A curve with rotation number one that is not universal for beacon routing

Florestan Brunck *Maarten Löffler* *Rodrigo I. Silveira*

Keywords: Beacon Routing, Plane Curves, Navigation, Universal Curves

Abstract

Beacon routing is the study of how to indirectly route objects through various domains. Consider two points a—which we think of as the *attractor*—and b—which we think of as the *ball*—lying on a smooth closed curve in the plane. In curve-restricted beacon routing, b moves along the curve as long as its Euclidean straight-line distance to a decreases, until this distance is locally minimal. Assuming b moves infinitely faster than a, the goal is to move a along the curve in such a way that a ends up meeting b. We say that a curve is universal if there always exists a strategy to catch the ball from every initial configuration of the attractor and the ball. Recent work of Abrahamsen et al. has shown that every simple curve is universal. The authors also conjectured that all curves with rotation number one are universal. In this note, we disprove their conjecture and present a curve with rotation number one that is not universal.

1 Introduction

Beacon routing. *Beacon routing* [1] is the study of how to indirectly route objects through various domains. Let $\gamma : S^1 \to \mathbb{R}^2$ be a smooth closed curve in the plane, and let a and b be two points on γ. In curve-restricted beacon routing [2], b always moves closer to a, if possible, while staying on the curve. More precisely, b moves as long as its distance to a decreases, until a point where the distance is locally minimal. As soon as the position of a allows b to get closer to a, b moves again. We may think of a as the *attractor* and b as the *ball*. We can directly control the position a, but we have no direct control over b. We assume that b moves infinitely faster than a. Our goal is to move a along the curve in such a way that a meets b (see Figure 1).

A bit of care is required to properly define this problem and make sure γ is generic enough for the problem to be well-defined; we discuss the formal requirements in Section 2.

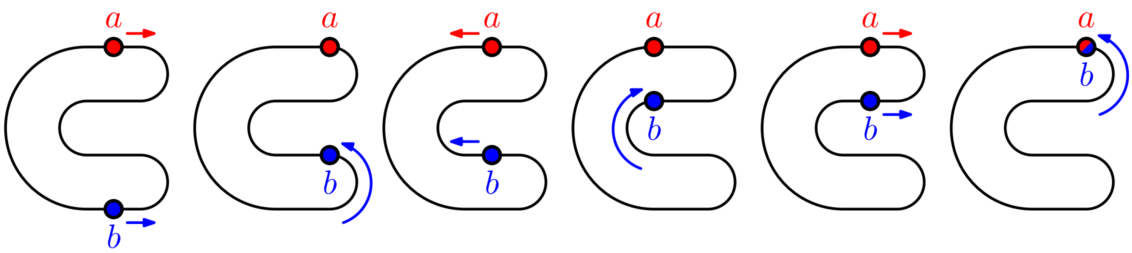

Fig. 1. An example of a strategy for the attractor to catch the ball on a c-shaped closed simple curve, by walking back and forth on a short section of the curve. (Figure taken from [2].)

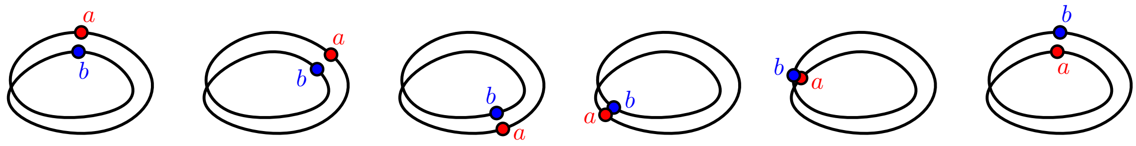

FIG. 2. An example of a non-universal curve: from a starting position in which a and b are close to each other, but on different loops, we can never reach a configuration where they meet.

Universal curves. A given curve γ is *universal* if there always exists a strategy to catch the ball from every initial configuration [2]. Biro et al. [1] ask whether every *simple* (i.e., not self-intersecting) curve is universal; Abrahamsen et al. [2] show that the answer is Yes. The condition of the curves being simple is necessary, and it is easy to find non-universal curves that may have self-intersections. The simplest example is the "doubled circle" (see Figure 2), where a and b start on different loops. Note that for curves with self-intersections, we consider the position of a and b at an intersection point to be different if their location (parameter value) at the curve is different, and we do not consider a and b to be meeting each other even if they are both at the same location in the plane, but at different locations on the curve.

In fact, the example of Figure 2 can be generalized: any curve that is "doubled" or "multiplied" is an example of a non-universal curve, since there are initial positions that cause a and b to always stay close to each other on different "copies" of the curve, without ever meeting each other. Refer to Figure 4 for some examples; for proper definitions and proofs see [2].

Rotation number. The *rotation number* (also called the *index* or *Whitney index*) [3, 4] of a closed curve is defined as the number of revolutions a tangent vector completes as it traverses the curve once. Inspired by the fact that with "multiplied curves", as in Figure 4, we can create examples of non-universal curves with any rotation number $k > 1$, as well as $k = 0$, but *not* for $k = 1$, Abrahamsen et al. [2] conjectured that, perhaps, every curve with rotation number one is universal.

Contribution. In this work, we disprove their conjecture by presenting a curve with rotation number one that is not universal. More precisely, we describe a smooth curve with rotation number one (Figure 5), as well as a polygonal analogue, and we show that these curves are not universal by specifying an initial configuration and arguing that there is no strategy to move a along the curve so that a and b eventually meet.

1.1 Related Work

The problem of classifying universal curves, graphs, or other domains, has recently received significant attention in computational geometry. This activity was initiated and inspired by *beacon-based geometric routing*, a framework that generalizes both greedy geometric routing and the art gallery problem. Versions of the framework were presented in the early 2010s [5, 6], and further developed in Michael Biro's PhD thesis [7] and subsequent papers [8, 9]. In early work in this framework, there is not a single moving attractor but rather a set of stationary *beacons* which can be activated and deactivated to create a "magnetic pull" on the moving point b towards itself, and the task is to route b through a polygonal domain by activating different beacons one at a time.

The first work that considers a moving attractor is due to Kouhestani and Rappaport [10], who study a scenario in which a moving point b is restricted to the interior of a (simple, polygonal) domain P, and a moving attractor a is restricted to the *boundary* of P. Kouhestani and Rappaport described a polynomial-time algorithm that finds a strategy for a to meet b, if such a strategy exists, given a simple polygon as input; they also conjectured that a capturing strategy always exists; that is, that simple polygons are always universal in this model. This conjecture was subsequently proven to be false by Abel et al. [11], by constructing non-universal polygons and their starting configurations. Their counter-examples even work in the less restrictive model where a can also move through the exterior of P. Their simplest counter-example is an orthogonal polygon with about 50 vertices.

The setting where a and b move in the same domain was first considered by Michael Biro in the open problem session of the 25th Canadian Conference on Computational Geometry (CCCG 2013), when he asked the question of whether every simple curve is universal. This question was answered in the affirmative by Abrahamsen et al. [2]. The setting where the domain is not a single curve, but a *embedded graph*, was recently considered by Ockenfels, Okamoto, and Schnider [12]. They show that every orthogonal planar embedding is universal. They also conjecture that *every* planar embedded graph is universal. Their proof heavily uses the orthogonality of the graph segments though, and is unlikely to be extensible to this more general setting. A different, very natural, scenario is where the domain for both a and b is a simply-connected region in the plane. Whether every such domain is universal remains an open question to this day.

2 Preliminaries

We noted earlier that while the setting should be very intuitive, a bit more care is required to carefully ensure that the problem is well-defined. We think of the attractor a and the ball b as each being defined by points on the circle S^1, the domain of the curve, or *track*, $\gamma : S^1 \hookrightarrow \mathbb{R}^2$. They thus correspond to two points in the Euclidean plane, $\gamma(a)$ and $\gamma(b)$. Much like we did in our introduction, we will sometimes abuse notation and confound a for $\gamma(a)$ and b for $\gamma(b)$. For any configuration (a, b) of an attractor and a ball, we write $D(a, b)$ to denote the Euclidean distance between the points $\gamma(a)$ and $\gamma(b)$. The rule of beacon-routing is that the ball moves in order to decrease its distance to the attractor, in a greedy fashion:

- If $D(a, b+\epsilon) < D(a, b)$ for all sufficiently small $\epsilon > 0$: the ball moves forward along the track.
- If $D(a, b-\epsilon) < D(a, b)$ for all sufficiently small $\epsilon > 0$: the ball moves backwards along the track.
- If both conditions hold, the ball moves in an arbitrary direction.
- If neither condition holds, the ball does not move and we say that the configuration is *stable*.

While the ball is moving, the attractor remains stationary. In other words, in our model, we assume that the ball moves infinitely faster than the attractor. If the configuration is stable, the attractor can be moved in any chosen direction along the curve, and the ball must then again conform to these four exhaustive rules.

For the conjecture to be meaningful, we shall restrict the curves under consideration to regular, piecewise C^1 curves, for which the rotation number is well-defined, as follows. For simplicity, we will further assume that the curves are piecewise C^2.

We say that a curve is regular if its derivative γ' is non vanishing everywhere. We say that it is piecewise C^2 if its signed curvature κ is continuous everywhere except at finitely many points $s_1, s_2, \ldots, s_n \in S^1$, at which we assume that both one-sided derivative vectors $\tau_\pm(s_i) = \lim_{s \to s_i^\pm} \frac{\gamma'(s)}{\|\gamma'(s)\|}$ exist (see Figure 3). Intuitively, the rotation number of γ measures the number of revolutions the tangent vector completes as we traverse the curve once. Formally, the rotation number of γ is then defined as the quantity:

$$\mathrm{rot}(\gamma) := \frac{1}{2\pi} \left(\sum_{i=1}^{n} \int_{s_i}^{s_{i+1}} \kappa(s) ds + \theta_i \right)$$

where θ_i is the clockwise oriented angle between $\tau_-(s_i)$ and $\tau_+(s_i)$ (see Figure 3) and s_{n+1} is defined as s_1. It is a classical theorem that this quantity is always an integer (see [3, 4]).

Remark 1: Due to their proof techniques, Abrahamsen et al ([2]) required the curve γ to satisfy stronger genericity conditions. Namely, that the curve be $\mathcal{C}^3$-continuous and intersects its evolute

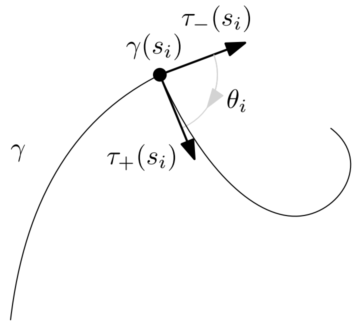

FIG. 3. A regular, piecewise C^2 curve has a well-defined rotation number.

transversely. It will be clear from our construction that our counter-example satisfies these additional conditions too. Their stronger genericity conditions also allow for a nice classification of the types of pairs possible: the so called forward, backward and critical configurations, where critical configurations are further classified into stable, unstable and (forward and backward) pivot configurations. This classification and the pivot configurations play an essential role in their proof (see [2]) but are not required to present our results.

3 Counter-example: smooth version

In this section, we describe a first intuitive smooth counter-example and give a sketch of a proof. In Section 4 we present a polygonal analogue of the same construction, and provide a more formal proof of our main result: a curve of rotation number one that is not universal; that is, a counter-example to the conjecture by Abrahamsen et al. [2]. We present the continuous version first, as we believe that it is easier to follow and contains the core ideas. The next theorem summarizes the main result of this paper.

Theorem 2: There exists a curve γ of rotation number one and an initial position for the pair (a, b) such that, no matter how the attractor a moves on γ, the ball b will never reach a.

3.1 Construction

We first describe how to construct the smooth version of our curve, globally illustrated in Figure 5.

Our curve is the concatenation of seven simpler curves, that we call *paths*. We begin by defining the family of paths $c_1, c_2, \ldots, c_7 : [0, 4] \to \mathbb{R}^2$ such that:

1. For all $i \in [7]$ and for choices of constants $C, \delta, \epsilon > 0$, we have:
   - $c_i(0) = (-1, -i\epsilon)$ and $c_i{}'(0) = (1, 0)$.
   - $c_i(1) = (1, i\epsilon + \delta)$ and $c_i{}'(1) = (1, 0)$, choosing δ so that $c_3(0) = c_3(1)$, i.e. $\delta = \epsilon(i + 1)$.
   - $c_i(2) = (1, i\epsilon - C)$ and $c_i{}'(2) = (-1, 0)$.
   - $c_i(3) = (-1, i\epsilon - C)$ and $c_i{}'(3) = (-1, 0)$.
   - $c_i(4) = c_{i+1}(0)$ and $c_i{}'(4) = (1, 0)$, with indices wrapping around so that $c_8 \coloneqq c_1$.
2. For all $t \in (0, 1) \cup (2, 3)$ and $i \in [7]$, $c_i(t)$ is contained in the strip $\{(x, y) | -1 < x < 1\}$.
3. For all $i \in [7] - \{2\}$, the tangent vector of each path c_i does not accomplish a single revolution, as t ranges from 0 to 1. Note, for instance, that the tangent vector never assumes the value $(-1, 0)$. The tangent vector of c_2 is the exception, and makes a single positive revolution.

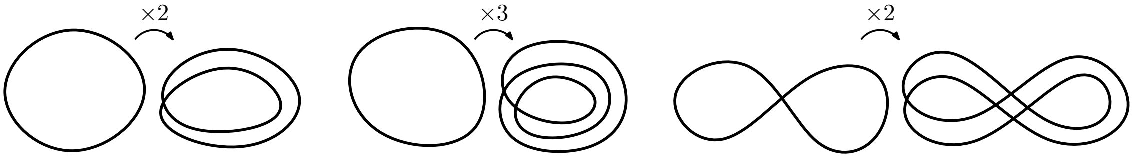

FIG. 4. Any curve can by "multiplied" by taking multiple copies of it at (almost) the same location, and introducing a switch-over point, to obtain a non-universal curve.

4. For all $i \in [7]$, c_i traces a counter-clockwise loop between $c_i(1)$ and $c_i(2)$, so that the tangent vector of each path c_i continuously changes from $(1, 0)$ to $(-1, 0)$ while avoiding $(0, -1)$ and thus "backtracks" instead of completing a full revolution, so that the tangent vector of each path c_i accomplishes a single negative revolution as t ranges from 1 to 2.
5. For all $i \in [7]$, c_i traces a horizontal line segment between $c_i(2)$ and $c_i(3)$.
6. For all $i \in [7]$, c_i smoothly interpolates between $c_i(3)$ and $c_i(4)$ without adding to the rotation number.
7. For each fixed $i \in [7]$, the geometric image of the intervals $(0, 1)$, $(1, 2)$, $(2, 3)$ and $(3, 4)$ under c_i are all pairwise disjoint.

By construction, the union of such a collection of 7 paths, $\cup_i c_i$, is a closed curve γ with rotation number one.

3.2 Non-universality

The crucial desired property that we would like to ensure with such a construction is the following:

Lemma 3: The paths $c_1, c_2, \ldots, c_7$ can be chosen such that, for all $i \in [7]$, if we position a and b either respectively at starting positions $c_i(0)$ and $c_{i+1}(0)$ or starting positions $c_i(1)$ and $c_{i+1}(1)$, then no matter how we move a along c_i:

- If we reach a point where a is at $c_i(1)$, then b is at $c_{i+1}(1)$.
- If we reach a point where a is at $c_i(0)$, then b is at $c_{i+1}(0)$.

Proving this lemma is an easier task in the polygonal case (see Section 4), nevertheless we believe the most important intuition is more easily communicated with the smooth construction, which is why we provide an extensive sketch of the proof.

Remark 4: Note that the lemma is clear for all parameters $t \in [1, 4]$ so we may consider only the restriction of curves c_i to the interval $(0, 1)$.

Sketch of Proof of Lemma 3. To motivate the lemma and our construction, recall that, as discussed earlier, the double circle curve is a simple example of a non-universal curve. One way to look at our construction and to think of Lemma 3 is to take the necessary steps to alter the double circle into a curve of rotation one, while keeping the desired non-universality. To that effect, let us picture the double circle as being obtained from two copies of a circle at almost the same location, one blue (call it c_2) and the other turquoise (call it c_3), cutting both at an arbitrary point, and switching their endpoints before reconnecting. Suppose that the outer loop ends up being the blue loop c_2. Since the double circle has rotation number two, we add a counter clockwise loop to subtract one to the rotation number, and get again a curve with rotation number one (see Figure 6).

Taking a closer look at this curve, we see that, unfortunately, this construction only works in one direction. If the attractor a is on the blue curve c_2, then b follows on the turquoise curve c_3. However if a leads on the turquoise curve, then b gets "stuck" in the blue loop we introduced, after which a need only continue its path until it meets with b (see Figure 7). To prevent this situation,

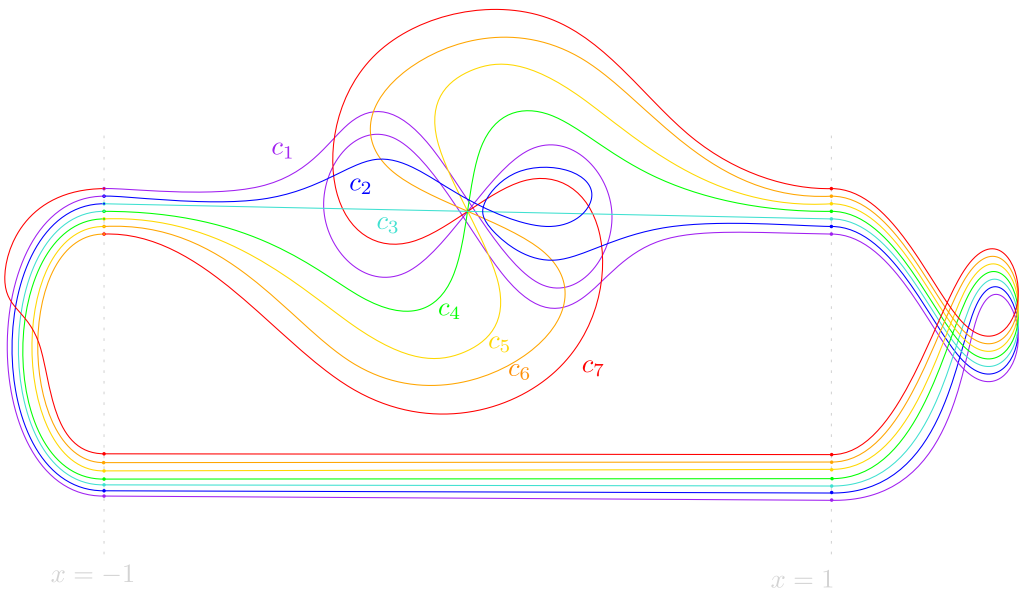

FIG. 5. Our construction γ has rotation number one and is not universal: there exist configurations from which the ball can never be caught. For example, start with a at one of the lowest points of the curve (on the purple strand) and b just above it (on the blue strand).

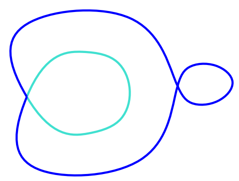

FIG. 6. The doubled circle can be thought of as a prototype example with two paths (blue and turquoise). We modify the doubled circle with a counter clockwise loop (on the right) to achieve rotation number one. However these two paths fail to verify Lemma 3.

we introduce a sequence of five "intermediate" curves that will ensure that, no matter how a moves, it will not be able to meet b.

In light of the previous explanation and the remark, we start by defining c_2 (the blue curve on Figure 7) to trace a single counter-clockwise loop between $c_2(0)$ and $c_2(1)$, while c_3 traces a line segment between $c_3(0)$ and $c_3(1)$ (the turquoise curve on Figure 7).

Remark 5: Note that in our analysis, for the sake of simplicity and economy of diagram making, we shall assume that a moves from the left to the right in the forward direction. For all the pairs of curves we introduce, a direct symmetry argument solves the backwards direction. The only exception is the pair c_1 and c_2, for which we detail the backwards case explicitly in Figure 8.

The first intermediate curve is the curve c_1 (in purple on Figure 8), which is essentially a multiplied copy of a curve "shaped like an 8" and thus has rotation number 0 (see the rightmost

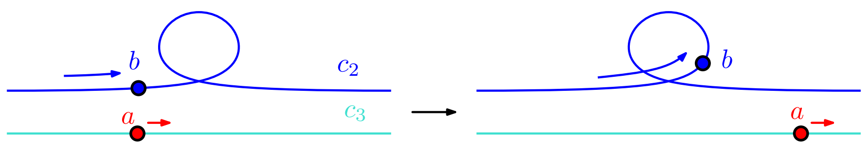

FIG. 7. If the attractor a leads on the turquoise curve c_3, the ball b gets trapped in the loop on the blue curve c_2.

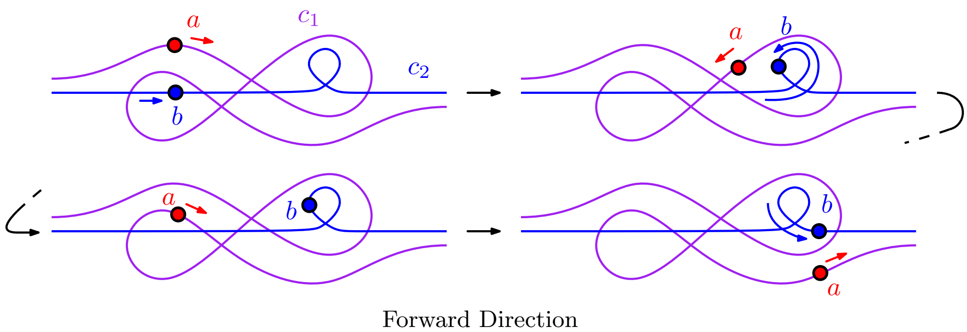

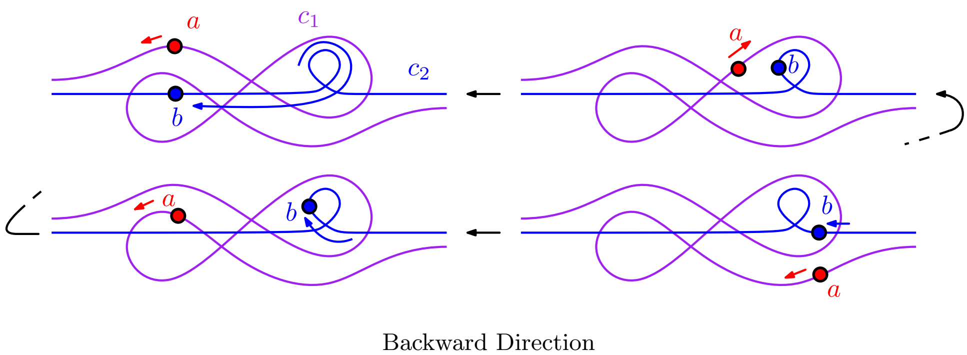

FIG. 8. Adding an intermediate purple curve c_1 to act as c_3 for c_2 fixes our previous problem, where if a leads on the purple c_1, b now follows on the blue c_2 instead of getting stuck in its loop. However, we simply postponed the problem to the last pair of curves c_3 and c_1.

diagram of Figure 4). In the forward case, as seen on Figure 8 (top), this curve has the property that if a leads on c_1, then b follows on c_2 and escapes the added loop. We thus now have curves c_1, c_2 and c_3 with the desired property of Lemma 3 with the exception of the last pair c_3 and c_1. Namely, if a leads on c_3 (turquoise), b does not follow on c_1 (purple). The backwards case is illustrated on Figure 8 (bottom).

To fix this last problem, we introduce the curve c_7 (pictured in red), which is designed as a spiral to ensure that if a leads on c_7, b follows along on c_1 (purple) and is guided to escape the double eight (see Figure 9). The last issue to fix now is that if a leads on c_3 (turquoise), b will not quite follow on this new red curve c_7. However note that c_3 and c_7 have the same topology and are both a simple line with no loops. So we need only introduce a fine enough sequence of intermediate curves that interpolate between the flat line curve c_3 (turquoise) and the spiraling c_7 (red). As can be seen on the original Figure 5 depicting our counter-example, three intermediate curves c_4, c_5, c_6 that each progressively spiral with less than a right-angled bend from each other are enough to accomplish this task, leading to the claim in the lemma. Note that we implictly also rely on the following remark, which should be clear from the description of our process.

Remark 6: At any point and for all the curve pairs, moving a back and forth instead of simply moving a monotonously along the curve does not introduce any new dynamics, and we can safely assume without loss of generality that a moves monotonously along the curve it is on.

∎

4 Counter-example: polygonal version

The original proof of Abrahamsen et al. treats the polygonal case separately because of its additional intricacies; in contrast, we note here that our counter-example can be turned into a polygonal curve in a straightforward way. The advantage of this discretization is that it allows

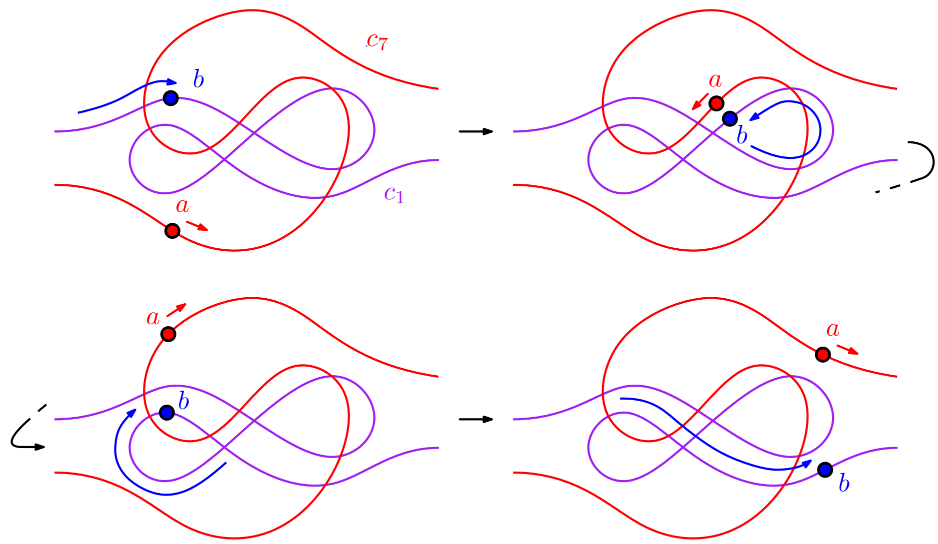

FIG. 9. Introducing the red curve c_7 to act as c_3 for c_1 solves the previous pathology: if a leads on c_7, b now follows on c_1. We have thus moved the problem to finding intermediate curves from the much simpler c_7 back to the flat turquoise curve c_3.

us to dispense with the more technical machinery that would be required to provide a rigorous proof of our smooth counter-example. Instead, the polygonal setting gives us a clean and simple algorithm to rigorously check the validity of our counter-example. This algorithmic check boils down to a single local lemma we discuss below (see Lemma 7). Note that, conceptually, it should be the case that we need only approximate the continuous curve with sufficiently many points to obtain a polygonal version of our previous counter-example. In practice, the polygonal setting allows us to distill the continuous one to give curves with very few vertices that behave similarly, albeit perhaps at the detriment of aesthetics considerations and readability.

In the polygonal setting, we consider the particular case where the curve γ is piecewise linear and the image of γ is a (self-crossing) oriented polygon P with finitely many oriented edges $e_1, e_2, \ldots, e_n$. Whenever we increment indices in what follows the resulting indices are considered to wrap around (i.e., $e_{n+1} = e_1$). The endpoints of the edge e_i will be denoted by v_i and v_{i+1}, following the orientation of P. For each edge e_i, let us partition the plane into the following three sets: $S_i^- := \{x \in \mathbb{R}^2 | \langle x - v_i, e_i \rangle \leq 0\}$ and $S_i^+ := \{x \in \mathbb{R}^2 | \langle x - v_{i+1}, e_i \rangle \geq 0\}$, and $S_i = (S_i^- \cup S_i^+)^c$ (see Fig. 10). To rigorously check our polygonal counter-example, we can use the following local lemma.

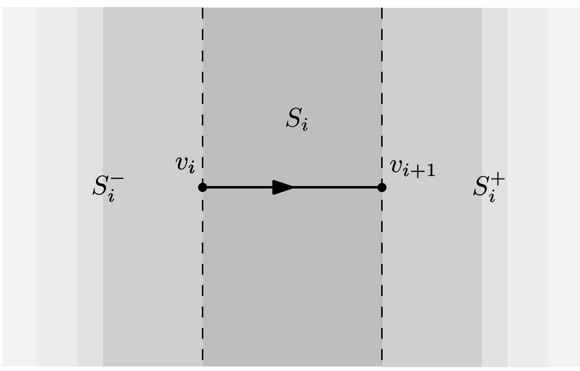

FIG. 10. Each edge e_i defines three regions S_i^-, S_i and S_i^+.

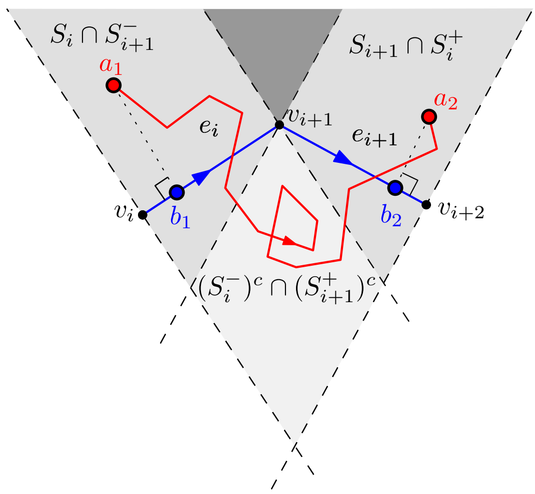

FIG. 11. Lemma 7 allows us to perform a local check on the successive pairs of curves that need to be accounted for.

Lemma 7: Suppose we are given a stable configuration (a_1, b_1) such that $\gamma(a_1)$ lies in $S_i \cap S_{i+1}^-$ and $\gamma(b_1)$ lies on the open edge e_i. Suppose further that the attractor describes a path along γ from the start position $\gamma(a_1)$ to the end position $\gamma(a_2)$ such that:

1. $\gamma(a_2) \in S_{i+1} \cap S_i^+$.
2. For all $t \in [a_1, a_2]$, $\gamma(t) \in (S_i^-)^c \cap (S_{i+1}^+)^c$ (see Fig. 11).
3. The path traced by the attractor crosses the boundary of S_i exactly once, and the boundary of S_{i+1} exactly once.

Then the ball describes a trajectory from the start position $\gamma(b_1)$ to the end position $\gamma(b_2)$ such that:

1. For all $t \in [b_1, b_2]$, $\gamma(t) \in e_i \cup e_{i+1}$ and $\gamma(b_2) \in e_{i+1}$.
2. (a_2, b_2) is a stable pair.

Proof.Observe that since the pair (a_1, b_1) is stable and the attractor starts in the region S_i, it must be that b_1 is the orthogonal projection of a on the open edge e_i. The triangle inequality then dictates that as long as a stays inside S_i, b follows on the edge e_i as the orthogonal projection of a. Thus as a moves to the common boundary of S_i and S_i^+, b approaches the vertex v_{i+1}. When a lies exactly on the boundary line, b coincides with v_{i+1} for a brief moment; then, (i) if a crossed this boundary inside S_{i+1}, the triangular inequality applied to the edge e_i dictates that b must move to the orthogonal projection of a onto the edge e_{i+1} to minimize the distance; (ii) if a crossed this boundary line in $S_i^+ \cap S_{i+1}^-$ (the darker shaded area on Fig. 11), the closest point from a to the union of the two edges is v_{i+1}, so the ball cannot decrease its distance locally, and remains there until a crosses the boundary line between S_{i+1}^- and S_{i+1}, where the triangle inequality applied to the edge e_{i+1} now dictates that b must move as the orthogonal projection of a onto that edge, forming a new stable pair. ■

Remark 8: Note that the previous lemma is completely symmetrical if we swap the orientation of the polygon P, i.e., if we simultaneously swap (a_1, b_1) with (a_2, b_2) together with the indices i and $i+1$ and the $+$ and $-$ superscripts. This shows that this move is fully reversible and the behavior of the attractor along the path from $\gamma(a_1)$ to $\gamma(a_2)$ does not matter.

Remark 9: We can approximate any $\mathcal{C}^1$ curve with inscribed polygons of increasingly small edge-lengths. In the limit, the gray region of the previous lemma becomes a ray starting from the center of curvature at $\gamma(b)$ in the direction of $\gamma(b)$. This would give a straightforward discretization of the smooth counter-example save for the fact that, given a pair of successive curves, we are not guaranteed that the successor curve stays between the curve and its evolute. In fact, this is a behaviour exhibited by curves c_4 through c_7, where for example c_6 intersects the evolute of curve c_7 (see Fig. 12). As we will see, we are able to eliminate this behaviour and spare one curve in the process of doing so, although arguably at the cost of immediate readability.

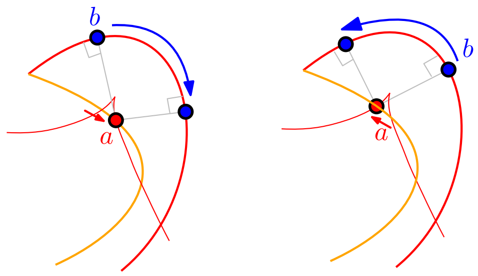

FIG. 12The curve c_6 (orange) intersects the evolute (thin red) of the curve c_7 (bold red).

Our polygonal curve γ is modeled after our previous smooth counter-example and can be seen as a “minimalist” polygonal approximation of it, see Figure 13 for a global view and Figure 15 for a more readable zoomed-in view on the successive pairs of curves. Analogously to our smooth counter-example, we name the different portions of our curve and its six sections $c_1, c_2, \dots, c_6$ (we were able to spare one curve). We denote the edges of γ_i by $e^i_1, e^i_2, \dots, e^i_{n_i}$ and its vertices by $v^i_1, \dots . v^i_{n_i}$, with $v^i_{n_i} = v^{i+1}_0$. The slabs previously defined are notated accordingly, i.e., S^i_j for the slab based on the edge e^i_j.

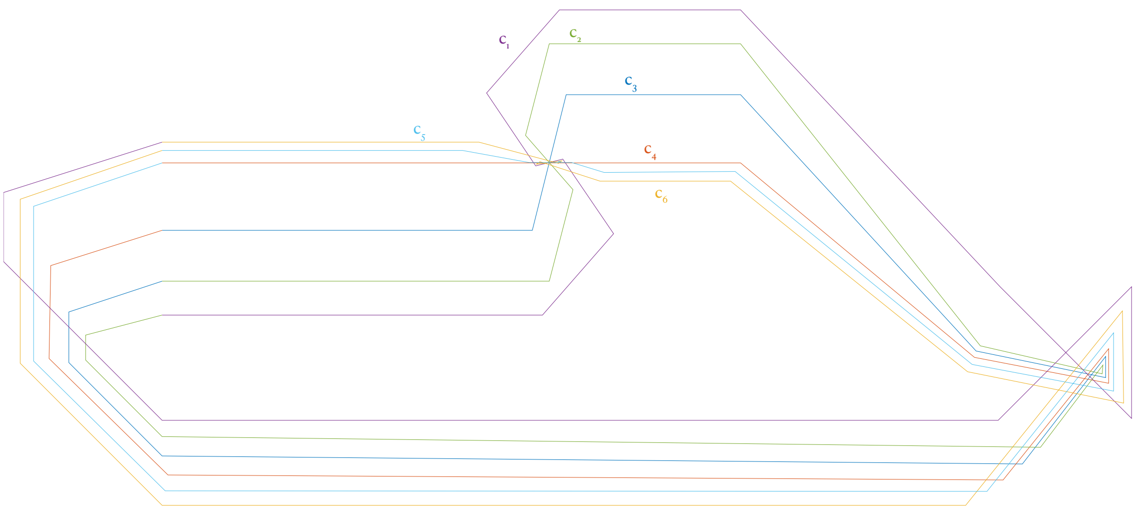

FIG. 13. The polygonal counter-example, a close-up of the central region is given on Figure 15.

Building on top of Lemma 7, the algorithm to verify our polygonal counter-example is straightforward: for each successive pairs of curves c_i, c_{i+1}, we position the attractor on v^i_0 and the ball on v^{i+1}_0 and sequentially verify whether the conditions of Lemma 7 are met, for two successive edges of c_i at a time, moving the attractor and the ball accordingly. Pseudocode for this procedure **check** $\boldsymbol{(c_i, c_{i+1})}$ is given in the appendix.

The following is then a direct consequence of applying Lemma 7 in sequence along the successive pairs of curves.

Corollary 10: For any index $i \in [6]$, the algorithm **check** (c_i, c_{i+1}) terminates and returns TRUE if Lemma 3 holds for the polygonal curves c_i and c_{i+1}, and FALSE otherwise.

The proof of Theorem 2 then follows from the following claim.

Claim 1: For every $i \in [6]$, we can always add additional vertices to the segments of c_{i+1} to obtain a subdivision c'_{i+1} of the curve c_{i+1} such that **check** (c_i, c'_{i+1}) terminates and does not return False.

Proof of Claim 1. The proof of this claim is given in a series of diagrams provided on Figure 15. The regions S^i_k are highlighted for each successive pair of curves (c_i, c_{i+1}). It is then a matter of following step by step the algorithm **check** (c_i, c'_{i+1}). We take as an example the pair (c_1, c_2), seen on the top-left part of Figure 15. The attractor leads on the purple c_1, starting at the bottom left of the curve, while the ball follows on the green c_2. The first green vertex lies as expected by Lemma 7 in the region S^1_1. Following the edge towards the second vertex, we leave S^1_1 and enter S^1_2 while staying in the desired region $((S^1_1)^-)^c \cap ((S^1_2)^+)^c$. The boundaries of both slabs is traversed exactly once, as needed. We then follow along each successive vertex along each successive pairs of slabs, diligently checking that the conditions for Lemma 7 are met, until we exit the central portion of the curve to the right of the diagram.

In the interest of space and following Remark 4, we omitted the portions of the curve outside of the $x \in [-1, 1]$ strip, as these are straightforward and less complicated. We mention however the rightmost loop, which is non trivial in the polygonal case and requires nesting each curve, as exemplified on Figure 14.

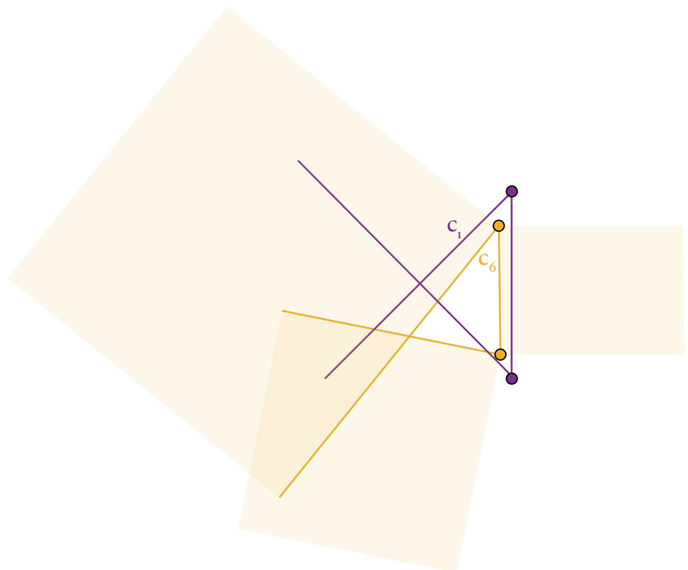

FIG. 14: A close-up of the right-most loop for a selected pair of successive curves.

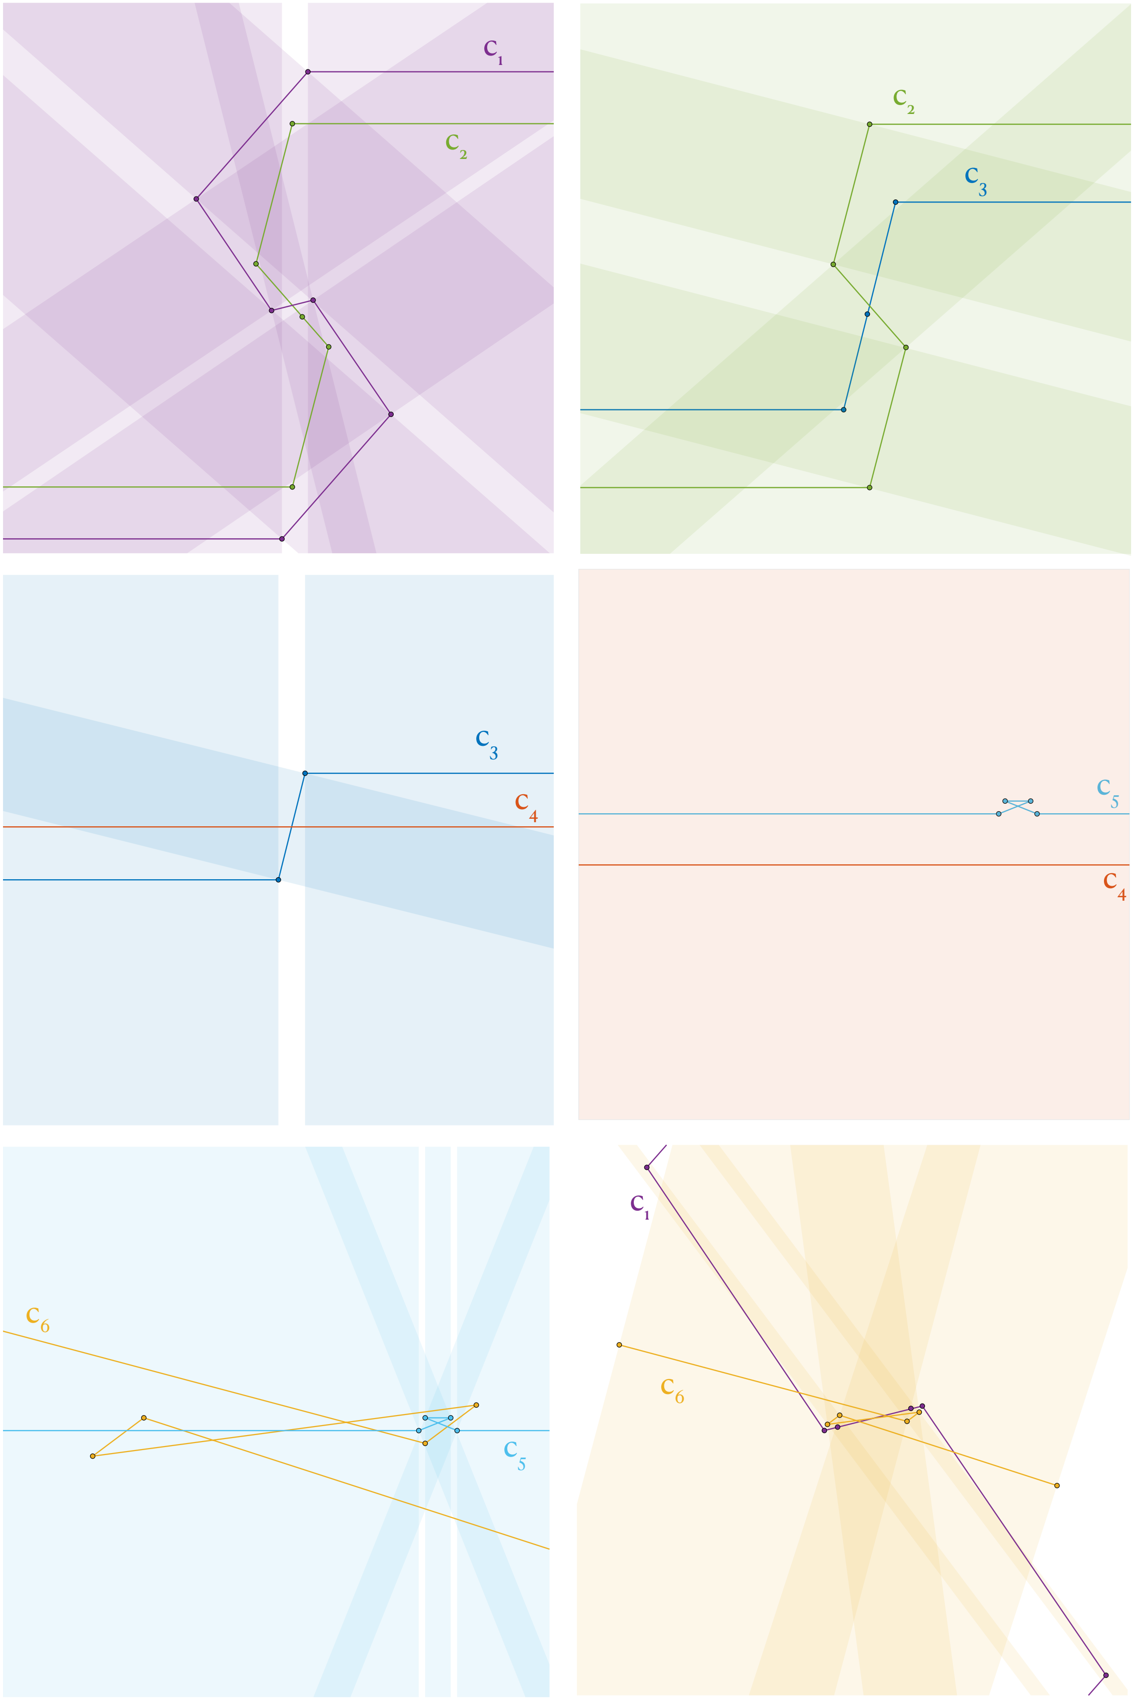

FIG. 15: The central region of interest for each successive pair of curves in our polygonal counter-example.

■

5 Discussion

We have presented two curves that have rotation number one, but are not universal, answering in the negative the question posed by Abrahamsen et al. [2]. This contrasts with their universality result for simple closed curves (which have rotation number one). We thus showed that one cannot extend their topological method to every other curve of rotation number 1, as they conjectured. It then remains an open question whether universality of curves can be classified or related to other properties, such as the number of self-intersections.

In particular, our curves are arguably not very close to being simple, since they are relatively complicated. In the smooth version, our curve has 108 self-intersections. The polygonal counter-example has 94 self-intersections. Note that we did not try to optimize neither number of vertices nor number of self-intersections. However, we observed that there seems to be a non-trivial relation between the number of self-intersections, the number of vertices of the curve, and the algebraic degree of the pieces of any counter-example. For example, taking a closer look at the right-most part of the polygonal counter-example shows that using fewer vertices leads to acute angles which force the nested structure seen on Figure 13 and many additional self-intersections. Using more vertices would allow obtuse angles and thus less self-intersections. Understanding this relationship better and finding the "simplest" counter-example would be of interest.

Beyond closed curves, Ockenfels, Okamoto and Schnider initiated the study of the same problem when the underlying domain is a *graph* [12]. They show that every orthogonal straight-line plane graph is universal, and they conjecture the same is true for any straight-line plane graph. The same question for graphs with curves edges may also be asked.

Even more generally, it would be interesting to understand beacon routing on general embedded graphs, that are not necessarily planar. Extensions of this problem where both the attractor and ball are instead constrained to a two-dimensional domain also remain open and unexplored.

Acknowledgements. This work was initiated at the Dagstuhl Seminar 24072: Triangulations in Geometry and Topology. We thank the organizers and the participants for the fruitful atmosphere. A preliminary version of this work was presented at the European Workshop on Computational Geometry (EuroCG) 2025. R.I. Silveira was partially supported by grant PID2023-150725NB-I00 funded by MICIU/AEI/10.13039/501100011033.

References

[1] 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, Waterloo, Ontario, Canada, August 8-10, 2013.* Carleton University, Ottawa, Canada, 2013. URL: `http://cccg.ca/proceedings/2013/papers/paper%5C_74.pdf`.

[2] Mikkel Abrahamsen, Jeff Erickson, Irina Kostitsyna, Maarten Löffler, Tillmann Miltzow, Jérôme Urhausen, Jordi Vermeulen, and Giovanni Viglietta. "Chasing Puppies: Mobile Beacon Routing on Closed Curves". In: *Journal of Computational Geometry* 13.2 (2022), pp. 115–150.

[3] Hassler Whitney. "On regular closed curves in the plane". In: *Compositio Mathematica* 4 (1937), pp. 276–284.

[4] Damián Wesenberg. *A new formula for rotation number.* 2020. arXiv: `2010.01422 [math.GT]`. URL: `https://arxiv.org/abs/2010.01422`.

[5] 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.

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

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

[8] 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)*. 2013, pp. 158–169. DOI: 10.1007/978-3-642-40104-6_14.

[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] Bahram Kouhestani and David Rappaport. "Edge patrolling beacon". In: *Japanese Conference on Discrete and Computational Geometry, Graphs, and Games (JCDCGGG 2017)*. 2017, pp. 101–102.

[11] 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". In: *Japanese Conference on Discrete and Computational Geometry, Graphs, and Games (JCDCGGG 2018)*. 2018, pp. 28–35. DOI: 10.1007/978-3-030-90048-9_3.

[12] Johanna Ockenfels, Yoshio Okamoto, and Patrick Schnider. "Chasing puppies on orthogonal straight-line plane graphs". In: *Proc. 41st European Workshop on Computational Geometry, EuroCG 2025 (to appear)*.

6 Appendix

check (c_i, c_{i+1})

Position a on v_0^i and b on v_0^{i+1}. // By construction of γ, this is a stable pair
$l \leftarrow 0$
bool1 $\leftarrow 0$ // Keep track of whether a point crossed the boundary of S_k^i
bool2 $\leftarrow 0$ // Keep track of whether a point crossed the boundary of S_{k+1}^i
for $k = 0, \ldots, n_i$ **do**
    if $v_l^{i+1} \in S_k^i \cap (S_{k+1}^i)^-$ **then**// Is the point correctly positioned in the starting region?
        $l \leftarrow l + 1$
        bool1 $\leftarrow 1$
    else
        Return False
    end if
    while $v_l^{i+1} \notin S_{k+1}^i \cap (S_k^i)^+$ **do** // Has the point reached the correct end region?
        if $v_l^{i+1} \in ((S_k^i)^-)^c \cap ((S_{k+1}^i)^+)^c$ **then** // Does point stay in the correct V-shaped region?
            if $v_l^{i+1} \in S_k^i$ **then** // Has the curve crossed the boundary of S_k^i more than once?
                if bool1 $= 0$ **then**
                    Return False
                end if
            else
                bool1 $\leftarrow 0$
            end if
            if $v_l^{i+1} \in S_{k+1}^i$ **then**
                bool2 $\leftarrow 1$
            else
                if bool2 $= 1$ **then** // Has curve crossed the boundary of S_{k+1}^i more than once?
                    Return False
                end if
            end if
            $l \leftarrow l + 1$
        else
            Return False
        end if
    end while
end for
Return True