WALKING YOUR DOG IN THE WOODS IN POLYNOMIAL TIME

ERIN WOLF CHAMBERS, ÉRIC COLIN DE VERDIÈRE, JEFF ERICKSON,
SYLVAIN LAZARD, FRANCIS LAZARUS, AND SHRIPAD THITE

ABSTRACT. The Fréchet distance between two curves in the plane is the minimum length of a leash that allows a dog and its owner to walk along their respective curves, from one end to the other, without backtracking. We propose a natural extension of Fréchet distance to more general metric spaces, which requires the leash itself to move continuously over time. For example, for curves in the punctured plane, the leash cannot pass through or jump over the obstacles (“trees”). We describe a polynomial-time algorithm to compute the *homotopic* Fréchet distance between two given polygonal curves in the plane minus a given set of obstacles, which are either points or polygons.

1. INTRODUCTION

Given two input curves, a natural question that arises is how similar the curves are to each other. One common measure is the Hausdorff distance, which simply takes the minimum distance between any two points, one from each curve. While the Hausdorff metric does measure closeness in space, it does not take into account the flow of the curves, which in many applications, such as morphing in computer graphics, is an important property of the curves.

The *Fréchet distance*, sometimes called the *dog-leash distance*, is defined as the minimum length of a leash required to connect a dog and its owner as they walk without backtracking along their respective curves from one endpoint to the other. The Fréchet metric takes into account the flow of the two curves because the pairs of points whose distance contributes to the Fréchet metric sweep continuously along their respective curves. It is therefore possible for two curves to have small Hausdorff distance but large Fréchet distance. Fréchet distance is used as a more accurate measure of similarity in many different applications [AG95, AB05].

When the two curves are embedded in a general metric space, the distance between two points on the curves (the length of the shortest leash joining them) is not the Euclidean distance but a geodesic distance. For instance, this is the case if the two curves lie on a terrain (or any surface) [MY05] or if the leash is constrained to the interior of a simple polygon [CW07, EGHP⁺02]. The definition of the ordinary Fréchet distance allows the leash to switch discontinuously, without penalty, from one side of an obstacle or a mountain to another.

In this paper, we introduce a continuity requirement on the motion of the leash. We require that the leash cannot switch, discontinuously, from one geodesic to another; in particular, the leash cannot jump over obstacles and can sweep over a mountain only if it is long enough. We define the *homotopic Fréchet distance* between two curves as the Fréchet distance with this additional continuity requirement. This continuity requirement is satisfied automatically for curves inside a simple polygon, but not in more general environments like convex polyhedra.

The motion of the leash defines a correspondence between the two curves that can be used to morph between the two curves—two points joined by a leash morph into each other. The homotopic Fréchet distance can thus be thought as the minimal amount of deformation needed to transform one curve into the other.

In robotics, the two curves being compared may be two motion sequences in the configuration space of a robot system. When the configuration space has obstacle regions, the similarity between the two curves is more accurately measured by the homotopic Fréchet distance rather than the ordinary Fréchet distance that ignores obstacles.

Efficiently computing the homotopic Fréchet distance in general metric spaces is a new open problem. We present a polynomial-time algorithm for a special case of the general problem, which is to compute the homotopic Fréchet distance between two polygonal curves in the plane with point obstacles.

2. DEFINITIONS

Let S be a fixed Hausdorff metric space. A **curve** in S is a continuous function from the unit interval $[0, 1]$ to S . We will sometimes abuse notation by using the same symbol to denote a curve $A: [0, 1] \rightarrow S$ and its image in S . A **reparameterization of $[0, 1]$** is a continuous, non-decreasing, surjection $\alpha: [0, 1] \rightarrow [0, 1]$. A reparameterization of a curve $A: [0, 1] \rightarrow S$ is any curve $A \circ \alpha$, where α is a reparameterization of $[0, 1]$. The **length** of any curve A , denoted $\text{len}(A)$, is defined by the metric of S ; in particular, two reparameterizations of the same curve have the same length.

A **leash** between two curves A and B is another curve $\lambda: [0, 1] \rightarrow S$ such that $\lambda(0) = A(s)$ and $\lambda(1) = B(t)$ for some parameters s and t . A **homotopy** between curves A and B is a continuous map $h: [0, 1] \times [0, 1] \rightarrow S$ such that $h(\cdot, 0) = A$ and $h(\cdot, 1) = B$. For any $t \in [0, 1]$, the one-parameter function $h(t, \cdot)$ is a leash from A to B . A **leash map** between curves A and B is a homotopy

between some reparameterization of A and some reparameterization of B . Intuitively, a leash map describes the *continuous* motion of a leash between a dog walking along A and its owner walking along B . The **length** of a leash map ℓ , denoted $\text{len}(\ell)$, is the maximum length of any leash $\ell(t, \cdot)$. Finally, the **homotopic Fréchet distance** between two curves A and B , denoted $\overline{\mathcal{F}}(A, B)$, is the infimum, over all leash maps ℓ between A and B , of the length of ℓ :

$$\overline{\mathcal{F}}(A, B) := \inf_{\text{leash map } \ell: [0,1]^2 \rightarrow S} \left(\max_{0 \leq t \leq 1} \text{len}(\ell(t, \cdot)) \right).$$

In contrast, the classical or **leashless** Fréchet distance is defined directly in terms of reparameterizations and distances:

$$\mathcal{F}(A, B) := \inf_{\alpha, \beta: [0,1] \rightarrow [0,1]} \left(\max_{0 \leq t \leq 1} d(A(\alpha(t)), B(\beta(t))) \right).$$

In spaces where shortest paths vary continuously as their endpoints move, such as Euclidean space or the interior of a simple polygon, the two definitions are equivalent. In general, however, homotopic Fréchet distance could be larger (but never smaller) than leashless Fréchet distance.

A **homotopy relative to A and B** , or simply **relative homotopy**, is a continuous function $h: [0, 1] \times [0, 1] \rightarrow S$, such that $h(\cdot, 0)$ and $h(\cdot, 1)$ are respectively of the form $A(u(\cdot))$ and $B(v(\cdot))$, where u and v are continuous functions from $[0, 1]$ to $[0, 1]$.

Two leashes λ and λ' are **relatively homotopic**, denoted $\lambda \simeq \lambda'$, if there is a relative homotopy h such that $h(0, \cdot) = \lambda$ and $h(1, \cdot) = \lambda'$. It is easy to prove that $\simeq$ is an equivalence relation over the set of leashes. Any leash map is (the transpose of) a relative homotopy; thus, all leashes $\ell(t, \cdot)$ determined by a leash map ℓ lie in the same relative homotopy class.

3. HOMOTOPY IN THE PUNCTURED PLANE

In this paper, we develop a polynomial-time algorithm to compute the homotopic Fréchet distance between two polygonal paths in $\mathcal{E} = \mathbb{E}^2 \setminus P$, for some set P of closed polygons, where the underlying metric is geodesic distance. The polygons P act as obstacles; in any leash map in $\mathcal{E}$, the moving leash can neither touch nor jump over any obstacle.

Specifically, the input to our problem consists of two polygonal curves A and B and a set P of polygons in the Euclidean plane. Curves A and B may (self-)intersect, but neither curve intersects any obstacle in P . To simplify our exposition, we assume that no three vertices of the input (vertices of polygons in P or vertices of A and B) are collinear; this assumption can be enforced algorithmically using standard perturbation techniques [Sei98].

Let $a_0, a_1, \dots, a_m$ denote the sequence of vertices of A ; these points define a standard parameterization $A: [0, m] \rightarrow \mathcal{E}$ whose restriction to any integer range $[i-1, i]$ is an affine map onto the corresponding edge $a_{i-1}a_i$. Similarly, the vertices $b_0, b_1, \dots, b_n$ of B define a standard piecewise-affine parameterization $B: [0, n] \rightarrow \mathcal{E}$. Let $P_1, P_2, \dots, P_o$ denote the obstacle polygons in P , and let k denote the total number of vertices in all obstacle polygons. In the special case where every obstacle is a single point, we obviously have $k = o = |P|$. Finally, let $N = n + m + k + 2$ denote the total complexity of the input.

Figure 1 illustrates optimum leash maps for a few sample inputs where P is a discrete set of points.

3.1. Geodesics and Geodesic Leash Maps. To simplify our presentation, we will allow ‘paths’ in $\mathcal{E}$ to touch obstacles in P . Specifically, we consider **geodesics**: piecewise-linear curves in the *closure* of $\mathcal{E}$, whose interior vertices are vertices in P . Although geodesics may run along obstacle boundaries, they do not intersect the interior of any obstacle.

In the special case where the obstacles are points (so the closure of $\mathcal{E}$ is the entire plane), we need some additional information to ensure that each geodesic lies in a unique homotopy class.

FIGURE 1. Optimum leash maps for three inputs. Dashed curves between matching numbers represent intermediate leashes.

Specifically, we associate a **turning angle** with each obstacle point that a geodesic touches. Let C_ε be a circle centered at obstacle point p and radius ε , small enough to exclude every other obstacles in P . A turning angle of θ at an obstacle point p indicates that replacing the portion of γ inside C_ε with a counterclockwise arc of length $\theta\varepsilon$ around C_ε yields a new path homotopic to γ . For example, a path with turning angle zero makes a U-turn at p without enclosing p ; a path that goes straight through p with p on its left (resp. right) has turning angle π (resp. $-\pi$); a turning angle of 10π means the path makes a U-turn after winding around the point five times counterclockwise. A geodesic could meet the same obstacle point more than once; we associate a different turning angle with each incidence.

It can be shown that for any two points x and y in $\mathcal{E}$, the shortest path from x to y in any relative homotopy class is a unique geodesic, in which every turning angle at a point obstacle is either at least π or at most $-\pi$. Conversely, every geodesic in which every turning angle is either at least π or at most $-\pi$ is a shortest path in some homotopy class.

A **geodesic leash map** is a leash map $\ell: [0, 1] \times [0, 1] \rightarrow \mathbb{E}^2$ in which every leash $\ell(t, \cdot)$ is a geodesic, and all these geodesics are in the same relative homotopy class. We next prove that for any leash map ℓ , there is a shorter geodesic leash map ℓ' in the same homotopy class, with the same parameterizations of A and B .

Lemma 3.1. *Suppose ℓ is a leash map between two curves A and B . There is a geodesic leash map ℓ' between A and B such that, for all $t \in [0, 1]$, the leash $\ell'(t, \cdot)$ is the shortest path homotopic to $\ell(t, \cdot)$ with the same endpoints. Additionally, the length of ℓ' is at most the length of ℓ .*

Proof. We lift ℓ to the universal cover $\hat{\mathcal{E}}$ of $\mathcal{E}$, obtaining a leash map $\hat{\ell}$ between the lifts $\hat{A}$ and $\hat{B}$ of A and B respectively. For each $t \in [0, 1]$, let $\hat{\ell}'(\cdot, t)$ be the shortest path between the endpoints of $\hat{\ell}(t, \cdot)$ in the same homotopy class. The universal cover $\hat{\mathcal{E}}$ is a simply-connected space with a locally Euclidean metric, so shortest paths in $\hat{\mathcal{E}}$ vary continuously as the endpoints move continuously. It follows that $\hat{\ell}'$ is a continuous function in both arguments, and therefore a (geodesic) leash map in $\hat{\mathcal{E}}$. The projection ℓ' of $\hat{\ell}'$ back to S is a (geodesic) leash map between A and B . For all t , the leash $\ell'(t, \cdot)$ is the shortest path homotopic to $\ell(t, \cdot)$, so $\max_t \text{len}(\ell'(t, \cdot)) \leq \max_t \text{len}(\ell(t, \cdot))$. $\square$

This lemma implies that the homotopic Fréchet distance between A and B is the length of a geodesic leash map in some homotopy class determined by some reparameterizations of A and B . Thus, the homotopic Fréchet distance can be redefined as the minimum, over all homotopy classes h , of the classical Fréchet distance, where distances are defined by shortest paths in relative homotopy class h :

$$\mathcal{F}_h(A, B) := \min_{\alpha, \beta: [0, 1] \rightarrow [0, 1]} \left(\max_{t \in [0, 1]} D_h(A(\alpha(t)), B(\beta(t))) \right)$$

$$\overline{\mathcal{F}}(A, B) := \min_{\text{homotopy class } h} \mathcal{F}_h(A, B)$$

(Here, $D_h(u, v)$ denotes the length of the shortest path from u to v in relative homotopy class h .)

We call a relative homotopy class h **optimal** if $\overline{\mathcal{F}}(A, B) = \mathcal{F}_h(A, B)$.

For the rest of the paper, we restrict ourselves to geodesic leashes and geodesic leash maps. In Section 4, we provide a characterization of an optimal homotopy class, and we use this characterization to enumerate possible optimal homotopy classes in polynomial time. In Section 5, we describe a polynomial-time algorithm to compute the Fréchet distance within a particular homotopy class. Combining these two subroutines gives us a polynomial-time algorithm to compute homotopic Fréchet distance.

3.2. Maintaining Homotopic Shortest Paths. Our algorithm relies on observations by Hersberger and Snoeyink [HS94] about shortest homotopic paths in the punctured plane; see also [GS98, CLMS04, EKL06, Bes03, Bes04]. Suppose we already know a shortest path (a leash) λ between points $a \in A$ and $b \in B$, such as, for instance, a straight-line segment ab . To compute the geodesic leash between some other pair of points in the same homotopy class as λ , we follow the continuous evolution of the geodesic as the points a and b move along their respective curves. The sequence of obstacle vertices on the leash behaves like a *double-ended queue* or *deque*. A new vertex is pushed onto one end of the deque whenever the first or last segment of λ collides with an obstacle vertex. Conversely, a vertex is popped off one end of the deque when the first or last two segments of λ become collinear, and if their common vertex is a point obstacle, the turning angle at that point is either π or $-\pi$.

4. CHARACTERIZING OPTIMAL HOMOTOPY CLASSES

4.1. Minimal Homotopy Classes. Let $\text{len}(\lambda)$ denote the length of any geodesic leash λ , and let $\text{turn}(\lambda)$ denote the sum of the absolute values of the turning angles at any point obstacles on λ . (Again, if λ meets the same point obstacle more than once, each incidence separately contributes to $\text{turn}(\lambda)$. If none of the obstacles are points, then $\text{turn}(\lambda) = 0$.) For any pair of leashes λ and λ' , we write $\lambda \preceq \lambda'$ if and only if either (a) $\text{len}(\lambda) < \text{len}(\lambda')$, or (b) $\text{len}(\lambda) = \text{len}(\lambda')$ and $\text{turn}(\lambda) \leq \text{turn}(\lambda')$. We write $\lambda \prec \lambda'$ whenever $\lambda \preceq \lambda'$ but $\lambda' \not\preceq \lambda$.

We can extend this partial order to homotopy classes as follows. For any relative homotopy class h and any $s, t \in [0, 1]$, let $\sigma_h(s, t)$ denote the shortest path in h between points $A(s)$ and $B(t)$. For any two homotopy classes h and h' , we write $h \preceq h'$ if and only if $\sigma_h(s, t) \preceq \sigma_{h'}(s, t)$ for all parameters s and t . We write $h \prec h'$ whenever $h \preceq h'$ but $h' \not\preceq h$.

Lemma 4.1. *For any relative homotopy classes h and h' , if $h \preceq h'$, then $\mathcal{F}_h(A, B) \leq \mathcal{F}_{h'}(A, B)$.*

Proof. Let ℓ' be an optimum leash map in homotopy class h' , so that $\text{len}(\ell') = \mathcal{F}_{h'}(A, B)$. For some reparameterizations α and β , we have $\ell'(t, \cdot) = \sigma_{h'}(\alpha(t), \beta(t))$ for all t . Let ℓ be the geodesic leash map in homotopy class h defined by the same reparameterizations: $\ell(t, \cdot) = \sigma_h(\alpha(t), \beta(t))$ for all t . The definition of $\preceq$ implies that $\text{len}(\ell(t, \cdot)) \leq \text{len}(\ell'(t, \cdot))$ for all t . It follows that $\mathcal{F}_h(A, B) \leq \text{len}(\ell) \leq \text{len}(\ell') = \mathcal{F}_{h'}(A, B)$. $\square$

A relative homotopy class h is *minimal* if $h' \preceq h$ implies $h \preceq h'$.

Lemma 4.2. *For any relative homotopy class h , there is a minimal relative homotopy class h' such that $h' \preceq h$.*

Proof. Assume, for the sake of contradiction, that there is no minimal relative homotopy class h' such that $h' \preceq h$. Then, by induction, we can construct an infinite descending chain of relative homotopy classes $h = h_0 \succ h_1 \succ h_2 \succ \dots$. To simplify notation, let $\sigma_n = \sigma_{h_n}(0, 0)$.

Consider the ordered list of obstacle points on each path σ_n . There are finitely many such ordered lists, because $\text{len}(\sigma_n) \leq \text{len}(\sigma_0)$ for each n . Thus, up to taking a subsequence, we may assume that every path σ_n visits the same sequence of obstacle points. This assumption implies that all paths σ_n are geometrically equivalent and thus have equal length. Thus, by definition of $\preceq$,

we have $\text{turn}(\sigma_n) < \text{turn}(\sigma_0)$ for all n . There are finitely many relative homotopy classes with a given ordered list of vertices and with bounded total absolute turning angle. (Specifically, since $\text{turn}(\sigma_n) - \text{turn}(\sigma_0)$ is always a multiple of 2π , there are at most $\lfloor \text{turn}(\sigma_0)/2\pi \rfloor + 1$ such classes.) $\square$

The two previous lemmas immediately imply that there is a *minimal* optimal homotopy class.

In the next two subsections, we characterize minimal homotopy classes and describe how to enumerate them efficiently, first for point obstacles and then for polygonal obstacles.

4.2. Point Obstacles. Suppose P is a fixed finite set of points. A *proper line segment* is a line segment in $\mathcal{E}$ joining a point in A to a point in B .

Proposition 4.3. *A relative homotopy class is minimal if and only if it contains a proper line segment.*

Proof. One direction of the proof is straightforward. Let h be the relative homotopy class of the proper line segment σ from $A(s)$ to $B(t)$. For any relative homotopy class $h' \neq h$, the shortest path $\sigma_{h'}(s, t)$ must be longer than σ , so $\sigma_{h'}(s, t) \not\leq \sigma = \sigma_h(s, t)$, which implies that $h' \not\leq h$. We conclude that h is minimal.

Now let h be an arbitrary minimal homotopy class. Let $\hat{A}$ and $\hat{B}$ be lifts of A and B in the universal cover $\hat{\mathcal{E}}$, such that for all s and t , the shortest path $\hat{\sigma}_h(s, t)$ between $\hat{A}(s)$ and $\hat{B}(t)$ is a lift of $\sigma_h(s, t)$. Let $\hat{P}$ denote the set of *all* lifts of points in P ; these lifted obstacle points lie on the boundary of $\hat{\mathcal{E}}$.

We prove that h contains a proper line segment in two stages. First, we prove that no lifted obstacle point $\hat{p} \in \hat{P}$ lies on *every* path $\hat{\sigma}_h(s, t)$. Next, we construct a relative homotopy from the initial leash $\sigma_h(0, 0)$ to a proper line segment.

STAGE 1: NO COMMON CORNER. For the sake of deriving a contradiction, suppose there is a lifted obstacle point $\hat{p} \in \hat{P}$ such that for all s and t , the path $\hat{\sigma}_h(s, t)$ passes through $\hat{p}$. For all s and t , the path $\hat{\sigma}_h(s, t)$ is a shortest path, so its turning angle at $\hat{p}$ must lie outside the open interval $(-\pi, \pi)$. This turning angle is a continuous function of s and t , so we can assume without loss of generality that it is always greater than π . In other words, we assume that every path $\hat{\sigma}_h(s, t)$ winds counterclockwise around $\hat{p}$.

Now $\hat{p}$ is a lift of some obstacle $p \in P$, and $\hat{\sigma}_h(s, t)$ similarly projects to a geodesic $\sigma_h(s, t)$. For each s and t , let $\tau(s, t)$ denote the path with the same vertices and turning angles as $\sigma_h(s, t)$, except that the turning angle at p is reduced by 2π . All paths $\tau(s, t)$ belong to a single relative homotopy class, which we denote h' .

Fix parameters s and t , and consider the turning angles of $\sigma_h(s, t)$ and $\tau(s, t)$ at p . If the turning angle of $\sigma_h(s, t)$ at p is strictly between π and 3π , then the turning angle of $\tau(s, t)$ at p is strictly between $-\pi$ and π . In this case, $\tau(s, t)$ cannot be the shortest path from s to t in this homotopy class, so $\text{len}(\sigma_{h'}(s, t)) < \text{len}(\tau(s, t)) = \text{len}(\sigma_h(s, t))$.

On the other hand, if the turning angle of $\sigma_h(s, t)$ at p is at least 3π , then the turning angle of $\tau(s, t)$ at p is at least π , which implies that $\tau(s, t)$ is the shortest path from s to t in h' . In this case $\sigma_h(s, t)$ and $\sigma_{h'}(s, t) = \tau(s, t)$ are geometrically equivalent and thus have equal length, but $\text{turn}(\sigma_{h'}(s, t)) = \text{turn}(\sigma_h(s, t)) - 2\pi < \text{turn}(\sigma_h(s, t))$.

Hence $\sigma_{h'}(s, t) \prec \sigma_h(s, t)$ for all s and t , which contradicts our assumption that h is a minimal homotopy class. We conclude that no lifted obstacle point $\hat{p}$ lies on every shortest path $\hat{\sigma}_h(s, t)$.

STAGE 2: HOMOTOPY CONSTRUCTION. If the shortest path $\hat{\sigma}_h(0, 0)$ is a proper line segment, then the geodesic $\sigma_h(0, 0)$ is also a proper line segment, and the proof is complete. Thus, we assume that $\hat{\sigma}_h(0, 0)$ passes through at least one point in $\hat{P}$.

Let $\hat{p}_1, \dots, \hat{p}_k$ be the sequence of lifted obstacle points on the shortest path $\hat{\sigma}_h(0, 0)$. (The points $\hat{p}_i$ are distinct, although their projections back into $\mathcal{E}$ might not be.) Our previous argument implies that for each i , there is a pair of parameters (s_i, t_i) such that $\hat{\sigma}_h(s_i, t_i)$ does not pass through $\hat{p}_i$.

We consider a continuous motion of the parameter point (s, t) , starting at $(s, t) = (0, 0)$ and then moving successively to each point (s_i, t_i) . Specifically, we define two continuous functions $s: [0, k] \rightarrow [0, m]$ and $t: [0, k] \rightarrow [0, n]$ such that $s(0) = t(0) = 0$, and for any integer i , we have $s(i) = s_i$ and $t(i) = t_i$. To simplify our notation, we write $\hat{\sigma}(\tau)$ to denote the shortest path $\hat{\sigma}_h(s(\tau), t(\tau))$.

As the parameter τ ('time') increases, points in $\hat{P}$ are inserted into and deleted from the deque of vertices of $\hat{\sigma}(\tau)$. If the deque is empty at any time τ , then the shortest path $\hat{\sigma}(\tau)$ is a proper line segment, which implies that the projected path $\sigma(\tau)$ is a proper line segment as well, concluding the proof. Thus, we assume to the contrary that the deque is never empty. Each vertex $\hat{p}_1, \dots, \hat{p}_k$ must be deleted from the deque at some time during the motion (but may be reinserted later).

Suppose $\hat{p}$ is the *last* point among $\hat{p}_1, \dots, \hat{p}_k$ to be removed from the deque for the *first* time. Without loss of generality, we assume $\hat{p}$ is first removed from the front of the deque at time τ_1 . Let $\hat{q}$ denote the second point in the deque just before $\hat{p}$ is removed; this point must exist, because the deque is never empty. The point $\hat{p}$ lies on the first segment $\hat{a}\hat{q}$ of $\hat{\sigma}(\tau_1)$, where $\hat{a} = \hat{A}(s(\tau_1))$.

By definition of $\hat{p}$, point $\hat{q}$ must have been pushed onto the *back* of the deque at some earlier time $\tau_2 < \tau_1$. Just after $\hat{q}$ is inserted, the last two points in the deque are $\hat{p}$ and $\hat{q}$, in that order. Moreover, $\hat{q}$ lies on the last segment $\hat{p}\hat{b}$ of $\hat{\sigma}(\tau_2)$, where $\hat{b} = \hat{B}(t(\tau_2))$.

Thus, there is an *improper* line segment $\hat{a}\hat{b}$ between a point in $\hat{A}$ and a point in $\hat{B}$. Since all line segments in $\hat{\mathcal{E}}$ are shortest paths, $\hat{a}\hat{b}$ is the shortest path $\hat{\sigma}_h(\tau_1, \tau_2)$. Thus, the path $\sigma_h(\tau_1, \tau_2)$ in $\mathcal{E}$ is an *improper* line segment in relative homotopy class h . Finally, for sufficiently small $\varepsilon > 0$, one of the four paths $\sigma_h(\tau_1 \pm \varepsilon, \tau_2 \pm \varepsilon)$ is a proper line segment (because no three vertices of the input are collinear). $\square$

Proposition 4.3 implies that we can enumerate the set of minimal relative homotopy classes in polynomial time as follows. For every pair of points $p, q \in P$, we find all intersections of the line $\overrightarrow{pq}$ with A and B , in $O(m + n)$ time by brute force. For each pair of intersection points $a \in A$ and $b \in B$, we obtain four proper line segments arbitrarily close to ab . (Alternately, we obtain four different straight geodesics from a to b , by assigning turning angles π or $-\pi$ at p and q .) Altogether, we find $O(mnk^2)$ proper line segments, at least one in each minimal homotopy class, in $O(mnk^2)$ time.

There are polygonal curves and point sets that admit $\Omega(mnk^2)$ distinct minimal relative homotopy classes; see Figure 2 for an example. Thus, any improvement in this portion of the algorithm will require a finer characterization of optimal relative homotopy classes.

FIGURE 2. Curves and points with $\Omega(N^4)$ minimal relative homotopy classes.

4.3. Polygonal Obstacles. The preceding argument breaks down if we allow polygonal obstacles; indeed, it is easy to construct instances where there are *no* proper line segments. Even if the instance allows proper line segments, the optimal homotopy class may not include one. Thus, we require a more complex characterization of minimal homotopy classes in this case.

We say that obstacle vertices p and q *pin* a geodesic γ if the globally shortest path from p to q (in the closure of $\mathcal{E}$) is a subpath of γ ; note that p and q may be the same point. A relative homotopy

class h is *pinned* if some pair of obstacle vertices pins every geodesic in h . The intersection of all geodesics within a pinned homotopy class is a shortest path between obstacle vertices, which we call the **pinned subpath** of h .

Lemma 4.4. *Every minimal relative homotopy class either contains a proper line segment or is pinned.*

Proof. As in the proof of Proposition 4.3, let h be an arbitrary minimal homotopy class. Let $\hat{A}$ and $\hat{B}$ be lifts of A and B in the universal cover $\hat{\mathcal{E}}$, such that for all s and t , the shortest path $\hat{\sigma}_h(s, t)$ between $\hat{A}(s)$ and $\hat{B}(t)$ is a lift of $\sigma_h(s, t)$. Let $\hat{P}$ denote the set of all lifts of the vertices of obstacles in P ; again, every point in $\hat{P}$ lies on the boundary of $\hat{\mathcal{E}}$.

Let $\hat{\pi}_h$ denote the intersection of all shortest paths $\hat{\sigma}_h(s, t)$. If $\hat{\pi} = \emptyset$, then by Stage 2 in the proof of Proposition 4.3, h contains a proper line segment.

Otherwise, $\hat{\pi}_h$ is a shortest path between some pair of lifted obstacle vertices $\hat{p}$ and $\hat{q}$. (In the special case where $\hat{\pi}_h$ is a single point, we have $\hat{p} = \hat{q} = \hat{\pi}_h$.) Now $\hat{p}$ and $\hat{q}$ are lifts of obstacle vertices p and q (which may be the same point, even if $\hat{p}$ and $\hat{q}$ are not), and $\hat{\pi}_h$ is similarly a lift of some path π_h from p to q .

Let σ_h denote the globally shortest path from p to q , and suppose that $\pi_h \neq \sigma_h$. For each s and t , let $\tau(s, t)$ be the curve obtained from $\sigma_h(s, t)$ by replacing π_h with σ_h . All paths $\tau(s, t)$ belong to the same relative homotopy class, which we will denote h' . We now easily confirm that $h' \prec h$, contradicting our assumption that h is minimal. We conclude that π_h is the shortest path from p to q , which implies that h is a pinned relative homotopy class, and π_h is its pinned subpath. $\square$

We call a geodesic γ **direct** if it consists of a proper line segment from A to some obstacle vertex p , the shortest path from p to an obstacle vertex q , and a proper line segment from q to B .

Lemma 4.5. *Every pinned relative homotopy class contains a direct geodesic.*

Proof. Let h be a pinned homotopy class, and let p and q denote the start and end of the pinned subpath π_h . Any geodesic $\sigma_h(s, t)$ consists of a geodesic $\alpha_h(s)$ from A to p , the pinned subpath π_h , and a geodesic $\beta_h(t)$ from q to B .

If $\alpha_h(0)$ is a proper segment, our claim is proved. Thus, we assume that $\alpha_h(0)$ is not a direct geodesic, which implies that the lifted path $\hat{\alpha}_h(0)$ passes through at least one obstacle vertex other than its endpoint $\hat{p}$. Let $\hat{p}^-$ be the last lifted obstacle vertex on $\hat{\alpha}_h(0)$ before $\hat{p}$. Let s_0 be the largest value such that $\hat{\alpha}_h(s)$ contains $\hat{p}^-$ for all $0 \leq s \leq s_0$. Because $\hat{p}^-$ is not on the pinned subpath, it is not on every geodesic $\hat{\alpha}_h(s)$, which implies that $s_0 < m$. For sufficiently small $\varepsilon > 0$, the geodesic $\alpha_h(s_0 + \varepsilon)$ is a proper line segment.

A similar argument implies that $\beta_h(t)$ is a proper line segment for some t . $\square$

We can compute all homotopy classes of proper line segments using only a small modification our earlier algorithm. In fact, we begin by running the earlier enumerate algorithm using only the vertices of polygons in P as obstacles. We then discard any line segment that intersects the interior of any obstacle polygon.

Lemma 4.5 implies that we can enumerate a superset of the minimal relative homotopy classes in polynomial time. We begin by computing the shortest paths between every pair of obstacle vertices [HS99]. Next, for every pair p and q of obstacle vertices, we extend the ray $\vec{pq}$ until it reaches the interior of an obstacle (or infinity), and then compute all $O(m+n)$ intersections between the resulting line segment and the curves A and B . Finally, we concatenate all $O(mk^2)$ initial segments, $O(k^2)$ shortest paths, and $O(nk^2)$ final segments to obtain $O(mnk^6)$ pinned paths in $O(mnk^6) = O(N^8)$ time. Every pinned relative homotopy class contains one of these paths.

Unlike Proposition 4.3, we do not have an exact characterization of minimal homotopy classes for polygonal obstacles—direct geodesics do not necessarily lie in pinned homotopy classes, and not all

pinned homotopy classes are minimal. However, there are problem instances with $\Omega(mnk^6)$ pinned homotopy classes; for example, replace each obstacle point in Figure 2 with a small triangle.

4.4. Non-Polygonal Obstacles. We can further generalize our characterization of minimal homotopy classes to arbitrary non-polygonal obstacles. If we replace ‘obstacle vertex’ with ‘obstacle boundary point’ in the definitions, Lemmas 4.4 and 4.5 are still true in this more general setting, although the proofs are somewhat more involved. We omit details from this extended abstract.

5. COMPUTING FRÉCHET DISTANCE IN ONE HOMOTOPY CLASS

In this section, we describe an algorithm to compute the Fréchet distance $\mathcal{F}_h(A, B)$ in some relative homotopy class h . Our algorithm is a direct adaptation of Alt and Godau’s algorithm for computing the classical Fréchet distance between polygonal paths in the plane [AG95].

As in the previous section, for any $s \in [0, m]$ and $t \in [0, n]$, let $\sigma_h(s, t)$ denote the shortest path from $A(s)$ to $B(t)$ in homotopy class h , and let $d_h(s, t) = \text{len}(\sigma_h(s, t))$. For any $\varepsilon > 0$, let $F_\varepsilon \subseteq [0, m] \times [0, n]$ denote the **free space** $\{(s, t) \mid d_h(s, t) \leq \varepsilon\}$. Our goal is to compute the smallest value of ε such that F_ε contains a monotone path from $(0, 0)$ to (m, n) ; this is precisely the Fréchet distance $\mathcal{F}_h(A, B)$.

The parameter space $[0, m] \times [0, n]$ decomposes naturally into an $m \times n$ grid; let $C_{i,j} = [i-1, i] \times [j-1, j]$ denote the grid cell representing paths from the i th edge of A to the j th edge of B . Our generalization of Alt and Godau’s algorithm requires that the restriction of the function d_h to any grid cell $C_{i,j}$ is convex. We prove this fact in Appendix A (Proposition A.5).

As input to our problem, we are given a path $\sigma_h(s_0, t_0)$ in relative homotopy class h ; based on the results of the previous section, this is either a proper line segment or a direct geodesic. Without loss of generality, we assume that the endpoints $A(s_0)$ and $B(t_0)$ are vertices of A and B ; otherwise, we insert them as new vertices and reparameterize.

5.1. Preprocessing for Distance Queries. The only significant difference between our algorithm and Alt and Godau’s is that we require additional preprocessing to compute several *critical distances* and an auxiliary data structure to answer certain *distance queries*. (If there are no obstacles, each critical distance can be computed, and each distance query can be answered, in constant time.)

There are three types of critical distances:

- **endpoint distances** $d_h(0, 0)$ and $d_h(m, n)$,
- **vertex-edge distances** $d_h(i, [j-1, j]) = \min\{d_h(i, t) \mid t \in [j-1, j]\}$ for all integers $i \in [0, m]$ and $j \in [1, n]$, and
- **edge-vertex distances** $d_h([i-1, i], j) = \min\{d_h(s, j) \mid s \in [i-1, i]\}$ for all integers $i \in [1, m]$ and $j \in [0, n]$.

Given integers i and j and any real value ε , a **horizontal distance query** asks for all values of $t \in [j-1, j]$ such that $d_h(i, t) = \varepsilon$, and a **vertical distance query** asks for all values of $s \in [i-1, i]$ such that $d_h(s, j) = \varepsilon$. The convexity of d_h within any grid cell implies that any distance query returns at most two values.

We first describe how to preprocess a single vertical edge in the parameter grid to answer distance queries; computation of the critical values will be done automatically during the preprocessing. Obviously a similar result applies to horizontal grid edges.

Lemma 5.1. *Suppose we are given a point p and a line segment $\ell = \overline{xy}$, parameterized over $[0, 1]$, as well as the geodesic $\sigma_h(p, x)$ and its length $d_h(p, x)$. In $O(k \log k)$ time, we can build a data structure of size $O(k)$ such that for any ε , all values $t \in [0, 1]$ such that $d_h(p, \ell(t)) = \varepsilon$ can be computed in $O(\log k)$ time. We also report the critical vertex-edge distance $d_h(p, \ell)$, the path $\sigma_h(p, y)$, and its length $d_h(p, y)$.*

Proof. We first compute a constrained Delaunay triangulation of the polygons P , the segment ℓ , and point p in time $O(k \log k)$. This triangulation includes ℓ and the edges of polygons in P as edges.

We apply the following observations used in the *funnel* algorithm for computing shortest homotopic paths [Cha82, LP84, HS94]. The shortest homotopic paths $\sigma_h(p, x)$ and $\sigma_h(p, y)$ may share a common subpath and then split at some vertex v ; v is then the apex of two concave chains that form a funnel with base xy . Each concave chain has complexity at most k and intersects a given edge of the triangulation at most twice.

The geodesic from p to x may have complexity greater than $O(k)$, but (as observed above) the concave chain from v to x will have at most $O(k)$ segments. Our goal is to find a vertex w on $\sigma_h(p, x)$ such that the path from w to x contains v . In other words, the chain from w to x along $\sigma_h(p, x)$ will be of complexity $O(k)$ and will contain the concave funnel path.

To find w , walk along the path from x to p . If we find a vertex where the chain is not concave, we must have passed v , so we mark the non-concave vertex as w . If we ever re-cross a segment of the triangulation a second time, we again must have passed the funnel apex v so we can mark the second crossing as w . (We walked along $O(k)$ edges of the chain to find w .) Let π be the portion of $\sigma_h(p, x)$ between p and w , and τ_1 be the portion of $\sigma_h(p, x)$ between w and x .

We know that π is contained in $\sigma_h(p, y)$, since w is before the apex of the funnel v . Let τ_2 be the portion of $\sigma_h(p, y)$ between w and y ; this can be computed in $O(k)$ time using the funnel algorithm. Given τ_2 , we can then find the apex of the funnel v in $O(k)$ time.

Imagine extending each line segment on the concave chains until it intersects ℓ , the line connecting x and y . Between the two concave chains, the combinatorial description of the distance function changes only at points where the extended lines meet ℓ . To answer distance queries, we will record the $O(k)$ intersections of the extended lines with ℓ . For each of the resulting intervals, record the (fixed) length of the geodesic up to the first vertex in the extended line, as well as the equations of the two lines that bracket the interval. In constant time per interval, we can also compute and store the value $t^* \in [0, 1]$ such that $d_h(p, \ell(t^*))$ is minimized, along with the path $\sigma_h(p, \ell(t^*))$; this gives the desired value $d_h(p, \ell)$.

Now given this data structure, we answer distance queries as follows. If the distance queried is smaller than $d_h(p, \ell)$, we return the empty set. If it is equal to $d_h(p, \ell)$, we return $\ell(t^*)$. If it is larger than $d_h(p, \ell)$, we do two binary searches, one on the intervals between x and $\ell(t^*)$ and the other on the intervals between $\ell(t^*)$ and y . $\square$

Lemma 5.2. *In $O(mnk \log k)$ time, we can compute all critical distances, as well as a data structure of size $O(mnk)$ that can answer any horizontal or vertical distance query in $O(\log k)$ time.*

Proof. We preprocess each edge of the parameter grid as described in Lemma 5.1. We start from the vertex (i, j) that is our given input, either a straight line segment or a direct geodesic. We then walk on the edges of the grid, visiting each edge at least once and at most $O(1)$ times. During this walk, at each current vertex (i, j) , we maintain the shortest homotopic path $\sigma_h(i, j)$ and its length $d_h(i, j)$. Each time we walk along an edge, we apply Lemma 5.1 to preprocess it and to compute the shortest homotopic path corresponding to the target vertex of that edge. Each step takes $O(k \log k)$ time, and there are $O(mn)$ edges, whence the running-time. $\square$

5.2. Decision Procedure. Like Alt and Godau, we first consider the following *decision problem*: Is $\mathcal{F}_h(A, B)$ at least some given value ε ? Equivalently, is there a monotone path in the free space F_ε from $(0, 0)$ to (m, n) ? Our algorithm to solve this decision problem is identical to Alt and Godau's, except for the $O(\log k)$ -factor penalty for distance queries; we briefly sketch it here for completeness.

For any integers i and j , let $h_{i,j}$ denote the intersection of the free space F_ε with the horizontal edge $([i-1, i], j)$, and let $v_{i,j}$ denote the intersection of F_ε with the vertical edge $(i, [j-1, j])$. In the first phase of the decision procedure, we compute $h_{i,j}$ and $v_{i,j}$ for all i and j , using one distance query (and $O(\log k)$ time) for each edge of the parameter grid.

In the second phase of the decision procedure, we propagate in lexicographic order from $C_{1,1}$ to $C_{m,n}$ and determine which $h_{i,j}$ and $v_{i,j}$ are reachable via a monotone path from $C_{1,1}$. Since the free space in each $C_{i,j}$ is convex, we can propagate through each cell in constant time.

Our decision algorithm returns true if and only if there is a monotone path that reaches (m, n) . The total running time is $O(mn \log k)$.

5.3. Computing Fréchet Distance. Finally, we describe how to use our decision procedure to compute the optimum value $\varepsilon^* = \min\{\varepsilon \mid (m, n) \in R_\varepsilon\}$; this is the Fréchet distance $\mathcal{F}_h(A, B)$.

We start by computing the critical distances and the distance-query data structure in time $O(mnk \log k)$, as described in Lemma 5.2. We then sort the $O(mn)$ critical distances. Using the decision procedure, we can compare the optimal distance ε^* with any critical distance ε in $O(mn \log k)$ time. By binary search, we can, repeating this step $O(\log mn)$ times, compute an interval $[\varepsilon^-, \varepsilon^+]$ that contains ε^* but no critical distances.

We then apply Megiddo's *parametric search* technique [Meg83]; see also [vOV04]. Parametric search combines our decision procedure with a 'generic' parallel algorithm whose combinatorial behavior changes at the optimum value ε^* . Alt and Godau observe that one of two events occurs when $\varepsilon = \varepsilon^*$:

- For some integers i, i', j , the bottom endpoint of $v_{i,j}$ and the top endpoint of $v_{i',j}$ lie on the same horizontal line.
- For some integers i, j, j' , the left endpoint of $h_{i,j}$ and the right endpoint of $h_{i,j'}$ lie on the same vertical line.

Thus, it suffices to use a 'generic' algorithm that sorts the $O(mn)$ endpoint *values* of all non-empty segments $h_{i,j}$ and $v_{i,j}$, where the value of an endpoint (s, j) of $h_{i,j}$ is s , and the value of an endpoint (i, t) of $v_{i,j}$ is t .

As suggested by Cole [Col87], we use the parallel sorting algorithm of Ajtai et al. [AKS83] ('the AKS network'). Each parallel step of the sorting algorithm needs to compare $O(mn)$ endpoints. The graph of an endpoint, considered as a function of ε , is convex, monotone, and made of $O(k)$ pieces, each having a simple closed form (see proof of Lemma 5.1). It follows that the sign of a comparison between two endpoints may change at $O(k)$ different values of ε that can be computed in $O(k)$ time. Applying the parametric search paradigm requires the following operations for each parallel step of the sorting algorithm:

- Compute the $O(mnk)$ values of ε corresponding to the changes of sign of the $O(mn)$ comparisons. This can be done in $O(mnk)$ time.
- Apply binary search to these values by median finding, calling the decision procedure to discard half of them at each step of the search. This takes $O(mnk + T_d \log(mnk))$ time, where the first term in this sum stands for the computation of the medians during binary search and $T_d = O(mn \log k)$ is the complexity of the decision procedure. We obtain this way an interval for ε where each of the $O(mn)$ comparisons has a determined sign.
- Deduce in $O(mn \log k)$ time the sign of each of the $O(mn)$ comparisons in the previously computed interval.

Taking into account the $O(\log mn)$ parallel steps of the sorting algorithm, the resulting parametric search algorithm runs in time $O(mn \log(mn)(k + \log k \log(mnk))) = O(N^3 \log^2 N)$.

Lemma 5.3. *Given a proper line segment or a direct geodesic in relative homotopy class h , the Fréchet distance $\mathcal{F}_h(A, B)$ can be computed in time $O(N^3 \log^2 N)$.*

5.4. Summary. Finally, to compute the homotopic Fréchet distance $\overline{\mathcal{F}}(A, B)$ in the plane minus a set of point obstacles, we compute the Fréchet distance $\mathcal{F}_h(A, B)$ in each of the $O(mnk^2)$ homotopy classes h that contain a line segment. We conclude:

Theorem 5.4. *The homotopic Fréchet distance between two polygonal curves in the plane minus a set of points can be computed in $O(N^7 \log^2 N)$ time.*

For polygonal obstacles, we obtain a slower algorithm, because we must check a larger number of relative homotopy classes.

Theorem 5.5. *The homotopic Fréchet distance between two polygonal curves in the plane minus a set of polygons can be computed in $O(N^{11} \log^2 N)$ time.*

6. CONCLUSION

In this paper, we introduced a natural generalization of the Fréchet distance between curves to more general metric spaces, called the homotopic Fréchet distance. We gave a polynomial-time algorithm to compute the homotopic Fréchet distance between polygonal curves in the plane with obstacles.

Improving the running time of our algorithms is the most immediate outstanding open problem. For point obstacles, we conjecture that the running time can be improved by at least a factor of N by optimizing leash maps in every minimal homotopy class simultaneously. Since shortest paths between the same endpoints but belonging to different homotopy classes are related, we expect to (partially) reuse the results of shortest path computations going from one homotopy class to another. For polygonal obstacles, an exact characterization of minimal homotopy classes would almost certainly lead to a significantly faster algorithm.

Cook and Wenk [CW07] describe an algorithm for computing geodesic Fréchet distance between two curves within a simple polygon, generalizing earlier results of Efrat *et al.* [EGHP⁺02]. Their algorithm is faster than ours by roughly a factor of N , in part because they use a randomized strategy in place of parametric search. Unfortunately, we have not been able to apply their technique to our more general problem, because it seems to require a simply-connected environment, but similar ideas may simplify and improve our algorithm.

The *weak* Fréchet distance is a variant of the ordinary dog-leash distance without the requirement that the endpoints move monotonically along their respective curves—the dog and its owner are allowed to backtrack to keep the leash between them short. Alt and Godau [AG95] gave a simpler algorithm for computing the weak Fréchet distance, using a graph shortest path algorithm instead of parametric search. A similarly simple algorithm computes the weak *homotopic* Fréchet distance in polynomial time. We omit further details from this extended abstract.

It would be interesting to compute optimum leash maps in more general spaces. In particular, we are interested in computing the homotopic Fréchet distance between two curves on a convex polyhedron, generalizing the algorithm of Maheshwari and Yi for leashless Fréchet distance [MY05]. The vertices of the polyhedron are ‘mountains’ over which the leash can pass only if it is long enough. Shortest paths on the surface of a convex polyhedron do *not* vary continuously as the endpoints move, because of the positive curvature at the vertices, so we cannot consider only geodesic leash maps.

Finally, it would also be interesting to consider the homotopic Fréchet distance between higher-dimensional manifolds; such problems arise with respect to surfaces in configuration spaces of robot systems. However, even the ordinary Fréchet distance is difficult to compute in higher dimensions [AB05].

ACKNOWLEDGMENTS

This research was initiated during a visit to INRIA Lorraine in Nancy, made possible by a UIUC/CNRS/INRIA travel grant. Research by Erin Chambers and Jeff Erickson was also partially supported by NSF grant DMS-0528086; Erin Chambers was additionally supported by an NSF graduate research fellowship. Research by Shripad Thite was partially supported by the Netherlands Organisation for Scientific Research (NWO) under project number 639.023.301 and by INRIA Lorraine. We thank Hazel Everett and Sylvain Petitjean for useful discussions and great company, and Kira and Nori for several walks in the woods.

REFERENCES

- [AB05] Helmut Alt and Maike Buchin. Semi-computability of the Fréchet distance between surfaces. In *Proc. 21st European Workshop on Computational Geometry*, pages 45–48, March 2005.
- [AG95] Helmut Alt and Michael Godau. Computing the Fréchet distance between two polygonal curves. *IJCGA*, 5(1–2):75–91, 1995.
- [AKS83] Miklós Ajtai, János Komlós, and Endre Szemerédi. Sorting in $c \log n$ parallel steps. *Combinatorica*, 3(1):1–19, 1983.
- [Bes03] Sergei Bespamyatnikh. Computing homotopic shortest paths in the plane. In *Proc. 14th Annu. ACM-SIAM Sympos. Discrete Algorithms*, pages 609–617, 2003.
- [Bes04] Sergei Bespamyatnikh. Encoding homotopy of paths in the plane. In *Proc. LATIN 2004: Theoretical Informatics*, volume 2976 of *Lect. Notes Comput. Sci.*, pages 329–338. Springer-Verlag, 2004.
- [Cha82] Bernard Chazelle. A theorem on polygon cutting with applications. In *Proc. 23rd Annu. IEEE Sympos. Found. Comput. Sci.*, pages 339–349, 1982.
- [CLMS04] Sergio Cabello, Yuanxin Liu, Andrea Mantler, and Jack Snoeyink. Testing homotopy for paths in the plane. *Discrete Comput. Geom.*, 31(1):61–81, 2004.
- [Col87] Richard Cole. Slowing down sorting networks to obtain faster sorting algorithms. *JACM*, 34(1):200–208, January 1987.
- [CW07] Altas Cook and Carola Wenk. Geodesic Fréchet and Hausdorff distance inside a simple polygon. Tech. Rep. CS-TR-2007-004, U. Texas San Antonio, 2007.
- [EGHP⁺02] 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 Comput. Geom.*, 28:535–569, 2002.
- [EKL06] Alon Efrat, Stephen G. Kobourov, and Anna Lubiw. Computing homotopic shortest paths efficiently. *Comput. Geom. Theory Appl.*, 35(3):162–172, 2006.
- [GS98] Dima Grigoriev and Anatol Slissenko. Polytime algorithm for the shortest path in a homotopy class amidst semi-algebraic obstacles in the plane. In *Proc. Internat. Sympos. Symbolic and Algebraic Computation*, pages 17–24, 1998.
- [HS94] John Hershberger and Jack Snoeyink. Computing minimum length paths of a given homotopy class. *Computational Geometry: Theory and Applications*, 4:63–67, 1994.
- [HS99] John Hershberger and Subhash Suri. An optimal algorithm for euclidean shortest paths in the plane. *SIAM J. Comput.*, 28(6):2215–2256, 1999.
- [LP84] D. T. Lee and F. P. Preparata. Euclidean shortest paths in the presence of rectilinear barriers. *Networks*, 14:393–410, 1984.
- [Meg83] N. Megiddo. Applying parallel computation algorithms in the design of serial algorithms. *J. ACM*, 30:852–866, 1983.
- [MY05] Anil Maheshwari and Jiehua Yi. On computing Fréchet distance of two paths on a convex polyhedron. In *Proc. 21st European Workshop on Computational Geometry*, pages 41–44, 2005.
- [Sei98] Raimund Seidel. The nature and meaning of perturbations in geometric computing. *Discrete Comput. Geom.*, 19:1–17, 1998.
- [vOV04] René van Oostrum and Remco C. Veltkamp. Parametric search made practical. *Computational Geometry: Theory and Applications*, 28:75–88, 2004.

APPENDIX A. CONVEXITY OF GEODESIC DISTANCE

Let A and B be arbitrary line segments in the universal cover $\hat{\mathcal{E}}$ of the punctured plane, and let $d: [0, 1] \rightarrow \mathbb{R}^+$ be such that $d(t)$ is the length of the shortest path in $\hat{\mathcal{E}}$ between $\hat{A}(t)$ and $\hat{B}(t)$. The goal of this section is to prove that this distance function is convex.

We first give some elementary lemmas concerning convex functions. Let $f: \mathbb{R} \rightarrow \mathbb{R}^2$ be an arbitrary affine function. Let $\|\cdot\|$ denote the Euclidean norm.

Lemma A.1. *The function $t \mapsto \|f(t)\|$ is convex.*

Proof. For any $s, t \in \mathbb{R}$ and for any $\alpha \in [0, 1]$, we have

$$\|f((1 - \alpha)s + \alpha t)\| = \|(1 - \alpha)f(s) + \alpha f(t)\|$$

From the triangle inequality and the homogeneity of the norm, we get

$$\|f((1 - \alpha)s + \alpha t)\| \leq (1 - \alpha)\|f(s)\| + \alpha\|f(t)\|,$$

which proves convexity. $\square$

Lemma A.2. *Let f and g be two real convex functions such that $f(t) \geq g(t)$ for all $t \in \mathbb{R}$, and let $\tau \in \mathbb{R}$ be such that $f(\tau) = g(\tau)$. Then the two functions $u(t)$ and $v(t)$ defined below are convex:*

$$u(t) = \begin{cases} f(t), & \text{if } t < \tau \\ g(t), & \text{otherwise} \end{cases} \quad v(t) = \begin{cases} g(t), & \text{if } t < \tau \\ f(t), & \text{otherwise} \end{cases}$$

Proof. We prove the convexity of u ; the convexity of v can be proved similarly. We need to show:

$$(1) \quad \forall s, t \text{ such that } s < \tau < t, \forall \alpha \in [0, 1] : u((1 - \alpha)s + \alpha t) \leq (1 - \alpha)u(s) + \alpha u(t)$$

Let $t = (1 - \alpha)s + \alpha t$. If $t \geq \tau$ then $u(t) = g(t)$, and the inequality follows from the convexity of g and because $g(t) \leq u(t)$.

If $t < \tau$ Equation 1 becomes

$$f(t) \leq f(s) + \alpha(g(t) - f(s))$$

Let

$$\rho_g(x, y) = \frac{g(y) - g(x)}{y - x}$$

Since g is convex, it is well-known that ρ_g is an increasing function of each variable when the other one is fixed.

Let $t = s + \beta(\tau - s)$. By convexity of f , we have

$$f(t) \leq f(s) + \beta(f(\tau) - f(s))$$

whence we conclude

$$(2) \quad \begin{aligned} f(t) &\leq f(s) + \beta(g(\tau) - g(s)) - \beta(f(s) - g(s)) \\ &\leq f(s) + \beta(g(\tau) - g(s)) - \alpha(f(s) - g(s)) \end{aligned}$$

The last inequality is implied by $\alpha < \beta$. Since $\beta(\tau - s) = \alpha(t - s)$ we can also write

$$(3) \quad \begin{aligned} \beta(g(\tau) - g(s)) &= \beta(\tau - s)\rho_g(\tau, s) \\ &\leq \alpha(t - s)\rho_g(t, s) \end{aligned}$$

Substituting Equation 3 in Equation 2, we get

$$f(t) \leq f(s) + \alpha(t - s)\rho_g(t, s) - \alpha(f(s) - g(s)),$$

which is the required inequality after expanding the right-hand side. $\square$

Corollary A.3. *Let f, g, h be three real convex functions such that $f(t) \geq h(t)$ and $g(t) \geq h(t)$ for all $t \in \mathbb{R}$, and let $\tau \in \mathbb{R}$ be such that $f(\tau) = g(\tau) = h(\tau)$. Then the function w defined below is convex:*

$$w(t) = \begin{cases} f(t), & \text{if } t < \tau \\ g(t), & \text{otherwise} \end{cases}$$

Proof. Define

$$u(t) = \begin{cases} f(t), & \text{if } t < \tau \\ h(t), & \text{otherwise} \end{cases} \quad v(t) = \begin{cases} h(t), & \text{if } t < \tau \\ g(t), & \text{otherwise} \end{cases}$$

Then, we have $w(t) = \max\{u(t), v(t)\}$ and the convexity of w follows from the convexity of u and v by Lemma A.2. $\square$

A function real f is *locally convex* if every $t \in \mathbb{R}$ has an open neighborhood $N(t)$ such that the restriction of f to $N(t)$ is convex.

Lemma A.4. *Every locally convex function is convex.*

Proof. Fix $r, t \in \mathbb{R}$ and $\alpha \in [0, 1]$. Consider a finite family of intervals covering $[r, t]$ such that f is convex over each of the intervals. Such a family exists by compactness of $[r, t]$. Let 2ε be a Lebesgues number for this covering: for all $s \in [r, t]$ the interval $[s - \varepsilon, s + \varepsilon]$ is contained in one of the intervals of the covering. Consider a sequence $r = s_1 < s_2 < \dots < s_n = t$ such that $s_{i+1} - s_i \leq \varepsilon$ for all $i \in [1, n - 1]$. Let g be the piecewise-linear function over $[s, t]$ defined by $g(s_i) = f(s_i)$ for all $i \in [1, n]$. The graph of g is convex and above the graph of f . It follows that $f((1 - \alpha)r + \alpha t) \leq g((1 - \alpha)r + \alpha t) \leq (1 - \alpha)f(r) + \alpha f(t)$. $\square$

Proposition A.5. *The function d is convex.*

Proof. Let $\sigma(t)$ be the shortest path from $A(t)$ to $B(t)$, and let $\hat{P}$ denote the set of lifts of obstacles points in P to the universal cover $\hat{\mathcal{E}}$. Consider the sequence $V(t)$ of points in $\hat{P}$ where $\sigma(t)$ bends. As t increases from 0 to 1, the sequence $V(t)$ evolves as follows.

- (1) If $V(t)$ is empty, then $\sigma(t)$ is a line segment and $V(t)$ remains empty as t increases until $\sigma(t)$ hits a point in $\hat{P}$ (or possibly two points simultaneously). Just after this event, $V(t)$ contains that point (or points).
- (2) If $V(t)$ is non-empty, then $\sigma(t)$ comprises a first segment $r(t)$ pivoting around the first vertex in $V(t)$ and a last segment $s(t)$ pivoting around the last vertex in $V(t)$. The rest of $\sigma(t)$ remains fixed until one of the following events occurs:
 - either $r(t)$ or $s(t)$ hits a vertex of T ;
 - $r(t)$ becomes aligned with the next segment along $\sigma(t)$; or
 - $s(t)$ becomes aligned with the previous segment along $\sigma(t)$.

Each of these events changes the sequence $V(t)$ by adding or removing at most two vertices; at all other times, $V(t)$ is constant.

The first type of event can happen only once during the motion. The same can be said of the reverse event when $r(t)$ becomes aligned with $s(t)$ and $V(t)$ becomes empty. All other events correspond to $r(t)$ or $s(t)$ becoming aligned with a segment defined by two points in $\hat{P}$. For example, if $r(t)$ hits $\hat{p} \in \hat{P}$, then $r(t)$ is aligned with $\hat{p}$ and the first vertex in $V(t)$. It follows that the number of events is finite.

Between two consecutive events, the length of $\sigma(t)$ is a sum of functions that are either constant or of the type $\|A(t) - B(t)\|$, $\|A(t) - v\|$, or $\|B(t) - w\|$ for some pivoting vertices v and w . Each of these functions is convex by Lemma A.1, hence d is convex between two consecutive events. Due to Lemma A.4 it suffices to prove that d is locally convex at each event time to prove that it is convex over the whole unit interval.

Consider an event time τ . If $V(t)$ contains the same vertex $\hat{p}$ just before and just after τ , then we can split $\sigma(t)$ into two parts $q(t)$ from $A(t)$ to $\hat{p}$ and $q'(t)$ from $\hat{p}$ to $B(t)$. The first segment $s(t)$ of $q(t)$ will pivot around a certain vertex $\hat{p}_1$ for $t < \tau$ and around a possibly different vertex $\hat{p}_2$ for $t > \tau$. If $\hat{p}_1 = \hat{p}_2$, then the length of $q(t)$ can be decomposed into the sum of closed forms has in Lemma A.1, hence it is locally convex at τ . If $\hat{p}_1 \neq \hat{p}_2$, then we can extend the two closed forms of the length $q(t)$ before and after τ and apply Lemma A.2 to see that the length of $q(t)$ is locally convex at τ . The same can be said of the length of $q'(t)$, which implies the local convexity of d at τ . If $V(t)$ has no vertex in common before and after τ it must be that $\sigma(\tau)$ is the straight line segment between $A(\tau)$ and $B(\tau)$. The local convexity of d then follows from Corollary A.3 where $h(t)$ is the length of the line segment between $A(t)$ and $B(t)$. $\square$

Proposition A.5 implies that the bivariate shortest-path distance function $D(u, v)$ between $A(u)$ and $B(v)$ is also convex, as follows. For any $u, u', v, v', t \in [0, 1]$ we denote by $d_{u,v,u',v'}(t)$ the shortest-path distance between $A((1-t)u + tu')$ and $B((1-t)v + tv')$. Said differently, we put $d_{u,v,u',v'}(t) = D((1-t)(u, v) + t(u', v'))$. Proposition A.5 implies that the univariate function $d_{u,v,u',v'}$ is convex. It follows that

$$\begin{aligned} D((1-t)(u, v) + t(u', v')) &= d_{u,v,u',v'}((1-t) \cdot 0 + t \cdot 1) \\ &\leq (1-t)d_{u,v,u',v'}(0) + td_{u,v,u',v'}(1) \\ &= (1-t)D(u, v) + tD(u', v'), \end{aligned}$$

which expresses the convexity of D .

The convexity of D immediately implies the following corollary, which is the last ingredient necessary for correctness of the algorithm described in Section 5. (See Section 5 for an explanation of the notation.)

Corollary A.6. *For all integers i and j , the restriction of d_h to any grid cell $C_{i,j}$ is convex.*

ERIN WOLF CHAMBERS, DEPARTMENT OF COMPUTER SCIENCE, UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN, USA

E-mail address: erinwolf@uiuc.edu

ÉRIC COLIN DE VERDIÈRE, CNRS, DÉPARTEMENT D'INFORMATIQUE, ÉCOLE NORMALE SUPÉRIEURE, PARIS, FRANCE

E-mail address: Eric.Colin.de.Verdiere@ens.fr

JEFF ERICKSON, DEPARTMENT OF COMPUTER SCIENCE, UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN, USA

E-mail address: jeffe@cs.uiuc.edu

SYLVAIN LAZARD, INRIA LORRAINE, LORIA, NANCY, FRANCE

E-mail address: lazard@loria.fr

FRANCIS LAZARUS, CNRS, LABORATOIRE LIS, INSTITUT NATIONAL POLYTECHNIQUE DE GRENOBLE GRENOBLE, FRANCE

E-mail address: Francis.Lazarus@lis.inpg.fr

SHRIPAD THITE, CENTER FOR THE MATHEMATICS OF INFORMATION, CALIFORNIA INSTITUTE OF TECHNOLOGY, USA

E-mail address: sthite@win.tue.nl