

Computing Optimal Homotopies over a Spiked Plane with Polygonal Boundary*

Benjamin Burton¹, Erin Chambers², Marc van Kreveld³,
Wouter Meulemans⁴, Tim Ophelders⁵, and Bettina Speckmann⁶

- 1 School of Mathematics and Physics, University of Queensland, Brisbane, Australia
bab@maths.uq.edu.au
- 2 Dept. of Computer Science, Saint Louis University, Saint Louis, MO, USA
echambe5@slu.edu
- 3 Dept. of Information and Computing Sciences, Utrecht University, The Netherlands
m.j.vankreveld@uu.nl
- 4 Dept. of Mathematics and Computer Science, TU Eindhoven, The Netherlands
w.meulemans@tue.nl
- 5 Dept. of Mathematics and Computer Science, TU Eindhoven, The Netherlands
t.a.e.ophelders@tue.nl
- 6 Dept. of Mathematics and Computer Science, TU Eindhoven, The Netherlands
b.speckmann@tue.nl

Abstract

Computing optimal deformations between two curves is a fundamental question with various applications, and has recently received much attention in both computational topology and in mathematics in the form of homotopies of disks and annular regions. In this paper, we examine this problem in a geometric setting, where we consider the boundary of a polygonal domain with *spikes*, point obstacles that can be crossed at an additive cost. We aim to continuously morph from one part of the boundary to another, necessarily passing over all spikes, such that the most expensive intermediate curve is minimized, where the cost of a curve is its geometric length plus the cost of any spikes it crosses.

We first investigate the general setting where each spike may have a different cost. For the number of inflection points in an intermediate curve, we present a lower bound that is linear in the number of spikes, even if the domain is convex and the two boundaries for which we seek a morph share an endpoint. We describe a 2-approximation algorithm for the general case, and an optimal algorithm for the case that the two boundaries for which we seek a morph share both endpoints, thereby representing the entire boundary of the domain.

We then consider the setting where all spikes have the same unit cost and we describe a polynomial-time exact algorithm. The algorithm combines structural properties of homotopies arising from the geometry with methodology for computing Fréchet distances.


1998 ACM Subject Classification I.3.5 Computational Geometry and Object Modeling

Keywords and phrases Fréchet distance, polygonal domain, homotopy, geodesic, obstacle

Digital Object Identifier 10.4230/LIPIcs.ESA.2017.23

* This research was initiated at Dagstuhl seminar 17072, “Applications of Topology to the Analysis of 1-Dimensional Objects”. E. Chambers is supported by NSF CCF-1614562 and CCF-1054779; W. Meulemans by NLeSC grant 027.015.G02; and T. Ophelders and B. Speckmann by NWO grant 639.023.208.





■ **Figure 1** Polygonal domain R bounded by polylines $f \cup g \cup \gamma_0 \cup \gamma_1$. K in red squares.

■ **Figure 2** Disk D is convex relative to R . γ is backwards- but not forwards-convex.

■ **Figure 3** The convex hull $\text{gh}_R(\gamma)$ of a curve $\gamma \in \Gamma$.

1 Introduction

Computing optimal deformations between two input curves is a fundamental building block in many application areas, such as robotics, motion planning, geographic information systems, and graphics. Such deformations or morphs have several salient properties, for example, the maximum necessary length of intermediate curves. The properties of an optimal morph between two curves also serve as a natural measure of similarity between these curves. Examples of such measures include the Fréchet distance and its variants.

In this paper we consider the following scenario: our input is a polygonal domain with point obstacles, similar to [6]. We aim to continuously morph from one part of the boundary to another. The intermediate curves (or leashes) are allowed to pass over the point obstacles, albeit for a fixed cost. We hence refer to the point obstacles as *spikes* to reinforce the intuition that encountering them is costly, but that they do not form impassable barriers. Our goal is to minimize the cost of the most expensive leash during the morph, where the cost of a leash is defined as its length plus the cost of the spikes it encounters. We consider both variable-cost spikes and unit-cost spikes, and describe several structural results as well as algorithms. In the following we first introduce some necessary definitions which allow us to state our results more precisely.

Definitions and problem statement. Let $f, g, \gamma_0, \gamma_1: [0, 1] \rightarrow \mathbb{R}^2$ be four interior-disjoint simple polylines (possibly of length 0) in the plane, with $f(0) = \gamma_0(0)$, $f(1) = \gamma_1(0)$, $g(0) = \gamma_0(1)$, and $g(1) = \gamma_1(1)$, whose union bounds a polygonal domain R of n vertices. Let $K \subset R$ be a finite set of k point-obstacles which we call *spikes* (see Figure 1). The *cost* of a spike is given by a function $w: K \rightarrow \mathbb{R}^{\geq 0}$. We assume that spikes lie in general position in the sense that no three spikes lie on the same geodesic in R . Standard perturbation techniques (as also described below) can lift this assumption.

Let Γ be the family of simple curves in R from f to g ; that is, all curves $\gamma: [0, 1] \rightarrow R$ with $\gamma(0) \in \text{Im}(f)$ and $\gamma(1) \in \text{Im}(g)$. A *homotopy* on R from γ_0 to γ_1 is a continuous map $h: [0, 1] \times [0, 1] \rightarrow R$ with $h(0, \cdot) = \gamma_0$, $h(1, \cdot) = \gamma_1$. All homotopies we consider have $h(t, \cdot) \in \Gamma$ for all $t \in [0, 1]$. With slight abuse of terminology, we refer to these simply as homotopies. We call a homotopy *monotone* if it is injective after infinitesimal perturbation. We refer to each curve $\gamma_t: p \mapsto h(t, p)$ as the *leash* of h at time t , and define the *cost* of a leash as its length plus the total cost of spikes in K (with multiplicity) it encounters. The *cost* of a homotopy is the cost of its maximum-cost leash. We are interested in the minimum-cost homotopy on R from γ_0 to γ_1 and refer to the cost of this homotopy as the *homotopy height* from γ_0 to γ_1 on (R, K, w) .

For two curves γ and γ' from a monotone homotopy, denote the region between them by $R(\gamma, \gamma')$; that is, $R(\gamma, \gamma')$ is the (possibly degenerate) topological disk bounded by γ, γ' , the arc of f between $\gamma(0)$ and $\gamma'(0)$, and the arc of g between $\gamma(1)$ and $\gamma'(1)$. For a simple curve $\gamma \in \Gamma$, define its *swept region* as $R(\gamma_0, \gamma)$ and symmetrically define its *unswept region* as $R(\gamma, \gamma_1)$. We call a region $D \subseteq R$ *convex relative to R* if for any two points in D , the shortest path in R connecting those points also lies in D [20]; for this definition we do not charge the additional cost of any spikes. A curve $\gamma \in \Gamma$ is *forwards-convex* if it lies on the boundary of a region D convex relative to R , and D is contained in the swept region of γ . Symmetrically, γ is *backwards-convex* if it lies on the boundary of a region D convex relative to R , and D is contained in the unswept region of γ (see Figure 2). A curve may be both forwards- and backwards-convex; a shortest path is always both.

Region R is convex relative to itself, and the intersection of any two convex sets relative to R is also convex relative to R . For $\gamma \in \Gamma$ define its convex hull (also known as geodesic hull) $\text{gh}_R(\gamma)$ relative to R as the unique minimal region containing γ which is convex relative to R (see Figure 3).

Consider a homotopy h that contains a leash γ_t that crosses several spikes. Infinitesimal perturbation of the leash at the spikes ensures that γ_t no longer crosses a spike, but is then forced into some homotopy class. In particular, this has two implications: (1) as the perturbation tends to zero, this has no effect on the length of the leash and thus, strictly speaking, the optimal homotopy is an infimum rather than a minimum if the given leash is the maximum-cost one in the optimal homotopy; (2) we can decompose a single leash crossing a number of spikes into a homotopy crossing each spike separately, essentially holding the leash fixed—hence an optimal homotopy exists that crosses spikes one at the time.

Results and organization. We consider various settings of a spiked plane with polygonal boundary. In Section 2, we investigate the general setting where each spike may have a different cost. First, we consider the number of inflection points that the leash may need in an optimal homotopy and present a lower bound that is linear in the number of spikes, even if R is convex and only f has positive length. We then present a 2-approximation algorithm for the general case and an optimal algorithm for computing the homotopy height for the case that f and g have length 0 (i.e., γ_1 and γ_2 together form the boundary of R).

In Section 3, we consider the setting where all spikes have the same unit cost. Here we present our main result: an algorithm to compute the exact homotopy height in polynomial time. The algorithm combines structural properties of homotopies arising from the geometry with methodology for computing Fréchet distances. To the best of our knowledge, these are the first polynomial-time algorithms to compute the exact homotopy height in any setting.

Related work. Chambers *et al.* [5] recently proved that there always is a minimum-cost homotopy between the boundaries of an annular surface that is an isotopy and monotone, that is, the intermediate leashes never move “backwards”. This proof readily transfers to our setting and is indeed supporting our results as described in Section 3.

Homotopy height was introduced independently in the computational geometry community [7] and in the combinatorics community [4]. On triangulated surfaces, the best known algorithm gives an $O(\log n)$ approximation, where n is the complexity of the surface [18]. More recently, it has also been studied in more general settings, where instead of having point obstacles, obstacles are modeled by assigning a non-Euclidean metric to R [10, 11, 12].

Fréchet distance is a well studied metric. Also known as the dog-leash distance, the goal is to minimize the length of the longest leash connecting a man walking along one curve and a dog along the other, where the man and dog walk monotonically along the curves.

The classic algorithm computes this metric in $O(n^2 \log n)$ time [1], with many variants and approximation algorithms having been studied since. The geodesic and homotopic Fréchet distance are particularly related to our setting. The former is a variant where the leashes must stay inside a polygonal boundary and remain geodesic, but no obstacles are present inside the polygon [13]. For the latter, point or polygonal obstacles are given which no leash may cross: this case can be solved in polynomial time [6].

Optimal morphs have also been studied in a variety of other settings. From the topology end, minimum-area homotopies can be computed for planar or surface embedded curves in polynomial time [9, 17, 19]. Minimum-area homologies are a closely related similarity measure on curves that can be computed very quickly using linear algebra packages [8, 14], but do not yield intuitive deformations or morphings in the same way as homotopy.

From the geometry and graph theory communities, much work has been done on computing morphs between inputs; indeed, it is well known that any two drawings of the same planar graph can be morphed to each other. Optimal morphs between such graphs are still being studied, including work that bounds the complexity of the morph [2, 3]. However, none of these morphs bound the length of any of the "leashes" tracing the paths of vertices during the morph. Morphs based on geodesic width [16] force all intermediate curves to not cross the input curves (which are part of the boundary polygon in our setting); however, none of these have been considered in the presence of obstacles. Dynamic time warping and related concepts [15] also consider ways to match and morph curves, but again do not extend to more general settings.

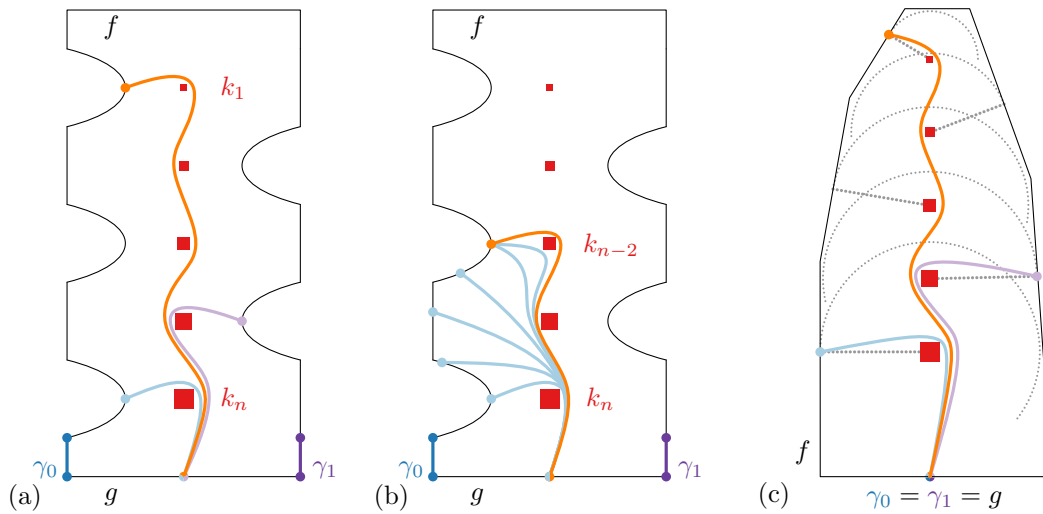
2 Variable-cost spikes

In this section, we consider the setting where each spike may have a different cost. As we prove, the variable costs have a profound effect on the leash complexity. Nonetheless, we obtain a general approximation algorithm as well as an optimal algorithm for a special case.


Leash complexity. We define the *complexity* of a leash as the number of inflection points that are not caused by the boundary of R . In particular, the leash complexity needed for an optimal homotopy is defined by the number of spikes that cause an inflection point. Unfortunately, in the general case, the complexity may be linear, even when R is convex. Correspondingly, we do not have a polynomial-time algorithm even when R is convex.

► **Lemma 1.** *In the worst case, the leash complexity for an optimal homotopy for (R, K, w) is $\Omega(|K|)$, even if γ_0, γ_1 and g have length 0 and f is convex.*

Proof. For ease of exposition, we first argue the general case, using the construction illustrated in Figure 4(a). We have n spikes, $K = k_1, \dots, k_n$, lined up in the middle, at vertical distance 1 from each other and f and g ; k_1 is the highest- and k_n the lowest-positioned spike. Moreover, the odd-numbered spikes have a closest point on the first half of f and the even-numbered spikes have a closest point on the second half of f ; these closest points are at distance 0.75. This implies that the optimal leash crossing k_i has cost $n - i + 1.75 + w(k_i)$. By setting $w(k_i) = c + i$, all these leashes have the same cost, for suitable constant $c > 0$, depending on the longest leash necessary that does not cross a spike. Now, the homotopy height of this instance is $1.75 + w(k_n) = 1.75 + c + n$. This requires the leashes to cross the spikes in the order of their closest points along f , that is, odd-numbered before even-numbered spikes (see Figure 4(b)). When crossing spike k_1 , the leash has crossed all other odd-numbered spikes, but none of the even-numbered spikes. Thus, the leash has a linear number of inflection points, if we perturb all odd-numbered spikes slightly.



■ **Figure 4** (a) A leash may need linear complexity when considering variable-cost spikes. (b) Part of the optimal homotopy, after crossing k_n up to crossing k_{n-2} . (c) This even holds in the convex case, with one boundary path and the two initial leashes having length 0. The closest point and corresponding distance circles are indicated for each spike.



■ **Figure 5** When f and g have length 0, we can apply a simple greedy strategy to shrink both γ_0 and γ_1 onto γ_{mid} .

A similar construction can be made when we require that f is convex and γ_0, γ_1 and g have length 0. This is illustrated in Figure 4(c). The same principle applies: we position spikes such that their closest point is alternately on the left half and right half of f . By setting the weights appropriately, we can again ensure that the optimal homotopy must cross the spikes in some order along f and force a linear number of inflection points. ◀

Algorithms. If f and g collapse onto a point, we can compute the homotopy height in polynomial time with a greedy algorithm. Interestingly, this contrasts the potential complexity of the problem if γ_0 and γ_1 and even g collapse onto a point, as suggested by the lower bound in Lemma 1.

► **Lemma 2.** *We can compute in polynomial time the homotopy height of (R, K, w) , if f and g have length 0.*

Proof. Consider the geodesic leash γ_{mid} between $f(0)$ and $g(0)$, ignoring any spikes; see Figure 5. As described below, we greedily shrink γ_0 until we reach γ_{mid} . We first compute the geodesic leash γ_t between $f(0)$ and $g(0)$ in the same homotopy class as γ_0 . By definition, γ_t cannot be longer than γ_0 . Then, we cross the minimal-cost spike $k \in K \cap \gamma_t$, resulting in a cost $\|\gamma_t\| + w(k)$. We repeat the process from γ_t , until we reach γ_{mid} . Analogously,

we shrink γ_1 to γ_{mid} . The maximum of $\|\gamma_0\|$, $\|\gamma_1\|$ and all intermediate $\|\gamma_t\| + w(k)$ is the homotopy height of (R, K, w) .

As all intermediate leashes from γ_0 to γ_{mid} are backwards-convex, these leashes grow only shorter. In particular, this implies that we cannot make a leash on some spike k shorter, by first crossing other spikes that are not on the leash but in the unswept area. In other words, we cannot improve the cost by crossing a spike k on a geodesic by first crossing spikes that are not on the geodesic. Hence, the greedy choice is optimal. ◀

The above proof readily implies that the longest leash in the optimal homotopy is determined by the initial and final leash, and thus results in the following lemma. Note that we are interested here only in the geometric length, excluding any spikes. We capture this in a separate lemma as it supports the unit-cost case, detailed in the next section.

► **Lemma 3.** *If $l \in \Gamma$ is backwards-convex and $r \in \Gamma$ is forwards-convex, with $l(0) = r(0)$, $l(1) = r(1)$, such that l and r together bound a region D convex relative to R , then there is a monotone homotopy from l to r consisting of only backwards- and forwards-convex leashes, and whose longest leash has length $\max\{\|l\|, \|r\|\}$.*

For the general case, there is also a simple 2-approximation achievable, by using the algorithm for the geodesic Fréchet distance.

► **Lemma 4.** *We can compute in $O(|R|^2 \log^2 |R|)$ time a 2-approximation of the homotopy height of (R, K, w) .*

Proof. The algorithm computes the geodesic Fréchet distance [13] in R , that is, ignoring the spikes. Consider the optimal geodesic Fréchet matching μ . We may extend μ into a homotopy h by infinitesimal perturbation to cross only one spike at once and by shortening γ_0 and γ_1 to the geodesics between $f(0)$ and $g(0)$ and between $f(1)$ and $g(1)$ respectively. We prove that h is a 2-approximation of the minimal-cost homotopy h^* including the spikes.


Let c be the longest leash in h and let w_{max} denote the maximal cost of a spike. The cost of h is bounded by $c + w_{\text{max}}$. Either c is defined by γ_0 or γ_1 (which must be in any homotopy) or c is defined by a leash in μ : in either case, c provides a lower bound on the maximal leash length in h' . Moreover, h' must also cross the maximal-cost spike. Hence, the cost of h' is at least $\max\{c, w_{\text{max}}\}$. We now have that $c + w_{\text{max}} \leq 2 \cdot \max\{c, w_{\text{max}}\}$, thus proving that h is a 2-approximation of h^* . ◀

3 Unit-cost spikes


In this section, we give an algorithm to compute the homotopy height in the case where all spikes have cost 1. We start by proving properties on the homotopy classes and lengths of curves in Γ . These properties allow us to construct for any homotopy, a homotopy of similar cost with a regular structure. Finally, we show how to decide the existence of such a regular homotopy cheaper than a given cost in polynomial time.

Shortcutting curves. Consider a curve $\gamma \in \Gamma$ and let $D = \text{gh}(\gamma)$ be its convex hull. Let l and r respectively be the backwards- and forwards-convex arcs of ∂D between the endpoints of γ . Consider an arc φ of $\gamma \setminus l$ or $\gamma \setminus r$. Let $\bar{\varphi}$ be the corresponding arc of l or r between the endpoints of φ . Then we refer to the disk bounded by $\varphi \cup \bar{\varphi}$ as a *pocket* of γ , and refer to $\bar{\varphi}$ as its *lid*, see Figure 6. Each lid is a shortest path in R , and the pockets of γ partition D .

► **Lemma 5.** *Let γ and $\gamma' \in \Gamma$ be two non-crossing simple curves. Each arc ψ of $\gamma' \cap \text{gh}(\gamma)$ has both endpoints on the same lid of the containing pocket of γ .*



■ **Figure 6** Pockets of γ with lids on r (shaded).



■ **Figure 7** (a) The curves r_i and l_j on ∂D_i and ∂D_j , respectively. (b) The curves r'_i and l'_j obtained after replacing arcs with the geodesic between x and x' . (c) An example curve γ' in the homotopy class of r'_i that is shorter than γ (dashed).


Proof. Since the endpoints of γ' lie on the boundary of R , ψ must have both endpoints on the boundary of its pocket. Since ψ does not cross γ , and the pocket is bounded by an arc of γ and a lid, the endpoints of ψ lie on the lid. ◀

Consider a forwards-convex and a backwards-convex curve in Γ . If these curves intersect, they intersect in at most two points or geodesics, and occur in the same order and direction along the curves. As such, their first and last point of intersection are naturally well defined.

► **Lemma 6.** *Let γ_i and $\gamma_j \in \Gamma$ be two non-intersecting simple curves, and assume γ_i lies in the swept region of γ_j . Let $D_i = \text{gh}(\gamma_i)$ and $D_j = \text{gh}(\gamma_j)$ be their convex hulls. Let $r_i \in \Gamma$ be the forwards-convex arc of ∂D_i and let $l_j \in \Gamma$ be the backwards-convex arc of ∂D_j . If r_i and l_j intersect, let x and x' be the first and last points of intersection of r_i and l_j . Let r'_i and l'_j be the curves obtained from r_i and l_j by replacing their arcs in $D_i \cap D_j$ by the geodesic between x and x' . If r_i and l_j do not intersect, let $r'_i = r_i$ and $l'_j = l_j$. Assume the region between γ_i and γ_j contains no spikes in its interior and consider a third simple curve γ in this region. There is a curve γ' with the same endpoints as γ and $\|\gamma'\| \leq \|\gamma\|$, such that γ' lies in the homotopy class of r'_i and l'_j .*

Proof. The setup is illustrated in Figure 7. We consider three cases illustrated in Figure 8, depending on the number of bends of r_i and l_j on $\partial(D_i \cap D_j)$ that are induced by γ_i and γ_j .

- (a) In the first case, assume there are no such bends on r_i or l_j , then the interior of $D_i \cap D_j$ is empty. Since D_i and D_j are disjoint, so are the pockets of γ_i and γ_j . If we replace all arcs of γ that lie in pockets of γ_i or γ_j by the geodesic between the endpoints on the corresponding lid, then we obtain a curve γ' between r'_i and l'_j with $\|\gamma'\| \leq \|\gamma\|$. Since there are no spikes between r'_i and l'_j , γ' lies in the same homotopy class as r'_i .
- (b) In the second case, assume either r_i or l_j has no such bend on $\partial(D_i \cap D_j)$, but the other has at least one bend. Without loss of generality, assume that r_i has at least one



■ **Figure 8** The three cases of Lemma 6. $D_i \cap D_j$ shaded.

bend. If l_j has one, a symmetric argument applies. We replace γ by a curve that passes through both x and x' . If γ does not pass through x already, then let $\bar{\varphi}$ and $\bar{\psi}$ be the lids of pockets of γ_i and γ_j , respectively, that intersect in x . It is also possible that $\bar{\varphi}$ is an edge of γ_i ; the proof is then similar. Denote by a the arc $\bar{\varphi} \setminus D_j$, by a' the arc $\bar{\varphi} \cap D_j$ and by c the arc $\bar{\psi} \setminus D_i$. For the lids $\bar{\varphi}'$ and $\bar{\psi}'$ of r_i and l_j crossing in x' , denote by b the arc $\bar{\varphi}' \setminus D_j$, by b' the arc $\bar{\varphi}' \cap D_j$ and by d the arc $\bar{\psi}' \setminus D_i$, see Figure 8 (b).


We claim that γ crosses either c , or both a and a' . If γ crosses c , we are done, so assume it does not. As $a \cup c$ connects γ_i and γ_j , γ must cross $a \cup c$ (and therefore a) an odd number of times to connect f and g . Lemma 5 implies that γ crosses $a \cup a'$ an even number of times. Hence, a' is crossed an odd number of times. We can thus find an arc of γ with endpoints on a and a' , and since this arc does not cross c , it lies in the pocket of $\bar{\varphi}$. Replace this arc by the arc of $\bar{\varphi}$ between those endpoints, which is a shortest path in R that passes through x . We now have a curve with the same endpoints as γ that passes through c or x , and this curve is not longer than γ . We allow the resulting curve to cross γ_i and γ_j along $\bar{\varphi}$, however the resulting curve contains a subcurve of γ that connects a , a' or c to g . Analogously, we can replace this subcurve by a curve that crosses either x' or d . This yields a curve from f to g that passes through c or x , and then through x' or d . Since l_j has no bends, c and d lie on the same lid, which passes through x and x' . Since this lid is a shortest path, we can replace the subpath between c or x and x' or d by a shortest path in R that passes through both x and x' .

The portions of the curve before x and after x' can be shortcut using the techniques of case (a) such that they lie between l_j and r_i and not in $D_i \cap D_j$. This yields a curve γ' in the homotopy class of r'_i with the same endpoints as γ , and $\|\gamma'\| \leq \|\gamma\|$.


- (c) In the final case, both r_i and l_j have bends on $\partial(D_i \cap D_j)$. As before, we shortcut γ such that it first passes through x and then through x' . Let the arcs a , a' , c be as before, and let c' be the arc $\bar{\psi} \cap D_i$. As $a \cup a'$, as well as $c \cup c'$ are crossed an even number of times, but $a \cup c$ is crossed an odd number of times, we have that either both a and a' , or both c and c' are crossed by γ . Replacing the arc between the crossings by the shortest path on the corresponding lid yields a path through x with length at most that of γ . Similarly, we can replace the remainder of the resulting path to also pass through x' . From here, we can use the same technique as in case (b) to obtain a curve γ' in the homotopy class of r'_i with the same endpoints as γ , and $\|\gamma'\| \leq \|\gamma\|$. ◀

3.1 Regular homotopies

Let $G(a, b)$ denote the geodesic in R between $f(a)$ and $g(b)$. Given a homotopy class σ , $G_\sigma(a, b)$ denotes the geodesic between $f(a)$ and $g(b)$ in σ . For $\gamma \in \Gamma$ denote its homotopy class by $\sigma(\gamma)$. If γ is a geodesic in R , we say that $\sigma(\gamma)$ is a *straight* homotopy class.



■ **Figure 9** A decomposed homotopy. Subhomotopies S_i shaded green, and B_i shaded orange.



■ **Figure 10** (a) The region between r_i and l_{i+1} obtained from the geodesic hulls of γ_{t_i} and $\gamma_{t_{i+1}}$. (b) The curves r'_i and l'_{i+1} , and the corresponding geodesics r''_i and l''_{i+1} in the same homotopy class. (c) For all $t_i \leq t \leq t_{i+1}$, the curve γ_t lies in the region between the curves γ_{t_i} and $\gamma_{t_{i+1}}$.

For an optimal homotopy, we may without loss of generality start by shortening γ_0 into the geodesic $G_{\sigma(\gamma_0)}(0, 0)$ in its homotopy class, use an optimal homotopy h from $G_{\sigma(\gamma_0)}(0, 0)$ to $G_{\sigma(\gamma_1)}(1, 1)$, and end by lengthening the geodesic $G_{\sigma(\gamma_1)}(1, 1)$ into γ_1 . The resulting homotopy has cost $\max\{\|\gamma_0\|, \text{cost}(h), \|\gamma_1\|\}$, and the computational challenge is to efficiently find an optimal homotopy h .

We call a homotopy from $G_{\sigma(\gamma_0)}(0, 0)$ to $G_{\sigma(\gamma_1)}(1, 1)$ a *regular* homotopy of order m if it can be decomposed into a sequence of homotopies $S_0, B_1, S_1, \dots, B_i, S_i, \dots, B_m, S_m$, subject to the following constraints, see also Figure 9.

- $S_i(1) = B_{i+1}(0)$ for $0 \leq i \leq m - 1$ and $B_i(1) = S_i(0)$ for $1 \leq i \leq m$, that is, the last leash of a homotopy matches the first leash of the next homotopy.
- For each homotopy B_i , the leashes all have the same endpoints on f and g , but the leashes can move over spikes. Moreover, the longest leash in B_i is either $B_i(0)$ or $B_i(1)$.
- For each homotopy S_i , all leashes are geodesics in the same homotopy class σ_i , but the endpoints of leashes can move along f and g .
- The respective homotopy classes of leashes in S_0 and S_m are $\sigma_0 = \sigma(\gamma_0)$ and $\sigma_m = \sigma(\gamma_1)$.

In Lemma 7, we show that there exists a minimum-cost homotopy h between $G_{\sigma(\gamma_0)}(0, 0)$ and $G_{\sigma(\gamma_1)}(1, 1)$ that is a regular homotopy of order at most k . In Lemma 8, we show that each σ_i (except possibly σ_0 and σ_m) can be assumed to be a straight homotopy class. For each homotopy B_i , some leash can be assumed be the geodesic in R between its endpoints.

For a homotopy h , denote by $\alpha_h(t) \in [0, 1]$, the position of the start of leash $h(t)$ in the parameter space of f (such that $f(\alpha_h(t)) = h(t, 0)$). Symmetrically, denote by $\beta_h(t) \in [0, 1]$ the position of the end of that leash in the parameter space of g (such that $g(\beta_h(t)) = h(t, 1)$). If h is monotone, then α_h and $\beta_h : [0, 1] \rightarrow [0, 1]$ are continuous nondecreasing surjections.

► **Lemma 7.** *Let h be a monotone homotopy from $G_{\sigma(\gamma_0)}(0, 0)$ to $G_{\sigma(\gamma_1)}(1, 1)$ of cost less than L , then there exists a regular homotopy of cost at most L .*

Proof. By monotonicity, as t increases, the number of spikes in the swept region of γ_t cannot decrease. Let $\{t_1, \dots, t_k\}$ be the minimum value t_i for which the swept region of γ_{t_i} contains at least i spikes. For each leash γ_{t_i} of h , let $D_i = \text{gh}_R(\gamma_{t_i})$ be its geodesic hull, and let l_i and $r_i \in \Gamma$ respectively be the backwards-convex and forwards-convex curves on the boundary of D_i connecting $\gamma_{t_i}(0)$ with $\gamma_{t_i}(1)$. For $1 \leq i \leq k$, define r'_i and l'_{i+1} respectively to be the curves obtained by replacing the arcs of r_i and respectively l_{i+1} inside $D_i \cap D_{i+1}$ by the geodesic between the crossings of l_{i+1} and r_i (if any), dashed in Figure 10(a). Then l'_i lies in the swept region of r'_i , and r'_i lies in that of l'_{i+1} . Let r''_i and l''_i be the geodesics with the same endpoints and homotopy class as r'_i and l'_i , respectively (see Figure 10(b)).

Lemma 3 gives us a homotopy B_i from l''_i to r''_i whose leashes have length at most $\max\{\|l''_i\|, \|r''_i\|\}$. Including the cost of spikes, the cost of B_i is at most $\max\{\|l''_i\|, \|r''_i\|\} + 1 + \varepsilon$ (for any $\varepsilon > 0$) by perturbing leashes to each encounter at most one spike. Because l''_i is a geodesic in the same homotopy class as l'_i , we have $\|l''_i\| \leq \|l'_i\|$. Moreover, l'_i is a copy of l_i with a subpath replaced by a shortest path, we have $\|l'_i\| \leq \|l_i\|$. Finally, because l_i lies on arcs of the convex hull of γ_{t_i} , we have $\|l_i\| \leq \|\gamma_{t_i}\|$. Therefore, we have $\|l''_i\| \leq \|\gamma_{t_i}\|$, and by symmetry, $\|r''_i\| \leq \|\gamma_{t_i}\|$. Since γ_{t_i} encounters a spike, and the cost of h is less than L , we have $\|\gamma_{t_i}\| < L - 1$. Hence, $\text{cost}(B_i) \leq \max\{\|l''_i\|, \|r''_i\|\} + 1 + \varepsilon \leq \|\gamma_{t_i}\| + 1 + \varepsilon \leq L$.

Define $r''_0 = G_{\sigma(\gamma_0)}(0, 0)$ and $l''_{k+1} = G_{\sigma(\gamma_1)}(1, 1)$. Moreover, let $t_0 = 0$ and $t_{k+1} = 1$. It remains to construct homotopies S_i between r''_i and l''_{i+1} for $0 \leq i \leq k$. Since there are no spikes interior to the region between r'_i and l'_{i+1} , they lie in the same homotopy class, which we denote by σ_i . To construct a homotopy from r''_i to l''_{i+1} , we consider the curves γ_t with $t_i \leq t \leq t_{i+1}$, and replace them by the geodesic $\gamma''_t = G_{\sigma_i}(\alpha_h(t), \beta_h(t))$ in homotopy class σ_i (see Figure 10(c)). These geodesics move continuously with t , so it remains to show that $\|\gamma''_t\| \leq \|\gamma_t\|$. This is not immediate since γ''_t may lie in a different homotopy class than γ_t . Instead, we use Lemma 6, which tells us that there is a curve with the same endpoints in the homotopy class of r'_i with length at most that of γ_t . Because γ''_t is a geodesic in the same homotopy class, its length is also at most that of γ_t . ◀

The straight homotopy classes in R can be enumerated by taking the geodesic between any two spikes in R , and extending it to the two points on the boundary of R . It is at these points where the geodesic hits the boundary of R that the homotopy class of the geodesic between points on the boundary of R changes: one can slide these points clockwise or counter-clockwise such the geodesic between them ends up in a different straight homotopy class. There are $O(k^2)$ straight homotopy classes.

► **Lemma 8.** *For $1 \leq i \leq k - 1$, we can assume σ_i to be a straight homotopy class without increasing its cost.*

Proof. Curve r_i is forwards-convex and l_{i+1} is backwards-convex, and the endpoints of r_i on f and g are not ahead of those of l_{i+1} . If r_i and l_{i+1} are disjoint, then we can find a geodesic in R separating r_i and l_{i+1} , and hence r''_i and l''_{i+1} .

If they are not disjoint, then the shortest path between their points of intersection lies on a geodesic between f and g , separating r'_i and l'_{i+1} , and hence r''_i and l''_{i+1} . ◀

3.2 Computation

A tool that is commonly used to compute the Fréchet distance is the *free space diagram* [1]. This tool captures between which points of f and g the geodesic is sufficiently short to be used as a leash in a homotopy of a given cost L . Formally, the free space diagram is defined as $\mathcal{F}(L) = \{(a, b) \in [0, 1] \times [0, 1] \mid \|G(a, b)\| \leq L\}$. More generally, for a given homotopy class σ , we define $\mathcal{F}_\sigma(L) = \{(a, b) \in [0, 1] \times [0, 1] \mid \|G_\sigma(a, b)\| \leq L\}$ to capture the geodesics in σ of length at most L .

Let Σ be the set of homotopy classes consisting of $\sigma(\gamma_0)$, $\sigma(\gamma_1)$, and all straight homotopy classes. There are $2 + O(k^2) = O(k^2)$ such homotopy classes, assuming $k \geq 1$. Let h be a regular homotopy of cost at most L , and let t_i and t'_i be the values of t in h at which the constituent homotopy S_i starts and stops, respectively. For all $t \in [t_i, t'_i]$, we have $\|h(t)\| \leq \text{cost}(S_i) \leq L$, so $(\alpha_h(t), \beta_h(t)) \in \mathcal{F}_{\sigma_i}(L)$, where σ_i is the homotopy class of the leashes in S_i . For $t \in [t'_{i-1}, t_i]$, the leashes $h(t)$ are part of homotopy B_i , and we even have $\|h(t)\| + 1 \leq L$, such that accounting for the spikes the leash passes over, we have $\text{cost}(B_i) \leq \max\{\|h(t'_{i-1})\| + 1, \|h(t_i)\| + 1\} \leq L$ by the construction of Lemma 7. Recall that the endpoints of leashes do not move throughout any homotopy B_i , so $\alpha_h(t) = \alpha_h(t_i)$ and $\beta_h(t) = \beta_h(t_i)$ for all $t \in [t'_{i-1}, t_i]$. Additionally, as the construction of Lemma 7 preserves monotonicity, we can assume α_h and β_h to both be continuous nondecreasing surjections. By Lemma 8, we can also assume that each σ_i lies in Σ . For the sake of presentation, since α_h and β_h are constant in the intervals $[t'_{i-1}, t_i]$, we assume from now on that $t'_{i-1} = t_i$, and prove that any homotopy with the structure imposed by Lemma 9 can be turned into a regular homotopy of cost L .

► **Lemma 9.** *We can construct a regular homotopy of cost at most L if we can find appropriate α_h, β_h, t_i and t'_i and values of $\sigma_i \in \Sigma$, with the following conditions. Let α and $\beta: [0, 1] \rightarrow [0, 1]$ be continuous nondecreasing surjections. Let $\sigma_0 = \sigma(\gamma_0)$, $\sigma_m = \sigma(\gamma_1)$, and $\sigma_i \in \Sigma$ for $i \in \{1, \dots, m-1\}$. Let $0 = t_0 \leq t_1 \leq \dots \leq t_{m+1} = 1$. Then, if $(\alpha(t), \beta(t)) \in \mathcal{F}_{\sigma_i}(L)$ for each $t \in [t_i, t_{i+1}]$, and additionally $(\alpha(t), \beta(t)) \in \mathcal{F}_{\sigma_i}(L-1) \cap \mathcal{F}_{\sigma_{i+1}}(L-1)$ for each t_i with $i \in \{1, \dots, m\}$, this corresponds to a regular homotopy of cost at most L .*

Proof. We use geodesics of σ_i for $t \in [t_i, t_{i+1}]$, and they move continuously. By Lemma 3, we can find a homotopy B_i of cost at most L if the geodesics of σ_i and σ_{i+1} based at $\alpha(t)$ and $\beta(t)$ both have length at most $L-1$. Furthermore, since $(\alpha(t), \beta(t)) \in \mathcal{F}_{\sigma_i}(L)$ for each $t \in [t_i, t_{i+1}]$, we can find a homotopy S_i of cost at most L between σ_i and σ_{i+1} . ◀

Computing free space diagrams. To compute the free space diagram in our setting, we subdivide the edges of f and g in such a way that for each pair of (subdivided) edges, the length of the geodesic can be described as a quadratic function in two parameters a and b . This subdivision is based on the lines through any pair of spikes and vertices of ∂R , and finding their projection onto f or g , if any. In total, this yields subdivided curves f' and g' of $O((n+k)^2)$ vertices. Using a standard rotating sweep around every spike and vertex, we can compute the projections in $O((n+k)^2 \log(n+k))$ time and sort them along every edge of f and g in the same time, giving the subdivided curves f' and g' .

Now, given any straight homotopy class, or the homotopy class of γ_0 or γ_1 , we compute the quadratic function for each pair of edges $(e_{f'}, e_{g'})$ of f' and g' . To this end, we take a straight homotopy class σ and determine the induced partition of K into K_1 and K_2 . Then we compute the convex hulls $\text{gh}_R(K_1)$ and $\text{gh}_R(K_2)$ of K_1 and K_2 relative to the domain R . Using the common inner tangents of $\text{gh}_R(K_1)$ and $\text{gh}_R(K_2)$ in R we can find all pairs of edges $(e_{f'}, e_{g'})$ of f' and g' for which the geodesic in σ is straight, and determine the corresponding quadratic functions (which are ellipses). For other pairs of edges of f' and g' , the geodesic contains vertices of $\text{gh}_R(K_1)$, $\text{gh}_R(K_2)$, and R itself, and their lengths are determined by a hyperbolic part in a , a hyperbolic part in b , and a constant part (between the first and last vertices not on f' and g'). We fix an edge $e_{g'}$ of g' and traverse all edges of f' sequentially, updating the three parts of the quadratic function when needed. Updates of the constant part happen only at the ends of the geodesic, and amortized we can do all updates in time linear in the number of parts of f' , that is, $O((n+k)^2)$.

Hence, we can compute all quadratic functions for all straight homotopy classes in time $O(k^2)$ (for the straight homotopy classes) times $O((n+k)^2)$ (for the number of segments of g') times $O((n+k)^2)$ (for the number of segments of f'). In total, this is $O((n+k)^4 k^2) = O(n^4 k^2 + k^6)$ time. The $O(n^4 + k^4)$ cells of the free space diagram each have $O(k^2)$ quadratic functions, at most one for each of the homotopy classes.

Decision algorithm. For a parameter L , we define the *reachable free space* as the set of coordinates $(\sigma, a, b) \in \Sigma \times [0, 1] \times [0, 1]$, such that there exist continuous nondecreasing surjections $\alpha: [0, 1] \rightarrow [0, a]$ and $\beta: [0, 1] \rightarrow [0, b]$, a value m , values $0 = t_0 \leq t_1 \leq \dots \leq t_{m+1} = 1$, and homotopy classes $\sigma_i \in \Sigma$ with $\sigma_0 = \sigma(\gamma_0)$ and $\sigma_m = \sigma$, such that for each $t \in [t_i, t_{i+1}]$, we have $(\alpha(t), \beta(t)) \in \mathcal{F}_{\sigma_i}(L)$ and for each $i \in \{1, \dots, m\}$, we have $(\alpha(t_i), \beta(t_i)) \in \mathcal{F}_{\sigma_i}(L-1) \cap \mathcal{F}_{\sigma_{i+1}}(L-1)$. The reachable free space corresponds to the classes σ and points $f(a)$ and $g(b)$ that have a monotone regular homotopy from $G_{\sigma_0}(0, 0)$ to $G_{\sigma}(f(a), g(b))$ of cost at most L . Deciding whether a regular homotopy of at most a certain cost L exists is then equivalent to testing whether $(\sigma(\gamma_1), 1, 1)$ lies in the reachable free space for parameter L . We can compute the reachable free space using dynamic programming. In contrast algorithms for most variants of the Fréchet distance, which need only information about the free space on the boundary of cells, we also need information about their interiors. In our dynamic program, we compute the reachable free space on the boundary of each cell.

The restriction of the free space to any cell and homotopy class is convex. Therefore, if a point (σ, a, b) lies in the reachable free space, then for all $a' \geq a$ and $b' \geq b$, in the same cell as (a, b) , if $(a', b') \in \mathcal{F}_{\sigma}(e)(L)$, then (σ, a', b') also lies in the reachable free space. Moreover, for σ and $\sigma' \in \Sigma$, if $(a, b) \in \mathcal{F}_{\sigma}(L-1)$ and $(a, b) \in \mathcal{F}_{\sigma'}(L-1)$, then (σ, a, b) lies in the reachable free space if and only if (σ', a, b) lies in the reachable free space. Our dynamic program starts as follows: for a homotopy of cost at most L to exist, check whether $(0, 0) \in \mathcal{F}_{\sigma_0}(L)$. If so, $(\sigma_0, 0, 0)$ lies in the reachable free space, and otherwise the reachable free space is empty. Now we propagate the reachable free space on a cell-by-cell basis, maintaining for each cell $[a, a'] \times [b, b']$ and for each homotopy class σ , two pieces of information. First, the minimum $a^* \in [a, a']$ for which (σ, a^*, b) lies in the reachable free space (if any); and second, the minimum $b^* \in [b, b']$ for which (σ, a, b^*) lies in the reachable free space (if any). The first piece of information can be propagated to a neighboring cell $[a, a'] \times [b', b'']$, and the second piece can be propagated to a neighboring cell $[a', a''] \times [b, b']$.

To propagate this information, we use a horizontal sweep line that maintains the reachable free space intersecting the sweep line for the cell $[a, a'] \times [b, b']$ in each of the $O(k^2)$ relevant homotopy classes, based on the coordinates (σ, a, b^*) and (σ, a^*, b) in each of those homotopy classes. Naively, we can propagate this information in $O(k^6)$ time per cell, using the coordinates of the $O(k^4)$ intersections of the boundary of free space cells from different homotopy classes as events for the sweep line.

After propagating the information through all $O(n^4 + k^4)$ cells of the free space in $O(n^4 k^6 + k^{10})$ time, we can return whether $(\sigma(\gamma_1), 1, 1)$ lies in the reachable free space to decide whether there exists a homotopy of cost at most L .

Exact computation. The candidate values for the minimum-cost regular homotopy depend on the values of L where the a - or b -coordinates of different intersections align. There are $O(((n+k)^2 k^4)^2)$ intersections that can align in this way, which yields $O(n^4 k^8 + k^{12})$ critical values, which we can enumerate in $O(1)$ time per value. We perform a binary search over these critical values, using linear-time median finding and running the decision procedure $O(\log nk)$ times to find the minimum cost of a regular homotopy. Doing so, we compute the homotopy height in $O(n^4 k^6 \log n + n^4 k^8 + k^{12})$ time.

4 Conclusion

We have shown that in a spiked plane with polygonal boundary, we can compute the homotopy height between two curves on the boundary in polynomial time for various cases. In particular, this holds if all spikes have the same (unit) weight, or if the two curves together form the entire boundary of the domain. We also provide a 2-approximation algorithm for the general case. We have also shown that intermediate leashes may require many inflection points for an optimal homotopy, even if the polygonal domain is convex. This complexity of the leash has been preventing us from developing a polynomial-time algorithm, and thus it remains open whether the general case can be solved optimally in polynomial time.

Future work. Various other settings can also be studied. The case where f and g are not on the boundary of the polygonal domain is a natural first step. However, the monotonicity [5] that supports our results is not known to hold in this case, which is likely a premise for efficient optimal algorithms. Our approximation algorithm (Lemma 4) extends to deal with the case that γ_0 and γ_1 are still on the boundary, since the algorithm upon which it is based [13] does not require f and g to lie on the boundary. If the initial leashes γ_0 and γ_1 are no longer specified, we readily get a 2-approximation algorithm by using the algorithm by Chambers *et al.* [6] to solve the decision variant combined with an appropriate search.

References

- 1 Helmut Alt and Michael Godau. Computing the Fréchet distance between two polygonal curves. *International Journal on Computational Geometry and Application*, 5(1–2):75–91, 1995.
- 2 Patrizio Angelini, Giordano Da Lozzo, Giuseppe Di Battista, Fabrizio Frati, Maurizio Patrignani, and Vincenzo Roselli. Morphing planar graph drawings optimally. In *Automata, Languages, and Programming (ICALP)*, LNCS 8572, pages 126–137, 2014.
- 3 Patrizio Angelini, Fabrizio Frati, Maurizio Patrignani, and Vincenzo Roselli. Morphing planar graph drawings efficiently. In *Graph Drawing*, volume 8242 of LNCS 8242, pages 49–60, 2013.
- 4 Graham R. Brightwell and Peter Winkler. Submodular percolation. *SIAM Journal on Discrete Mathematics*, 23(3):1149–1178, 2009.
- 5 Erin W. Chambers, Gregory R. Chambers, Arnaud de Mesmay, Tim Ophelders, and Regina Rotman. Monotone contractions of the boundary of the disc. Computing Research Repository (arXiv):1704.06175, 2017.
- 6 Erin W. Chambers, Éric Colin de Verdière, Jeff Erickson, Sylvain Lazard, Francis Lazarus, and Shripad Thite. Homotopic Fréchet distance between curves or, walking your dog in the woods in polynomial time. *Computational Geometry: Theory and Applications*, 43(3):295–311, 2010.
- 7 Erin W. Chambers and David Letscher. On the height of a homotopy. In *Proceedings of the 21st Canadian Conference on Computational Geometry*, pages 103–106, 2009.
- 8 Erin W. Chambers and Mikael Vejdemo-Johansson. Computing minimum area homologies. *Computer Graphics Forum*, 34(6):13–21, 2015.
- 9 Erin W. Chambers and Yusu Wang. Measuring similarity between curves on 2-manifolds via homotopy area. In *Proceedings of the 29th Annual Symposium on Computational Geometry*, pages 425–434, 2013.
- 10 Gregory R. Chambers. *Optimal Homotopies of Curves on Surfaces*. PhD thesis, University of Toronto, Canada, 2014.

- 11 Gregory R. Chambers and Yevgeny Liokumovich. Converting homotopies to isotopies and dividing homotopies in half in an effective way. *Geometric and Functional Analysis*, 24(4):1080–1100, 2014.
- 12 Gregory R. Chambers and Regina Rotman. Monotone homotopies and contracting discs on Riemannian surfaces. *Journal of Topology and Analysis*, 2016.
- 13 Atlas F. Cook and Carola Wenk. Geodesic Fréchet distance inside a simple polygon. *ACM Transactions on Algorithms*, 7(1):Art. 9, 2009.
- 14 Tamal K. Dey, Anil N. Hirani, and Bala Krishnamoorthy. Optimal homologous cycles, total unimodularity, and linear programming. *SIAM Journal on Computing*, 40(4):1026–1044, 2011.
- 15 Alon Efrat, Quanfu Fan, and Suresh Venkatasubramanian. Curve matching, time warping, and light fields: New algorithms for computing similarity between curves. *Journal of Mathematical Imaging and Vision*, 27(3):203–216, 2007. doi:10.1007/s10851-006-0647-0.
- 16 Alon Efrat, Leonidas J. Guibas, Sariel Har-Peled, Joseph S. B. Mitchell, and T. M. Murali. New similarity measures between polylines with applications to morphing and polygon sweeping. *Discrete & Computational Geometry*, 28(4):535–569, 2002.
- 17 Brittany Fasy, Selcuk Karakoç, and Carola Wenk. On minimum area homotopies. In *Computational Geometry: Young Researchers Forum*, pages 49–50, 2016.
- 18 Sariel Har-Peled, Amir Nayyeri, Mohammad Salavatipour, and Anastasios Sidiropoulos. How to walk your dog in the mountains with no magic leash. In *Proceedings of the 28th Annual Symposium on Computational Geometry*, pages 121–130, 2012.
- 19 Zipei Nie. On the minimum area of null homotopies of curves traced twice. Computing Research Repository (arXiv):1412.0101, 2014.
- 20 Godfried Toussaint. An optimal algorithm for computing the relative convex hull of a set of points in a polygon. School of Computer Science, McGill University, 1986.