

Fréchet Distance in Unweighted Planar Graphs



Ivor van der Hoog  

IT University of Copenhagen, Denmark

Thijs van der Horst  

Utrecht University, The Netherlands

TU Eindhoven, The Netherlands

Eva Rotenberg  

IT University of Copenhagen, Denmark

Lasse Wulf  

IT University of Copenhagen, Denmark

Abstract

The Fréchet distance is a distance measure between trajectories in \mathbb{R}^d or walks in a graph G . Given constant-time shortest path queries, the Discrete Fréchet distance $D_G(P, Q)$ between two walks P and Q can be computed in $O(|P| \cdot |Q|)$ time using a dynamic program. Driemel, van der Hoog, and Rotenberg [SoCG'22] show that for weighted planar graphs this approach is likely tight, as there can be no strongly-subquadratic algorithm to compute a 1.01-approximation of $D_G(P, Q)$ unless the Orthogonal Vector Hypothesis (OVH) fails.

Such quadratic-time conditional lower bounds are common to many Fréchet distance variants. However, they can be circumvented by assuming that the input comes from some well-behaved class: There exist $(1 + \varepsilon)$ -approximations, both in weighted graphs and in \mathbb{R}^d , that take near-linear time for c -packed or κ -straight walks in the graph. In \mathbb{R}^d there also exists a near-linear time algorithm to compute the Fréchet distance whenever all input edges are long compared to the distance.

We consider computing the Fréchet distance in unweighted planar graphs. We show that there exist no strongly-subquadratic 1.25-approximations of the discrete Fréchet distance between two disjoint simple paths in an unweighted planar graph in strongly subquadratic time, unless OVH fails. This improves the previous lower bound, both in terms of generality and approximation factor. We subsequently show that adding graph structure circumvents this lower bound: If the graph is a regular tiling with unit-weighted edges, then there exists an $\tilde{O}((|P| + |Q|)^{1.5})$ -time algorithm to compute $D_G(P, Q)$. Our result has natural implications in the plane, as it allows us to define a new class of well-behaved curves that facilitate $(1 + \varepsilon)$ -approximations of their discrete Fréchet distance in subquadratic time.

2012 ACM Subject Classification Theory of computation \rightarrow Computational geometry; Theory of computation \rightarrow Design and analysis of algorithms; Theory of computation \rightarrow Graph algorithms analysis

Keywords and phrases Fréchet distance, planar graphs, lower bounds, approximation algorithms

Digital Object Identifier 10.4230/LIPIcs.ESA.2025.24

Related Version *Full Version:* <https://arxiv.org/abs/2504.17342>

Funding This work was supported by the Independent Research Fund Denmark grant 2020-2023 (9131-00044B) “Dynamic Network Analysis”, the Carlsberg Foundation Young Researcher Fellowship CF21-0302 “Graph Algorithms with Geometric Applications”, the VILLUM Foundation grant (VIL37507) “Efficient Recomputations for Changeful Problems”, and the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 899987.



© Ivor van der Hoog, Thijs van der Horst, Eva Rotenberg, and Lasse Wulf;
licensed under Creative Commons License CC-BY 4.0

33rd Annual European Symposium on Algorithms (ESA 2025).

Editors: Anne Benoit, Haim Kaplan, Sebastian Wild, and Grzegorz Herman; Article No. 24; pp. 24:1–24:16

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

The Fréchet distance is a widely used metric for measuring the similarity between trajectories. It is often illustrated through a metaphor: imagine a person walking along one trajectory and their dog following another, connected by a leash. The Fréchet distance is the minimum possible leash length over all synchronized traversals of both trajectories.

This metric has numerous applications, particularly in movement data analysis [10, 18, 24, 15, 5, 13, 25, 33]. It is versatile, having been applied to handwriting recognition [29], coastlines [26], geometric shapes in geographic information systems [16], and trajectories of moving objects such as vehicles, animals, and athletes [28, 30, 6, 13]. We consider the *discrete* Fréchet distance, where trajectories are modeled as sequences of discrete points.

Alt and Godau [2] were the first to analyse the Fréchet distance from a computational perspective. They considered trajectories as polygonal curves in \mathbb{R}^d with n and m vertices, and compute the continuous Fréchet distance in $O(mn \log(n + m))$ time. Later, Buchin, Buchin, Meulemans and Mulzer [11] introduced a faster randomized algorithm, achieving a running time of $O(n^2 \sqrt{\log n} (\log \log n)^{3/2})$ on a real-valued pointer machine and $O(n^2 \log \log n)$ on a word RAM. For the discrete Fréchet distance, Eiter and Mannila [20] presented an $O(nm)$ time algorithm under constant-time distance computations. Agarwal, Avraham, Kaplan, and Sharir [1] later improved this to $O(nm(\log \log nm)/\log nm)$ in the word RAM.

Driemel, van der Hoog, and Rotenberg [19] initiate the study of Fréchet distance in graphs. They consider as input a (weighted) graph G where a trajectory is a walk in the graph. The distance between two vertices is the length of their shortest path. Given a constant-time distance oracle, the discrete Fréchet distance $D_G(P, Q)$ between any two walks in G can then be computed in $O(nm)$ time using the algorithm by Eiter and Mannila [20].

Conditional Lower Bounds. Several conditional lower bounds exist for computing the Fréchet distance or its constant-factor approximations. These results rely on complexity assumptions such as the Orthogonal Vector Hypothesis (OVH) and the Strong Exponential Time Hypothesis (SETH) [32]. Bringmann [7] established that no algorithm can compute the (discrete or continuous) Fréchet distance between two polygonal curves of n vertices in $O(n^{2-\delta})$ time for any $\delta > 0$, unless OVH fails. The same lower bound holds for small constant-factor approximations. Bringmann’s proof originally involved self-intersecting curves in the plane, but later work by Bringmann and Mulzer [9] extended the result to intersecting curves in \mathbb{R}^1 . Buchin, Ophelders, and Speckmann [12] further demonstrated that the same lower bound applies for any better-than-3-approximation algorithm for pairwise disjoint planar curves in \mathbb{R}^2 , or intersecting curves in \mathbb{R}^1 .

Bringmann’s lower bound [7] holds also for unbalanced inputs: given two curves with n and m vertices in the plane, no algorithm can compute the Fréchet distance in $O((nm)^{1-\delta})$ time assuming OVH. Driemel, van der Hoog and Rotenberg [19] give a similar lower bound for paths in a weighted planar graph. They prove that, unless OVH fails, no $O((nm)^{1-\delta})$ -time algorithm can 1.01-approximate the discrete Fréchet distance between arbitrary paths in a weighted planar graph – even if the ratio between smallest and highest weight is bounded by a constant. This lower bound applies only to weighted graphs and does therefore not exclude the existence of a fast algorithm on unweighted planar graphs. A closer inspection of their argument reveals that their proof cannot simply be adapted to the unweighted case by subdividing long edges often enough. In this case, new vertices get introduced to the graph and their arguments break down.

Avoiding lower bounds through well-behaved curves. These lower bounds can be circumvented if the input is well-behaved. Driemel, Har-peled and Wenk [17] consider three classes of curves in \mathbb{R}^d which are well-behaved as long as their corresponding behavioural parameter is low. These are c -packed curves [17, 8], ϕ -low density curves [17], and κ -bounded curves [3, 4, 17]. They show algorithms to compute a $(1 + \varepsilon)$ -approximation between two well-behaved curves whose running time is near-linear in n and m . Gudmundsson, Mirzanezhad, Mohades, and Wenk consider the special case where all edges of the input curves are long with respect to their Fréchet distance [23]. In this case, the Fréchet distance can be computed in $O((n + m) \log(n + m))$ time [23]. Driemel, van der Hoog, and Rotenberg [19] show that if the input are walks in a graph G and one of the walks is restricted to a c -packed or κ -straight path, then there exists $(1 + \varepsilon)$ -approximations that take near-linear time.

Contribution. Driemel, van der Hoog, and Rotenberg [19] ruled out a strongly subquadratic algorithm for a 1.01-approximation between paths in a weighted graph, while the construction by Buchin, Ophelders, and Speckmann [12] rules out a strongly subquadratic algorithm for a 3-approximation between walks in an unweighted planar graph. We consider the discrete Fréchet distance between paths in an unweighted planar graph. We prove (Theorem 5) that no strongly-subquadratic 1.25-approximation algorithm exists unless OVH fails, thereby improving both the generality (and also the approximation factor for planar graphs).

Next, we prove that adding graph structure circumvents this lower bound. If G is an unweighted regular tiling of the plane, then there exists an $\tilde{O}((n + m)^{1.5})$ -time algorithm to compute the Fréchet distance $D_G(P, Q)$, between curves of length n and m . Our results have natural implications in the plane, for a special class of well-behaved curves $((\varepsilon, \delta)$ -curves, to be defined). For those, we can compute the Fréchet distance under the L_1 metric (respectively, approximate it for any L_c metric) in $\tilde{O}(\frac{\sqrt{\delta}}{\varepsilon}(n + m)^{1.5})$ time (respectively, $\tilde{O}(\frac{\delta}{\varepsilon^2}(n + m)^{1.5})$ time). For completeness, we explain the continuous analogy to (ε, δ) -curves in the full version of this paper. We also compare this new curve class to existing well-behaved curve classes in the full version.

2 Preliminaries

Let G be an unweighted graph where $V(G)$ is its vertex set. We define $P = (p_1, \dots, p_n)$ and $Q = (q_1, \dots, q_m)$ as two paths in G . A path cannot visit a vertex twice. For integers $i, i' \in [n]$ with $i \leq i'$ we denote by $P[i : i']$ the subpath of P from p_i to $p_{i'}$. We define the discrete Fréchet distance $D_G(P, Q)$ between P and Q as follows [19]: A sequence $(a_i, b_i)_{i \in [k]}$ of pairs of indices is a *monotone walk* if for all $i \in [k - 1]$ we have $a_{i+1} \in \{a_i, a_i + 1\}$ and $b_{i+1} \in \{b_i, b_i + 1\}$. Let $\mathbb{F}_{n,m}$ be the set of all monotone walks from $(1, 1)$ to (n, m) . The cost of $F \in \mathbb{F}_{n,m}$ is the maximum of $\text{dist}(p_i, q_j)$ over $(i, j) \in F$, where $\text{dist}()$ denotes the length of the shortest path between vertices. $D_G(P, Q)$ is the minimum cost of a walk in $\mathbb{F}_{n,m}$.

$$D_G(P, Q) = \min_{F \in \mathbb{F}_{n,m}} \text{cost}(F) = \min_{F \in \mathbb{F}_{n,m}} \max_{(i,j) \in F} \text{dist}(p_i, q_j).$$




Figure 1 (a) The three regular plane tilings. (b) A hexagonal super-tiling of a triangular tiling.

Decision variant. We focus on the decision variant where the input includes some $\Delta \geq 0$ and the goal is to output whether $D_G(P, Q) \leq \Delta$. The Δ -free space matrix M_Δ is the n by m matrix where $M_\Delta[i, j]$ is 0 if $\text{dist}(p_i, q_j) \leq \Delta$ and 1 otherwise. We follow geometric convention, where i denotes the column of M_Δ (column i corresponds to some $p_i \in P$). Column entries with value 0 correspond to $q_j \in Q$ with $\text{dist}(p_i, q_j) \leq \Delta$. Eiter and Mannila [20] define a directed graph over M_Δ where there is an edge from $M_\Delta[a, b]$ to $M_\Delta[c, d]$ if and only if $c \in \{a, a+1\}$, $d \in \{b, b+1\}$ and $M_\Delta[a, b] = M_\Delta[c, d] = 0$. They prove that $D_G(P, Q) \leq \Delta$ if and only if there exists a directed path from $(1, 1)$ to (n, m) in this graph.

Orthogonal vectors. Our conditional lower bound reduces from the Orthogonal Vector Hypothesis. The input consists of two ordered sets of d -dimensional Boolean vectors U, W of sizes $|U| = n$ and $|W| = m$. We denote for $u \in U$ by $u_i \in \{0, 1\}$ its i 'th component. The *OV problem* is to find out if some vector of U is orthogonal to some vector of W .

► **Definition 1** (Williams [32]). *The orthogonal vector hypothesis (OVH) states that for all $\delta > 0$, there exists constants ω and $1 > \gamma > 0$ such that the OV problem for vectors of dimension $d = \omega \log n$ and $m = n^\gamma$ cannot be solved in $O((nm)^{1-\delta})$ time.*

We consider each vector in U (or W) as a string of d bits. We denote by $\text{str}(U)$ and $\text{str}(W)$ the concatenation of all strings in U and W , respectively. A *substring* of some string $s_1 s_2 \dots s_k$ is the sequence $s_a s_{a+1} \dots s_{b-1} s_b$ for some $1 \leq a \leq b \leq k$.


Regular tilings. A *regular tiling* T is an infinite plane graph where each face is a regular polygon and each face has the same degree. There exist exactly three regular tilings [22]: the triangular, square, and hexagonal tiling (Figure 1). Without loss of generality, we assume that our tilings are axis-aligned. A face F of T is an open region bounded by the edges of T , we denote by \bar{F} its closure. The *edge length* of T is the length of any edge. A tiling is *unit* if its edges have unit length. A path in T is a path in its corresponding graph. We always denote by T a unit tiling where $(0, 0)$ is a vertex of T . For any pair of vertices p, q of T , the term $\text{dist}(p, q)$ denotes the distance in the graph T . We assume that we can compute $\text{dist}(p, q)$ in constant time. Our algorithm considers what we will call a *super-tiling* of T and a natural partition of the edges of P and Q into subpaths.

► **Definition 2** (Figure 1 (b)). *Given a regular tiling T of the plane, a super-tiling \mathbb{T} of T is a regular tiling of the plane whose vertex set is a subset of the vertex set $V(T)$. The boundary vertices $B(\mathbb{T})$ are the vertices of T that lie on the boundary of some face of \mathbb{T} .*

► **Definition 3.** *Consider a regular tiling T , some super-tiling \mathbb{T} and a path P in T , we define the breakpoints $B(P, \mathbb{T})$ as all vertices in $P \cap B(\mathbb{T})$.*

Given a path P in a tiling T and a super-tiling \mathbb{T} of T , we want to consider how \mathbb{T} splits P into smaller snippets, each snippet “belonging” to a face of \mathbb{T} . More formally:

► **Definition 4** (Figure 2). *Consider a regular tiling T and a super-tiling \mathbb{T} of T . Given a path $P = (p_1, p_2, \dots, p_n)$ in T , let $\mathbb{I} := \{i \in [n] \mid p_i \in B(P, \mathbb{T}) \text{ or } i \in \{1, n\}\}$. We define the induced subpaths $S(P, \mathbb{T})$ as all subpaths $P[x : y]$, for consecutive $x, y \in \mathbb{I}$. Note that for each induced subpath, there exists a (not necessarily unique) face F in \mathbb{T} whose closure \bar{F} contains it; we say the induced subpath is assigned to one such face (breaking ties arbitrarily).*



■ **Figure 2** (a) For a face F in the super-tiling \mathbb{T} , we show the set $B(\mathbb{T}) \cap \overline{F}$. (b) Given a path P , we create a sequence of subpaths of P where consecutive subpaths overlap in a single vertex.

Submatrices of M_Δ . We show a subquadratic algorithm to decide whether $D_T(P, Q) \leq \Delta$ if P and Q are paths in a tiling T . We compute a monotone walk from $(1, 1)$ to (n, m) in M_Δ by efficiently propagating reachability information across *submatrices* of M_Δ . For integers $i, i' \in [n]$ and $j, j' \in [m]$ with $i \leq i'$ and $j \leq j'$ we denote by $M_\Delta[i : i', j : j']$ the submatrix of M_Δ that consist of all entries $M_\Delta[x, y]$ for $(x, y) \in [i, i'] \times [j, j']$. We define the *bottom-left* facets of $M_\Delta[i : i', j : j']$ as all entries $\{M_\Delta[i, y] \mid y \in [j, j']\}$ and $\{M_\Delta[x, j] \mid x \in [i, i']\}$. The *top-right* facets are defined as the entries $\{M_\Delta[i', y] \mid y \in [j, j']\}$ and $\{M_\Delta[x, j'] \mid x \in [i, i']\}$.

(ε, δ) -curves. Finally, we introduce a new class of well-behaved curves in the plane. Intuitively, an (ε, δ) -curve, for small δ , is a curve composed of short edges that does not revisit the same region of the plane too frequently. Formally, fix a universal constant $\gamma \in O(1)$. Let P be a curve in the plane, and let B_ε denote the set of all balls (in Euclidean metric) of radius ε centred at the vertices of P . We say that P is an (ε, δ) -curve if all its edges have length at most γ , and any point in the plane lies in at most δ balls from B_ε .

For any given pair (P, ε) , δ is determined. As ε decreases, so may δ . If δ remains large for small ε , then many vertices of P are densely clustered – i.e., effectively, the curve frequently revisits the same locations. We present an exact algorithm for the Discrete Fréchet distance under the L_1 metric with a linear dependency on $\frac{\sqrt{\delta}}{\varepsilon}$. We assume that the input specifies a desirable value of ε , and that this yields a favourable corresponding δ . Imposing that δ is constant amounts to requiring that the curve does not revisit any ε -ball more than a constant number of times – a restriction that somewhat resembles the well-studied c -packed curves [17, 19, 8]. In the full version of this paper, we discuss in detail how our new class relates to the prior well-behaved curve classes.

3 A lower bound for paths in an unweighted planar graph

Let G be an unweighted planar graph and P and Q be paths in G . We show a lower bound, conditioned on OVH, that excludes any strongly subquadratic algorithm to compute a 1.25-approximation of $D_G(P, Q)$. Since we are space-restricted, we present our argument only in the full version of this paper. Here, we assume that the reader is familiar with how these lower bounds are typically constructed [7, 12, 19, 9] and only mention how our construction differs from prior works. We start with some OV instance (U, W) . We build an unweighted planar graph G and two paths P and Q in G such that $D_G(P, Q) \leq 4$ if and only if the OV instance is a yes-instance and $D_G(P, Q) \geq 5$ otherwise. One novelty of our approach is a preprocessing step, where we transform (U, W) into an equivalent instance with additional properties:

► **Lemma 1.** *An instance $U', W' \subseteq \{0, 1\}^{d'}$ of OV can be preprocessed in $O(d'(n + m))$ time, resulting in a new instance $U, W \subseteq \{0, 1\}^d$ with $d \in O(d')$ such that:*


■ *a yes-instance stays a yes-instance and a no-instance stays a no-instance,*

- for all $u \in U$, $u_1 = u_d = 0$ and for all $w \in W$, $w_1 = w_d = 1$,
- if the instance is a no-instance, then $\forall u \in U$ the vector u is not only non-orthogonal to every $w \in W$, but even non-orthogonal to all consecutive length- d substrings of $\text{str}(W)$.

This preprocessing significantly simplifies our proofs and may be of independent interest. Previous lower bound constructions [7, 19, 12, 9] fix some $\Delta \in \mathbb{R}$. Given an OV-instance (U, W) , they construct curves (P, Q) by constructing for each $u \in U$ a subcurve $f(u)$ and each $w \in W$ a subcurve $g(w)$ (see Figure 3). The construction has a very strong property: Consider a traversal of P and Q where the person and the dog remain within distance Δ of one another. If the person is in $f(u)$ for some $u \in U$ and the dog is in $g(w)$ for some $w \in W$ then neither the person or the dog can stand still if the other moves.

Requiring the planar graph to be unweighted significantly restricts the freedom we have in engineering the pairwise distance between vertices. We are unable to recreate an equally strong property and instead obtain a much weaker statement: If the person is traversing the subpath $f(u)$ for some $u \in U$ and the dog is traversing $g(w)$ for some $w \in W$ and the person moves at least two steps then the dog has to move at least one. We rely upon our preprocessing to then still argue that there exists a traversal of P and Q such that the leash length is at most 4 if and only if there is an orthogonal pair $(u, w) \in U \times W$.

► **Theorem 5.** *Let P and Q be paths in an unweighted planar graph. Let $|P| = n$ and $|Q| = m = n^\gamma$ for some constant $\gamma > 0$. Then for all $\delta \in (0, 1)$, one cannot approximate $D_G(P, Q)$ by a factor better than 1.25 in $O((nm)^{1-\delta})$ time unless OVH fails.*



■ **Figure 3** We illustrate the construction for planar curves by Bringmann [7]. Each $u \in U$ becomes a subcurve $f(u)$ of P and each $w \in W$ becomes a subcurve $g(w)$ of Q . Our table shows a green cell if the distance between the vertices is less than Δ . If the person is in a subcurve $f(u)$ and the dog is in $g(w)$ and the person moves from its vertex then the dog *must* also move.

4 Upper bound

We show that subquadratic algorithms *do* exist for unweighted planar graphs that exhibit additional structure. Specifically, our main result is as follows:

► **Theorem 6.** *Let T be a regular unit tiling and (P, Q) be paths in T with $|P| = n$ and $|Q| = m$. Given $\Delta \in \mathbb{R}$, we can output whether $D_T(P, Q) \leq \Delta$ in $O((n + m)^{1.5} \log(n + m))$ time.*

We compute $D_T(P, Q)$ as follows. Given (P, Q) , we compute the smallest axis-aligned rectangles R_P and R_Q that contain P and Q , respectively, in $O(n + m)$ time. Let d_0 be the Euclidean length of the shortest segment (p, q) with $p \in R_P$ and $q \in R_Q$, and let d_1 be the Euclidean length of the longest such segment. Observe that $D_T(P, Q) \in [\frac{d_0}{2}, 2d_1]$, that this range contains $O(n + m)$ integers, and that $D_T(P, Q)$ is always integral. By performing a binary search over $[\frac{d_0}{2}, 2d_1] \cap \mathbb{Z}$, applying Theorem 6 at each step, we obtain:

► **Corollary 7.** *Let T be a regular unit tiling, and let P and Q be paths in T with $|P| = n$ and $|Q| = m$. Then $D_T(P, Q)$ can be computed in $O((n + m)^{1.5} \log^2(n + m))$ time.*

Additional definitions. To prove Theorem 6, we first introduce a few new concepts. Let T be a regular, unit, axis-aligned tiling. Based on T , we construct a super-tiling \mathbb{T} of the same type as T , that is, if T is a square/triangular/hexagonal tiling, then \mathbb{T} is also a square/triangular/hexagonal tiling. We state our results in general terms, but for ease of reading, the reader may assume T is a square tiling. For a (super) tiling \mathbb{T} , we define the *halfslabs* of any face F of \mathbb{T} by picture (Figure 4).

► **Definition 8.** *Let T be a regular tiling, and let \mathbb{T} be a super-tiling of T . We say that two faces F and F' of \mathbb{T} are well-separated by a vertex $v \in V(T)$ if, for all $p \in V(T) \cap \overline{F}$ and $q \in V(T) \cap \overline{F'}$, there exists a shortest path from p to q in T that passes through v .*

Having faces F, F' be well-separated proves to be very useful, since in such a pair the distance between some vertex $p \in V(T) \cap \overline{F}$ and $q \in V(T) \cap \overline{F'}$ can be well-understood. We provide an efficient algorithm to compute the pairwise Fréchet distance between subpaths of P and Q whenever their corresponding faces in \mathbb{T} are well-separated. We identify a sufficient condition that ensures that some pair F, F' is well-separated, based on the *alignment* of the two:


► **Definition 9.** *Let T be a regular tiling. Two faces F and F' of T are aligned if F' intersects a halfslab of F (or F intersects a halfslab of F').*

► **Lemma 10.** *Let T be an axis-aligned regular unit tiling, and let \mathbb{T} be a super-tiling of T of the same type as T . Any two faces F, F' of \mathbb{T} that are not aligned are well-separated by a vertex $v \in V(T)$.*

Proof. Consider first the case where T is a square tiling (and so \mathbb{T} is too). The situation is depicted in Figure 5 (a). If the faces F, F' are not aligned, then F' is contained in one of the four regions (R_1, R_2, R_3, R_4) of the plane formed by the halfslabs of F . Let v be the vertex at the apex of the region which contains F' . We claim that for any vertices $p \in V(T) \cap \overline{F}$ and $q \in V(T) \cap \overline{F'}$, there exists a shortest path between p and q that passes through v . Hence F and F' are well-separated. The claim is easy to see from the properties of shortest paths in grid graphs. We provide an alternative proof of the claim, with the advantage that it can be more easily generalized to the triangular and hexagonal grid. See Figure 5 (b) and consider for $i = 1, 2, \dots$ the circles of radius i in the plane graph T centred around p , and q .

$$h_i := \{x \in V(T) : \text{dist}(p, x) = i\}, \quad h'_i := \{x \in V(T) : \text{dist}(q, x) = i\}.$$


We can see that h_i, h'_i can be described as diamonds with four sides. Let us among the four sides of h_i denote by X_i the side that is facing towards F' . Likewise, let us among the



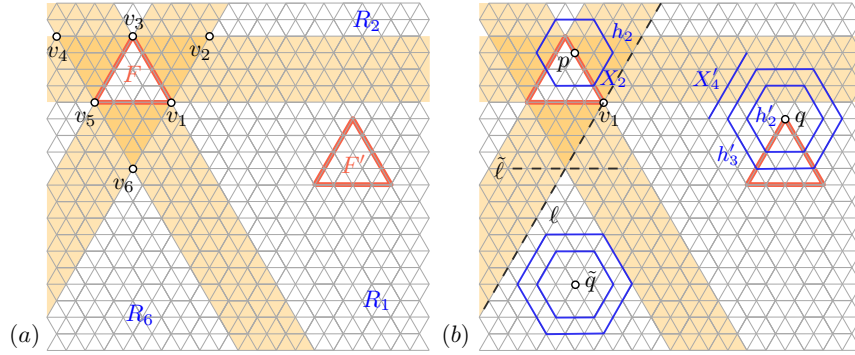
■ **Figure 4** For a face F in a square/triangular/hexagonal (super-)tiling, its halfslabs are the infinite slabs enclosing F along the directions depicted here (orange).

four sides of h'_i denote by X'_i the side that is facing towards F . Note that all of X_i, X'_i are parallel for $i \in \mathbb{N}$. Let ℓ denote the line through vertex v parallel to the X_i . Let $a \in \mathbb{N}$ be minimal such that h_a touches ℓ . Since X_a is parallel to ℓ , we have $X_a \subseteq \ell$. Furthermore, no matter which $p \in V(T) \cap \overline{F}$ was chosen as the centre of the circles h_i , we always have $v \in X_a$. This follows due to the way and the angle in which the circles h_i expand. Analogously, let $b \in \mathbb{N}$ be minimal such that h'_b touches ℓ . Due to the way the h'_i expand and since F and F' are not aligned, we have $v \in X'_b$. Now, the circle h_a is entirely on one side of ℓ , the circle h'_b is entirely on the other side, and $v \in h_a \cap h'_b$. By definition of h_a, h'_b , this implies $\text{dist}(p, q) = a + b$ and there is a shortest path between p, q that passes through v .

For the case of triangular and hexagonal grids, the situation is depicted in Figures 6 and 7. Here the halfslabs divide the plane into 6 cone-shaped regions R_1, \dots, R_6 . Let v_j be the point at the apex of R_j for $j = 1, \dots, 6$. Note that v_j is also a vertex of $V(T)$. For both the triangular and hexagonal grid, the definitions of $h_i, h'_i, X_i, X'_i, \ell, a, b$ can be made in an analogous matter. Analogously, it follows that $v_j \in h_a \cap h'_b$ which concludes the lemma. ◀



■ **Figure 5** (a) The two faces F, F' of the square super-tiling \mathbb{T} are non-aligned.




■ **Figure 6** (a) Vertex v_1 well-separates non-aligned faces in a triangular grid. (b) Proof.

Algorithm overview. Given T and paths P and Q , we first compute a suitable super-tiling:

► **Lemma 11.** *Let T be a regular unit tiling, and let P and Q be two paths in T of lengths n and m . In $O((n + m)^{1.5})$ time, we can compute a super-tiling \mathbb{T} such that:*

- *the edges of \mathbb{T} have a length in $\Theta(\sqrt{n + m})$, and*
- *the total number of induced subpaths in $S(P, \mathbb{T}) \cup S(Q, \mathbb{T})$ (see Definition 4) is $O(\sqrt{n + m})$.*

Proof. Our proof is inspired by that of Chan and Rahmati [14, Lemma 1], who prove an analogous statement for grids and curves in higher-dimensional real space (without graphs). For the proof, we assume without loss of generality that T is axis-aligned.




■ **Figure 7** (a) Vertex v_1 well-separates non-aligned faces in a hexagonal grid. (b) Proof.

Let $f = \lceil \sqrt{n+m} \rceil$, and let \mathbb{T}_0 be any axis-aligned super-tiling of T with an edge length in $\Theta(f)$. For each $i \in [f-1]$, define \mathbb{T}_i to be the super-tiling obtained by shifting \mathbb{T}_0 by the vector (i, i) if T is square, or by $(0, i\sqrt{3})$ if T is triangular or hexagonal. These shifted super-tilings remain axis-aligned super-tilings of T , as all their vertices align with those of T .

Since we shift in a direction not aligned with any edge of T , and since $i \leq f-1$, each edge or vertex of T lies on the boundary of $O(1)$ faces across the family $\{\mathbb{T}_i\}_{i=0}^{f-1}$. It follows that for any discrete vertex set $V \subseteq V(T)$ with $|V| = k$, there exists some \mathbb{T}_i such that only $O(k/f)$ vertices lie on boundaries of faces in \mathbb{T}_i . Therefore, there exists a super-tiling $\mathbb{T} \in \{\mathbb{T}_0, \dots, \mathbb{T}_{f-1}\}$ such that the total number of boundary vertices from P and Q is $O((n+m)/f)$, which implies $O((n+m)/f) = O(\sqrt{n+m})$ induced subpaths. To compute \mathbb{T} , we iterate over all vertices v of P and Q and, for each, determine the $O(1)$ super-tilings \mathbb{T}_i with $v \in B(\mathbb{T}_i)$ by iterating over all $\mathbb{T} \in \{\mathbb{T}_0, \dots, \mathbb{T}_{f-1}\}$. This takes $O((n+m)f)$ total time. We then select the super-tiling \mathbb{T}_i that minimizes $|(P \cup Q) \cap B(\mathbb{T}_i)|$. ◀

Let \mathbb{T} be the computed super-tiling of T . We partition P and Q into induced subpaths, denoted by $S(P, \mathbb{T})$ and $S(Q, \mathbb{T})$. These $O(\sqrt{n+m})$ subpaths naturally induce a subdivision of the free-space diagram M_Δ into $O(n+m)$ rectangular submatrices $M_\Delta[i:i', j:j']$, where $P[i:i'] \in S(P, \mathbb{T})$ and $Q[j:j'] \in S(Q, \mathbb{T})$ (see Figure 8(a)).

Our algorithm makes use of a wave-front method. We define a wave-front W as a monotone walk through the grid $[1, n] \times [1, m]$ that separates $(1, 1)$ from (n, m) . At each grid point $(i, j) \in W$, we store a Boolean value indicating whether there exists a monotone walk F from $(1, 1)$ to (i, j) such that all points on F lie in free space: that is, $M_\Delta[x, y] = 0$ for all $(x, y) \in F$ (Figure 8(b)).



■ **Figure 8** (a) For a given tiling T , we first compute a super-tiling \mathbb{T} which partitions the edges of P and Q . These subpaths divide the Δ -free space matrix M_Δ into submatrices that correspond to aligned and non-aligned curve pairs (green dashed versus neutral). (b) We maintain a wave-front W which is a monotone curve through M_Δ that has a Boolean value at each entry. An *update* selects a submatrix $M_\Delta[i:i', j:j']$ and removes from W its bottom-left facets (adding its top-right facets).

An *update* to the wave-front W selects a pair of subpaths $(P[i : i'], Q[j : j']) \in S(P, \mathbb{T}) \times S(Q, \mathbb{T})$ such that the bottom and left facets of the submatrix $M_\Delta[i : i', j : j']$ lie on the current wave-front. The update removes these facets from W and inserts the top and right facets instead. To perform the update, we consider the faces F_P and F_Q of \mathbb{T} that $P[i : i']$, respectively $Q[j : j']$, are assigned to, and proceed via a case distinction:

1. If F_P and F_Q are not aligned, then they are well-separated by a vertex v (by Lemma 10), and we show how to perform a fast update using v .
2. If F_P and F_Q are aligned, we perform a brute-force update and use a counting argument to bound the total running time of these updates.

We analyse both cases separately, observing that while some instances may involve only Case 1 or only Case 2 submatrices, we can independently bound the total cost for each case.

Case 1: Updating with $(P[i : i'], Q[j : j'])$ where F_P and F_Q are not aligned. We show an update algorithm that is near-linear in the perimeter of the submatrix $M_\Delta[i : i', j : j']$.

► **Lemma 12.** *Let $P[i : i'] \in S(P, \mathbb{T})$ and $Q[j : j'] \in S(Q, \mathbb{T})$ be subpaths whose faces F_P and F_Q in \mathbb{T} are not aligned. Then the corresponding wave-front update can be performed in $O(k \log k)$ time, where $k = |P[i : i']| + |Q[j : j']|$.*

Proof. By Lemma 10, the subcurves $P[i : i']$ and $Q[j : j']$ are well-separated by some vertex $v \in V(T)$, which we can identify in constant time. Given v , we compute two curves in \mathbb{R} whose Δ -free space matrix equals $M_\Delta[i : i', j : j']$. We define these curves as follows.

For a point $p \in P[i : i']$ we let $\bar{p} = \text{dist}(p, v)$. Likewise, for a point $q \in Q[j : j']$, we let $\bar{q} = -\text{dist}(q, v)$. Note the difference in sign. Because $P[i : i']$ and $Q[j : j']$ are well-separated by v , we have $\text{dist}(p, q) = \text{dist}(p, v) + \text{dist}(v, q) = |\bar{p} - \bar{q}|$ for any $p \in P[i : i']$ and $q \in Q[j : j']$. We let $\bar{P} = (\bar{p}_i, \dots, \bar{p}_{i'})$ and $\bar{Q} = (\bar{q}_j, \dots, \bar{q}_{j'})$.

The Δ -free-space matrix $M_\Delta[i : i', j : j']$ is equal to the Δ -free space matrix of the sequences \bar{P} and \bar{Q} over the real line under the absolute value metric. Let \bar{M}_Δ be this matrix. We compute the top and right facets of \bar{M}_Δ in $O(k \log k)$ time with existing fast algorithms for 1D separated sequences [8, 31]. ◀

Case 2: Updating with $(P[i : i'], Q[j : j'])$ where F_P and F_Q are aligned. For aligned face pairs, we fall back on a brute-force update strategy based on *switching cells*, a concept introduced by Har-Peled, Knauer, Wang and Wenk. [4].

► **Definition 13** (Switching cells). *A cell (i, j) in M_Δ is a switching cell if $M_\Delta[i, j] = 0$, but either $M_\Delta[i, j - 1] = 1$ or $M_\Delta[i, j + 1] = 1$ (or both).*

The switching cells in a submatrix $M_\Delta[i : i', j : j']$ allow for a direct update of the wave-front:

► **Lemma 14** ([4]). *Let $P[i : i']$ and $Q[j : j']$ be two subpaths. Given all s switching cells in $M_\Delta[i : i', j : j']$, a brute-force wave-front update takes $O(|P[i : i']| + |Q[j : j']| + s)$ time.*

We apply the brute-force update algorithm to all subpath pairs $(P[i : i'], Q[j : j']) \in S(P, \mathbb{T}) \times S(Q, \mathbb{T})$ whose associated faces F_P and F_Q in \mathbb{T} are aligned. We now show that the total number of switching cells encountered in these submatrices is not too large:

► **Lemma 15.** *Let T be a regular unit tiling, and let P and Q be two paths in T . Let \mathbb{T} be any super-tiling of T with edge length $\Theta(\sqrt{n+m})$. Then there are at most $O(n\sqrt{n+m})$ pairs (p_i, q_j) such that $M_\Delta[i, j]$ is a switching cell and p_i and q_j lie on aligned faces of \mathbb{T} . We can compute these pairs in $O(n\sqrt{n+m} \log m)$ time.*

Proof. Fix a vertex p_i and let F_p be a face of \mathbb{T} such that $p_i \in \overline{F_p}$. Since P and Q are paths in a unit tiling, a switching cell $M_\Delta[i, j]$ occurs only if q_j lies at distance exactly Δ from p_i .

Let $B_i(\Delta) \subseteq V(T)$ denote the set of all vertices at distance exactly Δ from p_i . This set lies on the boundary of a shape whose geometry depends on the tiling type:

- $B_i(\Delta)$ is the boundary of a diamond if T is a square tiling.
- $B_i(\Delta)$ is the boundary of a regular hexagon if T is a triangular tiling.
- $B_i(\Delta)$ is the boundary of a (non-regular) hexagon if T is a hexagonal tiling.

These regions are convex. Consider a face F' of \mathbb{T} that is aligned with F_p . The intersection of $B_i(\Delta)$ with $\overline{F'}$ contains at most $O(\sqrt{n+m})$ vertices. Each of these vertices could correspond to at most one vertex $q_j \in Q$ where (i, j) is a switching cell. There are only $O(1)$ faces of \mathbb{T} that are aligned with F_p , whose closures intersect $B_i(\Delta)$. Hence, there are at most $O(\sqrt{n+m})$ such q_j values per p_i . Summing over all n vertices of P gives the desired $O(n\sqrt{n+m})$ bound.

What remains is to compute these pairs (p_i, q_j) . We preprocess Q in a membership query data structure. For each $p_i \in P$, we select its assigned face F_p in constant time and we identify the $O(1)$ faces F' of \mathbb{T} that are aligned with F_p and intersect $B_i(\Delta)$ in constant time. For every such face F' , we iterate over the vertices v in $\overline{F'} \cap B_i(\Delta)$. We find if there exists a $q_j \in Q$ with $v = q_j$ in $O(\log m)$ time through a membership query. If so, then we compute $\text{dist}(p_i, q_{j+1})$ and $\text{dist}(p_i, q_{j-1})$ to determine whether $M_\Delta[i, j]$ is a switching cell. ◀

We are now ready to prove our main theorem:

► **Theorem 6.** *Let T be a regular unit tiling and (P, Q) be paths in T with $|P| = n$ and $|Q| = m$. Given $\Delta \in \mathbb{R}$, we can output whether $D_T(P, Q) \leq \Delta$ in $O((n+m)^{1.5} \log(n+m))$ time.*

Proof. Given T , P , and Q , we begin by applying Lemma 11 to compute in $O((n+m)^{1.5})$ time a super-tiling \mathbb{T} and the corresponding subpaths $S(P, \mathbb{T})$ and $S(Q, \mathbb{T})$. The number of subpath pairs $(P[i : i'], Q[j : j']) \in S(P, \mathbb{T}) \times S(Q, \mathbb{T})$ is $O(n+m)$. For each pair, we determine if the corresponding faces F_P and F_Q in \mathbb{T} are aligned, which takes $O(n+m)$ total time.

The product set $S(P, \mathbb{T}) \times S(Q, \mathbb{T})$ induces a grid of $O(n+m)$ rectangular submatrices of M_Δ with $O(\sqrt{n+m})$ rows and columns. Thus:

$$\sum_{\text{submatrices } M' \text{ of } M_\Delta} |\text{facets}(M')| = \sum_{(P[i:i'], Q[j:j']) \in S(P, \mathbb{T}) \times S(Q, \mathbb{T})} (i' - i + j' - j) \in O((n+m)^{1.5}). \quad (1)$$

Next, we initialize the wave-front W by computing the first row and first column of M_Δ , i.e., all entries $M_\Delta[i, 1]$ for $i \in [n]$ and $M_\Delta[1, j]$ for $j \in [m]$, in $O(n+m)$ time. We then iterate diagonally through all pairs $(P[i : i'], Q[j : j'])$ in $S(P, \mathbb{T}) \times S(Q, \mathbb{T})$ and perform a wave-front update for each pair.

If the two faces F_P and F_Q of \mathbb{T} assigned to $P[i : i']$ and $Q[j : j']$ are *not aligned* then we apply Lemma 12. The total time spent in this case equals:

$$\sum_{(P[i:i'], Q[j:j']) \text{ where } (F_P, F_Q) \text{ are not aligned}} O((i' - i + j' - j) \log(i' - i + j' - j))$$

By Equation 1, this sum is upper bound by $O((n+m)^{1.5} \log(n+m))$.

If the two faces F_P and F_Q of \mathbb{T} assigned to $P[i : i']$ and $Q[j : j']$ are aligned then we apply Lemma 14 instead. The total time spent is:

$$\sum_{P[i:i'], Q[j:j']) \text{ where } (F_P, F_Q) \text{ are aligned}} O((i' - i + j' - j) + \text{switching cells}(M_\Delta[i : i', j : j']))$$

The first term and second term are bounded by $O((n+m)^{1.5} \log(n+m))$ by Equation 1 and Lemmas 14+15, respectively. Once the wave-front W reaches the entry (n, m) , we have determined whether $D_T(P, Q) \leq \Delta$, proving the theorem. ◀

5 The discrete Fréchet distance under the L_1 metric

We extend our techniques for paths in tilings to curves in the plane. Fix some universal constant $\gamma > 0$ and fix some value $\varepsilon \in (0, 1)$. We consider \mathbb{R}^2 under the L_1 metric, and two polygonal curves P and Q with n and m vertices, respectively. In full generality, we only require that all edges in P have length at most γ and that Q is an (ε, δ) -curve. Denote their discrete Fréchet distance under the L_1 metric by $D_F(P, Q)$. We show that we can compute $D_F(P, Q)$ in $\tilde{O}\left(\frac{\sqrt{\delta}}{\varepsilon}(n+m)^{1.5}\right)$ time. We adapt this algorithm in the full version of this paper, to make it work for any L_c metric, by allowing some inaccuracy.

Our approach. Given some $\Delta \in \mathbb{R}$, we first show how to decide whether $D_F(P, Q) \leq \Delta$ using the wave-front propagation technique from Section 4, with suitable modifications for curves in the plane. Let $t > 1$ be a parameter to be determined later. We define an axis-aligned square tiling \mathbb{T} of the plane where each square has edge length t .

► **Definition 16.** We define the breakpoints $B(P, \mathbb{T})$ of P as the set of vertices $p \in V(P)$ such that some edge of P incident to p intersects an edge of \mathbb{T} .

Note that $B(P, \mathbb{T})$ may contain overlapping but distinct vertices of P (since an (ε, δ) -curve may visit the same point up to δ times). If we remove all edges from P which intersect some edge of \mathbb{T} , we see that the breakpoints induce a partition of the vertices of P into pairwise vertex-disjoint subcurves $P[i : i']$, each fully contained in the closure \bar{F} of a face $F \in \mathbb{T}$. We refer to these subcurves as the *induced subcurves* $S(P, \mathbb{T})$ and we assign each induced subcurve to a face in \mathbb{T} . We prove that for a suitable choice of \mathbb{T} , the total number of induced subcurves satisfies $|S(P, \mathbb{T}) \cup S(Q, \mathbb{T})| = O((n+m)/t)$.

We then apply our wave-front algorithm over the product set $S(P, \mathbb{T}) \times S(Q, \mathbb{T})$. For a pair of subcurves $(P[i : i'], Q[j : j'])$, let F_P and F_Q denote their respective containing faces in \mathbb{T} . We distinguish two cases:

- If F_P and F_Q are not aligned, we show that $(P[i : i'], Q[j : j'])$ is well-separated by a corner of either F_P or F_Q . We then apply Lemma 12 to perform a wave-front update in $O((i' - i + j' - j) \log(n+m))$ time.
- If F_P and F_Q are aligned, we apply the brute-force algorithm of Lemma 14, but with a more careful analysis of the number of switching cells.

Balancing the choice of t is critical: a larger t reduces the number of wave-front updates, while a smaller t reduces the number of switching cells encountered. We choose t to optimize the overall running time.

Choosing a tiling. We construct a tiling \mathbb{T} such that it creates few induced subcurves.

► **Lemma 17.** Given $t \geq 1$, we can compute in $O((n+m)t)$ time a square tiling \mathbb{T} of edge length t such that the number of induced subcurves satisfies $|S(P, \mathbb{T}) \cup S(Q, \mathbb{T})| = O((n+m)/t)$.

Proof. The proof mirrors that of Lemma 11. Let \mathbb{T}_0 be an axis-aligned square tiling with edge length t , and let \mathbb{T}_i denote the tiling obtained by shifting \mathbb{T}_0 by (i, i) for $i \in [t-1]$. We required that each edge e of P or Q has length at most $\gamma \in O(1)$. Thus each edge intersects the boundary of at most $O(1)$ of the shifted tilings. It follows that there exists an index i such that \mathbb{T}_i induces only $O((n+m)/t)$ breakpoints across P and Q .

To find \mathbb{T}_i , we iterate over all $O(n + m)$ edges of P and Q . For each edge, we test for intersection with each of the t candidate tilings in constant time. This gives an overall running time of $O((n + m)t)$ to find the tiling \mathbb{T}_i that minimizes $|B(P, \mathbb{T}_i) \cup B(Q, \mathbb{T}_i)|$. ◀

Wave-front updates. Unlike in Section 4, the subcurves in $S(P, \mathbb{T})$ are vertex-disjoint. Thus, we never have a submatrix $M_\Delta[i : i', j : j']$ whose bottom-left facet coincides with the current wave-front W . However, we can always find a submatrix $M' := M_\Delta[i : i', j : j']$ where each of the cells on the bottom-left facet of M' is adjacent to a cell in W . We extend W to cover it in $O(i' - i + j' - j)$ time via brute force and thereby define wave-front updates analogously.

Case 1: Updating with $(P[i : i'], Q[j : j'])$ where F_P and F_Q are not aligned. Our key observation is that, even though the edges of $P[i : i']$ and $Q[j : j']$ do not coincide with \mathbb{T} , a pair of faces (F_P, F_Q) of \mathbb{T} are still well-separated whenever they are not aligned:

► **Lemma 18.** *Let $P[i : i'] \in S(P, \mathbb{T})$ and $Q[j : j'] \in S(Q, \mathbb{T})$ be subcurves whose faces F_P and F_Q in \mathbb{T} are not aligned. Then the corresponding wave-front update can be performed in $O(k \log k)$ time, where $k = |P[i : i']| + |Q[j : j']|$.*

Proof. Assume, without loss of generality, that the bottom-left corner of F_P lies above and to the right of the top-right corner of F_Q . Then any shortest L_1 path from a point in $\overline{F_P}$ to one in $\overline{F_Q}$ may be rerouted through that corner. The rest follows from Lemma 12. ◀

Case 2: Updating with $(P[i : i'], Q[j : j'])$ where F_P and F_Q are aligned. We upper bound the number of switching cells $M_\delta[i, j]$ where p_i and q_j lie in aligned faces of \mathbb{T} :

► **Lemma 19.** *Let \mathbb{T} be the tiling from Lemma 17. Then there are at most $O(\frac{\delta n(t+1)}{\varepsilon^2})$ pairs (p_i, q_j) such that $M_\Delta[i, j]$ is a switching cell and p_i and q_j lie in aligned faces of \mathbb{T} . We can compute these pairs in $O(\frac{\delta n(t+1)}{\varepsilon^2} \log m)$ time.*

Proof. We preprocess Q by snapping its vertices to a square grid G of edge length ε^{-1} , forming Q' . Each vertex in Q' corresponds to $O(\delta)$ original vertices since Q is an (ε, δ) -curve. We store Q' in a spatial data structure (e.g., a lexicographically sorted balanced binary tree), where a query point z returns all $q_j \in Q$ corresponding to z in $O(\delta + \log n)$ time.

Recall that all edges of Q have length at most $\gamma \in O(1)$. Fix a vertex p_i and let F_p be a face of \mathbb{T} that contains p_i . Let $B_i(\Delta)$ be the metric circle (under the L_1 metric) of radius Δ centred at p_i . Let Γ denote a ball with radius 2γ and $B_i^*(\Delta)$ be the metric annulus that is the Minkowski-sum $\Gamma \oplus B_i(\Delta)$. Observe that an edge (q_j, q_{j+1}) of Q intersects $B_i(\Delta)$ only if one of its corresponding vertices q'_j and q'_{j+1} of Q' lies in $B_i^*(\Delta)$.

There are $O(1)$ faces F' of \mathbb{T} that are aligned with F_p and that intersect $B_i^*(\Delta)$. Let $A := B_i^*(\Delta) \cap F'$. Distinguishing between $\Delta > 1$ and $\Delta \leq 1$, we see that A has an area of $O(\max\{1, t\}\gamma)$ and therefore A contains at most $O(\frac{\gamma \max\{1, t\}}{\varepsilon^2}) = O(\frac{t+1}{\varepsilon^2})$ vertices of G . We iterate over the vertices in $B_i^*(\Delta) \cap \overline{F'} \cap V(G)$ and for each we perform a membership query on Q' . This returns $O(\frac{\delta(t+1)}{\varepsilon^2})$ vertices of Q' . For each reported vertex q'_j we test if $M_\Delta[i, j]$ is a switching cell in constant time. Thus, for any $p_i \in P$, we compute all switching cells $M_\Delta[i, j]$ that correspond to the lemma statement in $O(\frac{\delta(t+1)}{\varepsilon^2} \log m)$ time. ◀

► **Theorem 20.** *Let P and Q be two curves in \mathbb{R}^2 under the L_1 metric with $|P| = n$ and $|Q| = m$. Fix a universal constant $\gamma \in O(1)$. For $\varepsilon > 0$ (with $\varepsilon \in \omega(\sqrt{\frac{\delta}{n}})$) and $\Delta \in \mathbb{R}$, if P has edge length at most γ and Q is an (ε, δ) -curve, then we can compute whether $D_F(P, Q) \leq \Delta$ in $O(\frac{\sqrt{\delta}}{\varepsilon}(n + m)^{1.5} \log(n + m))$ time.*

Proof. Let $t \geq 1$ be a value that we set later. Given P and Q , we apply Lemma 17 to compute in $O((n+m)t)$ time our tiling \mathbb{T} and the $O((n+m)/t)$ induced subcurves $S(P, \mathbb{T})$ and $S(Q, \mathbb{T})$. The product set $S(P, \mathbb{T}) \times S(Q, \mathbb{T})$ induces a grid over M_Δ with $O((n+m)/t)$ rows and columns. Thus:

$$\sum_{\text{submatrices } M' \text{ of } M_\Delta} |\text{facets}(M')| = \sum_{(P[i:i'], Q[j:j']) \in S(P, \mathbb{T}) \times S(Q, \mathbb{T})} (i' - i + j' - j) \in O((n+m)^2/t). \quad (2)$$

We iteratively perform wave-front updates with the input $(P[i:i'], Q[j:j'])$. Let F_P and F_Q denote the corresponding two faces of \mathbb{T} . If F_P and F_Q are *not aligned* then we apply Lemma 14. The total time spent in this case equals:

$$\sum_{(P[i:i'], Q[j:j']) \text{ where } (F_P, F_Q) \text{ are not aligned}} O((i' - i + j' - j)) \log(n+m)$$

By Equation 2, this sum is upper bound by $O(\frac{(n+m)^2}{t} \log(n+m))$.

If F_P and F_Q are aligned then we apply Lemma 15 instead. The total time spent is:

$$\sum_{P[i:i'], Q[j:j'] \text{ where } (F_P, F_Q) \text{ are aligned}} O((i' - i + j' - j) + \text{switching cells}(M_\Delta[i:i', j:j']))$$

The first term is upper bounded by $O(\frac{(n+m)^2}{t})$ by Equation 2. The second term is upper bounded by $O(\frac{\delta n(t+1)}{\varepsilon^2} \log m)$ by combining Lemma 14 and 19. It follows, assuming that $\frac{\varepsilon}{\sqrt{\delta}} \sqrt{n}$ goes to infinity, that if we choose $t \in \Theta(\frac{\varepsilon \cdot (n+m)}{\sqrt{\delta n}})$ the total running time is upper bounded by $O(\frac{\sqrt{\delta}}{\varepsilon} (n+m)^{1.5} \log(n+m))$. ◀

Computing $D_F(P, Q)$. Given P and Q , there is a set of $O(nm)$ values S , which can be computed in $O(nm)$ time, such that $D_F(P, Q) \in S$. The classical algorithm to compute $D_F(P, Q)$ performs binary search over S . Since we want to use subquadratic time, we cannot compute the set S explicitly. Instead, we represent S as the *Cartesian sum* $X \oplus Y$ of two sets of $O(n+m)$ values. The Cartesian sum of X and Y is the multiset $X \oplus Y = \{x + y \mid x \in X, y \in Y\}$. We can use existing techniques [21, 27] for selection in Cartesian sums to binary search over the Cartesian sum $X \oplus Y$ without explicitly computing its $O(|X| \cdot |Y|)$ elements.

Specifically, we define X and Y as follows. The set X contains, for each point $(x, y) \in P$, the four values $x + y$, $-x + y$, $x - y$ and $-x - y$. The set Y is defined analogously for points in Q . For all $(p, q) \in P \times Q$, the distance between p and q is the sum of an element in X and an element in Y . That is, their distance is an element of $X \oplus Y$. Hence, $D_F(P, Q) \in X \oplus Y$.

After sorting X and Y , we can compute the i^{th} smallest value in $X \oplus Y$, for any given i , in $O(|X| + |Y|) = O(n+m)$ time [21, 27]. We binary search over the integers $1, \dots, |X| \cdot |Y|$, and for each considered integer i , we compute the i^{th} smallest value Δ in $X \oplus Y$. Then we use our decision algorithm (Theorem 20) to decide whether $D_F(P, Q) \leq \Delta$ to guide the search to $D_F(P, Q)$. Thus, we obtain the following result:

► **Corollary 21.** *Given $\gamma > 0$, let P and Q be curves in \mathbb{R}^2 under L_1 with $|P| + |Q| = n + m$. For any $\varepsilon > 0$, if P has edge length at most γ and Q is an (ε, δ) -curve, then $D_F(P, Q)$ can be computed in $O(\frac{\sqrt{\delta}}{\varepsilon} (n+m)^{1.5} \log^2(n+m))$ time.*

References

- 1 Pankaj K Agarwal, Rinat Ben Avraham, Haim Kaplan, and Micha Sharir. Computing the discrete Fréchet distance in subquadratic time. *SIAM Journal on Computing*, 43(2):429–449, 2014. doi:10.1137/130920526.

- 2 Helmut Alt and Michael Godau. Computing the Fréchet distance between two polygonal curves. *International Journal of Computational Geometry & Applications*, 5(01n02):75–91, 1995. doi:10.1142/S0218195995000064.
- 3 Helmut Alt, Christian Knauer, and Carola Wenk. Comparison of distance measures for planar curves. *Algorithmica*, 38(1):45–58, 2004. doi:10.1007/s00453-003-1042-5.
- 4 Boris Aronov, Sariel Har-Peled, Christian Knauer, Yusu Wang, and Carola Wenk. Fréchet distance for curves, revisited. In *European Symposium on Algorithms (ESA)*, pages 52–63, 2006. doi:10.1007/11841036_8.
- 5 Yoonsik Bang, Jiyoung Kim, and Kiyun Yu. An improved map-matching technique based on the Fréchet distance approach for pedestrian navigation services. *Sensors*, 16(10), 2016. doi:10.3390/s16101768.
- 6 Sotiris Brakatsoulas, Dieter Pfoser, Randall Salas, and Carola Wenk. On map-matching vehicle tracking data. In *International Conference on Very Large Data Bases (VLDB)*, pages 853–864, 2005. doi:10.5555/1083592.1083691.
- 7 Karl Bringmann. Why walking the dog takes time: Fréchet distance has no strongly subquadratic algorithms unless SETH fails. In *IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 661–670, 2014. doi:10.1109/FOCS.2014.76.
- 8 Karl Bringmann and Marvin Künnemann. Improved approximation for Fréchet distance on c-packed curves matching conditional lower bounds. *International Journal of Computational Geometry & Applications*, 27(1-2):85–120, 2017. doi:10.1142/S0218195917600056.
- 9 Karl Bringmann and Wolfgang Mulzer. Approximability of the discrete Fréchet distance. *Journal of Computational Geometry*, 7(2):46–76, 2016. doi:10.4230/LIPIcs.SOCG.2015.739.
- 10 Kevin Buchin, Maike Buchin, David Duran, Brittany Terese Fasy, Roel Jacobs, Vera Sacristan, Rodrigo I Silveira, Frank Staals, and Carola Wenk. Clustering trajectories for map construction. In *ACM International Conference on Advances in Geographic Information Systems (SIGSPATIAL)*, pages 1–10, 2017. doi:10.1145/3139958.3139964.
- 11 Kevin Buchin, Maike Buchin, Wouter Meulemans, and Wolfgang Mulzer. Four soviets walk the dog: Improved bounds for computing the Fréchet distance. *Discrete & Computational Geometry*, 58(1):180–216, 2017. doi:10.1137/1.9781611973402.103.
- 12 Kevin Buchin, Tim Ophelders, and Bettina Speckmann. SETH says: Weak Fréchet distance is faster, but only if it is continuous and in one dimension. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2887–2901, 2019. doi:10.5555/3310435.3310614.
- 13 Maike Buchin, Bernhard Kilgus, and Andrea Kölzsch. Group diagrams for representing trajectories. *International Journal of Geographical Information Science*, 34(12):2401–2433, 2020. doi:10.1145/3283207.3283208.
- 14 Timothy Chan and Zahed Rahmati. An improved approximation algorithm for the discrete Fréchet distance. *Information Processing Letters (IPL)*, 138:72–74, 2018. doi:10.1016/J.IPL.2018.06.011.
- 15 Daniel Chen, Anne Driemel, Leonidas Guibas, Andy Nguyen, and Carola Wenk. Approximate map matching with respect to the Fréchet distance. In *ACM Workshop on Algorithm Engineering and Experiments (ALENEX)*, pages 75–83, 2011. doi:10.1137/1.9781611972917.8.
- 16 Thomas Devogele. A new merging process for data integration based on the discrete Fréchet distance. In *Advances in spatial data handling*, pages 167–181, 2002. doi:10.1007/978-3-642-56094-1_13.
- 17 Anne Driemel, Sariel Har-Peled, and Carola Wenk. Approximating the Fréchet distance for realistic curves in near linear time. *Discrete Computational Geometry*, 48(1):94–127, 2012. doi:10.1007/s00454-012-9402-z.
- 18 Anne Driemel, Amer Krivošija, and Christian Sohler. Clustering time series under the Fréchet distance. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 766–785. SIAM, 2016. doi:10.5555/2884435.2884490.

- 19 Anne Driemel, Ivor van der Hoog, and Eva Rotenberg. On the discrete Fréchet distance in a graph. In *International Symposium on Computational Geometry (SoCG)*, pages 36:1–36:18, 2022. doi:10.4230/LIPICS.SOCG.2022.36.
- 20 Thomas Eiter and Heikki Mannila. Computing discrete Fréchet distance. Technical Report CD-TR 94/64, Christian Doppler Laboratory for Expert Systems, TU Vienna, Austria, 1994.
- 21 Greg N. Frederickson and Donald B. Johnson. Generalized selection and ranking: Sorted matrices. *SIAM Journal of Computing*, 13(1):14–30, 1984. doi:10.1137/0213002.
- 22 Branko Grünbaum and Geoffrey Colin Shephard. *Tilings and patterns*. Courier Dover Publications, 1987. doi:10.5555/19304.
- 23 Joachim Gudmundsson, Majid Mirzanezhad, Ali Mohades, and Carola Wenk. Fast Fréchet distance between curves with long edges. *International Journal of Computational Geometry & Applications*, 29(02):161–187, 2019. doi:10.1145/3191801.3191811.
- 24 Richard Kenefic. Track clustering using Fréchet distance and minimum description length. *Journal of Aerospace Information Systems*, 11(8):512–524, 2014. doi:10.2514/1.I010170.
- 25 Maximilian Konzack, Thomas McKetterick, Tim Ophelders, Maike Buchin, Luca Giuggioli, Jed Long, Trisalyn Nelson, Michel A Westenberg, and Kevin Buchin. Visual analytics of delays and interaction in movement data. *International Journal of Geographical Information Science*, 31(2):320–345, 2017. doi:10.5555/3048386.3048392.
- 26 Ariane Mascaret, Thomas Devogele, Iwan Le Berre, and Alain Hénaff. Coastline matching process based on the discrete Fréchet distance. In *Progress in Spatial Data Handling*, pages 383–400. Springer, 2006. doi:10.1007/3-540-35589-8_25.
- 27 Andranik Mirzaian and Eshrat Arjomandi. Selection in X+Y and matrices with sorted rows and columns. *Information Processing Letters*, 20(1):13–17, 1985. doi:10.1016/0020-0190(85)90123-1.
- 28 Roniel Sousa, Azzedine Boukerche, and Antonio Loureiro. Vehicle trajectory similarity: Models, methods, and applications. *ACM Computing Surveys*, 53(5), 2020. doi:10.1145/3406096.
- 29 E Sriraghavendra, K Karthik, and Chiranjib Bhattacharyya. Fréchet distance based approach for searching online handwritten documents. In *Ninth International Conference on Document Analysis and Recognition (ICDAR)*, pages 461–465, 2007. doi:10.1109/ICDAR.2007.4378752.
- 30 Han Su, Shuncheng Liu, Bolong Zheng, Xiaofang Zhou, and Kai Zheng. A survey of trajectory distance measures and performance evaluation. *Proceedings of the VLDB Endowment*, 29(1):3–32, 2020. doi:10.1007/s00778-019-00574-9.
- 31 Thijs van der Horst, Marc van Kreveld, Tim Ophelders, and Bettina Speckmann. The geodesic Fréchet distance between two curves bounding a simple polygon. *CoRR*, abs/2501.03834, 2025. doi:10.48550/arXiv.2501.03834.
- 32 Ryan Williams. A new algorithm for optimal 2-constraint satisfaction and its implications. *Theoretical Computer Science (TCS)*, 348(2-3):357–365, 2005. doi:10.1016/J.TCS.2005.09.023.
- 33 Dong Xie, Feifei Li, and Jeff M Phillips. Distributed trajectory similarity search. *Proceedings of the VLDB Endowment*, 10(11):1478–1489, 2017. doi:10.14778/3137628.3137655.