Improved space bounds for Fréchet distance queries

Maike Buchin[1], Ivor van der Hoog[2], Tim Ophelders[3], Rodrigo I. Silveira[4], Lena Schlipf[5], and Frank Staals[2]

1 **Ruhr University Bochum**
`maike.buchin@rub.de`
2 **Utrecht University**
`[i.d.vanderhoog,f.staals]@uu.nl`
3 **Michigan State University**
`ophelder@egr.msu.edu`
4 **Universitat Politècnica de Catalunya**
`rodrigo.silveira@upc.edu`
5 **Universität Tübingen**
`schlipf@informatik.uni-tuebingen.de`

Abstract

We revisit a data structure from de Berg, Mehrabi and Ophelders that can store a polygonal curve P, such that for any directed horizontal query segment pq one can compute the Fréchet distance between P and pq in polylogarithmic time. We extend their analysis of the geometric constructions that can realize the Fréchet distance between P and pq and prove that in fact, their data structure only requires $O(n^{3/2})$ space, as opposed to the $O(n^2)$ space originally believed.

1 Introduction

Comparing the shapes of polygonal curves is an important task that arises in many contexts such as GIS applications [2, 4], protein classification [8], curve simplification [3], curve clustering [1] and even speech recognition [9]. Within computational geometry, there are two well studied distance measures for polygonal curves: the Hausdorff and the Fréchet distance. In this paper, we consider the problem of preprocessing a polygonal curve P of n edges in the plane, such that given a query segment pq traversed from p to q, the Fréchet distance between pq and P can be computed in sublinear time. The curve may self-intersect and pq may intersect P. For this version, the proofs have been omitted.

We give an overview of recent results that preprocess a polygonal chain P in order to compute the Fréchet distance between P and a query segment pq. Driemel and Har-Peled [6] studied how to process a polygonal chain P, such that given a query segment pq one can compute a $(1+\epsilon)$-approximation of the Fréchet distance between P and pq in $O(\epsilon^{-2} \log n \log \log n)$ time. Gudmundsson, Mirzanezhad, Mohades and Wenk [7] consider the Fréchet distance between polygonal curves where each curve contains only edges which are long when compared to the Fréchet distance between the two curves. A corollary of their result is that they can preprocess a curve P such that given a query segment pq one can compute the exact Fréchet distance between P and pq in $O(\log^2 n)$ time, provided that the length of pq and each edge of P is relatively large compared to this distance. Recently [5], de Berg, Mehrabi and Ophelders presented a paper in which they preprocess a curve P in $O(n \log^2 n)$ time, using $O(n^2)$ space, such that for any *horizontal* query segment pq one can compute the Fréchet distance between P and pq in $O(\log^2 n)$ time. In this paper we extend these results, by showing, via a more involved analysis, that the data structure by de Berg, Mehrabi and Ophelders requires only $O(n^{3/2})$ space.

36th European Workshop on Computational Geometry, Würzburg, Germany, March 16–18, 2020.
This is an extended abstract of a presentation given at EuroCG'20. It has been made public for the benefit of the community and should be considered a preprint rather than a formally reviewed paper. Thus, this work is expected to appear eventually in more final form at a conference with formal proceedings and/or in a journal.

1.1 Preliminaries.

The *directed Hausdorff distance* is a distance measure between any two point sets. Let A and C be two point sets, we define the *directed Hausdorff distance* from A to C as $D_{\overrightarrow{H}}(A, C) := \sup_{a \in A} \inf_{c \in C} ||a - c||$. The Fréchet distance is a distance measure between any two curves, which is commonly explained with the following "leash" analogy: consider two curves in the plane P and Q where a person walks along curve P and a dog walks along curve Q, and neither of them is allowed to walk backwards. Then what is the minimum length that a leash between the person and the dog needs to have? Formally, we denote by $\alpha : [0,1] \to P$ a (non-strict) monotone traversal of P in the time interval $[0,1]$ and we denote by $\beta : [0,1] \to Q$ an identical traversal of Q. The Fréchet distance between P and Q is the infimum over all choices of α and β, of the maximal distance realized during the traversal:

$$D_F(P, Q) = \inf_{\substack{\alpha:[0,1]\to P \\ \beta:[0,1]\to Q}} \left\{ \max_{t \in [0,1]} ||\alpha(t) - \beta(t)|| \right\}$$

De Berg, Mehrabi and Ophelders consider the scenario where P is a polygonal curve $P = (v_0, v_1, \ldots, v_n)$, where each vertex v_i is a point in the plane, and Q is a horizontal segment pq in the plane, at height y with p left of q. Their data structure uses the notion of *backward pairs*: any ordered pair of vertices (v_i, v_j) with $i \leq j$ in P form a *backward pair* if v_j lies further to the left than v_i. Note that a vertex of P can be be part of many backward pairs, even if its outgoing edge is pointed along the directed edge from p to q (Figure 1). There are $O(n^2)$ backward pairs in total. We denote the set of backward pairs by $\mathcal{B}(P)$. De Berg, Mehrabi and Ophelders note that a backward pair v_i, v_j has the following effect on the Fréchet distance. For the point on the query segment minimizing the distance to the farthest of v_i and v_j, that distance is a lower bound on the Fréchet distance. This point on the query segment lies either on the bisector B_{v_i,v_j} between v_i and v_j, or it is the point on the query segment closest to v_i or v_j. We define the distance function $F(y, v_i, v_j) := ||v_i - \ell_y \cap B_{v_i,v_j}||$ from v_i to the closest point that lies both on the bisector of v_i and v_j, and on the horizontal line ℓ_y at height y. De Berg, Mehrabi and Ophelders observe the following:

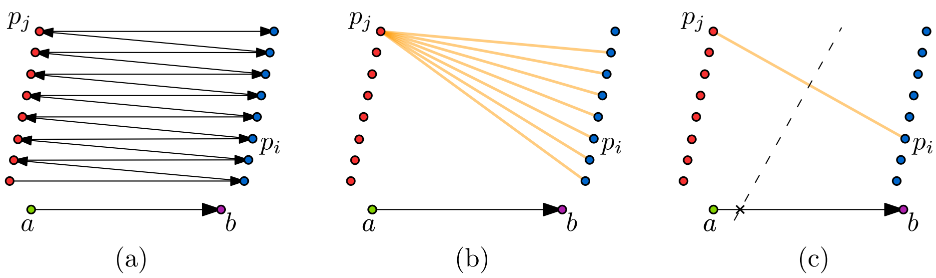

Figure 1 (Left) a polygonal curve that zigzags and a query segment from left to right. (Middle) The red vertex v_j forms a backward pair with all but one blue vertex. (Right) For a fixed backward pair (v_i, v_j), we consider the point of intersection between their bisector and pq (cross) and we are interested in the distance between that point and either v_i or v_j.

▶ Observation 1 (From [5]). For all $(v_i, v_j) \in \mathcal{B}(P)$, for any y, if the intersection between ℓ_y and B_{v_i,v_j} lies in the rectangle spanned by v_i and v_j, then $F(y, v_i, v_j)$ is a hyperbolic segment with absolute slope smaller than 1. Otherwise, it is a line with slope 1 or -1 (Fig 2).

They prove that for any backward pair (v_i, v_j), the value $F(y, v_i, v_j)$ is a lower bound for the Fréchet distance between pq and P if pq has height y. Note that the Fréchet distance is also lower-bounded by the distance between (1) p and the start of P, (2) q and the end of P and (3) by the directed Hausdorff distance from P to pq. Specifically, de Berg, Mehrabi and Ophelders prove that the Fréchet distance is the maximum of any of these lower bounds:

$$D_F(P, pq) = \max \left\{ ||v_0 - p||, \quad ||v_n - q||, \quad d_{\overrightarrow{H}}(P, pq), \quad \max_{(v_i,v_j) \in \mathcal{B}(P)} F(y, v_i, v_j) \right\}$$

In this paper, we perform a deeper analysis on the data structure of de Berg, Mehrabi and Ophelders that computes these four terms, and give better bounds on its space complexity.

2 A data structure for horizontal segments

For any polygonal curve $P = (v_0, v_1, \ldots, v_n)$ and for any segment pq the distance $||v_0 - p||$ and $||v_n - q||$ can be computed in constant time. De Berg, Mehrabi and Ophelders present a linear space data structure that can compute the directed Hausdorff distance from P to pq in $O(\log^2 n)$ time. To compute the remaining component of the lower bound on the Fréchet distance they provide a data structure with $O(\log^2 n)$ query time whose space is linear in the number of backward pairs. They obtained this as follows: they consider the function $F(y, v_i, v_j)$ for every backward pair $(v_i, v_j) \in \mathcal{B}(P)$ and compute the upper envelope of all these functions. They argue that the upper envelope is linear in the number of backward pairs, which gives a quadratic upper bound on the space of the data structure.

We extend their analysis with the following observation: consider a vertex $v_i \in P$, the set of vertices $V_i := \{v' \mid (v_i, v') \in \mathcal{B}(P)\}$ and the upper envelope of all $\{F(y, v', v_i) \mid v' \in V_i\}$ (Figure 3). We define $L_i(y) := \min_{v_j \in V_i} (\ell_y \cap B_{v_{ij}, v_i})_x$ as the *left chain* of V_i and the function $F_i^L(y) := ||\ell_y \cap L_i(y) - v_i||$ as the distance from a point on L_i at height y to v_i. Note that the points $(L_i(y), y)$, that for simplicity we will denote L_i, correspond to the "left envelope" in the arrangement of bisectors. We use the term *chain* to avoid confusion with the upper envelope of the distances functions F_i^L, which we denote by $F^L(y) := \max_i \{F_i^L(y)\}$. Analogously, we define the right chain R_i, its corresponding distance function F_i^R, and the upper envelope $F^R(y) = \max_i F_i^R(y)$. It then follows that $F(y) = \max\{F^L(y), F^R(y)\}$ and thus:

▶ Observation 2. If the complexity of $F^L(y)$ and $F^R(y)$ are both upper bounded by $O(n^{3/2})$ then the complexity of $F(y)$ is upper bound by $O(n^{3/2})$.

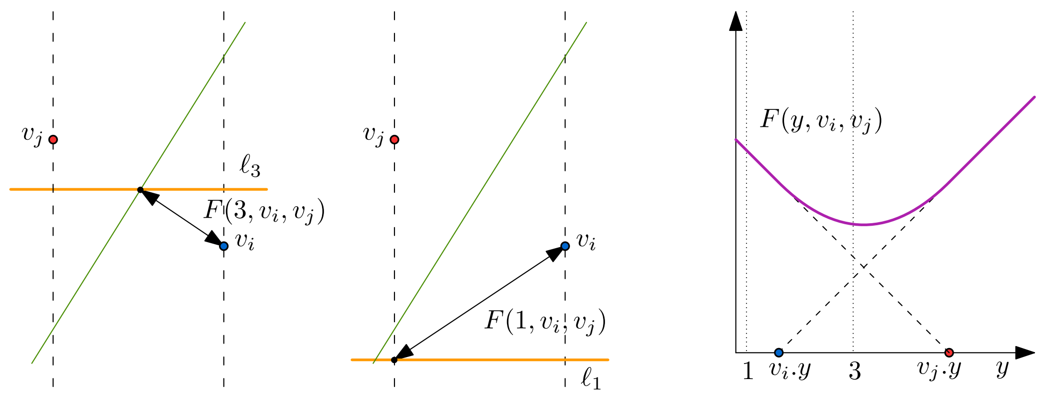

Figure 2 (Left) The backward pair (v_i, v_j), a line at height 3 and the distance $F(3, v_i, v_j)$. (Middle) A line at height 1 and the distance $F(1, v_i, v_j)$. (Right) The function $F(y, v_i, v_j)$.

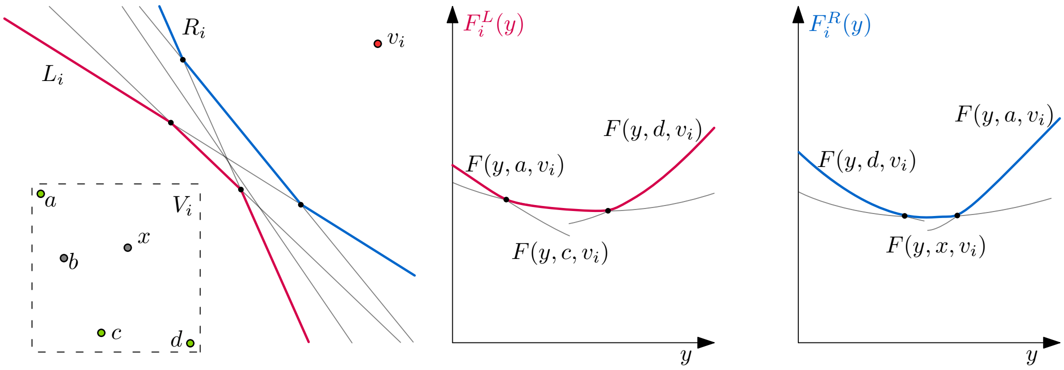

Figure 3 (Left) A point v_i and the set of vertices V_i that form a backward pair with v_i, together with the left and the right chains; vertices of V_i^* in green. (Middle) The envelope $F_i^L(y)$ is a piecewise curve with three pieces. (Right) The envelope $F_i^R(y)$ is also a piecewise curve with three pieces.

In the remainder of this section, we bound the complexity of $F^L(y)$, the proof for $F^R(y)$ is analogous. We denote by V_i^* the subset of V_i, such that $v' \in V_i^*$ if and only if a piece of the bisector B_{v',v_i} appears on L_i, and the distance function F_i^L is maximal there, i.e., $F_i^L(y) = F^L(y)$. The proofs of the following observations are deferred to the appendix.

▶ **Lemma 2.1.** *For any vertex v_i, the vertices V_i^* lie in convex position and the left chain L_i is a convex chain where the clockwise ordering of the bisectors on L_i is identical to the clockwise ordering of the vertices of V_i^*.*

▶ **Lemma 2.2.** *Let $v \in V_i^* \cap V_j^*$ be a point that forms a backward pair with both v_i and v_j. The bisectors B_{v,v_i} and B_{v,v_j} intersect B_{v_i,v_j} at a single point.*

▶ **Lemma 2.3.** *For any v_i, v_j, the chains L_i and L_j intersect B_{v_i,v_j} in a common point c. Moreover, the bisectors that intersect in c correspond to the same point $v \in V_i^* \cap V_j^*$.*

▶ **Corollary 2.4.** *For any v_i, v_j, for any horizontal line ℓ_y of height y, the line ℓ_y intersects L_i and L_j on the same side of the bisector B_{v_i,v_j}.*

▶ **Lemma 2.5.** *(Illustrated by Figure 4) For any two vertices v_i, v_j consider their Voronoi diagram. If the set $V_i^* \cap V_j^*$ contains at least four elements, then there is an edge (which is not a halfline) on L_i that is entirely contained in the Voronoi cell of v_i, or there is an edge (which is not a halfline) on L_j that is entirely contained in the Voronoi cell of v_j.*

Recall that the upper envelope $F^L(y)$ begins and ends with halflines of slope -1/1. All other bisectors generate at most one hyperbolic segment on $F^L(y)$. In the following lemma, we consider all pairs $v_i, v_j \in P$ that do not participate in these two halflines.

▶ **Lemma 2.6.** *For any v_i, v_j that do not participate in a halfline of $F^L(y)$, the set $V_i^* \cap V_j^*$ contains at most three vertices.*

Proof. Suppose for the sake of contradiction that $V_i^* \cap V_j^*$ contains four vertices (a, b, c, d) in counter-clockwise order, with v_i below the bisector B_{v_i,v_j}. By Lemma 2.5, we can assume without loss of generality that B_{a,v_i} and B_{a,v_j} are entirely contained in the Voronoi cell of v_i. Because of the assumption of the lemma, and because $a \in V_i^*$, the bisector B_{a,v_i} generates a hyperbolic segment on $F^L(y)$. Next we prove that this hyperbolic segment cannot appear on the upper envelope of $\{F_i^L(y), F_j^L(y)\}$ which contradicts the assumption that a is in V_i^*.

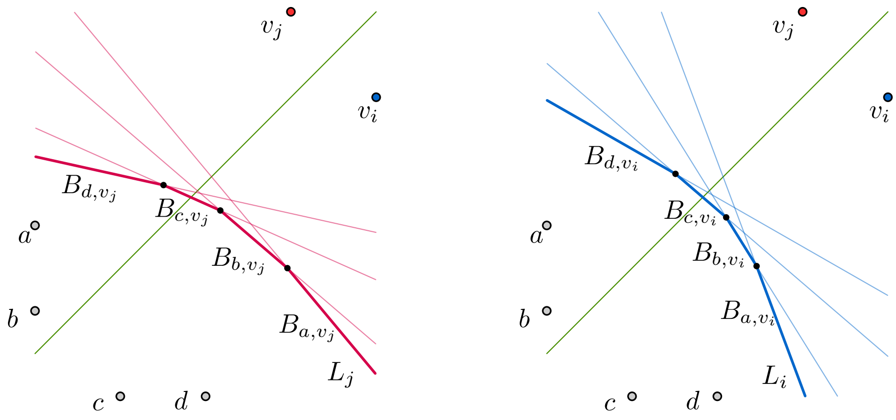

Figure 4 Two vertices v_i and v_j and their Voronoi diagram. We drew four points $a, b, c, d \in V_i^* \cap V_j^*$. Left we see the bisectors between these points and V_j^*, and their left chain L_j. Right we see the bisectors between these points and V_i^*, and their left chain L_i.

We denote the point of intersection between ℓ_y and B_{a,v_i} by S_i. Consider the point of intersection between ℓ_y and L_j and denote it by S_j. Observe that S_j must also lie in the Voronoi cell of v_i (Corollary 2.4). We split the proof in three cases depending on S_i. We prove that the first case cannot exists. The remaining cases are illustrated in Figure 5.

Case S_j lies left of S_i. Following Observation 1 and the assumption that v_i and v_j are both the rightmost points of the backward pairs, S_i must lie left of v_i and S_j must lie left of v_j. Thus the distance $\|v_i - S_i\|$ is smaller than the distance $\|v_i - S_j\|$. But since S_j lies in the Voronoi cell of v_i, the distance $\|v_i - S_j\|$ is smaller than the distance $\|v_j - S_j\|$. This implies that the distance $\|v_i - S_i\|$ is smaller than the distance $\|v_j - S_j\|$ which contradicts the assumption that for this y-coordinate $F(y, v_i, a)$ lies on $F_i^L(y)$.

Case S_j lies right of S_i and on the bisector B_{a,v_j}. We note, that since S_i and S_j lie on bisectors, that $\|v_i - S_i\| = \|a - S_i\|$ and $\|v_j - S_j\| = \|a - S_j\|$ and thus per assumption $\|a - S_j\| > \|a - S_i\|$. However, S_j and S_i both lie to the right of a and they have the same y-coordinate. Since S_j lies further to the right than S_i, we know that $\|a - S_i\| < \|a - S_j\|$ which is a contradiction.

Case S_j lies right of S_i and not on the bisector B_{a,v_j}. We say that S_j lies on B_{b,v_j} but the argument works for any bisector further than a in the ordering, the argument is illustrated by Figure 6. We pinpoint the location of the point a with three observations:

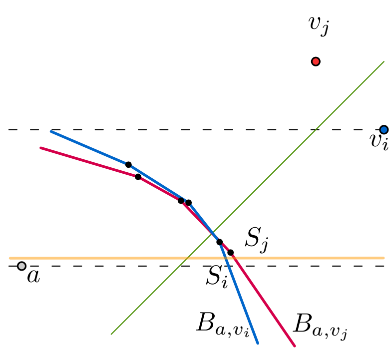

Figure 5 A horizontal line at a y-coordinate in yellow. The points of intersection S_j lies right of S_i and the intersection points originate from the bisectors B_{a,v_i} and B_{a,v_j}.

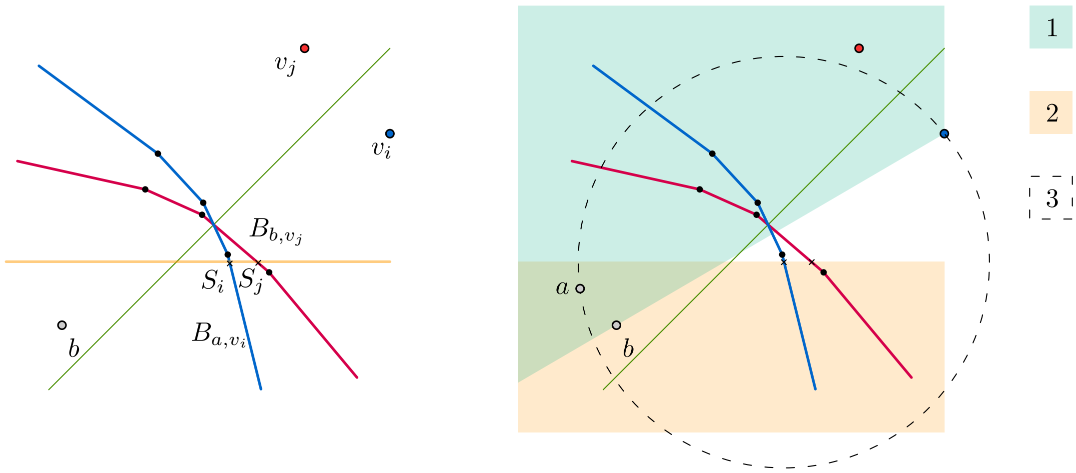

Figure 6 (left) Case S_j left of S_i *and* S_j lies on B_{b,v_j}. (right) Three regions where a can lie.

1. The bisector B_{a,v_i} is the last bisector in the clockwise ordering of the left chain L_i, therefore the point a must lie above the line through b and v_i (Lemma 2.1).
2. The point a must lie on the opposite side of ℓ_y with respect to v_i since B_{a,v_i} realizes a hyperbolic segment on the upper envelope $F^L(y)$ and all hyperbolic segments come from intersection points that lie in the rectangle bound by a and v_i (Observation 1).
3. Per assumption, $\|b - S_j\| < \|a - S_i\|$. The point a must lie on the circle centered at S_i with radius $\|a - S_i\|$. Combining our assumption with the fact that S_j lies right of S_i, we know that the point b cannot lie left of this circle. Since a lies clockwise of b with respect to V_i, it now follows that a lies left of b.

These three observations imply that the bisector B_{a,v_j} intersects the line ℓ_y left of S_j which contradicts the assumption that S_j lies on L_j. ◀

▶ **Lemma 2.7.** *The function F^L has complexity $O(n^{3/2})$.*

Proof. We upper bound the number of hyperbolic segments on F^L by bounding the number of elements in $\cup_i V_i^*$. Let v_a, v_b be the (at most) two rightmost vertices that participate in a backward pair whose bisector is a halfline of slope 1/-1 on $F^L(y)$. The sets V_a^* and V_b^* contain at most $O(n)$ elements. Now consider the bipartite graph $G = (L \cup R \backslash \{v_a, v_b\}, E)$ in which L and R are two copies of the vertices in $P \backslash \{v_a, v_b\}$. There is an edge between $v_j \in L$ and $v_i \in R$ if and only if $v_j \in V_i^*$. By Lemma 2.6, the graph G is $K_{4,2}$-free, and thus has at most $O(n^{3/2})$ edges [10, Theorem 4.5.2]. Since every edge corresponds to a (relevant) backward pair, it follows that the number of elements in $\cup_i V_i^*$ and therefore the number of hyperbolic segments on F^L is bounded by $O(n^{3/2})$. ◀

By Observation 2 the function F thus has complexity at most $O(n^{3/2})$ as well. Therefore, the data structure of de Berg et al. [5] uses at most $O(n^{3/2})$ space. We conclude:

▶ **Theorem 2.8.** *Given a polygonal curve P in $\mathbb{R}^2$ with n vertices, there is a data structure of size $O(n^{3/2})$, that can be built in $O(n^2 \log^2 n)$ time, that can report the Fréchet distance between P and a horizontal query segment in $O(\log^2 n)$ time.*

Acknowledgements. This research was initiated during the Dagstuhl seminar 19352.

References

1 Pankaj K Agarwal, Sariel Har-Peled, Nabil H Mustafa, and Yusu Wang. Near-linear time approximation algorithms for curve simplification. *Algorithmica*, 42(3-4):203–219, 2005.

2 H. Alt, A. Efrat, G. Rote, and C. Wenk. Matching planar maps. *Journal of algorithms*, 49(2):262–283, 2003.

3 K. Buchin, M. Buchin, J. Gudmundsson, M. Löffler, and J. Luo. Detecting commuting patterns by clustering subtrajectories. *International Journal of Computational Geometry & Applications*, 21(03):253–282, 2011.

4 Kevin Buchin, Maike Buchin, Marc Van Kreveld, Maarten Löffler, Rodrigo I Silveira, Carola Wenk, and Lionov Wiratma. Median trajectories. *Algorithmica*, 66(3):595–614, 2013.

5 Mark de Berg, Ali D Mehrabi, and Tim Ophelders. Data structures for fréchet queries in trajectory data. In *CCCG*, pages 214–219, 2017.

6 A. Driemel and S. Har-Peled. Jaywalking your dog: Computing the fréchet distance with shortcuts. *SIAM Journal on Computing*, 42(5):1830–1866, 2013.

7 J. Gudmundsson, M. Mirzanezhad, A. Mohades, and C. Wenk. Fast fréchet distance between curves with long edges. In *Proceedings of the 3rd International Workshop on Interactive and Spatial Computing*, IWISC '18, pages 52–58. ACM, 2018.

8 M. Jiang, Y. Xu, and B. Zhu. Protein structure–structure alignment with discrete Fréchet distance. *Journal of bioinformatics and computational biology*, 6(01):51–64, 2008.

9 S. Kwong, QH He, K. Man, KS Tang, and CW Chau. Parallel genetic-based hybrid pattern matching algorithm for isolated word recognition. *International Journal of Pattern Recognition and Artificial Intelligence*, 12(05):573–594, 1998.

10 J. Matoušek. *Lectures on discrete geometry*, volume 108. Springer, 2002.

A Omitted proofs

Lemma 2.1: For any vertex v_i, the vertices V_i^* lie in convex position and the left chain L_i is a convex chain where the clockwise ordering of the bisectors on L_i is identical to the clockwise ordering of the vertices of V_i^*.

Proof of Lemma 2.1. Any face in an arrangement of lines is convex. So in particular L_i is a convex chain. Consider traversing L_i from top to bottom, and let $p_1, .., p_k$ be the corresponding points in V_i. At every vertex of L_i we make a right turn, going from B_{p_j,v_i} to B_{p_{j+1},v_i}. Therefore, p_{j+1} lies above the line through v_i and p_j. Hence, $p_1, .., p_j$ corresponds to a clockwise traversal of the points, ordered around v_i. All that remains to argue is that p_{j+1} lies left of the line through p_j and p_{j-1}. Assume, by contradiction that p_{j+1} lies right of this line. Since p_{j+1} contributes an edge to L_i, the intersection point c of B_{p_j,v_i} and B_{p_{j+2},v_i} lies right of the bisector B_{p_{j+1},v_i}. Hence, the distance $\|p_{j+1}-c\| > \|v-c\| = \|p_j-c\| = \|p_{j+2}-c\|$. However, since p_{j+1} lies right of the line through p_j and p_{j+2}, and in between p_j and p_{j+2} in the clockwise ordering around v_i, it follows that p_{j+1} is closer to c than p_j or p_{j+2}, and thus $\|p_{j+1} - c\| < \|v - c\|$. Contradiction. The lemma follows. ◀

Lemma: 2.2: Let $v \in V_i^* \cap V_j^*$ be a point that forms a backward pair with both v_i and v_j. The bisectors $B_{v,i}$ and $B_{v,j}$ intersect B_{v_i,v_j} at a single point.

Proof of Lemma: 2.2. Let $B_{v,v_i} \cap B_{v_i,v_j} = x$. Then $\|x - v\| = \|x - v_i\| = \|x - v_j\|$, and hence $x \in B_{v_i,v_j}$. So all three lines B_{v,v_i}, B_{v,v_j}, and B_{v_i,v_j} intersect at a single point. ◀

Lemma 2.3: For any v_i, v_j, the chains L_i and L_j intersect B_{v_i,v_j} in a common point c. Moreover, the bisectors that intersect in c correspond to the same point $v \in V_i^* \cap V_j^*$.

Proof of Lemma 2.3. Denote by c_v the intersection point of the bisector B_{v,v_i} and B_{v_i,v_j}, for $v \in V_i^* \cap V_j^*$. Since (v, v_i) is a backward pair, we observe that all points on the bisectors B_{v_i,v_j} right of c_v must be closer to v_i than to v. Consider the leftmost such intersection point c_v among all points $v \in V_i^* \cap V_j^*$. It follows that all other points $c_{v'}$, with $v' \in V_i^* \cap V_j^*$ lie right of B_{v,v_i}, and hence these points cannot appear in the leftmost chain L_i. Hence, L_i intersects B_{v_i,v_j} at most once in c_v. Similarly, we obtain that L_j intersects B_{v_i,v_j} at most once, namely at point c_k. By Lemma 2.2 it follows that B_{v,v_i} intersects B_{v_i,v_j} at the same point as B_{v,v_j}. It then follows that $c_v = c_k$, and thus L_i and L_j intersect B_{v_i,v_j} at the same point $c := c_v = c_k$. ◀

Corollary 2.4: For any v_i, v_j. for any horizontal line ℓ_y of height y, the line ℓ_y intersects L_i and L_j on the same side of the bisector B_{v_i,v_j}.

Proof of Corollary 2.4. Suppose that the corollary does not hold and assume without loss of generality that L_j intersects ℓ_y to the left of L_i. Then L_j must intersect B_{v_i,v_j} below y. Similarly the left chain L_i must intersect B_{v_i,v_j} above y, contradicting Lemma 2.3. ◀

Lemma 2.5: For any two vertices v_i, v_j consider their Voronoi diagram. If the set $V_i^* \cap V_j^*$ contains at least four elements, then there is an edge (which is not a halfline) on L_i that is entirely contained in the Voronoi cell of v_i, or there is an edge (which is not a halfline) on L_j that is entirely contained in the Voronoi cell of v_j.

Proof of Lemma 2.5. The bisector B_{v_i,v_j} intersects only one edge s of L_i. So if we consider the four edges of L_i ordered by (a, b, c, d), then following Lemma 2.1 there is a vertex s in that ordering such that the edges of L_i corresponding to $[a, \ldots, s)$ lie on one side of B_{v_i,v_j} and $(s, \ldots d]$ on the other. Following Lemma 2.3, L_j also intersects B_{v_i,v_j} in s and thus following Lemma 2.1 the edges of L_j corresponding to $[a, \ldots, s)$ lie on one side of B_{v_i,v_j} and $(s, \ldots d]$ on the other. It follows that either the edge given by the bisector-segments B_{b,v_i}, B_{b,v_j} lie in the same Voronoi cell or the edge corresponding to the bisector-segments B_{c,v_i}, B_{c,v_j} lie in the same Voronoi cell. ◀