Four Soviets Walk the Dog—with an Application to Alt's Conjecture

Kevin Buchin*   Maike Buchin†   Wouter Meulemans*   Wolfgang Mulzer‡

Abstract

Given two polygonal curves in the plane, there are many ways to define a notion of similarity between them. One measure that is extremely popular is the Fréchet distance. Since it has been proposed by Alt and Godau in 1992, many variants and extensions have been studied. Nonetheless, even more than 20 years later, the original $O(n^2 \log n)$ algorithm by Alt and Godau for computing the Fréchet distance remains the state of the art (here n denotes the number of vertices on each curve). This has led Helmut Alt to conjecture that the associated decision problem is 3SUM-hard.

In recent work, Agarwal *et al.* show how to break the quadratic barrier for the *discrete* version of the Fréchet distance, where one considers sequences of points instead of polygonal curves. Building on their work, we give a randomized algorithm to compute the Fréchet distance between two polygonal curves in time $O(n^2\sqrt{\log n}(\log \log n)^{3/2})$ on a pointer machine and in time $O(n^2(\log \log n)^2)$ on a word RAM. Furthermore, we show that there exists an algebraic decision tree for the decision problem of depth $O(n^{2-\varepsilon})$, for some $\varepsilon > 0$. This provides evidence that the decision problem may not be 3SUM-hard after all and reveals an intriguing new aspect of this well-studied problem.

1 Introduction

Shape matching is a fundamental problem in computational geometry, computer vision, and image processing. A simple version can be stated as follows: given a database $\mathcal{D}$ of shapes (or images) and a query shape S, find the shape in $\mathcal{D}$ that most resembles S. How-ever, before we can solve this problem, we first need to address a much more fundamental issue: what does it mean for two shapes to be similar? In the mathematical literature, there are many different notions of distance between two sets, a prominent example being the *Hausdorff distance*. Informally, the Hausdorff distance is defined as the maximal distance between two elements when every element of one set is mapped to the closest element in the other. It has the advantage of being simple to describe and easy to compute for discrete sets. In the context of shape matching, however, the Hausdorff distance often turns out to be unsatisfactory: it does not take the continuity of the shapes into account. There are well known examples where the distance fails to capture the similarity of shapes as perceived by human observers [4].

In order to address this issue, Alt and Godau introduced the Fréchet distance into the computational geometry literature [6, 37]. They argued that the Fréchet distance is better suited as a similarity measure, and they described an $O(n^2 \log n)$ time algorithm to compute it on a real RAM or pointer machine.[1] Since Alt and Godau's seminal paper, there has been a wealth of research in various directions, such as extensions to higher dimensions [5, 19, 22, 24, 29, 38], approximation algorithms [7, 8, 32], the geodesic and the homotopic Fréchet distance [25, 30, 33, 40], and much more [2, 12, 18, 21, 31, 42, 44, 45]. All known approximation algorithms make further assumptions on the curves, and only an $O(n^2)$-time approximation algorithm is known for arbitrary polygonal curves [20]. The Fréchet distance and its variants, such as dynamic time-warping [11], have found various applications, with recent work particularly focusing on geographic applications such as map-matching tracking data [13, 49] and moving objects analysis [15, 16, 39].

Despite the large amount of published research, the original algorithm by Alt and Godau has not been improved, and the quadratic barrier on the running time of the associated decision problem remains unbroken. If we cannot improve on a quadratic bound for a geometric problem despite many efforts, a possible culprit may

---

[1]For a brief overview of the different computational models in this paper, refer to Appendix A.

be the underlying 3SUM-hardness [36]. This situation induced Helmut Alt to make the following conjecture.[2]

CONJECTURE 1.1. (ALT'S CONJECTURE) *Let P, Q be two polygonal curves in the plane. Then it is* 3SUM-*hard to decide whether the Fréchet distance between P and Q is at most* 1.

Here, 1 can be considered as an arbitrary constant, which can be changed to any other bound by scaling the curves. So far, we know only that this problem takes $\Omega(n \log n)$ steps in the algebraic computation tree model [17].

Recently, Agarwal *et al.* [1] showed how to achieve a subquadratic running time for the *discrete* version of the Fréchet distance. Their approach relies on reusing small parts of the solution. We follow a similar approach based on the so-called Four-Russian-trick which precomputes small recurring parts of the solution and uses table-lookup to speed up the whole computation.[3] The result by Agarwal *et al.* is stated in the word RAM model of computation. They ask whether their result can be generalized to the case of the original (continuous) Fréchet distance.

Our contribution. We address the question by Agarwal *et al.* and show how to extend their approach to the Fréchet distance between two polygonal curves. Our algorithm requires total expected time $O(n^2 \sqrt{\log n} (\log \log n)^{3/2})$. This is the first algorithm with a running time of $o(n^2 \log n)$ and constitutes the first improvement for the general case since the original paper by Alt and Godau [6]. To achieve this running time we give the first subquadratic algorithm for the decision problem of the Fréchet distance. We emphasize that these algorithms run on a real RAM/pointer machine and do not require any bit-manipulation tricks. Therefore, our results are more in the line of Chan's recent subcubic-time algorithms for all-pairs-shortest paths [26,27] or recent subquadratic-time algorithms for min-plus convolution [14] than the subquadratic-time algorithms for 3SUM [10].

If we relax the model to allow constant time table-lookups, the running time can be improved to be almost quadratic, up to $O(\log \log n)$ factors. As in Agarwal *et al.*, our results are achieved by first giving a faster algorithm for the decision version, and then performing an appropriate search over the critical values to solve the optimization problem.

Finally, we address Alt's conjecture by showing that *non-uniformly*, the Fréchet distance can be computed in subquadratic time. More precisely, we prove that the decision version of the problem can be solved by an algebraic decision tree [9] of depth $O(n^{2-\varepsilon})$, for some fixed $\varepsilon > 0$. It is conjectured that no such decision tree exists for 3SUM [46] and an $\Omega(n^2)$ lower bound is known in a restricted linear decision tree model [3,34]. Our result therefore provides strong evidence that the Fréchet distance problem is not 3SUM-hard.

It is, however, not clear how to implement this decision tree in subquadratic time, which hints at a discrepancy between the decision tree and the uniform complexity of the Fréchet problem. This puts it into the illustrious company of such notorious problems as SORTING $X + Y$ [35], MIN-PLUS-CONVOLUTION [14], or finding the Delaunay triangulation for a point set that has been sorted in two orthogonal directions [23]. We find that this aspect of the Fréchet distance is highly intriguing and deserves further study.

2 Preliminaries and Basic Definitions

Let P and Q be two polygonal curves in the plane, defined by their vertices $p_0, p_1, \ldots, p_n$ and $q_0, q_1, \ldots, q_n$.[4] Depending on the context, we interpret P and Q either as sequences of n edges, or as continuous functions $P, Q : [0, n] \to \mathbb{R}^2$. In the latter case, we have $P(i + \lambda) = (1 - \lambda) p_i + \lambda p_{i+1}$ for $i = 0, \ldots, n - 1$ and $\lambda \in [0, 1]$, and similarly for Q. Let Ψ be the set of all continuous and nondecreasing functions $\sigma : [0, n] \to [0, n]$ with $\sigma(0) = 0$ and $\sigma(n) = n$. The *Fréchet distance* between P and Q is defined as

$$d_F(P, Q) := \inf_{\sigma \in \Psi} \max_{x \in [0,n]} \| P(x) - Q(\sigma(x)) \|,$$

where $\| \cdot \|$ denotes the Euclidean distance.

The classic approach to computing $d_F(P, Q)$ uses the *free-space diagram* $\mathrm{FSD}(P, Q)$. It is defined as

$$\mathrm{FSD}(P, Q) := \{ (x, y) \in [0, n]^2 \mid \| P(x) - Q(y) \| \le 1 \}.$$

In other words, $\mathrm{FSD}(P, Q)$ is the subset of the joint parameter space for P and Q where the corresponding points on the curves have distance at most 1, see Fig. 1.

The structure of $\mathrm{FSD}(P, Q)$ is easy to describe. Let $R := [0, n]^2$ be the ground set. We subdivide R into n^2 *cells* $C(i, j) = [i, i+1] \times [j, j+1]$, for $i, j = 0, \ldots, n-1$. The cell $C(i, j)$ corresponds to the edge pair e_{i+1} and f_{j+1}, where e_{i+1} is the $(i+1)^{\text{th}}$ edge of P and f_{j+1} is the $(j+1)^{\text{th}}$ edge of Q. Then the set $F(i, j) := \mathrm{FSD}(P, Q) \cap C(i, j)$ represents all pairs of points on $e_{i+1} \times f_{j+1}$ with


Figure 1: Two polygonal curves P and Q, together with their associated free-space diagram. The reachable region reach(P, Q) is shown in blue.

distance at most 1. Elementary geometry shows that $F(i, j)$ is the intersection of $C(i, j)$ with an ellipse [6]. In particular, the set $F(i, j)$ is convex, and the intersection of FSD(P, Q) with the boundary of $C(i, j)$ consists of four (possibly empty) intervals, one on each side of $\partial C(i, j)$. We call these intervals the *doors* of $C(i, j)$ in FSD(P, Q). A door is said to be *closed* if the interval is empty, and *open* otherwise.

A path π in FSD(P, Q) is *bimonotone* if it is both x- and y-monotone, i.e., every vertical and every horizontal line intersects π in at most one connected component. Alt and Godau observed that it suffices to determine whether there exists a bimonotone path from $(0, 0)$ to (n, n) inside FSD(P, Q). We define reach(P, Q) as the set of points in FSD(P, Q) that are reachable from $(0, 0)$ on a bimonotone path. Then $d_F(P, Q) \leq 1$ if and only if $(n, n) \in$ reach(P, Q). It is not necessary to compute all of reach(P, Q): since FSD(P, Q) is convex inside each cell, we need just the intersections reach$(P, Q) \cap \partial C(i, j)$. The sets defined by reach$(P, Q) \cap \partial C(i, j)$ are subintervals of the doors of the free-space diagram; they are defined by endpoints of doors in the free-space diagram in the same row or column. We call the intersection of a door with reach(P, Q) a *reach-door*. The intersections can be found in $O(n^2)$ time through a simple traversal of the cells [6]. In the next sections, we show how to obtain the crucial information, i.e. whether $(n, n) \in$ reach(P, Q), in $o(n^2)$ instead.

Basic approach and intuition. In our algorithm for the decision problem, we basically want to compute reach(P, Q). But instead of propagating the reachability information cell by cell, we always group τ by τ cells (with $1 \ll \tau \ll n$) into an *elementary box* of cells. When processing a box, we can assume that we know

which parts of the left and the bottom boundary of the box are reachable. That is, we know the reach-doors on the bottom and left boundary, and we need to compute the reach-doors on the top and right boundary of the elementary box. These reach-doors are determined by the combinatorial structure of the box. More specifically, if we know for every row and column the order of the door endpoints (including the reach-doors on the left and bottom boundary), we know which door boundaries determine the reach-doors on the top and right boundary. We call the sequence of these orders, the *(full) signature* of the box.

The total number of possible signatures is bounded by an expression in terms of τ. Thus, if we pick τ sufficiently small compared to n, we can pre-compute for all possible signatures the reach-doors on the top and right boundary, and build a data structure to query these quickly (Section 3). Since the reach-doors on the bottom and left boundary are required to make the signature, we initially have only partial signatures. In Section 4, we describe how to compute these efficiently. The partial signatures are then used to preprocess the data structure such that we can quickly find the full signature once we know the reach-doors of an elementary box. After building and preprocessing the data structure, it is possible to determine $d_F(P, Q) \leq 1$ efficiently by traversing the free-space diagram elementary box by elementary box, as explained in Section 5.

3 Building a Lookup Table

3.1 Preprocessing an elementary box Before it considers the input, our algorithm builds a lookup table. As mentioned above, the purpose of this table is to speed up the computation of small parts of the free-

Figure 2: The elementary box. The cell $D(2,1)$ is shown white. Its boundaries—$l(2,1)$, $l(3,1)$, $b(2,1)$, $b(2,2)$—are indicated.

space diagram.

Let $\tau \in \mathbb{N}$ be a parameter.[5] The *elementary box* is a subdivision of $[0,\tau]^2$ into τ columns and rows, thus τ^2 cells.[6] For $i,j = 0,\ldots,\tau-1$, we denote the cell $[i, i+1] \times [j, j+1]$ with $D(i,j)$. We denote the left side of the boundary $\partial D(i,j)$ by $l(i,j)$ and the bottom side by $b(i,j)$. Note that $l(i,j)$ coincides with the right side of $\partial D(i-1,j)$ and $b(i,j)$ with the top of $\partial D(i,j-1)$. Thus, we write $l(\tau+1, j)$ for the *right* side of $D(\tau, j)$ and $b(i, \tau+1)$ for the *top* side of $D(i, \tau)$. Fig. 2 shows the elementary box.

The *door-order* σ_j^r for a row j is a permutation of $\{s_0, t_0, \ldots, s_\tau, t_\tau\}$, thus having $2\tau + 2$ elements. For $i = 1, \ldots, \tau$, the element s_i represents the lower endpoint of the door on $l(i,j)$, and t_i represents the upper endpoint. The elements s_0 and t_0 are an exception: they describe the reach-door on the boundary $l(0,j)$ (i.e. its intersection with $\text{reach}(P,Q)$). The door-order σ_j^r represents the combinatorial order of these endpoints, as projected onto a vertical line, i.e., they are sorted into their vertical order. Some door-orders may encode the same combinatorial structure. In particular when door i is closed, the exact position of s_i and t_i in a door-order is irrelevant, up to t_i being before s_i. For a closed door i ($i > 0$), we assign s_i to the upper endpoint of $l(i,j)$ and t_i to the lower endpoint. The values of s_0 and t_0 are defined by the reach-door and their relative order is thus a result of computation. We break ties between s_i and $t_{i'}$ by placing s_i before $t_{i'}$, and any other ties are resolved by index. A door-order σ_i^c is defined

---

[5]A preview for the impatient reader: we later set $\tau = \Theta(\sqrt{\log n / \log \log n})$.

[6]For now, the elementary box is a combinatorial concept. In the next section, we overlay these boxes on the free-space diagram to obtain "concrete" elementary boxes.

analogously for a column i. We write $x <_i^c y$ if x comes before y in σ_i^c, and $x <_j^r y$ if x comes before y in σ_j^r. A *partial door-order* is a door-order in which s_0 and t_0 are omitted (i.e. the intersection of $\text{reach}(P,Q)$ with the door is still unknown); see Fig. 3.

We can now define the *(full) signature* of the elementary box as the aggregation of the door-orders of its rows and columns. Therefore, a signature $\Sigma = (\sigma_1^c, \ldots, \sigma_\tau^c, \sigma_1^r, \ldots, \sigma_\tau^r)$ consists of 2τ door-orders: one door-order σ_i^c for each column i and one door-order σ_i^r for each row j of the elementary box. Similarly, a *partial signature* is the aggregation of partial door-orders.

For a given signature, we define the *combinatorial reachability structure* of the elementary box as follows. For each column i and for each row j, the combinatorial reachability structure indicates which door boundaries in the respective column or row define the reach-door of $b(i, \tau)$ or $l(\tau, j)$.

LEMMA 3.1. *Let Σ be a signature for the elementary box. Then we can determine the combinatorial reachability structure of Σ in total time $O(\tau^2)$.*

Proof. We use dynamic programming, very similar to the algorithm by Alt and Godau [6]. For each vertical edge $l(i,j)$ we define a variable $\widehat{l}(i,j)$, and for each horizontal edge $b(i,j)$ we define a variable $\widehat{b}(i,j)$. The $\widehat{l}(i,j)$ are pairs of the form (s_u, t_v), representing the reach-door $\text{reach}(P,Q) \cap l(i,j)$. If this reach-door is closed, then $t_v <_j^r s_u$ holds. If the reach-door is open, then it is bounded by the lower endpoint of the door on $l(u,j)$ and by the upper endpoint of the door on $l(v,j)$. (Note that in this case we have $v = i$.) Once again s_0 and t_0 are special and represent the reach-door on $l(0,j)$. The variables $\widehat{b}(i,j)$ are defined analogously.

Now we can compute $\widehat{l}(i,j)$ and $\widehat{b}(i,j)$ recursively as follows: first, we set

$$\widehat{l}(0,j) = \widehat{b}(i,0) = (s_0, t_0), \quad \text{for } i,j = 0, \ldots, \tau-1.$$

Next, we describe how to find $\widehat{l}(i,j)$ given $\widehat{l}(i-1,j)$ and $\widehat{b}(i-1,j)$, see Fig. 4.

Case 1: Suppose $\widehat{b}(i-1,j)$ is open. This means that $b(i-1,j)$ intersects $\text{reach}(P,Q)$, so $\text{reach}(P,Q) \cap l(i,j)$ is limited only by the door on $l(i,j)$, and we can set $\widehat{l}(i,j) := (s_i, t_i)$.

Case 2: If both $\widehat{b}(i-1,j)$ and $\widehat{l}(i-1,j)$ are closed, it is impossible to reach $l(i,j)$ and thus we set $\widehat{l}(i,j) := \widehat{l}(i-1,j)$.

Case 3: If $\widehat{b}(i-1,j)$ is closed and $\widehat{l}(i-1,j)$ is open, we may be able to reach $l(i,j)$ via $l(i-1,j)$. Let s_u be the lower endpoint of $\widehat{l}(i-1,j)$. We need to pass $l(i,j)$ above s_u and s_i and below t_i, and therefore set

Figure 4: The three cases for the recursive definition of $\widehat{l}(i,j)$. If the lower boundary is reachable, we can reach the whole right door (left). If neither the lower nor the left boundary is reachable, the right door is not reachable either (middle). Otherwise, the lower boundary is the maximum of $\widehat{l}(i-1,j)$ and the lower boundary of the right door (right).

$\widehat{l}(i,j) := (\max(s_u, s_i), t_i)$, where the maximum is taken according to the order $<^r_j$.

The recursion for the variable $\widehat{b}(i,j)$ is defined similarly. We can implement the recursion in time $O(\tau^2)$ for any given signature, for example by traversing the elementary box column by column, while processing each column from bottom to top. □

There are at most $((2\tau + 2)!)^{2\tau} = \tau^{O(\tau^2)}$ distinct signatures for the elementary box. We choose $\tau = \lambda\sqrt{\log n / \log\log n}$ for a sufficiently small constant $\lambda > 0$, so that this number becomes $o(n)$. Thus, during the preprocessing stage we have time to enumerate all possible signatures and determine the corresponding combinatorial reachability structure inside the elementary box. This information is then stored in an appropriate data structure.

3.2 Building the data structure

Before we describe this data structure, we first explain how the door-orders are represented. This depends on the computational model. By our choice of τ, there are $o(n)$ distinct door-orders. On the word RAM, we represent each door-order and partial door-order by an integer between 1 and $(2\tau)!$. This fits into a word of $\log n$ bits. On the pointer machine, we create a record for each door-order and partial door-order; we represent an order by a pointer to the corresponding record.

The data structure has two *stages*, as schematized in Fig. 5. In the first stage (Fig. 5 (a-b)), we assume we know the partial door-order for each row and for each column of the elementary box[7], and we wish to determine the partial signature. In the second stage (Fig. 5 (c-d)), we have obtained the reach-doors for the left and bottom sides of the elementary box, and we are looking for the full signature. The details of our method depend on the computational model. One way uses table lookup and requires the word RAM; the other way works on the pointer machine, but is a bit more involved.

Word RAM. We organize the lookup table as a large tree T. In the first stage, each level of T corresponds to a row or column of the elementary box. Thus, there are 2τ levels. Each node has $(2\tau)!$ children, representing the possible partial door-orders for the next row or column. Since we represent door-orders by positive integers, each node of T may store an array for its children; we can choose the appropriate child for a given partial door-order in constant time. Thus, determining the partial signature for an elementary box requires $O(\tau)$ steps on a word RAM.

For the second stage, we again use a tree structure. Now the tree has $O(\tau)$ *layers*, each with $O(\log \tau)$ levels. Again, each layer corresponds to a row or column of the elementary box. The levels inside each layer then implement a balanced binary search tree that allows us to locate the endpoints of the reach-door within the partial signature. Since there are 2τ endpoints, this requires $O(\log \tau)$ levels. Thus, it takes $O(\tau \log \tau)$ time to find the full signature of a given elementary box.

Pointer model. Unlike in the word RAM model, we are not allowed to store a lookup table on every level of the tree T, and there is no way to quickly find the appropriate child for a given door-order. Instead, we must rely on batch processing to achieve a reasonable

---
[7]In the next section, we describe how to determine the partial door-orders efficiently.


Figure 3: The door-order of a row (the vertical order of the points) encodes the combinatorial structure of the doors. The door-order for the row in the figure is $s_1s_3s_4t_5t_3t_0s_2t_4s_0s_5t_1t_2$. Note that s_0 and t_0 represent the reach-door, which is empty in this case. These are omitted in the partial door-order.

Figure 5: Overview of the data structure. a–b) Using partial door-orders, we find the partial signature of an elementary box. c–d) Using the reach-doors on the bottom and left boundary of a box, we find the full signature and the combinatorial reachability.

running time.

Thus, suppose that during the first stage we want to find the partial signatures for a set B of m elementary boxes, where again for each box in B we know the partial door-order for each row and each column. Recall that we represent the door-order by a pointer to the corresponding record. With each such record, we store a queue of elementary boxes that is empty initially.

We now simultaneously propagate the boxes in B through T, proceeding level by level. In the first level, all of B is assigned to the root of T. Then, we go through the nodes of one level of T, from left to right. Let v be the current node of T. We consider each elementary box b assigned to v. We determine the next partial door-order for b, and we append b to the queue for this partial door-order—the queue is addressed through the corresponding record, so all elementary boxes with the same next partial door-order end up in the same queue. Next, we go through the nodes of the next level, again from left to right. Let v' be the current node. The node v' corresponds to a next partial door-order σ that extends the known signature of its parents. We consider the queue stored at the record for σ. By construction, the elementary boxes that should be assigned to v' appear consecutively at the beginning of this queue. We remove these boxes from the queue and assign them to v'. After this, all the queues are empty, and we can continue by propagating the boxes to the next level. During this procedure, we traverse each node of T a constant number of times, and in each level of the T we consider all the boxes in B. Since T has $o(n)$ nodes, the total running time is $O(n + m\tau)$.

For the second stage, the data structure works just as in the word RAM case, because no table lookup is necessary. Again, we need $O(\tau \log \tau)$ steps to process one box.

After the second stage we obtain the combinatorial reachability structure of the box in constant time since

we precomputed this information for each box. Thus, we have shown the following lemma, independently of the computational model.

LEMMA 3.2. *For* $\tau = \lambda\sqrt{\log n / \log \log n}$ *with a sufficiently small constant* $\lambda > 0$*, we can construct in* $o(n)$ *time a data structure of size* $o(n)$ *such that*

- *given a set of m elementary boxes where the partial door-orders are known, we can find the partial signature of each box in total time* $O(n + m\tau)$*;*
- *given the partial signature and the reach-doors on the bottom and left boundary of an elementary box, we can find the full signature in* $O(\tau \log \tau)$ *time;*
- *given the full signature of an elementary box, we can find the combinatorial reachability structure of the box in constant time.*

4 Preprocessing a Given Input

Next, we perform a second preprocessing phase that considers the input curves P and Q. Our eventual goal is to compute the intersection of reach(P, Q) with the cell boundaries, taking advantage of the data structure from Section 3. For this, we aggregate the cells of FSD(P, Q) into (concrete) elementary boxes consisting of $\tau \times \tau$ cells. There are n^2/τ^2 such boxes. We may avoid rounding issues by either duplicating vertices or handling a small part of FSD(P, Q) without lookup tables.

The goal is to determine the signature for each elementary box S. At this point this is not quite possible yet, since the signature depends on the intersection of reach(P, Q) with the lower and left boundary of S. Nonetheless, we can find the partial signature, in which the positions of s_0, t_0 (the reach-door) in the (partial) door-orders σ_i^r, σ_j^c are still to be determined.

We aggregate the columns of FSD(P, Q) into vertical *strips*, each corresponding to a single column of elementary boxes (i.e. τ consecutive columns of cells in FSD(P, Q)). See Fig. 6.

Figure 6: $\mathrm{FSD}(P,Q)$ is subdivided into n^2/τ^2 elementary boxes of size $\tau \times \tau$. The free-space diagram is subdivided into n^2/τ^2 elementary boxes of size $\tau \times \tau$. A strip is a column of elementary boxes: it corresponds to a subcurve of P with τ edges.

Let A be such a strip. It corresponds to a subcurve P' of P with τ edges. The following lemma implies that we can build a data structure for A such that, given any segment of Q, we can efficiently find its partial door-order within the elementary box in A.

LEMMA 4.1. *There exists a constant c such that the following holds: given a subcurve P' with τ edges, we can compute in $O(\tau^c)$ time a data structure that requires $O(\tau^c)$ space and that allows us to determine the partial door-order of any line segment on Q in time $O(\log \tau)$.*

Proof. Consider the arrangement $\mathcal{A}$ of unit circles whose centers are the vertices of P' (see Fig. 7). The partial door-order of a line segment s is determined by the intersections of s with the arcs of $\mathcal{A}$ (and for a circle not intersecting s by whether s lies inside or outside of the circle). Let ℓ_s be the line spanned by line segment s. Suppose we wiggle ℓ_s. The order of intersections of ℓ_s and the arcs of $\mathcal{A}$ changes only when ℓ_s moves over a vertex of $\mathcal{A}$ or if ℓ_s leaves or enters a circle.

We use the standard duality transform that maps a line $\ell : y = ax + b$ to the point $\ell^* : (a, -b)$, and vice versa. Consider a unit circle C in $\mathcal{A}$ with center (c_x, c_y). Elementary geometry shows that the set of all lines that are tangent to C from above dualizes to the curve $t_a^*(C) : y = c_x x - c_y - \sqrt{1+x^2}$. Similarly, the lines that are tangent to C from below dualize to the curve $t_b^*(C) : y = c_x x - c_y + \sqrt{1+x^2}$. Define $C^* := \{t_a^*(C), t_b^*(C) \mid C \in \mathcal{A}\}$. Since any pair of distinct circles C_1, C_2 has at most four common tangents, one for each choice of above/below C_1 and above/below C_2, it follows that any two curves in C^* intersect at most once.

Let V be the set of vertices in $\mathcal{A}$, and let V^* be the lines dual to the points in V (note that $|V| = O(\tau^2)$).

Since for any vertex $v \in V$ and any circle $C \in \mathcal{A}$ there are at most two tangents through v on C, each line in V^* intersects each curve in C^* at most once. Thus, the arrangement $\mathcal{B}$ of the curves in $V^* \cup C^*$ is an arrangement of *pseudolines* with complexity $O(\tau^4)$. Furthermore, it can be constructed in the same expected time, together with a point location structure that finds the containing cell in $\mathcal{B}$ of any given point in time $O(\log \tau)$ [47, Chapter 6.6.1].

Now consider a line segment s and the supporting line ℓ_s. As observed in the first paragraph, the combinatorial structure of the intersection between ℓ_s and $\mathcal{A}$ is completely determined by the cell of $\mathcal{B}$ that contains the dual point ℓ_s^*. Thus, for every cell $f(s) \in \mathcal{B}$, we construct a list $L_{f(s)}$ that represents the combinatorial structure of $\ell_s \cap \mathcal{A}$. There are $O(\tau^4)$ such lists, each having size $O(\tau)$. We can compute $L_{f(s)}$ by traversing the zone of ℓ_s in $\mathcal{A}$. Since circles intersect at most twice and also a line intersects any circle at most twice, the zone has complexity $O(\tau 2^{\alpha(\tau)})$, where $\alpha(\cdot)$ denotes the inverse Ackermann function [47, Theorem 5.11]. Since $O(\tau 2^{\alpha(\tau)}) \subset O(\tau^2)$, we can compute all lists in $O(\tau^6)$ time.

Given the list $L_{f(s)}$, the partial door-order of s is determined by the position of the endpoints of s in $L_{f(s)}$. There are $O(\tau^2)$ possible ways for this, and we build a table $T_{f(s)}$ that represents them. For each entry in $T_{f(s)}$, we store a representative for the corresponding partial door-order. As described in the previous section, the representative is a positive integer in the word RAM model and a pointer to the appropriate record on a pointer machine.

The total size of the data structure is $O(\tau^6)$ and it can be constructed in the same time. A query works as follows: given s, we can compute ℓ_s^* in constant time.

Figure 7: By using the arrangement $\mathcal{A}$ defined by unit circles centered at vertices of P', we can determine the partial door-order of each segment s on Q. This is done by locating the dual point of ℓ_s in the dual arrangement $\mathcal{B}$. The dual arrangement also contains pseudolines to determine when ℓ_s leaves a circle of $\mathcal{A}$.

Then we use the point location structure of $\mathcal{B}$ to find $f(s)$ in $O(\log \tau)$ time. Using binary search on $T_{f(s)}$ (or an appropriate tree structure in the case of a point machine), we can then determine the position of the endpoints of s in the list $L_{f(s)}$ in $O(\log \tau)$ time. This bound holds both on the word RAM and on the pointer machine. $\qquad \square$

LEMMA 4.2. *Given the data structure of Lemma 3.2, the partial signature for each elementary box can be determined in time $O(n\tau^{c-1} + n^2(\log \tau)/\tau)$ for some constant c.*

Proof. By building and using the data structure from Lemma 4.1, we determine the partial door-order for each row in each vertical τ-strip in $O(\frac{n}{\tau}(\tau^c + n \log \tau)) = O(n\tau^{c-1} + n^2 \log \tau/\tau)$ time. We repeat the procedure with the horizontal strips. Now we know for each elementary box in $\mathrm{FSD}(P, Q)$ the partial door-order for each row and each column. We use the data structure of Lemma 3.2 to combine these. As there are n^2/τ^2 boxes, the number of steps is $O(n^2/\tau + n) = O(n^2/\tau)$. Hence, the partial signature for each elementary box is computed in $O(n\tau^{c-1} + n^2(\log \tau)/\tau)$. $\qquad \square$

5 Solving the Decision Problem

With the data structures and preprocessing from the previous sections, we have all elements in place to determine whether $d_F(P, Q) \leq 1$. We know for each elementary box its partial signature and we have a data structure to derive its full signature (and with it, the combinatorial reachability structure) when its reach-doors are known. What remains to be shown is that we can efficiently process the free-space diagram to determine whether $(n, n) \in \mathrm{reach}(P, Q)$. This is captured in the following lemma.

LEMMA 5.1. *If the partial signature for each elementary box is known, we can determine whether $(n, n) \in \mathrm{reach}(P, Q)$ in time $O(n^2(\log \tau)/\tau)$.*

Proof. We go through all of the elementary boxes of $\mathrm{FSD}(P, Q)$, processing them one column at a time, going from bottom to top in each column. Initially, we know the full signature for the box S in the lower left corner of $\mathrm{FSD}(P, Q)$. We use the signature to determine the intersections of $\mathrm{reach}(P, Q)$ with the upper and right boundary of S. There is a subtlety here: the signature gives us only the combinatorial reachability structure, and we need to map the resulting s_i, t_j back to the corresponding vertices on the curves. On the word RAM, this can be done easily through table lookups. On the pointer machine, we use representative records for the s_i, t_i elements and use $O(\tau)$ time before processing the box to store a pointer from each representative record to the appropriate vertices on P and Q.

We proceed similarly for the other boxes. By the choice of the processing order of the elementary boxes we always know the incoming reach-doors on the bottom and left boundary when processing a box. Given the incoming reach-doors, we can determine the full signature and find the structure of the outgoing reach-doors in total time $O(\tau \log \tau)$, using Lemma 3.2. Again, we need $O(\tau)$ additional time on the pointer machine to establish the mapping from the abstract s_i, t_i elements to the concrete vertices of P and Q. In total, we spend $O(\tau \log \tau)$ time per box. Thus, it takes time $O(n^2(\log \tau)/\tau)$ to process all boxes, as claimed. $\qquad \square$

As a result, we obtain the following theorem for a pointer machine (and by extension, for the real RAM model). For the word RAM model, we obtain an even faster algorithm (see Section 6).

THEOREM 5.1. *The decision version of the Fréchet problem can be solved in $O(n^2(\log\log n)^{3/2}/\sqrt{\log n})$ time on a pointer machine.*

Proof. Set $\tau = \lambda\sqrt{\log n/\log\log n}$ for a sufficiently small constant $\lambda > 0$. The theorem follows by applying Lemmas 3.2, 4.2, and 5.1 in sequence. $\qquad\square$

6 Improved Bound on Word RAM

We now explain how the running time of our algorithm can be improved if our computational model allows for constant time table-lookup. We use the same τ as above (up to a constant factor). However, we change a number of things. "Signatures" are represented differently and the data structure to obtain combinatorial reachability structures is changed accordingly. Furthermore, we aggregate elementary boxes into clusters and determine "partial door-orders" for multiple boxes at the same time. Finally, we walk the free-space diagram based on the clusters to decide $d_F(P,Q) \leq 1$.

Clusters and extended signatures. We introduce a second level of aggregation in the free-space diagram: a *cluster* is a collection of $\tau \times \tau$ elementary boxes, that is, $\tau^2 \times \tau^2$ cells in $\mathrm{FSD}(P,Q)$. Let R be a row of cells in $\mathrm{FSD}(P,Q)$ of a certain cluster. As before, the row R corresponds to an edge e on Q and a subcurve P' of P with τ^2 edges. We associate with R an ordered set $Z = \langle e_0, z'_0, z_1, z'_1, z_2, z'_2, \ldots, z_k, z'_k, e_1 \rangle$ with $2 \cdot k + 3$ elements. Here k is the number of intersections of e with the unit circles centered at the τ vertices of P' (all but the very first). Hence, k is bounded by 2τ and $|Z|$ is bounded by $4\tau+3$. The order of Z indicates the order of these intersections with e directed along Q. Elements e_0 and e_1 represent the endpoints of e and take a special role. In particular, these are used to represent closed doors and snap open doors to the edge e. The elements z'_i are placeholders for the positions of the endpoints of the reach-doors: z'_0 represents a possible reach-door endpoint between e_0 and z_1, the element z'_1 is an endpoint between z_1 and z_2, etc.

Consider a row R' of an elementary box inside the row R of a cluster, corresponding to an edge e of Q. The *door-index* of R' is an ordered set $\langle s_0, t_0, \ldots, s_\tau, t_\tau \rangle$ of size $2\tau + 2$. Similar to a door-order, elements s_0 and t_0 represent the reach-door at the leftmost boundary of R'; the elements s_i and t_i ($1 \leq i \leq \tau$) represent the door at the right boundary of the i^{th} cell in R'. However, instead of rearranging the set to indicate relative positions, the elements s_i and t_i simply refer to an element in Z. If the door is open, they refer to the intersections with e (possibly snapped to e_0 or e_1). If the door is closed, s_i is set to e_1 and t_i is set to e_0. The elements s_0 and t_0 are special, representing

the reach-door and refer to one of the elements z'_i. A *partial door-index* is a door-index without s_0 and t_0. The advantage of a door-index over a door-order is that the reach-door is always at the start. Hence, completing a partial door-index to a full door-index can be done in constant time. Since a door-index has size $2\tau + 2$, the number of possible door-indices for R' is $\tau^{O(\tau)}$.

We define the door-indices for the columns analogously. We concatenate the door-indices for the rows and the columns to obtain the *indexed signature* for an elementary box. Similarly, we define the *partial indexed signature*. The total number of possible indexed signatures remains $\tau^{O(\tau^2)}$.

For each possible partial indexed signature Σ we build a lookup table T_Σ as follows: the input is a word with 4τ fields of $O(\log\tau)$ bits each. Each field stores the positions in Z of the endpoints of the ingoing reach-doors for the elementary box: 2τ fields for the left side, 2τ fields for the lower side. The output consists of a word that represents the indices for the elements in Z that represent the outgoing reach-doors for the upper and right boundary of the box. Thus, the input of T_Σ is a word of $O(\tau\log\tau)$ bits, and T_Σ has size $\tau^{O(\tau)}$. Hence, for all partial indexed signatures combined, the size is $\tau^{O(\tau^2)} = o(n)$ by our choice of τ.

Preprocessing a given input. During the preprocessing for a given input P, Q, we use *superstrips* consisting of τ strips. That is, a superstrip is a column of clusters and consists of τ^2 columns of the free-space diagram. Lemma 4.1 still holds, albeit with a larger constant c. The data structure gets as input a query edge e, and it returns in $O(\log\tau)$ time a word that contains τ fields. Each field represents the partial door-index for e in the corresponding elementary box and thus consists of $O(\tau\log\tau)$ bits. Hence, the word size is $O(\tau^2\log\tau) = O(\log n)$ by our choice of τ. Thus, the total time for building a data structure for each superstrip and for processing all rows is $O(n/\tau^2\,(\tau^c + n\log\tau)) = O(n^2(\log\tau)/\tau^2)$. We now have parts of the partial indexed signature for each elementary box packed into different words. To obtain the partial indexed signature, we need to rearrange the information such that the partial door-indices of the rows in one elementary box are in a single word. This corresponds to computing a transposition of a matrix, as is illustrated in Fig. 8. For this, we need the following lemma, which can be found—in slightly different form—in Thorup [48, Lemma 9].

LEMMA 6.1. *Let X be a sequence of τ words that contain τ fields each, so that X can be interpreted as a $\tau \times \tau$ matrix. Then we can compute in time $O(\tau\log\tau)$ on a word RAM a sequence Y of τ words with τ fields each that represents the* transposition *of X.*

Figure 8: (left) Every field represents the partial door-index of a row in an elementary box. (center) The fields are grouped into words per row in a cluster. (right) Transposition yields the desired organization, where a word represents the partial door-index of the rows in an elementary box.

Proof. The algorithm is recursive and solves a more general problem: let X be a sequence of a words that represents a sequence M of b different $a \times a$ matrices, such that the i^{th} word in X contains the fields of the i^{th} row of each matrix in M from left to right. Compute a sequence of words Y that represents the sequence M' of the transposed matrices in M.

The recursion works as follows: if $a = 1$, there is nothing to be done. Otherwise, we split X into the sequence X_1 of the first $a/2$ words and the sequence X_2 of the remaining words. X_1 and X_2 now represent a sequence of $2b$ $(a/2) \times (a/2)$ matrices, which we transpose recursively. After the recursion, we put the $(a/2) \times (a/2)$ submatrices back together in the obvious way. To finish, we need to transpose the off-diagonal submatrices. This can be done simultaneously for all matrices in time $O(a)$, by using appropriate bit-operations (or table lookup). Hence, the running time obeys a recursion of the form $T(a) = 2T(a/2) + O(a)$, giving $T(a) = O(a \log a)$, as desired. $\square$

By applying the lemma to the words that represent τ consecutive rows in a superstrip, we obtain the partial door-indices of the rows for each elementary box. This takes total time $O((n/\tau^2) \cdot (n/\tau) \cdot \tau \log \tau) = O(n^2 (\log \tau)/\tau^2)$. We repeat this procedure for the horizontal superstrips. By using an appropriate lookup table to combine the partial door-indices of the rows and columns, we obtain the partial indexed signature for each elementary box in total time $O(n^2 (\log \tau)/\tau^2)$.

The actual computation. We traverse the free-space diagram cluster by cluster (recall that a cluster consists of $\tau \times \tau$ elementary boxes). The clusters are processed column by column from left to right, and inside each column from bottom to top. Before processing a cluster, we walk along the left and lower boundary of the cluster to determine the incoming reach-doors. This is done by performing a binary search for each box on the boundary, and determining the appropriate elements z_i' which correspond to the incoming reach-doors. Using this information, we assemble the appropriate words that represent the incoming information for each elementary box. Since there are n^2/τ^4 clusters, this step requires time $O((n^2/\tau^4)\tau^2 \log \tau) = O(n^2 (\log \tau)/\tau^2)$.

We then process the elementary boxes inside the cluster, in a similar fashion. Now, however, we can process each elementary box in constant time through a single table lookup, so the total time is $O(n^2/\tau^2)$. Hence, the total running time of our algorithm is $O(n^2 (\log \tau)/\tau^2)$. By our choice of $\tau = \lambda\sqrt{\log n / \log \log n}$ for a sufficiently small $\lambda > 0$, we obtain the following theorem.

THEOREM 6.1. *The decision version of the Fréchet problem can be solved in $O(n^2 (\log \log n)^2 / \log n)$ time on a word RAM machine.*

7 Computing the Fréchet Distance

The optimization version of the Fréchet problem, i.e., computing the Fréchet distance, can be done in $O(n^2 \log n)$ time using parametric search with the decision version as a subroutine [6]. We showed that the decision problem can be solved in $o(n^2)$ time. However, this does not directly yield a faster algorithm for the optimization problem: if the running time of the decision problem is $T(n)$, parametric search gives an $O((T(n) + n^2) \log n)$ time algorithm [6]. There is an alternative randomized algorithm by Raichel and Har-Peled [41]. Their algorithm also needs $O((T(n) + n^2) \log n)$ time, but below we adapt it to obtain the following lemma.

LEMMA 7.1. *The Fréchet distance of two polygonal curves with n vertices each can be computed by a randomized algorithm in $O(n^2 2^{\alpha(n)} + T(n) \log n)$ expected time, where $T(n)$ is the time for the decision problem.*

Before we prove the lemma, we recall that possible values of the Fréchet distance are limited to a certain set of *critical values* [6]:

1. the distance between a vertex of one curve and a vertex of the other curve (**vertex-vertex**);

2. the distance between a vertex of one curve and an edge of the other curve (**vertex-edge**); and

3. for two vertices of one curve and an edge of the other curve, the distance between one of the vertices and the intersection of e with the bisector of the two vertices (if this intersection exists) (**vertex-vertex-edge**).

If we also include vertex-vertex-edge tuples with no intersection, we can sample a critical value uniformly at random in constant time. The algorithm now works as follows (see Har-Peled and Raichel [41] for more details): first, we sample a set S of $K = 4n^2$ critical values uniformly at random. Next, we find $a, b \in S$ such that the Fréchet distance lies between a and b and such that $[a, b]$ contains no other value from S. In the original algorithm this is done by sorting S and performing a binary search using the decision version. Using median-finding instead, this step can be done in $O(K + T(n) \log K)$ time. Alternatively, the running time of this step could be reduced by picking a smaller K. However, this does not improve the final bound, since it is dominated by a $O(n^2 2^{\alpha(n)})$ component. The interval $[a, b]$ with high probability contains only a small number of the remaining critical values. More precisely, for $K = 4n^2$ the probability that $[a, b]$ has more than $2cn \ln n$ critical values is at most $1/n^c$ [41, Lemma 6.2].

The remainder of the algorithm determines the K' critical values in the interval $[a, b]$, sorts them, and performs a binary search. Excluding the time to determine the critical values, this takes $O(K' + T(n) \log K')$ time with median-finding. Thus the crucial part is to determine the K' critical values. In $O(n^2)$ time we can check for each vertex-vertex and vertex-edge pair whether the corresponding critical value lies in $[a, b]$. It remains to determine the critical values corresponding to vertex-vertex-edge tuples. These critical values are found by a standard sweepline algorithm. For this, take an edge e of P and the vertices of Q. The sweep starts with circles of radius a around the vertices of Q, and it increases the radii until they reach b. During this sweep, the algorithm maintains the order in which the circle arcs intersect e. A critical value of the vertex-vertex-edge type corresponds to the event that two different circles intersect e in the same point. Besides these events the sweepline algorithm requires the following events: (a) a circle intersects e for the first time, and (b) a circle intersects one of the vertices of e. Both event types correspond to critical values involving e or a vertex of e. Thus if we perform a sweep for each edge of P (and similarly for Q), the total number of events is $O(K')$. Thus, the overall running time of all sweeps ignoring the time for initialization is $O(K' \log n)$.

It remains to show that we can quickly find the initial order in which the circle arcs intersect e. First, compute the arrangement $\mathcal{A}$ of circles with radius a around the vertices of Q. This takes $O(n^2)$ time [28]. To find the intersection order, traverse in $\mathcal{A}$ the zone of the line ℓ spanned by e. The time for the traversal is bounded by the complexity of the zone. Since circles pairwise intersect at most twice and ℓ intersects

each circle only twice, the complexity of the zone is $O(n 2^{\alpha(n)})$ [47, Theorem 5.11]. Summing over all edges e, this adds a total of $O(n^2 2^{\alpha(n)})$ to the running time. Thus the overall time is $O(T(n) \log(n) + n^2 2^{\alpha(n)} + K' \log n)$. The event $K' > 8n \ln n$ has probability less than $1/n^4$, and we always have $K' = O(n^3)$. Thus, this case adds $o(1)$ to the expected running time. Given $K' \leq 8n \ln n$, the running time is $O(n \log^2 n)$. Lemma 7.1 follows. Theorem 7.1 now results from Lemma 7.1, Theorem 5.1, and Theorem 6.1.

THEOREM 7.1. *The Fréchet distance of two polygonal curves with n vertices each can be computed by a randomized algorithm in time $O(n^2 \sqrt{\log n} (\log \log n)^{3/2})$ on a pointer machine and in time $O(n^2 (\log \log n)^2)$ on a word RAM.*

8 Decision Trees

Our results also have implications for the decision-tree complexity of the Fréchet problem. Since in that model we account only for comparisons between the input elements, the preprocessing comes for free, and hence the size of the elementary boxes can be increased.

Before we consider the continuous Fréchet problem, we first note that a similar result can be obtained for the *discrete* Fréchet problem: suppose we have two sequences $P = \langle p_1, p_2, \ldots, p_n \rangle$ and $Q = \langle q_1, q_2, \ldots, q_n \rangle$. For $\delta > 0$, we define a directed graph G_δ with vertex set $P \times Q$. In G_δ, there is an edge between two vertices (p_i, q_j), (p_i, q_{j+1}) if and only if both $d(p_i, q_j) \leq \delta$ and $d(p_i, q_{j+1}) \leq \delta$. The condition is similar for an edge between vertices (p_i, q_j) and (p_{i+1}, q_j), and vertices (p_i, q_j) and (p_{i+1}, q_{j+1}). There are no further edges in G_δ. Now the problem is to find the smallest δ for which G_δ has a path from (p_1, q_1) to (p_n, q_n). For the discrete Fréchet problem we obtain the following bound.[8]

THEOREM 8.1. *There is an algebraic computation tree for the* discrete *Fréchet problem of depth $\widetilde{O}(n^{4/3})$.*

Proof. We first discuss the decision problem, where we are given P, Q and a δ, and we need to decide whether we can reach (p_n, q_n) from (p_1, q_1). For the discrete case, the analogue of the reachable free-space is just an $n \times n$ boolean matrix M, where the bit in the i^{th} row and the j^{th} column indicates whether the pair (p_i, q_j) can be reached from (p_1, q_1) in G_δ.

As shown by Katz and Sharir [43], we can compute a representation of the set of points (p_i, q_j) with $\|p_i - q_j\| \leq \delta$ in $\widetilde{O}(n^{4/3})$. This information suffices to complete M without further comparisons. As shown

---

[8]The notation $\widetilde{O}(\cdot)$ (pronounced "soft-Oh") stands for $O(\cdot)$ up to poly-logarithmic factors.

by Agarwal *et al.*, one can then solve the optimization problem at the cost of another $O(\log n)$-factor, which is absorbed into the $\widetilde{O}$-notation. $\qquad\square$

Given our results above, we prove an analogous statement for the continuous Fréchet distance.

THEOREM 8.2. *There exists an algebraic decision tree for the Fréchet problem (decision version) of depth $O(n^{2-\varepsilon})$, for a fixed constant $\varepsilon > 0$.*

Proof. We reconsider the steps of our algorithm. The only phases that actually involve the input are the second preprocessing phase and the traversal of the elementary boxes. The reason of our choice for τ was to keep the time for the first preprocessing phase polynomial. This is no longer a problem. By Lemmas 4.2 and 5.1, the remaining cost is bounded by $O(n\tau^{c-1} + n^2(\log\tau)/\tau)$, where c is the constant from Lemma 4.1. Choosing $\tau = n^{1/c}$, we get a decision tree of depth $n \cdot n^{\frac{c-1}{c}} + n^2 \log n/n^{1/c}$. This is $O(n^{2-(1/c)}\log n) = O(n^{2-\varepsilon})$ for, say, $\varepsilon = 1/2c$. $\qquad\square$

Assuming linear time reductions, this has the following consequence for Alt's conjecture.

COROLLARY 8.1. *If the decision version of the Fréchet problem is 3SUM-hard, then 3SUM has an algebraic decision tree of depth $O(n^{2-\varepsilon})$.*

9 Conclusion

In this paper we have broken the long-standing quadratic upper bound for the decision version of the Fréchet problem. Moreover, we have shown that this problem has an algebraic decision tree of depth $O(n^{2-\varepsilon})$, for some $\varepsilon > 0$ and where n is the number of vertices of the polygonal curves. This strongly indicates that the problem is not 3SUM-hard after all. We have shown how our faster algorithm for the decision version can be used for a faster algorithm to compute the Fréchet distance. If we allow constant-time table-lookup, we obtain a running time in close reach of $O(n^2)$.

This leaves us with intriguing open research questions. Can we reduce the time needed for the decision version to $O(n^{2-\varepsilon})$, the bound we obtain from the algebraic decision tree? Can we devise a quadratic or even subquadratic algorithm for the optimization version? Can we devise such an algorithm on the word RAM, that is, with constant-time table-lookup? Or, on the other hand, can we establish a connection between the Fréchet distance and other problems which exhibit a discrepancy between the decision tree and the uniform complexity, such as MIN-PLUS-CONVOLUTION?

References

[1] P. K. Agarwal, R. Ben Avraham, H. Kaplan, and M. Sharir. Computing the discrete Fréchet distance in subquadratic time. In *Proc. 24th Annu. ACM-SIAM Sympos. Discrete Algorithms (SODA)*, pages 156–167, 2013.

[2] P. K. Agarwal, S. Har-Peled, N. H. Mustafa, and Y. Wang. Near-linear time approximation algorithms for curve simplification. *Algorithmica*, 42(3–4):203–219, 2005.

[3] N. Ailon and B. Chazelle. Lower bounds for linear degeneracy testing. *J. ACM*, 52(2):157–171, 2005.

[4] H. Alt. The computational geometry of comparing shapes. In *Efficient Algorithms*, volume 5760 of *Lecture Notes in Computer Science*, pages 235–248. Springer-Verlag, 2009.

[5] H. Alt and M. Buchin. Can we compute the similarity between surfaces? *Discrete Comput. Geom.*, 43(1):78–99, 2010.

[6] H. Alt and M. Godau. Computing the Fréchet distance between two polygonal curves. *Internat. J. Comput. Geom. Appl.*, 5(1–2):78–99, 1995.

[7] H. Alt, C. Knauer, and C. Wenk. Comparison of distance measures for planar curves. *Algorithmica*, 38(1):45–58, 2003.

[8] B. Aronov, S. Har-Peled, C. Knauer, Y. Wang, and C. Wenk. Fréchet Distances for Curves, Revisited. In *Proc. 14th Annu. European Sympos. Algorithms (ESA)*, volume 4168 of *Lecture Notes in Computer Science*, pages 52–63, 2006.

[9] S. Arora and B. Barak. *Computational complexity. A modern approach.* Cambridge University Press, Cambridge, 2009.

[10] I. Baran, E. D. Demaine, and M. Pătraşcu. Subquadratic algorithms for 3SUM. *Algorithmica*, 50(4):584–596, April 2008.

[11] R. Bellman and R. Kalaba. On adaptive control processes. *IRE Transactions on Automatic Control*, 4(2):1–9, 1959.

[12] M. de Berg, A. F. Cook IV, and J. Gudmundsson. Fast Fréchet queries. *Comput. Geom. Theory Appl.*, 46(6):747–755, 2013.

[13] S. Brakatsoulas, D. Pfoser, R. Salas, and C. Wenk. On map-matching vehicle tracking data. In *Proc. 31st Int. Conf. on Very Large Data Bases*, pages 853–864, 2005.

[14] D. Bremner, T. M. Chan, E. D. Demaine, J. Erickson, F. Hurtado, J. Iacono, S. Langerman, M. Pătraşcu, and P. Taslakian. Necklaces, convolutions, and $X + Y$. *Algorithmica*, pages 1–21, 2012.

[15] K. Buchin, M. Buchin, and J. Gudmundsson. Constrained free space diagrams: a tool for trajectory analysis. *Int. J. of GIS*, 24(7):1101–1125, 2010.

[16] K. Buchin, M. Buchin, J. Gudmundsson, M. Löffler, and J. Luo. Detecting commuting patterns by clustering subtrajectories. *Internat. J. Comput. Geom. Appl.*, 21(3):253–282, 2011.

[17] K. Buchin, M. Buchin, C. Knauer, G. Rote, and C. Wenk. How difficult is it to walk the dog? In *Proc. 23rd European Workshop Comput. Geom. (EWCG)*, pages 170–173, 2007.

[18] K. Buchin, M. Buchin, W. Meulemans, and B. Speckmann. Locally correct Fréchet matchings. In *Proc. 20th Annu. European Sympos. Algorithms (ESA)*, volume 7501 of *Lecture Notes in Computer Science*, pages 229–240, 2012.

[19] K. Buchin, M. Buchin, and A. Schulz. Fréchet distance of surfaces: Some simple hard cases. In *Proc. 18th Annu. European Sympos. Algorithms (ESA)*, volume 6347 of *Lecture Notes in Computer Science*, pages 63–74, 2010.

[20] K. Buchin, M. Buchin, R. van Leusden, W. Meulemans, and W. Mulzer. Computing the Fréchet distance with a retractable leash. In *Proc. 21st Annu. European Sympos. Algorithms (ESA)*, Lecture Notes in Computer Science, pages 241–252, 2013.

[21] K. Buchin, M. Buchin, and Y. Wang. Exact algorithms for partial curve matching via the Fréchet distance. In *Proc. 20th Annu. ACM-SIAM Sympos. Discrete Algorithms (SODA)*, pages 645–654, 2009.

[22] K. Buchin, M. Buchin, and C. Wenk. Computing the Fréchet distance between simple polygons. *Comput. Geom. Theory Appl.*, 41(1–2):2–20, 2008.

[23] K. Buchin and W. Mulzer. Delaunay triangulations in $O(\text{sort}(n))$ time and more. *J. ACM*, 58(2):Art. 6, 27, 2011.

[24] M. Buchin. *On the Computability of the Fréchet Distance Between Triangulated Surfaces*. PhD thesis, Free University Berlin, Institute of Computer Science, 2007.

[25] E. Chambers, É. de Verdière, J. Erickson, S. Lazard, F. Lazarus, and S. Thite. Homotopic Fréchet distance between curves or, walking your dog in the woods in polynomial time. *Comput. Geom. Theory Appl.*, 43(3):295–311, 2010.

[26] T. M. Chan. All-pairs shortest paths with real weights in $O(n^3/\log n)$ time. *Algorithmica*, 50(2):236–243, 2008.

[27] T. M. Chan. More algorithms for all-pairs shortest paths in weighted graphs. *SIAM J. Comput.*, 39(5):2075–2089, 2010.

[28] B. M. Chazelle and D. T. Lee. On a circle placement problem. *Computing*, 36(1–2):1–16, 1986.

[29] A. F. Cook, A. Driemel, S. Har-Peled, J. Sherette, and C. Wenk. Computing the Fréchet distance between folded polygons. In *Proc. 12th Symposium on Algorithms and Data Structures (WADS)*, pages 267–278, 2011.

[30] A. F. Cook and C. Wenk. Geodesic Fréchet distance inside a simple polygon. *ACM Transactions on Algorithms*, 7(1):Art. 9, 2010.

[31] A. Driemel and S. Har-Peled. Jaywalking your dog: computing the Fréchet distance with shortcuts. In *Proc. 23rd Annu. ACM-SIAM Sympos. Discrete Algorithms (SODA)*, pages 318–337, 2012.

[32] A. Driemel, S. Har-Peled, and C. Wenk. Approximating the Fréchet distance for realistic curves in near linear time. In *Proc. 26th Annu. ACM Sympos. Comput. Geom. (SoCG)*, pages 365–374, 2010.

[33] A. Efrat, L. Guibas, S. Har-Peled, J. Mitchell, and T. Murali. New similarity measures between polylines with applications to morphing and polygon sweeping. *Discrete Comput. Geom.*, 28(4):535–569, 2002.

[34] J. Erickson. Bounds for linear satisfiability problems. *Chicago J. Theor. Comput. Sci.*, page Article 8, 1999.

[35] M. L. Fredman. How good is the information theory bound in sorting? *Theoret. Comput. Sci.*, 1(4):355–361, 1975/76.

[36] A. Gajentaan and M. H. Overmars. On a class of $O(n^2)$ problems in computational geometry. *Comput. Geom. Theory Appl.*, 5(3):165–185, 1995.

[37] M. Godau. A natural metric for curves - computing the distance for polygonal chains and approximation algorithms. In *Proc. 8th Sympos. Theoret. Aspects Comput. Sci. (STACS)*, pages 127–136, 1991.

[38] M. Godau. *On the Complexity of Measuring the Similarity Between Geometric Objects in Higher Dimensions*. PhD thesis, Freie Universität Berlin, Germany, 1998.

[39] J. Gudmundsson and T. Wolle. Towards automated football analysis: Algorithms and data structures. In *Proc. 10th Australasian Conf. on Mathematics and Computers in Sport*, 2010.

[40] S. Har-Peled, A. Nayyeri, M. Salavatipour, and A. Sidiropoulos. How to walk your dog in the mountains with no magic leash. In *Proc. 28th Annu. ACM Sympos. Comput. Geom. (SoCG)*, pages 121–130, 2012.

[41] S. Har-Peled and B. Raichel. The Fréchet distance revisited and extended. In *Proc. 27th Annu. ACM Sympos. Comput. Geom. (SoCG)*, pages 448–457, 2011.

[42] P. Indyk. Approximate nearest neighbor algorithms for Frechet distance via product metrics. In *Proc. 18th Annu. ACM Sympos. Comput. Geom. (SoCG)*, pages 102–106, 2002.

[43] M. Katz and M. Sharir. An expander-based approach to geometric optimization. *SIAM J. Comput.*, 26(5):1384–1408, 1997.

[44] A. Maheshwari, J.-R. Sack, K. Shahbaz, and H. Zarrabi-Zadeh. Fréchet distance with speed limits. *Comput. Geom. Theory Appl.*, 44(2):110–120, 2011.

[45] A. Maheshwari, J.-R. Sack, K. Shahbaz, and H. Zarrabi-Zadeh. Improved algorithms for partial

curve matching. In *Proc. 19th Annu. European Sympos. Algorithms (ESA)*, volume 6942 of *Lecture Notes in Computer Science*, pages 518–529, 2011.

[46] M. Pătraşcu. Towards polynomial lower bounds for dynamic problems. In *Proc. 42nd Annu. ACM Sympos. Theory Comput. (STOC)*, pages 603–610, 2010.

[47] M. Sharir and P. K. Agarwal. *Davenport-Schinzel Sequences and Their Geometric Applications*. Cambridge University Press, 1995.

[48] M. Thorup. Randomized sorting in $O(n \log \log n)$ time and linear space using addition, shift, and bitwise boolean operations. *J. Algorithms*, 42(2):205–230, 2002.

[49] C. Wenk, R. Salas, and D. Pfoser. Addressing the need for map-matching speed: Localizing global curve-matching algorithms. In *Proc. 18th Int. Conf. on Sci. and Stat. Database Management*, pages 379–388, 2006.

A Computational Models

Real RAM. The standard machine model in computational geometry is the *real RAM*. Here, data is represented as an infinite sequence of storage cells. These cells can be of two different types: they can store real numbers or integers. The model supports standard operations on these numbers in constant time, including addition, multiplication, and elementary functions like square-root, sine or cosine. Furthermore, the integers can be used as indices to memory locations. Integers can be converted to real numbers in constant time, but we need to be careful about the reverse direction. The *floor* function can be used to truncate a real number to an integer, but if we were allowed to use it arbitrarily, the real RAM could solve PSPACE-complete problems in polynomial time. Therefore, we usually have only a restricted floor function at our disposal.

Word RAM. The *word RAM* is essentially a real RAM without support for real numbers. However, on a real RAM, the integers are usually treated as atomic, whereas the word RAM allows for powerful bit-manipulation tricks. More precisely, the word RAM represents the data as a sequence of w-bit words, where $w = \Omega(\log n)$. Data can be accessed arbitrarily, and standard operations, such as Boolean operations (`and`, `xor`, `shl`, ...), addition, or multiplication take constant time. There are many variants of the word RAM, depending on precisely which instructions are supported in constant time. The general consensus seems to be that any function in AC^0 is acceptable.[9] However, it is

always preferable to rely on a set of operations as small, and as non-exotic, as possible. Note that multiplication is not in AC^0, but nevertheless is usually included in the word RAM instruction set.

Pointer machine. The *pointer machine* model disallows the use of constant time table lookup, and is therefore a restriction of the (real) RAM model. The data structure is modeled as a directed graph G with bounded out-degree. Each node in G represents a *record*, with a bounded number of pointers to other records and a bounded number of (real or integer) data items. The algorithm can access data only by following pointers from the inputs (and a bounded number of global entry records); random access is not possible. The data can be manipulated through the usual real RAM operations, but without support for the floor function, for reasons mentioned above.

Algebraic computation tree. *Algebraic computation trees* (ACTs) [9] are the computational geometry analogue of binary decision trees, and like these they are mainly used for proving lower bounds. Let $x_1, \ldots, x_n \in \mathbb{R}$ be the inputs. An ACT is a binary tree with two kinds of nodes: *computation* and *branch* nodes. A computation node v has one child and is labeled with an expression of the type $y_v = y_u \oplus y_w$, where $\oplus \in \{+, -, *, /, \sqrt{\cdot}\}$ is a operator and y_u, y_w is either an input variable $x_1, \ldots, x_n$ or corresponds to a computation node that is an ancestor of v. A branch node has degree 2 and is labeled by $y_u = 0$ or $y_u > 0$, where again y_u is either an input or a variable for an ancestor. A family of algebraic computation trees $(T_n)_{n \in \mathbb{N}}$ solves a computational problem (like Delaunay triangulation or convex hull computation), if for each $n \in \mathbb{N}$, the tree T_n accepts inputs of size n, and if for any such input $x_1, \ldots, x_n$ the corresponding path in T_n (where the children of the branch nodes are determined according to the conditions they represent) constitutes a computation that represents the answer in the variables y_v encountered during the path.

Algebraic *decision* trees are defined as follows: we allow only branch nodes. Each branch node is labeled with a predicate of the form $p(x_1, \ldots, x_n) = 0$ or $p(x_1, \ldots, x_n) > 0$. The leaves are labeled *yes* or *no*. Fix some $r \in \{1, \ldots, n\}$. If p is restricted to be of the form $p(x_1, \ldots, x_n) = \sum_{i=1}^{n} a_i x_i - b$, with at most r coefficients $a_i \neq 0$, we call the decision tree r-*linear*. Erickson [34] showed that any 3-linear decision tree for 3SUM has depth $\Omega(n^2)$. However, this bound does not say anything about more general predicates (e.g., if p may include quadratic terms). This severely limits its

---

[9]AC^0 is the class of all functions $f \colon \{0,1\}^* \to \{0,1\}^*$ that can be computed by a family of circuits $(C_n)_{n \in \mathbb{N}}$ with the following properties: (i) each C_n has n inputs; (ii) there exist constants a, b, such that C_n has at most an^b gates, for $n \in \mathbb{N}$; (iii) there is a constant d such that for all n the length of the longest path from an input to an output in C_n is at most d (ie, the circuit family

has bounded depth); (iv) each gate has an arbitrary number of incoming edges (i.e., the fan-in is unbounded).

applicability to geometric problems. For example, there is no r-linear decision tree for the Fréchet problem, no matter the choice of r: with r-linear decision trees, we cannot even decide whether two given points p and q have Euclidean distance at most 1.