

Approximability of the Discrete Fréchet Distance

Karl Bringmann^{*1} and Wolfgang Mulzer^{†2}

1 Institute of Theoretical Computer Science, ETH Zurich, Switzerland
karlb@inf.ethz.ch

2 Institut für Informatik, Freie Universität Berlin, Germany
mulzer@inf.fu-berlin.de

Abstract

The Fréchet distance is a popular and widespread distance measure for point sequences and for curves. About two years ago, Agarwal et al [SIAM J. Comput. 2014] presented a new (mildly) subquadratic algorithm for the discrete version of the problem. This spawned a flurry of activity that has led to several new algorithms and lower bounds.

In this paper, we study the approximability of the discrete Fréchet distance. Building on a recent result by Bringmann [FOCS 2014], we present a new conditional lower bound that strongly subquadratic algorithms for the discrete Fréchet distance are unlikely to exist, even in the *one-dimensional* case and even if the solution may be approximated up to a factor of 1.399.

This raises the question of how well we can approximate the Fréchet distance (of two given d -dimensional point sequences of length n) in strongly subquadratic time. Previously, no general results were known. We present the first such algorithm by analysing the approximation ratio of a simple, linear-time greedy algorithm to be $2^{\Theta(n)}$. Moreover, we design an α -approximation algorithm that runs in time $O(n \log n + n^2/\alpha)$, for any $\alpha \in [1, n]$. Hence, an n^ε -approximation of the Fréchet distance can be computed in strongly subquadratic time, for any $\varepsilon > 0$.

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems—Geometrical problems and computations.

Keywords and phrases Fréchet distance, approximation, lower bounds, Strong Exponential Time Hypothesis

Digital Object Identifier 10.4230/LIPIcs.xxx.yyy.p

1 Introduction

Let P and Q be two polygonal curves with n vertices each. The *Fréchet distance* provides a meaningful way to define a distance between P and Q that overcomes some of the shortcomings of the classic Hausdorff distance [6]. Since its introduction to the computational geometry community by Alt and Godau [6], the concept of Fréchet distance has proven extremely useful and has found numerous applications (see [4, 6–10] and the references therein).

The Fréchet distance has two classic variants: *continuous* and *discrete* [6, 12]. In this paper, we focus on the discrete variant. In this case, the Fréchet distance between two sequences P, Q of n points in d dimensions is defined as follows: imagine two frogs traversing the sequences P and Q , respectively. In each time step, a frog can jump to the next vertex along its sequence, or it can stay where it is. The discrete Fréchet distance is the minimal length of a leash required to connect the two frogs while they traverse the two sequences from start to finish.

* Supported by an ETH Zurich Postdoctoral Fellowship.

† Supported in part by DFG Grants MU 3501/1 and MU 3501/2.



The original algorithm for the continuous Fréchet distance by Alt and Godau has running time $O(n^2 \log n)$ [6]; while the algorithm for the discrete Fréchet distance by Eiter and Mannila needs time $O(n^2)$ [12]. These algorithms have remained the state of the art until very recently: in 2013, Agarwal et al [4] presented a slightly subquadratic algorithm for the discrete Fréchet distance. Building on their work, Buchin et al [9] managed to find a slightly improved algorithm for the continuous Fréchet distance a year later. At the time, Buchin et al thought that their result provides evidence that computing the Fréchet distance may not be 3SUM-hard [13], as had previously been conjectured by Alt [5]. Even though a recent result by Grønlund and Pettie [15], showing that 3SUM has subquadratic decision trees, casts new doubt on the connection between 3SUM and the Fréchet distance, the conclusions of Buchin et al motivated Bringmann [7] to look for other explanations for the apparent difficulty of the Fréchet distance.

He found a possible explanation in the *Strong Exponential Time Hypothesis* (SETH) [16,17], which roughly speaking asserts that satisfiability cannot be decided in time¹ $O^*((2 - \varepsilon)^n)$ for any $\varepsilon > 0$ (see Section 2 for details). Since exhaustive search takes time $O^*(2^n)$ and since the fastest known algorithms are only slightly faster than that, SETH is a reasonable assumption that formalizes an algorithmic barrier. It has been shown that SETH can be used to prove conditional lower bounds even for polynomial time problems [1, 2, 18, 20]. In this line of research, Bringmann [7] showed, among other things, that there are no strongly subquadratic algorithms for the Fréchet distance unless SETH fails. Here, *strongly subquadratic* means any running time of the form $O(n^{2-\varepsilon})$, for constant $\varepsilon > 0$. Bringmann's lower bound works for two-dimensional curves and both classic variants of the Fréchet distance. Thus, it is unlikely that the algorithms by Agarwal et al and Buchin et al can be improved significantly, unless a major algorithmic breakthrough occurs.

1.1 Our Contributions

In this extended abstract we focus on the discrete Fréchet distance. In Section 6, we will discuss how far our results carry over to the continuous version. Our main results are as follows.

Conditional Lower Bound. We strengthen the result of Bringmann [7] by showing that even in the one-dimensional case computing the Fréchet distance remains hard. More precisely, we show that any 1.399-approximation algorithm in strongly subquadratic time for the one-dimensional discrete Fréchet distance violates the Strong Exponential Time Hypothesis. Previously, Bringmann [7] had shown that no strongly subquadratic algorithm approximates the two-dimensional Fréchet distance by a factor of 1.001, unless SETH fails.

One can embed any one-dimensional sequence into the two-dimensional plane by fixing some $\varepsilon > 0$ and by setting the y -coordinate of the i -th point of the sequence to $i \cdot \varepsilon$. For sufficiently small ε , this embedding roughly preserves the Fréchet distance. Thus, unless SETH fails, there is also no strongly subquadratic 1.399-approximation for the discrete Fréchet distance on (1) two-dimensional curves without self-intersections, and (2) two-dimensional x -monotone curves (also called time-series). These interesting special cases had been open.

Approximation: Greedy Algorithm. A simple greedy algorithm for the discrete Fréchet distance goes as follows: in every step, make the move that minimizes the current distance,

¹ The notation $O^*(\cdot)$ hides polynomial factors in the number of variables n and the number of clauses m .

where a “move” is a step in either one sequence or in both of them. This algorithm has a straightforward linear time implementation. We analyze the approximation ratio of the greedy algorithm, and we show that, given two sequences of n points in d dimensions, the maximal distance attained by the greedy algorithm is a $2^{\Theta(n)}$ -approximation for their discrete Fréchet distance. We emphasize that this approximation ratio is *bounded*, depending only on n , but not the coordinates of the vertices. This is surprising, since so far no bounded approximation algorithm that runs in strongly subquadratic time was known at all. Moreover, although an approximation ratio of $2^{\Theta(n)}$ is huge, the greedy algorithm is the best *linear time* approximation algorithm that we could come up with.

Approximation: Improved Algorithm. For the case that slightly more than linear time is acceptable, we provide a much better approximation algorithm: given two sequences P and Q of n points in d dimensions, we show how to find an α -approximation of the discrete Fréchet distance between P and Q in time $O(n \log n + n^2/\alpha)$, for any $1 \leq \alpha \leq n$. In particular, this yields an $n/\log n$ -approximation in time $O(n \log n)$, and an n^ε -approximation in strongly subquadratic time for any $\varepsilon > 0$. We leave it open whether these approximation ratios can be improved.

2 Preliminaries and Definitions

We call an algorithm an α -approximation for the Fréchet distance if, given curves P, Q , it returns a number that is at least the Fréchet distance of P, Q and at most α times the Fréchet distance of P, Q .

2.1 Discrete Fréchet Distance

Since we focus on the discrete Fréchet distance throughout, we will sometimes omit the term “discrete”. Let $P = \langle p_1, \dots, p_n \rangle$ and $Q = \langle q_1, \dots, q_n \rangle$ be two sequences of n points in d dimensions. A *traversal* β of P and Q is a sequence of pairs in $(p, q) \in P \times Q$ such that (i) the traversal β begins with the pair (p_1, q_1) and ends with the pair (p_n, q_n) ; and (ii) the pair $(p_i, q_j) \in \beta$ can be followed only by one of (p_{i+1}, q_j) , (p_i, q_{j+1}) , or (p_{i+1}, q_{j+1}) . We call β *simultaneous*, if it only makes steps of the third kind, i.e., if in each step β advances in both P and Q . We define the *distance* of the traversal β as $\delta(\beta) := \max_{(p,q) \in \beta} d(p, q)$, where $d(\cdot, \cdot)$ denotes the Euclidean distance. The *discrete Fréchet distance* of P and Q is now defined as $\delta_{\text{dF}}(P, Q) := \min_{\beta} \delta(\beta)$, where β ranges over all traversals of P and Q .

We review a simple $O(n^2 \log n)$ time algorithm to compute $\delta_{\text{dF}}(P, Q)$ that is the starting point of our second approximation algorithm. First, we describe a *decision procedure* that, given a value γ , decides whether $\delta_{\text{dF}}(P, Q) \leq \gamma$. For this, we define the *free-space matrix* F . This is a Boolean $n \times n$ matrix such that for $i, j = 1, \dots, n$, we set $F_{ij} = 1$ if $d(p_i, q_j) \leq \gamma$, and $F_{ij} = 0$, otherwise. Then $\delta_{\text{dF}}(P, Q) \leq \gamma$ if and only if F allows a *monotone traversal* from $(1, 1)$ to (n, n) , i.e., if we can go from entry F_{11} to F_{nn} while only going down, to the right, or diagonally, and while only using 1-entries. This is captured by the *reach matrix* R , which is again an $n \times n$ Boolean matrix. We set $R_{11} = F_{11}$, and for $i, j = 1, \dots, n$, $(i, j) \neq (1, 1)$, we set $R_{ij} = 1$ if $F_{ij} = 1$ and either one of $R_{(i-1)j}$, $R_{i(j-1)}$, or $R_{(i-1)(j-1)}$ equals 1 (we define any entry of the form $R_{(-1)j}$ or $R_{i(-1)}$ to be 0). Otherwise, we set $R_{ij} = 0$. From these definitions, it is straightforward to compute F and R in total time $O(n^2)$. Furthermore, by construction we have $\delta_{\text{dF}}(P, Q) \leq \gamma$ if and only if $R_{nn} = 1$.

With this decision procedure at hand, we can use binary search to compute $\delta_{\text{dF}}(P, Q)$ in total time $O(n^2 \log n)$ by observing that the optimum must be achieved for one of the n^2

distances $d(p_i, q_j)$, for $i, j = 1, \dots, n$. Through a more direct use of dynamic programming, the running time can be reduced to $O(n^2)$ [12].

2.2 Hardness Assumptions

Strong Exponential Time Hypothesis (SETH). As is well-known, the k -SAT problem is as follows: given a CNF-formula Φ over boolean variables x_1, \dots, x_n with clause width k , decide whether there is an assignment of x_1, \dots, x_n that satisfies Φ . Of course, k -SAT is NP-hard, and it is conjectured that no subexponential algorithm for the problem exists [14]. The Strong Exponential Time Hypothesis (SETH) goes a step further and basically states that the exhaustive search running time of $O^*(2^n)$ cannot be improved to $O^*(1.99^n)$ [16, 17].

► **Conjecture 2.1 (SETH).** For no $\varepsilon > 0$, k -SAT has an $O(2^{(1-\varepsilon)n})$ algorithm for all $k \geq 3$.

The fastest known algorithms for k -SAT take time $O(2^{(1-c/k)n})$ for some constant $c > 0$ [19]. Thus, SETH is reasonable and, due to lack of progress in the last decades, can be considered unlikely to fail. It is by now a standard assumption for conditional lower bounds.


Orthogonal Vectors (OV) is the following problem: Given $u_1, \dots, u_N, v_1, \dots, v_N \in \{0, 1\}^D$, decide whether there are $i, j \in [N]$ with $(u_i)_k \cdot (v_j)_k = 0$ for all $1 \leq k \leq D$. Here we denote by u_i the i -th vector and by $(u_i)_k$ its k -th entry. We write $u_i \perp v_j$ if $(u_i)_k \cdot (v_j)_k = 0$ for all $1 \leq k \leq D$. This problem has a trivial $O(DN^2)$ algorithm. The fastest known algorithm runs in time $N^{2-1/O(\log(D/\log N))}$ [3], which is only slightly subquadratic for $D \gg \log N$. It is known that OV has no strongly subquadratic time algorithms unless SETH fails [21]; we present a proof for completeness.

► **Lemma 2.2.** For no $\varepsilon > 0$, OV has an $D^{O(1)} \cdot N^{2-\varepsilon}$ algorithm, unless SETH fails.

Proof. Given a k -SAT formula Φ on variables x_1, \dots, x_n and clauses C_1, \dots, C_m , we build an equivalent OV-instance with $N = 2^{n/2}$ and $D = m$. Denote all possible assignments of true and false to the variables $x_1, \dots, x_{n/2}$ (the first half of the variables) by ϕ_1, \dots, ϕ_N , $N = 2^{n/2}$. For every such assignment ϕ_i we construct a vector u_i where $(u_i)_k$ is 0 if ϕ_i causes C_k to evaluate to true, and 1 otherwise. Similarly, we enumerate the assignments ψ_1, \dots, ψ_N of the variables $x_{n/2+1}, \dots, x_n$ (the second half of the variables), and for every ψ_j we construct a vector v_j where $(v_j)_k$ is 0 if ψ_j causes C_k to evaluate to true, and 1 otherwise. Then, $(u_i)_k \cdot (v_j)_k = 0$ if and only if one of ϕ_i and ψ_j satisfies clause C_k . Thus, we have $(u_i)_k \cdot (v_j)_k = 0$ for all $1 \leq k \leq D$ if and only if (ϕ_i, ψ_j) forms a satisfying assignment of the formula Φ . Hence, we constructed an equivalent OV-instance of the required size. The constructed OV instance can be computed in time $O(DN)$.

It follows that any algorithm for OV with running time $D^{O(1)} \cdot N^{2-\varepsilon}$ gives an algorithm for k -SAT with running time $m^{O(1)} 2^{(1-\varepsilon/2)n}$. Since $m \leq n^k \leq 2^{o(n)}$, this contradicts SETH. ◀

A problem P is *OV-hard* if there is a reduction that transforms an instance I of OV with parameters N, D , to an equivalent instance I' of P of size $n \leq D^{O(1)}N$, in time $D^{O(1)}N^{2-\varepsilon}$ for some $\varepsilon > 0$. A strongly subquadratic algorithm (i.e., $O(n^{2-\varepsilon'})$ for some $\varepsilon' > 0$) for an OV-hard problem P would yield an algorithm for OV with running time $D^{O(1)}N^{2-\min\{\varepsilon, \varepsilon'\}}$. Thus, by Lemma 2.2 an OV-hard problem does not have strongly subquadratic time algorithms unless SETH fails. Most known SETH-based lower bounds for polynomial time problems are actually OV-hardness results; our lower bound in the next section is no exception. Note that OV-hardness is potentially stronger than a SETH-based lower bound, since it is conceivable that SETH fails, but OV still has no strongly subquadratic algorithms.



■ **Figure 1** The point set \mathcal{P} constructed in the conditional lower bound.

3 Hardness of Approximation in One Dimension

We prove OV-hardness of the discrete Fréchet distance on one-dimensional curves. By Lemma 2.2, this also yields a SETH-based lower bound.

Let $u_1, \dots, u_N, v_1, \dots, v_N \in \{0, 1\}^D$ be an instance of the Orthogonal Vectors problem. Without loss of generality, we assume that D is even (if not, we duplicate a dimension). We show how to construct two sequences P and Q of $O(DN)$ points in \mathbb{R} in time $O(DN)$ such that there are $i, j \in \{1, \dots, N\}$ with $u_i \perp v_j$ if and only if $\delta_{\text{dF}}(P, Q) \leq 1$. Our sequences P and Q consist of elements from the following set \mathcal{P} of 13 points; see Figure 1.

- $a_0^o = -0.8, a_1^o = -0.4, a_1^e = 0.4, a_0^e = 0.8.$
- $b_1^o = -1.8, b_0^o = -1.4, b_0^e = 1.4, b_1^e = 1.8.$
- $s = 0, x_1 = -1.2, x_2 = 1.2$
- $w_1 = -0.2, w_2 = 0.2.$

We first construct *vector gadgets*. For each $u_i, i \in \{1, \dots, N\}$, we define a sequence A_i of D points from \mathcal{P} as follows: For $1 \leq k \leq D$ let $p \in \{o, e\}$ be the parity of k (odd or even). Then the k th point of A_i is $a_{(u_i)_k}^p$. Similarly, for each v_j , we define a sequence B_j of D points from \mathcal{P} . For B_j , we use the points b_*^p instead of a_*^p . The next claim shows how the vector gadgets encode orthogonality.

► **Claim 3.1.** Fix $i, j \in \{1, \dots, N\}$ and let β be a traversal of (A_i, B_j) . (i) If β is not a simultaneous traversal, then $\delta(\beta) \geq 1.8$; (ii) if β is a simultaneous traversal and $u_i \perp v_j$, then $\delta(\beta) \leq 1$; and (iii) if β is a simultaneous traversal and $u_i \not\perp v_j$, then $\delta(\beta) = 1.4$.

Proof. First, suppose that β is not a simultaneous traversal. Consider the first time when β makes a move on one sequence but not the other. Then, the current points on A_i and B_j lie on different sides of s , which forces $\delta(\beta) \geq \min\{d(a_1^o, b_0^e), d(a_1^e, b_0^o)\} = 1.8$.

Next, suppose that $u_i \perp v_j$. Then, the simultaneous traversal β of A_i and B_j has $\delta(\beta) \leq 1$. Indeed, for each dimension $1 \leq k \leq D$ at least one of $(u_i)_k$ and $(v_j)_k$ is 0. Thus, if we consider the k th point of A_i and the k th point of B_j , both of them lie on the same side of s , and at least one of them is in $\{a_0^o, a_0^e, b_0^o, b_0^e\}$. It follows that the distance between the k th points in β is at most 1, for all k .

Finally, suppose that $(u_i)_k = (v_j)_k = 1$ for some k . Let β be the simultaneous traversal of A_i and B_j , and consider the time when β reaches the k th points of A_i and B_j . These are either $\{a_1^o, b_1^o\}$ or $\{a_1^e, b_1^e\}$, so $\delta(\beta) = \min\{d(a_1^o, b_1^o), d(a_1^e, b_1^e)\} = 1.4$. ◀

Let W be the sequence of $D(N-1)$ points that alternates between a_0^o and a_0^e , starting with a_0^o . We set

$$P = W \cdot x_1 \cdot s \cdot A_1 \cdot s \cdot A_2 \cdot \dots \cdot s \cdot A_N \cdot s \cdot x_2 \cdot W$$

and

$$Q = w_1 \cdot B_1 \cdot w_2 \cdot w_1 \cdot B_2 \cdot w_2 \cdot \dots \cdot w_1 \cdot B_N \cdot w_2,$$

where \cdot denotes the concatenation of sequences. The idea is to implement an *or-gadget*. If there is a pair of orthogonal vectors, then P and Q should be able to reach the corresponding vector gadgets and traverse them simultaneously. If there is no such pair, it should not be possible to “cheat”. The purpose of the sequences W and the points w_1 and w_2 is to provide a buffer so that one sequence can wait while the other sequence catches up. The purpose of the points x_1 , x_2 , and s is to synchronize the traversal so that no cheating can occur. The next two claims make this precise. First, we show completeness.

► **Claim 3.2.** If there are $i, j \in \{1, \dots, N\}$ with $u_i \perp v_j$, then $\delta_{\text{dF}}(P, Q) \leq 1$.

Proof. Let u_i, v_j be orthogonal. We traverse P and Q as follows:

1. P goes through $D(N - j)$ points of W ; Q stays at w_1 .
2. For $k = 1, \dots, j - 1$: we perform a simultaneous traversal of B_k and the next portion of W starting with a_0^o and the first point on B_k . When the traversal reaches a_0^e and the last point of B_k , P stays at a_0^e while Q goes to w_2 and w_1 . If $k < j - 1$, the traversal continues with a_0^o on P and the first point of B_{k+1} on Q . If $k = j - 1$, we go to Step 3.
3. P proceeds to x_1 and walks until the point s before A_i , Q stays at w_1 before B_j .
4. P and Q go simultaneously through A_i and B_j , until the pair (s, w_2) after A_i and B_j .
5. P continues to x_2 while Q stays at w_2 .
6. For $k = j + 1, \dots, N$: P goes to the next a_0^o on W while Q goes to w_1 . We then perform a simultaneous traversal of B_k and the next portion of W . When the traversal reaches a_0^e and the last point of B_k , P stays at a_0^e while Q continues to w_2 . If $k < N$, the traversal continues with the next iteration, otherwise we go to Step 7.
7. P finishes the traversal of W while Q stays at w_2 .

We use the notation $\max\text{-}d(S, T) := \max_{s \in S, t \in T} d(s, t)$, and $\max\text{-}d(s, T) := \max\text{-}d(\{s\}, T)$, $\max\text{-}d(S, t) := \max\text{-}d(S, \{t\})$. The traversal maintains a maximum distance of 1: for Step 1, this is implied by $\max\text{-}d(\{a_0^o, a_0^e\}, w_1) = 1$. For Step 2, it follows from D being even and from

$$\max\text{-}d(a_0^o, \{b_1^o, b_0^o\}) = \max\text{-}d(a_0^e, \{b_1^e, b_0^e, w_1, w_2\}) = 1.$$

For Step 3, it is because $\max\text{-}d(\{x_1, a_0^o, a_1^o, s, a_1^e, a_0^e\}, w_1) = 1$. For Step 4, we use Claim 3.1 and $d(s, w_2) = 0.2$. In Step 5, it follows from $\max\text{-}d(\{a_0^o, a_1^o, s, a_1^e, a_0^e, x_2\}, w_2) = 1$. In Step 6, we again use that D is even and that

$$\max\text{-}d(a_0^o, \{b_1^o, b_0^o, w_1\}) = \max\text{-}d(a_0^e, \{b_1^e, b_0^e, w_2\}) = 1.$$

Step 7 uses $\max\text{-}d(\{a_0^o, a_0^e\}, w_2) = 1$. ◀

The second claim establishes the soundness of the construction.

► **Claim 3.3.** If there are no $i, j \in \{1, \dots, N\}$ with $u_i \perp v_j$, then $\delta_{\text{dF}}(P, Q) \geq 1.4$.

Proof. Let β be a traversal of (P, Q) . Consider the time when β reaches x_1 on P . If Q is not at either w_1 or at a point from $B^o = \{b_0^o, b_1^o\}$, then $\delta(\beta) \geq 1.4$ and we are done. Next, suppose that the current position is in $\{x_1\} \times B^o$. In the next step, β must advance P to s or Q to $\{b_0^e, b_1^e\}$ (or both).² In each case, we get $\delta(\beta) \geq 1.4$. From now on, suppose we reach x_1 in position (x_1, w_1) . After that, P must advance to s , because advancing Q to B^o would take us to a position in $\{x_1\} \times B^o$, implying $\delta(\beta) \geq 1.4$ as we saw above.

² Recall that we assumed D to be even.

Now consider the next step when Q leaves w_1 . Then Q must go to a point from B^o . At this time, P must be at a point from $A^o = \{a_0^o, a_1^o\}$ or we would get $\delta(\beta) \geq 1.4$ (note that P has already passed the point x_1). This point on P belongs to a vector gadget A_i or to the final gadget W (again because P is already past x_1). In the latter case, we have $\delta(\beta) \geq 1.4$, because in order to reach the final W , P must have gone through x_2 and $d(x_2, w_1) = 1.4$. Thus, P is at a point in A^o in a vector gadget A_i , and Q is at the starting point (from B^o) of a vector gadget B_j .

Now β must alternate simultaneously in P and Q among both sides of s , or again $\delta(\beta) \geq 1.4$, see Claim 3.1. Furthermore, if P does not start in the first point of A_i , then eventually P has to go to s while Q has to go to a point in B^o or stay in $\{b_0^e, b_1^e\}$, giving $\delta(\beta) \geq 1.4$. Thus, we may assume that β simultaneously reached the starting points of A_i and B_j and traverses A_i and B_j simultaneously. By assumption, the vectors u_i, v_j are not orthogonal, so Claim 3.1 gives $\delta(\beta) \geq 1.4$. ◀

► **Theorem 3.4.** *Fix $\alpha \in [1, 1.4)$. Computing an α -approximation of the discrete Fréchet distance in one dimension is OV-hard. In particular, the discrete Fréchet distance in one dimension has no strongly subquadratic α -approximation unless SETH fails.*

Proof. We use Claims 3.2 and 3.3 and the fact that P and Q can be computed in time $O(DN)$ from $u_1, \dots, u_N, v_1, \dots, v_N$: any $O(n^{2-\varepsilon})$ time α -approximation for the discrete Fréchet distance would yield an OV algorithm in time $D^{O(1)}N^{2-\varepsilon}$, which by Lemma 2.2 contradicts SETH. ◀

► **Remark.** The proofs of Claims 3.2 and 3.3 yield a system of linear inequalities that constrain the points in \mathcal{P} . Using this system, one can see that the inapproximability factor 1.4 in Theorem 3.4 is best possible for our current proof.

4 Approximation Quality of the Greedy Algorithm

In this section we study the following greedy algorithm. Let $P = \langle p_1, \dots, p_n \rangle$ and $Q = \langle q_1, \dots, q_n \rangle$ be two sequences of n points in \mathbb{R}^d . We construct a traversal $\beta_{\text{greedy}} = \beta_{\text{greedy}}(P, Q)$. We begin at (p_1, q_1) . If the current position is (p_i, q_j) , there are at most three possible successor configurations: (p_{i+1}, q_j) , (p_i, q_{j+1}) , and (p_{i+1}, q_{j+1}) (or fewer, if we have already reached the last point from P or Q). Among these, we pick the pair $(p_{i'}, q_{j'})$ that minimizes the distance $d(p_{i'}, q_{j'})$. We stop when we reach (p_n, q_n) . We denote the largest distance taken by the greedy traversal by $\delta_{\text{greedy}}(P, Q) := \delta(\beta_{\text{greedy}}(P, Q))$.

► **Theorem 4.1.** *Let P and Q be two sequences of n points in \mathbb{R}^d . Then, $\delta_{\text{dF}}(P, Q) \leq \delta_{\text{greedy}}(P, Q) \leq 2^{O(n)}\delta_{\text{dF}}(P, Q)$. Both inequalities are tight, i.e., there are polygonal curves P, Q with $\delta_{\text{greedy}}(P, Q) = \delta_{\text{dF}}(P, Q) > 0$ and $\delta_{\text{greedy}}(P, Q) = 2^{\Omega(n)}\delta_{\text{dF}}(P, Q) > 0$, respectively.*

The inequality $\delta_{\text{dF}}(P, Q) \leq \delta_{\text{greedy}}(P, Q)$ follows directly from the definition, since the traversal $\beta_{\text{greedy}}(P, Q)$ is a candidate for an optimal traversal. Furthermore, one can check that if P and Q are increasing one-dimensional sequences, then the greedy traversal is optimal (this is similar to the merge step in mergesort). Thus, there are examples where $\delta_{\text{greedy}}(P, Q) = \delta_{\text{dF}}(P, Q)$. It remains to show the upper bound $\delta_{\text{greedy}}(P, Q) \leq 2^{O(n)}\delta_{\text{dF}}(P, Q)$ and to provide an example where this inequality is tight.

4.1 Upper Bound

We call a pair $p_i p_{i+1}$ of consecutive points on P an *edge* of P , for $i = 1, \dots, n-1$, and similarly for Q . Let m be the total number of edges of P and Q , and let $\ell_1 \leq \ell_2 \leq \dots \leq \ell_m$ be the sorted sequence of the edge lengths. We pick $k^* \in \{0, \dots, m\}$ minimum such that

$$4\delta_{\text{dF}}(P, Q) + 2 \sum_{i=1}^{k^*} \ell_i < \ell_{k^*+1},$$

where we set $\ell_{m+1} = \infty$. We define δ^* as the left hand side, $\delta^* := 4\delta_{\text{dF}}(P, Q) + 2 \sum_{i=1}^{k^*} \ell_i$.

► **Lemma 4.2.** *We have (i) $\delta^* \geq 4\delta_{\text{dF}}(P, Q)$; (ii) $\sum_{i=1}^{k^*} \ell_i \leq \delta^*/2 - 2\delta_{\text{dF}}(P, Q)$; (iii) there is no edge with length in $(\delta^*/2 - 2\delta_{\text{dF}}(P, Q), \delta^*)$; and (iv) $\delta^* \leq 3^{k^*} 4\delta_{\text{dF}}(P, Q)$.*

Proof. Properties (i) and (ii) follow by definition. Property (iii) holds since for $i = 1, \dots, k^*$, we have $\ell_i \leq \delta^*/2 - 2\delta_{\text{dF}}(P, Q)$, by (ii), and for $i = k^* + 1, \dots, m$, we have $\ell_i \geq \delta^*$, by definition. It remains to prove (iv): for $k = 0, \dots, k^*$, we set $\delta_k = 4\delta_{\text{dF}}(P, Q) + 2 \sum_{i=1}^k \ell_i$, and we prove by induction that $\delta_k \leq 3^k 4\delta_{\text{dF}}(P, Q)$. For $k = 0$, this is immediate. Now suppose we know that $\delta_{k-1} \leq 3^{k-1} 4\delta_{\text{dF}}(P, Q)$, for some $k \in \{1, \dots, k^*\}$. Then, $k \leq k^*$ implies $\ell_k \leq \delta_{k-1}$, so $\delta_k = \delta_{k-1} + 2\ell_k \leq 3\delta_{k-1} \leq 3^k 4\delta_{\text{dF}}(P, Q)$, as desired. Now (iv) follows from $\delta^* = \delta_{k^*}$. ◀

We call an edge *long* if it has length at least δ^* , and *short* otherwise. In other words, the short edges have lengths $\ell_1, \dots, \ell_{k^*}$, and the long edges have lengths $\ell_{k^*+1}, \dots, \ell_m$. Let β be an optimal traversal of P and Q , i.e., $\delta(\beta) = \delta_{\text{dF}}(P, Q)$.

► **Lemma 4.3.** *The sequences P and Q have the same number of long edges. Furthermore, if $p_{i_1} p_{i_1+1}, \dots, p_{i_k} p_{i_k+1}$ and $q_{j_1} q_{j_1+1}, \dots, q_{j_k} q_{j_k+1}$ are the long edges of P and of Q , for $1 \leq i_1 < \dots < i_k < n$ and $1 \leq j_1 < \dots < j_k < n$, then both β and β_{greedy} contain the steps $(p_{i_1}, q_{j_1}) \rightarrow (p_{i_1+1}, q_{j_1+1}), \dots, (p_{i_k}, q_{j_k}) \rightarrow (p_{i_k+1}, q_{j_k+1})$.*

Proof. First, we show that for every long edge $p_i p_{i+1}$ of P , the optimal traversal β contains the step $(p_i, q_j) \rightarrow (p_{i+1}, q_{j+1})$, where q_j, q_{j+1} is a long edge of Q . Consider the step of β from p_i to p_{i+1} . This step has to be of the form $(p_i, q_j) \rightarrow (p_{i+1}, q_{j+1})$ for some $q_j \in Q$: since $\max\{d(p_i, q_j), d(p_{i+1}, q_j)\} \geq d(p_i, p_{i+1})/2 \geq \delta^*/2 \geq 2\delta_{\text{dF}}(P, Q)$, by Lemma 4.2(i), staying in q_j would result in $\delta(\beta) \geq 2\delta_{\text{dF}}(P, Q)$. Now, since $\max\{d(p_i, q_j), d(p_{i+1}, q_{j+1})\} \leq \delta(\beta) = \delta_{\text{dF}}(P, Q)$, the triangle inequality gives $d(q_j, q_{j+1}) \geq d(p_i, p_{i+1}) - 2\delta_{\text{dF}}(P, Q) \geq \delta^* - 2\delta_{\text{dF}}(P, Q)$. Lemma 4.2(iii) now implies $d(q_j, q_{j+1}) \geq \delta^*$, so the edge $q_j q_{j+1}$ is long.

Thus, β traverses every long edge of P simultaneously with a long edge of Q . A symmetric argument shows that β traverses every long edge of Q simultaneously with a long edge of P . Since β is monotone, it follows that P and Q have the same number of long edges, and that β traverses them simultaneously in their order of appearance along P and Q .

It remains to show that the greedy traversal β_{greedy} traverses the long edges of P and Q simultaneously. Set $i_0 = j_0 = 0$. We will prove for $a \in \{0, \dots, k-1\}$ that if β_{greedy} contains the position (p_{i_a+1}, q_{j_a+1}) , then it also contains the step $(p_{i_{a+1}}, q_{j_{a+1}}) \rightarrow (p_{i_{a+1}+1}, q_{j_{a+1}+1})$ and hence the position $(p_{i_{a+1}+1}, q_{j_{a+1}+1})$. The claim on β_{greedy} then follows by induction on a , since β_{greedy} contains the position (p_1, q_1) by definition. Thus, fix $a \in \{0, \dots, k-1\}$ and suppose that β_{greedy} contains (p_{i_a+1}, q_{j_a+1}) . We need to show that β_{greedy} also contains the step $(p_{i_{a+1}}, q_{j_{a+1}}) \rightarrow (p_{i_{a+1}+1}, q_{j_{a+1}+1})$. For better readability, we write i for i_a , j for j_a , i' for i_{a+1} , and j' for j_{a+1} . Consider the first position of β_{greedy} when β_{greedy} reaches either $p_{i'}$ or $q_{j'}$. Without loss of generality, this position is of the

from $(p_{i'}, q_l)$, for some $l \in \{j+1, \dots, j'\}$. Then, $d(p_{i'}, q_l) \leq \delta^*/2 - \delta_{\text{dF}}(P, Q)$, since we saw that $d(p_{i'}, q_{j'}) \leq \delta(\beta) = \delta_{\text{dF}}(P, Q)$ and since the remaining edges between q_l and $q_{j'}$ are short and thus have total length at most $\delta^*/2 - 2\delta_{\text{dF}}(P, Q)$, by Lemma 4.2(ii). The triangle inequality now gives $d(p_{i'+1}, q_l) \geq d(p_{i'}, p_{i'+1}) - d(p_{i'}, q_l) \geq \delta^*/2 + \delta_{\text{dF}}(P, Q)$. If $l < j'$, the same argument applied to q_{l+1} shows that $d(p_{i'}, q_{l+1}) \leq \delta^*/2 - \delta_{\text{dF}}(P, Q)$ and thus $d(p_{i'+1}, q_{l+1}) \geq \delta^*/2 + \delta_{\text{dF}}(P, Q)$. Thus, β_{greedy} moves to $(p_{i'}, q_{l+1})$. If $l = j'$, then β_{greedy} takes the step $(p_{i'}, q_{j'}) \rightarrow (p_{i'+1}, q_{j'+1})$, as $d(p_{i'+1}, q_{j'+1}) \leq \delta(\beta) = \delta_{\text{dF}}(P, Q)$, but $d(p_{i'}, q_{j'+1}), d(p_{i'+1}, q_{j'}) \geq \delta^* - \delta_{\text{dF}}(P, Q) \geq 3\delta_{\text{dF}}(P, Q)$, by Lemma 4.2(i). \blacktriangleleft

Finally, we can show the desired upper bound on the quality of the greedy algorithm.

► Lemma 4.4. *We have $\delta_{\text{greedy}}(P, Q) \leq \delta^*/2$.*

Proof. By Lemma 4.3, P and Q have the same number of long edges. Let $p_{i_1}p_{i_1+1}, \dots, p_{i_k}p_{i_k+1}$ and $q_{j_1}q_{j_1+1}, \dots, q_{j_k}q_{j_k+1}$ be the long edges of P and of Q , where $1 \leq i_1 < \dots < i_k < n$ and $1 \leq j_1 < \dots < j_k < n$. By Lemma 4.3, β_{greedy} contains the positions (p_{i_a}, q_{j_a}) and (p_{i_a+1}, q_{j_a+1}) for $a = 1, \dots, k$, and $d(p_{i_a}, q_{j_a}), d(p_{i_a+1}, q_{j_a+1}) \leq \delta_{\text{dF}}(P, Q)$ for $a = 1, \dots, k$. Thus, setting $i_0 = j_0 = 0$ and $i_{k+1} = j_{k+1} = n$, we can focus on the subtraversals $\beta_a = (p_{i_a+1}, q_{i_a+1}), \dots, (p_{i_{a+1}}, q_{i_{a+1}})$ of β_{greedy} , for $a = 0, \dots, k$. Now, since all edges traversed in β_a are short, and since $d(p_{i_a+1}, q_{i_a+1}) \leq \delta_{\text{dF}}(P, Q)$, we have $\delta(\beta_a) \leq \delta_{\text{dF}}(P, Q) + \delta^*/2 - 2\delta_{\text{dF}}(P, Q) \leq \delta^*/2$ by Lemma 4.2(iii) and the triangle inequality. Thus, $\delta(\beta_{\text{greedy}}) \leq \max\{\delta_{\text{dF}}(P, Q), \delta(\beta_1), \dots, \delta(\beta_k)\} \leq \delta^*/2$, as desired. \blacktriangleleft

Lemmas 4.2(iv) and 4.4 prove the desired inequality $\delta_{\text{greedy}}(P, Q) \leq 2^{O(n)}\delta_{\text{dF}}(P, Q)$, since $k^* \leq m = 2n - 2$.

4.2 Tight Example for the Upper Bound

Fix $1 < \alpha < 2$. Consider the sequence $P = \langle p_1, \dots, p_n \rangle$ with $p_i := (-\alpha)^i$ and the sequence $Q = \langle q_1, \dots, q_{n-2} \rangle$ with $q_i := (-\alpha)^{i+2}$. We show the following:

1. The greedy traversal $\beta_{\text{greedy}}(P, Q)$ makes $n - 2$ simultaneous steps in P and Q followed by 2 single steps in P . This results in a maximal distance of $\delta_{\text{greedy}}(P, Q) = \alpha^n + \alpha^{n-1}$.
2. The traversal which makes 2 single steps in P followed by $n - 2$ simultaneous steps in both P and Q has distance $\alpha^3 + \alpha^2$.

Together, this shows that $\delta_{\text{greedy}}(P, Q)/\delta_{\text{dF}}(P, Q) = \Omega(\alpha^n) = 2^{\Omega(n)}$, proving that the inequality $\delta_{\text{greedy}}(P, Q) \leq 2^{O(n)}\delta_{\text{dF}}(P, Q)$ is tight.

To see (1), assume that we are at position (p_i, q_i) . Moving to (p_i, q_{i+1}) would result in a distance of $d(p_i, q_{i+1}) = \alpha^{i+3} + \alpha^i$. Similarly, the other possible moves to (p_{i+1}, q_i) and to (p_{i+1}, q_{i+1}) would result in distances $\alpha^{i+2} + \alpha^{i+1}$, and $\alpha^{i+3} - \alpha^{i+1}$, respectively. It can be checked that for all $\alpha > 1$ we have $\alpha^{i+3} + \alpha^i > \alpha^{i+2} + \alpha^{i+1}$. Moreover, for all $\alpha < 2$ we have $\alpha^{i+2} + \alpha^{i+1} > \alpha^{i+3} - \alpha^{i+1}$. Thus, the greedy algorithm makes the move to (p_{i+1}, q_{i+1}) . Using induction, this shows that the greedy traversal starts with $n - 2$ simultaneous moves in P and Q . In the end, the greedy algorithm has to take two single moves in P . Thus, the greedy traversal contains the pair (p_{n-1}, q_{n-2}) , which is in distance $d(p_{n-1}, q_{n-2}) = \alpha^n + \alpha^{n-1} = 2^{\Omega(n)}$.

To see (2), note that the traversal which makes 2 single steps in P followed by $n - 2$ simultaneous moves in P and Q starts with (p_1, q_1) and (p_2, q_1) followed by (p_i, q_{i-2}) for $i = 2, \dots, n$. Note that $d(p_1, q_1) = \alpha^3 - \alpha$, $d(p_2, q_1) = \alpha^3 + \alpha^2$, and $p_i = q_{i-2}$, so that the remaining distances are 0. Thus, we have $\delta_{\text{dF}}(P, Q) \leq \alpha^3 + \alpha^2 = O(1)$.

5 Improved Approximation Algorithm

Let $P = \langle p_1, \dots, p_n \rangle$ and $Q = \langle q_1, \dots, q_n \rangle$ be two sequences of n points in \mathbb{R}^d , where d is constant. Let $1 \leq \alpha \leq n$. We show how to find a value δ^* with $\delta_{\text{dF}}(P, Q) \leq \delta^* \leq \alpha \delta_{\text{dF}}(P, Q)$ in time $O(n \log n + n^2/\alpha)$. For simplicity, we will assume that all points on P and Q are pairwise distinct. This can be achieved by an infinitesimal perturbation of the point set.

5.1 Decision Algorithm

We begin by describing an approximate decision procedure. For this, we prove the following theorem.

► **Theorem 5.1.** *Let P and Q be two sequences of n points in \mathbb{R}^d , and let $1 \leq \alpha \leq n$. Suppose that the points of P and Q have been sorted along each coordinate axis. There exists a decision algorithm with running time $O(n^2/\alpha)$ and the following properties: if $\delta_{\text{dF}}(P, Q) \leq 1$, the algorithm returns YES; if $\delta_{\text{dF}}(P, Q) \geq \alpha$, the algorithm returns NO; if $\delta_{\text{dF}}(P, Q) \in (1, \alpha)$, the algorithm may return either YES or NO. The running time depends exponentially on d .*

Consider the regular d -dimensional grid with diameter 1 (all cells are axis-parallel cubes with side length $1/\sqrt{d}$). The distance between two grid cells C and D , $d(C, D)$, is defined as the smallest distance between a point in C and a point in D . The distance between a point x and a grid cell C , $d(x, C)$, is the distance between x and the closest point in C . For a point $x \in \mathbb{R}^d$, we write B_x for the closed unit ball with center x and C_x for the grid cell that contains x (since we are interested in approximation algorithms, we may assume that all points of $P \cup Q$ lie strictly inside the cells). We compute for each point $r \in P \cup Q$ the grid cell C_r that contains it. We also record for each nonempty grid cell C the number of points from Q contained in C . This can be done in total linear time as follows: we scan the points from $P \cup Q$ in x_1 -order, and we group the points according to the grid intervals that contain them. Then we split the lists that represent the x_2, \dots, x_d -order correspondingly, and we recurse on each group to determine the grouping for the remaining coordinate axes. Each iteration takes linear time, and there are d iterations, resulting in a total time of $O(n)$. In the following, we will also need to know for each non-empty cell the neighborhood of all cells that have a certain constant distance from it. These neighborhoods can be found in linear time by modifying the above procedure as follows: before performing the grouping, we make $O(1)$ copies of each point $r \in P \cup Q$ that we translate suitably to hit all neighboring cells for r . By using appropriate cross-pointers, we can then identify the neighbors of each non-empty cell in total linear time. Afterwards, we perform a clean-up step, so that only the original points remain.

A grid cell C is *full* if $|C \cap Q| \geq 5n/\alpha$. Let \mathcal{F} be the set of full grid cells. Clearly, $|\mathcal{F}| \leq \alpha/5$. We say that two full cells $C, D \in \mathcal{F}$ are *adjacent* if $d(C, D) \leq 4$. This defines a graph H on \mathcal{F} of constant degree. Using the neighborhood finding procedure from above, we can determine H and its connected components L_1, \dots, L_k in time $O(n + \alpha)$. For $C \in \mathcal{F}$, the *label* L_C of C is the connected component of H containing C .

For each $q \in Q$, we search for a full cell $C \in \mathcal{F}$ with $d(q, C) \leq 2$. If such a cell exists, we label q with $L_q = L_C$; otherwise, we set $L_q = \perp$. Similarly, for each $p \in P$, we search a full cell $C \in \mathcal{F}$ with $d(p, C) \leq 1$. In case of success, we set $L_p = L_C$; otherwise, we set $L_p = \perp$. Using the neighborhood finding procedure from above, this takes linear time. Let $P' = \{p \in P \mid L_p \neq \perp\}$ and $Q' = \{q \in Q \mid L_q \neq \perp\}$. The labeling has the following properties.

► **Lemma 5.2.** *We have*

1. for every $r \in P \cup Q$, the label L_r is uniquely determined;
2. for every $x, y \in P' \cup Q'$ with $L_x = L_y$, we have $d(x, y) \leq \alpha$;
3. if $p \in P'$ and $q \in B_p \cap Q$, then $L_p = L_q$; and
4. if $p \in P \setminus P'$, there are $O(n/\alpha)$ points $q \in Q$ with $d(p, C_q) \leq 1$. Hence, $|B_p \cap Q| = O(n/\alpha)$.

Proof. Let $r \in P \cup Q$ and suppose there are $C, D \in \mathcal{F}$ with $d(r, C) \leq 2$ and $d(r, D) \leq 2$. Then $d(C, D) \leq d(C, r) + d(r, D) \leq 4$, so C and D are adjacent in H . It follows that $L_C = L_D$ and that L_r is determined uniquely.

Fix $x, y \in P' \cup Q'$ with $L_x = L_y$. By construction, there are $C, D \in \mathcal{F}$ with $d(x, C) \leq 2$, $d(y, D) \leq 2$ and $L_C = L_D$. This means that C and D are in the same component of H . Therefore, C and D are connected by a sequence of adjacent cells in \mathcal{F} . We have $|\mathcal{F}| \leq \alpha/5$, any two adjacent cells have distance at most 4, and each cell has diameter 1. Thus, the triangle inequality gives $d(x, y) \leq 2 + 4(|\mathcal{F}| - 1) + |\mathcal{F}| + 2 \leq \alpha$.

Let $p \in P'$ and $q \in B_p \cap Q$. Take $C \in \mathcal{F}$ with $d(p, C) \leq 1$. By the triangle inequality, $d(q, C) \leq d(q, p) + d(p, C) \leq 2$, so $L_q = L_p = L_C$.

Take $p \in P$ and suppose there is a grid cell C with $|C \cap Q| > 5n/\alpha$ and $d(p, C) \leq 1$. Then $C \in \mathcal{F}$, so $L_p \neq \perp$, which means that $p \in P'$. The contrapositive gives (4). ◀

Lemma 5.2 enables us to design an efficient approximation algorithm. For this, we define the *approximate free-space matrix* F . This is an $n \times n$ matrix with entries from $\{0, 1\}$. For $i, j \in \{1, \dots, n\}$, we set $F_{ij} = 1$ if either (i) $p_i \in P'$ and $L_{p_i} = L_{q_j}$; or (ii) $p_i \in P \setminus P'$ and $d(p_i, q_j) \leq 1$. Otherwise, we set $F_{ij} = 0$. The matrix F is approximate in the following sense:

► **Lemma 5.3.** *If $\delta_{\text{dF}}(P, Q) \leq 1$, then F allows a monotone traversal from $(1, 1)$ to (n, n) . Conversely, if F has a monotone traversal from $(1, 1)$ to (n, n) , then $\delta_{\text{dF}}(P, Q) \leq \alpha$.*

Proof. Suppose that $\delta_{\text{dF}}(P, Q) \leq 1$. Then there is a monotone traversal β of (P, Q) with $\delta(\beta) \leq 1$. By Lemma 5.2(3), β is also a traversal of F .

Now let β be a monotone traversal of F . By Lemma 5.2(2), we have $\delta(\beta) \leq \alpha$, as desired. ◀

Additionally, we define the *approximate reach matrix* R , which is an $n \times n$ matrix with entries from $\{0, 1\}$. We set $R_{ij} = 1$ if F allows a monotone traversal from $(1, 1)$ to (i, j) , and $R_{ij} = 0$, otherwise. By Lemma 5.3, R_{nn} is an α -approximate indicator for $\delta_{\text{dF}} \leq 1$. We describe how to compute the rows of R successively in total time $O(n^2/\alpha)$.

First, we perform the following preprocessing steps: we break Q into *intervals*, where an interval is a maximal consecutive subsequence of points $q \in Q$ with the same label $L_q \neq \perp$. For each point in an interval, we store pointers to the first and the last point of the interval. This takes linear time. Furthermore, for each $p_i \in P \setminus P'$, we compute a sparse representation T_i of the corresponding row of F , i.e., a sorted list of all the column indices j for which $F_{ij} = 1$. Using hashing and bucketing, this can be done in total time $O(n^2/\alpha)$, by Lemma 5.2(4).

Now we successively compute a sparse representation for each row i of R , i.e., a sorted list I_i of disjoint intervals $[a, b] \in I_i$ such that for $j = 1, \dots, n$, we have $R_{ij} = 1$ if and only if there is an interval $[a, b] \in I_i$ with $j \in [a, b]$. We initialize I_1 as follows: if $F_{11} = 0$, we set $I_1 = \emptyset$ and abort. Otherwise, if $p_1 \in P'$, then I_1 is initialized with the interval of q_1 (since $F_{11} = 1$, we have $L_{p_1} = L_{q_1}$ by Lemma 5.2(3)). If $p_1 \in P \setminus P'$, we determine the maximum b such that $F_{1j} = 1$ for all $j = 1, \dots, b$, and we initialize I_1 with the *singleton* intervals $[j, j]$ for $j = 1, \dots, b$. This can be done in time $O(n/\alpha)$, irrespective of whether p_i lies in P' or not.

Now suppose we already have the interval list I_i for some row i , and we want to compute the interval list I_{i+1} for the next row. We consider two cases.

Case 1: $p_{i+1} \in P'$. If $L_{p_{i+1}} = L_{p_i}$, we simply set $I_{i+1} = I_i$. Otherwise, we go through the intervals $[a, b] \in I_i$ in order. For each interval $[a, b]$, we check whether the label of q_b or the label of q_{b+1} equals the label of p_{i+1} . If so, we add the maximal interval $[b', c]$ to I_{i+1} with $b' = b$ or $b' = b + 1$ and $L_{p_{i+1}} = L_{q_j}$ for all $j = b', \dots, c$. With the information from the preprocessing phase, this takes $O(1)$ time per interval. The resulting set of intervals may not be disjoint (if $p_i \in P \setminus P'$), but any two overlapping intervals have the same endpoint. Also, intervals with the same endpoint appear consecutively in I_{i+1} . We next perform a clean-up pass through I_{i+1} : we partition the intervals into consecutive groups with the same endpoint, and in each group, we only keep the largest interval. All this takes time $O(|I_i| + |I_{i+1}|)$.

Case 2: $p_{i+1} \in P \setminus P'$. In this case, we have a sparse representation T_{i+1} of the corresponding row in F at our disposal. We simultaneously traverse I_i and T_{i+1} to compute I_{i+1} as follows: for each $j \in \{1, \dots, n\}$ with $F_{(i+1)j} = 1$, if I_i has an interval containing $j - 1$ or j or if $[j - 1, j - 1] \in I_{i+1}$, we add the singleton $[j, j]$ to I_{i+1} . This takes total time $O(|I_i| + |I_{i+1}| + n/\alpha)$.

The next lemma shows that the interval representation remains sparse throughout the execution of the algorithm, and that the intervals I_i indeed represent the approximate reach matrix R .

► **Lemma 5.4.** *We have $|I_i| = O(n/\alpha)$ for $i = 1, \dots, n$. Furthermore, the intervals in I_i correspond exactly to the 1-entries in the approximate reach matrix R .*

Proof. First, we prove that $|I_i| = O(n/\alpha)$ for $i = 1, \dots, n$. This is done by induction on i . We begin with $i = 1$. If $p_1 \in P'$, then $|I_1| = 1$. If $p_1 \in P \setminus P'$, then Lemma 5.2(4) shows that the first row of F contains at most $O(n/\alpha)$ 1-entries, so $|I_1| = O(n/\alpha)$. Next, suppose that we know by induction that $|I_i| = O(n/\alpha)$. We must argue that $|I_{i+1}| = O(n/\alpha)$. If $p_{i+1} \in P \setminus P'$, then the $(i + 1)$ -th row of F contains $O(n/\alpha)$ 1-entries by Lemma 5.2(4), and $|I_{i+1}| = O(n/\alpha)$ follows directly by construction. If $p_{i+1} \in P'$ and $L_{p_{i+1}} = L_{p_i}$, then $I_{i+1} = I_i$, and the claim follows by induction. Finally, if $p_{i+1} \in P'$ and $L_{p_{i+1}} \neq L_{p_i}$, then by construction, every interval in I_i gives rise to at most one new interval in I_{i+1} . Thus, by induction, $|I_{i+1}| \leq |I_i| = O(n/\alpha)$.

Second, we prove that I_i represents the i -th row of R , for $i = 1, \dots, n$. Again, the proof is by induction. For $i = 1$, the claim holds by construction, because the first row of R consists of the initial segment of 1s in F . Next, suppose we know that I_i represents the i -th row of R . We must argue that I_{i+1} represents the $(i + 1)$ th row of R . If $p_{i+1} \in P \setminus P'$, this follows directly by construction, because the algorithm explicitly checks the conditions for each possible 1-entry of R ($R_{(i+1)j}$ can only be 1 if $F_{(i+1)j} = 1$). If $p_{i+1} \in P'$ and $L_{p_{i+1}} = L_{p_i}$, then the $(i + 1)$ -th row of F is identical to the i -th row of F , and the same holds for R : there can be no new monotone paths, and all old monotone paths can be extended by one step along Q . Finally, consider the case $p_{i+1} \in P'$ and $L_{p_{i+1}} \neq L_{p_i}$. If $p_i \in P \setminus P'$, then every interval in I_i is a singleton $[b, b]$, from which a monotone path could potentially reach $(i + 1, b)$ and $(i + 1, b + 1)$, and from there walk to the right. We explicitly check both of these possibilities. If $p_i \in P'$, then for every interval $[a, b] \in I_i$ and for all $j \in [a, b]$ we have $L_{q_j} = L_{p_i} \neq L_{p_{i+1}}$. Thus, the only possible move is to $(i + 1, b + 1)$, and from there walk to the right, which is what we check. ◀

The first part of Lemma 5.4 implies that the total running time is $O(n^2/\alpha)$, since each row is processed in time $O(n/\alpha)$. By Lemma 5.3 and the second part of Lemma 5.4, if I_n has an interval containing n then $\delta_{\text{dF}}(P, Q) \leq \alpha$, and if $\delta_{\text{dF}}(P, Q) \leq 1$ then n appears in I_n .

Since the intervals in I_n are sorted, this condition can be checked in $O(1)$ time. Theorem 5.1 follows.

5.2 Optimization Procedure

We now leverage Theorem 5.1 to an optimization procedure.

► **Theorem 5.5.** *Let P and Q be two sequences of n points in \mathbb{R}^d , and let $1 \leq \alpha \leq n$. There is an algorithm with running time $O(n^2 \log n / \alpha)$ that computes a number δ^* with $\delta_{\text{dF}}(P, Q) \leq \delta^* \leq \alpha \delta_{\text{dF}}(P, Q)$. The running time depends exponentially on d .*

Proof. If $\alpha \leq 5$, we compute $\delta_{\text{dF}}(P, Q)$ directly in $O(n^2)$ time. Otherwise, we set $\alpha' = \alpha/5$. We sort the points of $P \cup Q$ according to the coordinate axes, and we compute a $(1/3)$ -well-separated pair decomposition $\mathcal{P} = \{(S_1, T_1), \dots, (S_k, T_k)\}$ for $P \cup Q$ in time $O(n \log n)$ [11]. Recall the properties of a well-separated pair decomposition: (i) for all pairs $(S, T) \in \mathcal{P}$, we have $S, T \subseteq P \cup Q$, $S \cap T = \emptyset$, and $\max\{\text{diam}(S), \text{diam}(T)\} \leq d(S, T)/3$ (here, $\text{diam}(S)$ denotes the maximum distance between any two points in S); (ii) the number of pairs is $k = O(n)$; and (iii) for every distinct $q, r \in P \cup Q$, there is exactly one pair $(S, T) \in \mathcal{P}$ with $q \in S$ and $r \in T$, or vice versa.

For each pair $(S_i, T_i) \in \mathcal{P}$, we pick arbitrary $s \in S_i$ and $t \in T_i$, and set $\delta_i = 3d(s, t)$. After sorting, we can assume that $\delta_1 \leq \dots \leq \delta_k$. We call δ_i a *YES-entry* if the algorithm from Theorem 5.1 on input α' and the point sets P and Q scaled by a factor of δ_i returns YES; otherwise, we call δ_i a *NO-entry*. First, we test whether δ_1 is a YES-entry. If so, we return $\delta^* = \alpha' \delta_1$. If δ_1 is a NO-entry, we perform a binary search on $\delta_1, \dots, \delta_k$: we set $l = 1$ and $r = k$. Below, we will prove that δ_k must be a YES-entry. We set $m = \lceil (l+r)/2 \rceil$. If δ_m is a NO-entry, we set $l = m$, otherwise, we set $r = m$. We repeat this until $r = l + 1$. In the end, we return $\delta^* = \alpha' \delta_r$. The total running time is $O(n \log n + n^2 \log n / \alpha)$. Our procedure works exactly like binary search, but we presented it in detail in order to emphasize that $\delta_1, \dots, \delta_k$ is not necessarily monotone: NO-entries and YES-entries may alternate.

We now argue correctness. The algorithm finds a YES-entry δ_r such that either $r = 1$ or δ_{r-1} is a NO-entry. By Theorem 5.1, any δ_i is a NO-entry if $\delta_i \leq \delta_{\text{dF}}(P, Q) / \alpha'$. Thus, we certainly have $\delta^* = \alpha' \delta_r > \delta_{\text{dF}}(P, Q)$. Now take a traversal β with $\delta(\beta) = \delta_{\text{dF}}(P, Q)$, and let $(p, q) \in P \times Q$ be a position in β that has $d(p, q) = \delta(\beta)$. There is a pair $(S_{r^*}, T_{r^*}) \in \mathcal{P}$ with $p \in S_{r^*}$ and $q \in T_{r^*}$, or vice versa. Let $s \in S_{r^*}$ and $t \in T_{r^*}$ be the points we used to define δ_{r^*} . Then

$$d(s, t) \geq d(p, q) - \text{diam}(S_{r^*}) - \text{diam}(T_{r^*}) \geq d(p, q) - 2d(S_{r^*}, T_{r^*})/3 \geq d(p, q)/3,$$

and

$$d(s, t) \leq d(p, q) + \text{diam}(S_{r^*}) + \text{diam}(T_{r^*}) \leq d(p, q) + 2d(S_{r^*}, T_{r^*})/3 \leq 5d(p, q)/3,$$

so $\delta_{r^*} = 3d(s, t) \in [\delta(\beta), 5\delta(\beta)]$. Since by Theorem 5.1 any δ_i is a YES-entry if $\delta_i \geq \delta_{\text{dF}}(P, Q)$, all δ_i with $i \geq r^*$ are YES-entries (in particular, δ_k is a YES-entry). Thus, $\delta^* \leq \alpha' \delta_{r^*} \leq 5\alpha' \delta_{\text{dF}}(P, Q) \leq \alpha \delta_{\text{dF}}(P, Q)$. ◀

The running time of Theorem 5.5 can be improved as follows.

► **Theorem 5.6.** *Let P and Q be two sequences of n points in \mathbb{R}^d , and let $1 \leq \alpha \leq n$. There is an algorithm with running time $O(n \log n + n^2 / \alpha)$ that computes a number δ^* with $\delta_{\text{dF}}(P, Q) \leq \delta^* \leq \alpha \delta_{\text{dF}}(P, Q)$. The running time depends exponentially on d .*

Proof. If $\alpha \leq 4$, we can compute $\delta_{\text{dF}}(P, Q)$ exactly. Otherwise, we use Theorem 5.5 to compute a number δ' with $\delta_{\text{dF}}(P, Q) \leq \delta' \leq n \cdot \delta_{\text{dF}}(P, Q)$, or, equivalently, $\delta_{\text{dF}}(P, Q) \in [\delta'/n, \delta']$. This takes time $O(n \log n)$. Set $i^* = \lceil \log(n/\alpha) \rceil + 1$ and for $i = 1, \dots, i^*$ let $\alpha_i = n/2^{i+1}$. Also, set $a_1 = \delta'/n$ and $b_1 = \delta'$.

We iteratively obtain better estimates for $\delta_{\text{dF}}(P, Q)$ by repeating the following for $i = 1, \dots, i^* - 1$. As an invariant, at the beginning of iteration i , we have $\delta_{\text{dF}}(P, Q) \in [a_i, b_i]$ with $b_i/a_i = 4\alpha_i$. We use the algorithm from Theorem 5.1 with inputs α_i and P and Q scaled by a factor $2a_i$ (since $\alpha_i \geq \alpha_{i^*-1} = n/2^{\lceil \log(n/\alpha) \rceil + 1} \geq \alpha/4$, the algorithm can be applied). If the answer is YES, it follows that $\delta_{\text{dF}}(P, Q) \leq \alpha_i 2a_i = b_i/2$, so we set $a_{i+1} = a_i$ and $b_{i+1} = b_i/2$. If the answer is NO, then $\delta_{\text{dF}}(P, Q) \geq 2a_i$, so we set $a_{i+1} = 2a_i$ and $b_{i+1} = b_i$. This needs time $O(n^2/\alpha_i)$ and maintains the invariant.

In the end, we return a_{i^*} . The invariant guarantees $\delta_{\text{dF}}(P, Q) \in [a_{i^*}, b_{i^*}]$ and $b_{i^*}/a_{i^*} = 4\alpha_{i^*} \leq \alpha$, as desired. The total running time is proportional to

$$n \log n + \sum_{i=1}^{i^*-1} n^2/\alpha_i = n \log n + \sum_{i=1}^{i^*-1} n2^{i+1} \leq n \log n + n2^{i^*+1} = O(n \log n + n^2/\alpha). \quad \blacktriangleleft$$

6 Conclusions

We have obtained several new results on the approximability of the discrete Fréchet distance. As our main results,

1. we showed a conditional lower bound for the *one-dimensional* case that there is no 1.399-approximation in strongly subquadratic time unless the Strong Exponential Time Hypothesis fails. This sheds further light on what makes the Fréchet distance a difficult problem.
2. we determined the approximation ratio of the *greedy* algorithm as $2^{\Theta(n)}$ in any dimension $d \geq 1$. This gives the first general linear time approximation algorithm for the problem; and
3. we designed an α -*approximation* algorithm running in time $O(n \log n + n^2/\alpha)$ for any $1 \leq \alpha \leq n$ in any constant dimension $d \geq 1$. This significantly improves the greedy algorithm, at the expense of a (slightly) worse running time.

Our lower bounds exclude only (too good) constant factor approximations with strongly subquadratic running time, while our best strongly subquadratic approximation algorithm has an approximation ratio of n^ϵ . It remains a challenging open problem to close this gap.

References

- 1 Amir Abboud and Virginia Vassilevska Williams. Popular conjectures imply strong lower bounds for dynamic problems. In *Proc. 55th Annu. IEEE Sympos. Found. Comput. Sci. (FOCS)*, pages 434–443, 2014.
- 2 Amir Abboud, Virginia Vassilevska Williams, and Oren Weimann. Consequences of faster alignment of sequences. In *Proc. 41st Internat. Colloq. Automata Lang. Program. (ICALP)*, volume 8572 of *LNCS*, pages 39–51, 2014.
- 3 Amir Abboud, Ryan Williams, and Huacheng Yu. More applications of the polynomial method to algorithm design. In *Proc. 26th Annu. ACM-SIAM Sympos. Discrete Algorithms (SODA)*, pages 218–230, 2015.
- 4 Pankaj K. Agarwal, Rinat Ben Avraham, Haim Kaplan, and Micha Sharir. Computing the discrete Fréchet distance in subquadratic time. *SIAM J. Comput.*, 43(2):429–449, 2014.

- 5 Helmut Alt. Personal communication. 2012.
- 6 Helmut Alt and Michael Godau. Computing the Fréchet distance between two polygonal curves. *Internat. J. Comput. Geom. Appl.*, 5(1–2):78–99, 1995.
- 7 Karl Bringmann. Why walking the dog takes time: Fréchet distance has no strongly subquadratic algorithms unless SETH fails. In *Proc. 55th Annu. IEEE Sympos. Found. Comput. Sci. (FOCS)*, pages 661–670, 2014.
- 8 Karl Bringmann and Marvin Künnemann. Improved approximation for Fréchet distance on c -packed curves matching conditional lower bounds. [arXiv:1408.1340](https://arxiv.org/abs/1408.1340), 2014.
- 9 Kevin Buchin, Maike Buchin, Wouter Meulemans, and Wolfgang Mulzer. Four soviets walk the dog - with an application to Alt’s conjecture. In *Proc. 25th Annu. ACM-SIAM Sympos. Discrete Algorithms (SODA)*, pages 1399–1413, 2014.
- 10 Kevin Buchin, Maike Buchin, Rolf van Leusden, Wouter Meulemans, and Wolfgang Mulzer. Computing the Fréchet distance with a retractable leash. In *Proc. 21st Annu. European Sympos. Algorithms (ESA)*, pages 241–252, 2013.
- 11 Paul B. Callahan and S. Rao Kosaraju. A decomposition of multidimensional point sets with applications to k -nearest-neighbors and n -body potential fields. *J. ACM*, 42(1):67–90, 1995.
- 12 Thomas Eiter and Heikki Mannila. Computing Discrete Fréchet Distance. Technical Report CD-TR 94/64, Christian Doppler Laboratory, 1994.
- 13 Anka Gajentaan and Mark H. Overmars. On a class of $O(n^2)$ problems in computational geometry. *Comput. Geom. Theory Appl.*, 5(3):165–185, 1995.
- 14 Michael R. Garey and David S. Johnson. *Computers and intractability. A guide to the theory of NP-completeness*. W. H. Freeman, 1979.
- 15 Allan Grønlund and Seth Pettie. Threesomes, degenerates, and love triangles. In *Proc. 55th Annu. IEEE Sympos. Found. Comput. Sci. (FOCS)*, pages 621–630, 2014.
- 16 Russell Impagliazzo and Ramamohan Paturi. On the complexity of k -SAT. *J. Comput. System Sci.*, 62(2):367–375, 2001.
- 17 Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which problems have strongly exponential complexity. *J. Comput. System Sci.*, 63(4):512–530, 2001.
- 18 Mihai Pătraşcu and Ryan Williams. On the possibility of faster SAT algorithms. In *Proc. 21st Annu. ACM-SIAM Sympos. Discrete Algorithms (SODA)*, pages 1065–1075, 2010.
- 19 Ramamohan Paturi, Pavel Pudlák, Michael E. Saks, and Francis Zane. An improved exponential-time algorithm for k -sat. *J. ACM*, 52(3):337–364, 2005.
- 20 Liam Roditty and Virginia Vassilevska Williams. Fast approximation algorithms for the diameter and radius of sparse graphs. In *Proc. 45th Annu. ACM Sympos. Theory Comput. (STOC)*, pages 515–524, 2013.
- 21 Ryan Williams. A new algorithm for optimal 2-constraint satisfaction and its implications. *Theoret. Comput. Sci.*, 348(2):357–365, 2005.