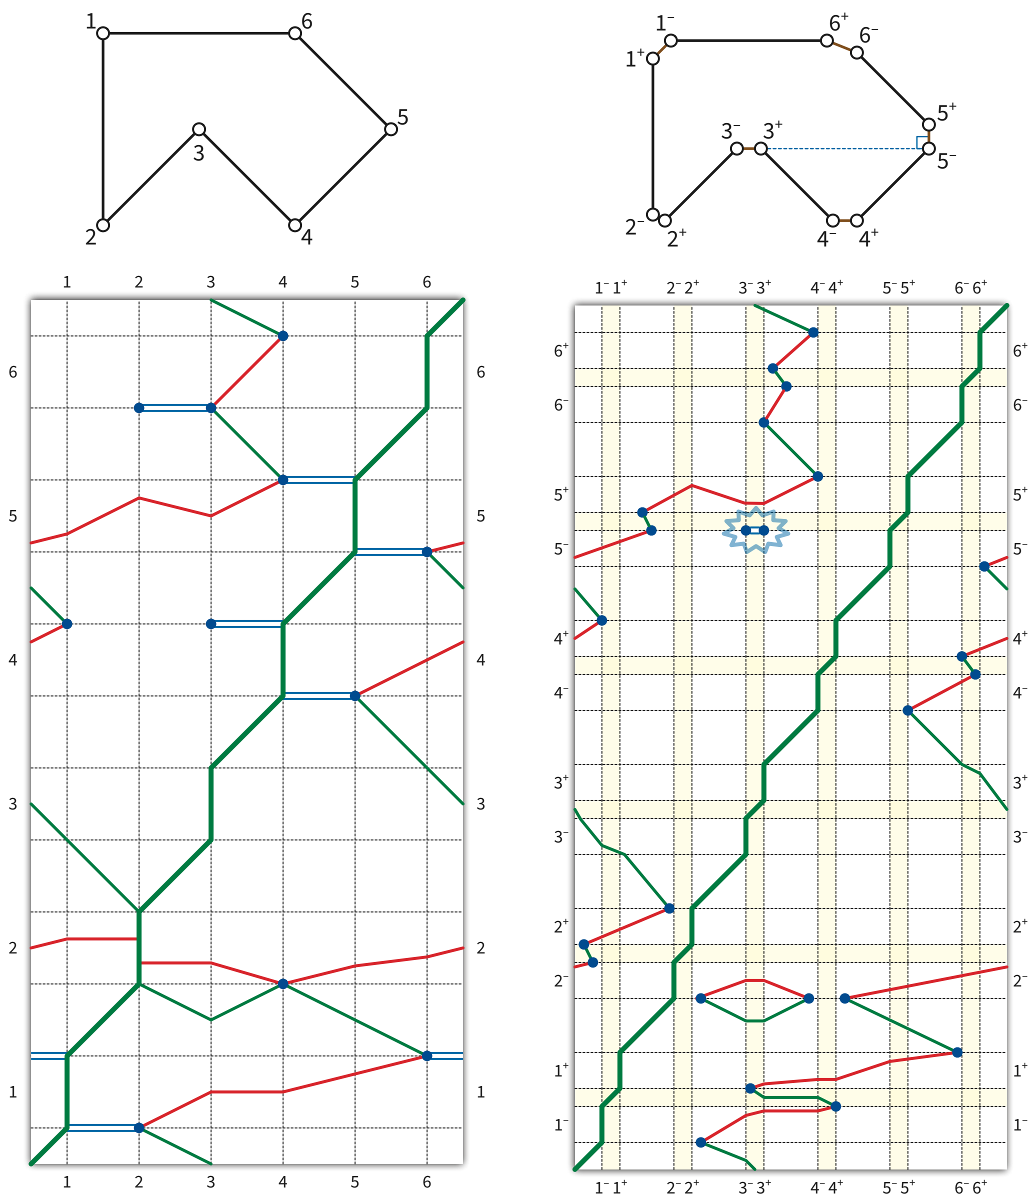

Figure 25: The attraction diagram of a degenerate polygon, before and after chamfering. All existing degeneracies disappeared in the chamfered polygon, which does have one new but harmless type-3b degeneracy.

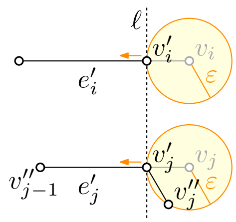

Figure 26: Potential new degenerate pivot configurations based on a (shortened) original edge e_i'. For ε small enough, there can be no degeneracy.

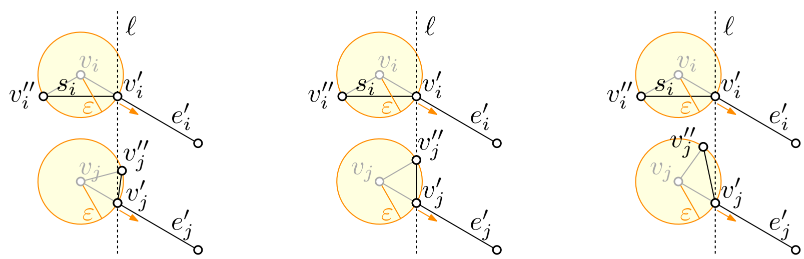

Figure 27: Potential new degenerate pivot configurations based on a short edge s_i. For any ε we may still have a new degeneracy of type 2b (left), 3b (middle), or no degeneracy (right).

Lemma 14. *Assuming ε is sufficiently small, for any valid path σ between two stable edgy configurations (x_1, y_1) and (x_2, y_2) in the attraction diagram of $\bar{P}$, there is a valid path σ' between (x_1, y_1) and (x_2, y_2) in the attraction diagram of P.*

Proof. We will describe how to obtain σ' by slightly deforming σ in the non-edgy configurations, assuming that ε is small enough. In fact, it will suffice to show that σ and σ' determine the same "qualitative behavior" of the puppy. That is, let ψ be a valid path in the attraction diagram of P or $\bar{P}$, and consider the ordered sequence of all configurations $((\tilde{x}_i, \tilde{y}_i))_{1\leq i\leq k}$ along ψ where the puppy's parameter value $\tilde{y}_i$ transitions from edgy to verty or vice versa. The *qualitative behavior* of the puppy determined by ψ is defined as the sequence $q_\psi = (\tilde{y}_i)_{1\leq i\leq k}$. We will show that $q_\sigma = q_{\sigma'}$, thus proving the lemma.

The intuition is that there is a direct correspondence between edgy configurations in the two diagrams, and we only have to ensure that the puppy has the correct behavior when the configuration is not edgy, i.e., the human or the puppy is in an ε-neighborhood of a vertex of P.

Let ρ be a maximal subpath of σ where the puppy's parameter y remains edgy except possibly at the endpoints, i.e., the puppy remains on some edge e_i' of $\bar{P}$ while the human walks along $\bar{P}$. We argue that, if the human moves in the same way along P, thus determining a path ρ' in the attraction diagram of P, then the puppy never leaves e_i. Moreover, if ρ terminates with the puppy on an endpoint of e_i', say v_i'', then ρ' terminates with the puppy in a verty position corresponding to v_i.

Observe that, if the projection of a vertex v_j on the line supporting e_i lies in the interior of e_i, then the projection of the short edge s_j on the same line lies in the interior of

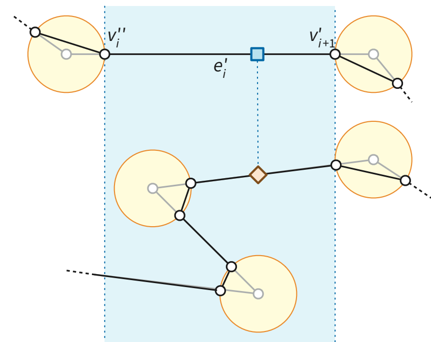

Figure 28: As long as the puppy stays on the chamfered edge e_i', its qualitative behavior is the same on the original and chamfered polygon.

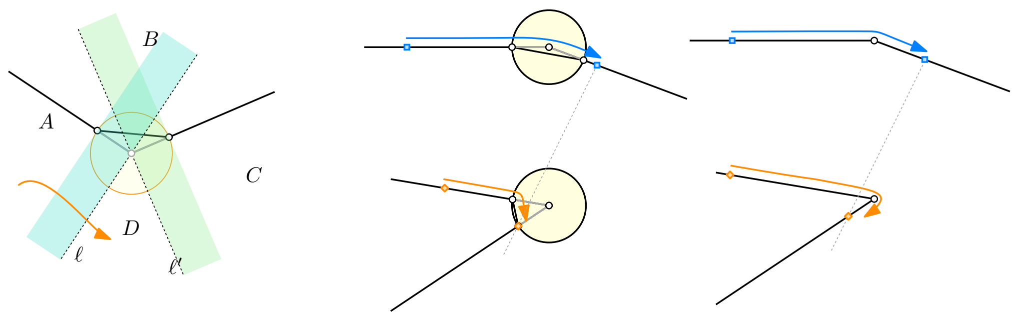

Figure 29: Left: When the puppy is around a vertex, its qualitative behavior is determined by the region where the human lies (either A, B, C, or D). Center and right: Once the human gets in a neighborhood of the lower vertex, the puppy makes a jump forward. This behavior can be replicated in P, as the configuration corresponds to a type-2a degeneracy.

e_i', assuming that ε is small enough. Thus, the puppy's behavior according to ρ' is the same as with ρ, except when the human reaches a neighborhood of a vertex v_j that projects on an endpoint of e_i, say v_i.

In the latter case, since the chamfering parameter ε is the same for both v_i and v_j, the human cannot reach the interior of the short edge s_j before the puppy reaches the interior of the short edge s_i. However, since ρ keeps the puppy on e_i', this is not possible. Thus, the puppy in ρ' behaves in the same way as in ρ in every case. See Figure 28.

Let us now consider a maximal subpath τ of σ where the puppy's parameter y remains verty. Furthermore, assume that both endpoints of τ have a puppy parameter at the boundary between verty and edgy (such is the situation when τ is between two subpaths of σ where the puppy parameter is edgy). As before, we will construct a path τ' in the attraction diagram of P such that the puppy has the same qualitative behavior as in τ. Refer to Figure 29.

By assumption, throughout τ, the puppy always remains on a short edge, say s_i, possibly rotating its direction vector while it is at a vertex of s_i. Let ℓ and ℓ' be the lines through v_i orthogonal to e_{i-1} and e_i, respectively. We say that v_i is *generic* if no other vertex lies on either ℓ or ℓ'. We denote by S the infinite strip of width ε bounded by ℓ and v_i'. Similarly, we denote the infinite strip bounded by ℓ' and v_i'' by S'.

If v_i is generic, then we can choose ε small enough such that the strips S and S' intersect no short edges of $\bar{P}$ other than s_i. Thus, whenever the human moves within one of the strips S or S', it stays within some edge e_j' of $\bar{P}$. It follows that, if the human in τ' replicates the identical behavior within S and S' as the human in τ, this determines the same qualitative behavior of the puppy (i.e., the puppy in τ leaves s_i from one of its endpoints if and only if the puppy in τ' moves to the corresponding edgy position in P).

Denote by A, B, C, D the four regions of the plane bounded by ℓ and ℓ', as in Figure 29 (left), and assume that the human in τ moves outside of S and S' within one of these four regions. In the case of D, the puppy never leaves s_i; replicating the human's movements in P (straightforwardly modified around the vertices to match the shape of P) causes the puppy to stay at v_i, thus having the same qualitative behavior. On the other hand, if the human is anywhere in $A \setminus S$ or in $C \setminus S'$, the puppy immediately moves to an edgy position, both in $\bar{P}$ and in P.

Suppose now that the human is in B, and consider the open strip S'' consisting of the union of all the lines perpendicular to s_i that intersect the interior of s_i. If the human moves within $B \setminus (S \cup S' \cup S'')$, we reason in the same way as with $A \setminus S$ and $C \setminus S'$. If the human is anywhere in $B \cap S''$, then the configuration stabilizes with the puppy in the interior of s_i. However, observe that, in order to reach this region, the human must have crossed $B \cap S''$, thus causing the puppy to move outside of s_i or never enter s_i in the first place. Hence, this case never occurs.

Finally, let us consider the case where v_i is not generic. We can argue in the same way as in the generic case, except when the human moves in a neighborhood of a vertex v_j that lies on, say, ℓ. In this case, we can choose ε small enough so that both S' and S'' (as defined above) are disjoint from the disk of radius ε centered at v_j. Now, if the human ever enters the region C while in a neighborhood of v_j, we can reason as above.

The only remaining case is the one where v_i and v_j give rise to a type-2a degeneracy in the attraction diagram of P, as illustrated in Figure 29 (center and right). Since the chamfering parameter ε is the same for both v_i and v_j, the short segment s_j lies entirely in the strip S. Also, by our choice of ε, s_j lies outside the open strip S''. Thus, if the human in τ ever reaches s_j, the puppy exits s_i from v_i''. This behavior can be replicated in P if the human moves to the vertex v_j, which causes the puppy to travel around vertex v_i.

We have proved that the path σ can be decomposed into subpaths ρ_1, τ_1, ρ_2, τ_2, $\dots$, ρ_k, each of which has a corresponding path ρ_i' or τ_i' in the attraction diagram of P which determines the same qualitative behavior of the puppy. By definition of "qualitative behavior", the ending point of any path in the sequence ρ_1', τ_1', ρ_2', τ_2', $\dots$, ρ_k' coincides with the starting point of the next path. Thus, the paths can be concatenated to form the desired path σ'. □

We are now ready to prove our main result.

Theorem 15. *Let P be a simple polygon. The human can catch the puppy on P, starting from any initial configuration.*

Proof. Let ε be so small as to satisfy both Lemma 13 and Lemma 14. Consider an arbitrary starting configuration on P. If the starting configuration is not stable, we let the puppy move until it is. If the resulting configuration is not edgy, we move the human along P until we reach an edgy configuration (x, y). (This must be possible, except if the puppy stays in an ε-neighborhood of a vertex for the entire time; in that case, we can catch the puppy trivially, by going to that vertex.)

By Lemma 13, the ε-chamfered polygon $\bar{P}$ has no degeneracies of type 1 or type 2a or type 3a. Thus, by Corollary 12, there exists a strategy for the human to catch the puppy on $\bar{P}$. If the end configuration of this strategy is not edgy, we may now simply move human and puppy together to an edgy final configuration (f, f). By Lemma 14, there is an equivalent strategy to reach (f, f) from (x, y) on P. Combined with the initial path to (x, y), this gives us a path from an arbitrary starting configuration to a final configuration on P. □

6 Further questions

For simple curves, we have only proved that a catching strategy exists. At least for polygonal tracks, it is straightforward to compute such a strategy in $O(n^2)$ time by searching the attraction diagram. In fact, we can compute a strategy that minimizes the total distance traveled by either the human or the puppy in $O(n^2)$ time, using fast algorithms for shortest paths in toroidal graphs [15, 17]. Unfortunately, this quadratic bound is tight in the worst case if the output strategy must be represented as an explicit path through the attraction diagram. We conjecture that an optimal strategy can be described in only $O(n)$ space by listing only the human's initial direction and the sequence of points where the human reverses direction. On the other hand, an algorithm to compute such an optimal strategy in subquadratic time seems unlikely.

If the track is a *smooth curve* of length ℓ whose attraction diagram has k pivot configurations, a trivial upper bound on the distance the human must walk to catch the puppy is $\ell \cdot k/2$. In any optimal strategy, the human walks straight to the point on the curve corresponding to a pivot located at one of the two endpoints of the current "stable sub-curve" of a critical curve (walking less than ℓ). Then the configuration moves to another stable sub-curve, and so on, never visiting the same stable sub-curve twice. Our question is whether a better upper bound can be proved.

In fact, if minimizing distance is not a concern, we conjecture that *no* reversals are necessary. That is, on *any* simple track, starting from *any* configuration, we conjecture that the human can catch the puppy *either* by walking only forward along the track *or* by walking only backward along the track. Figure 2 and its reflection show examples where each of these naïve strategies fails, but we have no examples where both fail. (Our proof of Theorem 2 implies that the human can always catch the puppy on an *orthogonal* polygon by walking *at most once* around the track in some direction, depending on the starting configuration.)

More ambitiously, we conjecture that the following *oblivious* strategy is always successful: walk twice around the track in one (arbitrary) direction, then walk twice around the track in the opposite direction.

Another interesting question is to what extent our result applies to self-intersecting curves in the plane, when we consider the two strands of the curve at an intersection point to be distinct. It is easy to see that the human cannot catch the puppy on a curve that traverses a circle twice; see Figure 4. Indeed, we know how to construct examples of bad curves with any rotation number *except* -1 and $+1$. We conjecture that Lemma 5, and therefore our main result, extends to all non-simple tracks with rotation number ± 1. Similarly, are there interesting families of curves in $\mathbb{R}^3$ there the human and puppy can always meet?

Finally, it is natural to consider similar pursuit-attraction problems in more general domains. Theorem 1 shows that the human can always catch the puppy in the interior of a simple polygon, by walking along the dual tree of any triangulation. Can the human always catch the puppy in any planar straight-line graph? Inside any polygon with holes?

Acknowledgements

The authors would like to thank Ivor van der Hoog, Marc van Kreveld, and Frank Staals for helpful discussions in the early stages of this work, and Joseph O'Rourke for clarifying the history of the problem. Portions of this work were done while the second author was visiting Utrecht University.

M.A. partially supported by the VILLUM Foundation grant 16582. M.L. partially supported by the Dutch Research Council (NWO) under project number 614.001.504 and 628.011.005. T.M. supported by the Dutch Research Council (NWO) under Veni grant EAGER J.U. supported by the Dutch Research Council (NWO) under project number 612.001.651. J.V. supported by the Dutch Research Council (NWO) under project number 612.001.651.

References

[1] Zachary Abel, Hugo A. Akitaya, Erik D. Demaine, Martin L. Demaine, Adam Hesterberg, Matias Korman, Jason S. Ku, and Jayson Lynch. Negative instance for the edge patrolling beacon problem. *Preprint,* *https://arxiv.org/abs/2006.01202*, 2020.

[2] Israel Aldana-Galván, Jose L. Alvarez-Rebollar, Juan Carlos Cataa-Salazar, Nestaly Marin-Nevárez, Erick Solís-Villareal, Jorge Urrutia, and Carlos Velarde. Beacon coverage in orthogonal polyhedra. In *Proceedings of the 29th Canadian Conference on Computational Geometry (CCCG 2017)*, pages 156–161, 2017.

[3] Israel Aldana-Galván, Jose L. Alvarez-Rebollar, Juan Carlos Cataa-Salazar, Nestaly Marin-Nevárez, Erick Solís-Villareal, Jorge Urrutia, and Carlos Velarde. Tight bounds for illuminating and covering of orthotrees with vertex lights and vertex beacons. *Graphs and Combinatorics*, 36:617–630, 2929.

[4] Helmut Alt and Michael Godau. Computing the Fréchet distance between two polygonal curves. *Int. J. Comput. Geom. Appl.*, 5:75–91, 1995. doi:10.1142/S0218195995000064.

[5] Sang Won Bae, Chan-Su Shin, and Antoine Vigneron. Tight bounds for beacon-based coverage in simple rectilinear polygons. *Computational Geometry*, 80:40–52, 2019. doi:10.1016/j.comgeo.2019.02.002.

[6] Michael Biro. *Beacon-based routing and guarding*. PhD thesis, State University of New York at Stony Brook, 2013.

[7] Michael Biro, Jie Gao, Justin Iwerks, Irina Kostitsyna, and Joseph S. B. Mitchell. Beacon-based routing and coverage. In *Proceedings of the 21st Fall Workshop on Computational Geometry*, 2011.

[8] Michael Biro, Jie Gao, Justin Iwerks, Irina Kostitsyna, and Joseph S. B. Mitchell. Beacon based structures in polygonal domains. In *Abstracts of the 1st Computa- tional Geometry: Young Researchers Forum*, 2012.

[9] Michael Biro, Jie Gao, Justin Iwerks, Irina Kostitsyna, and Joseph S. B. Mitchell. Combinatorics of beacon routing and coverage. In *Proceedings of the 25th Canadian Conference on Computational Geometry (CCCG 2013)*, 2013.

[10] Michael Biro, Justin Iwerks, Irina Kostitsyna, and Joseph S. B. Mitchell. Beacon-based algorithms for geometric routing. In *13th International Symposium on Algorithms and Data Structures (WADS 2013)*, pages 158–169, 2013. doi:10.1007/978-3-642-40104-6_14.

[11] Prosenjit Bose and Thomas C. Shermer. Gathering by repulsion. *Computational Geometry*, 90:101627, 2020. doi:10.1016/j.comgeo.2020.101627.

[12] Pierre Bouguer. Sur de nouvelles courbes ausquelle on peut donner le nom de linges de poursuite. *Mémoires de mathématique et de physique tirés des registres de l'Académie royale des sciences*, pages 1–14, 1732. URL: https://gallica.bnf.fr/ark:/12148/bpt6k35294.

[13] Johans Cleve and Wolfgang Mulzer. Combinatorics of beacon-based routing in three dimensions. *Computational Geometry*, 91:101667, 2020. doi:10.1016/j.comgeo.2020.101667.

[14] Pierre de Maupertuis. Sure les courbes de poursuite. *Mémoires de mathématique et de physique tirés des registres de l'Académie royale des sciences*, pages 15–17, 1732. URL: https://gallica.bnf.fr/ark:/12148/bpt6k35294.

[15] John R. Gilbert, Joan P. Hutchinson, and Robert Endre Tarjan. A separator theorem for graphs of bounded genus. *J. Algorithms*, 5(3):391–407, 1984. doi:10.1016/0196-6774(84)90019-1.

[16] Arthur S. Hathaway. Problems and solutions: Problem 2801. *American Mathematical Monthly*, 27(1):31, 1920. doi:10.2307/2973244.

[17] Monika R. Henzinger, Philip Klein, Satish Rao, and Sairam Subramanian. Faster shortest-path algorithms for planar graphs. *J. Comput. Syst. Sci.*, 55(1):3–23, 1997. doi:10.1006/jcss.1997.1493.

[18] Irina Kostitsyna, Bahram Kouhestani, Stefan Langerman, and David Rappaport. An Optimal Algorithm to Compute the Inverse Beacon Attraction Region. In *34th International Symposium on Computational Geometry (SoCG 2018)*, 2018. doi:10.4230/LIPIcs.SoCG.2018.55.

[19] Bahram Kouhestani and David Rappaport. Edge patrolling beacon. In *Abstracts from the 20th Japan Conference on Discrete and Computational Geometry, Graphs, and Games (JCDCGGG 2017)*, pages 101–102, 2017.

[20] Bahram Kouhestani, David Rappaport, and Kai Salomaa. On the inverse beacon attraction region of a point. In *Proceedings of the 27th Canadian Conference on Computational Geometry (CCCG 2015)*, 2015.

[21] Bahram Kouhestani, David Rappaport, and Kai Salomaa. The length of the beacon attraction trajectory. In *Proceedings of the 28th Canadian Conference on Computational Geometry (CCCG 2016)*, pages 69–74, 2016.

[22] Bahram Kouhestani, David Rappaport, and Kai Salomaa. Routing in a polygonal terrain with the shortest beacon watchtower. *Computational Geometry*, 68:34–47, 2018. doi:10.1016/j.comgeo.2017.05.005.

[23] John E. Littlewood. *Littlewood's Miscallany: edited by Béla Bollobás*. Cambridge University Press, 1986.

[24] Amirhossein Mozafari and Thomas C. Shermer. Transmitting particles in a polygonal domain by repulsion. In *12th International Conference on Combinatorial Optimization and Applications (COCOA 2018)*, pages 495–508, 2018. doi:10.1007/978-3-030-04651-4_33.

[25] Paul J. Nahin. *Chases and Escapes: The Mathematics of Pursuit and Evasion*. Princeton University Press, 2007.

[26] Thomas C. Shermer. A combinatorial bound for beacon-based routing in orthogonal polygons. In *Proceedings of the 27th Canadian Conference on Computational Geometry (CCCG 2015)*, 2015.