Why walking the dog takes time: Fréchet distance has no strongly subquadratic algorithms unless SETH fails

Karl Bringmann
Max Planck Institute for Informatics
Saarbrücken, Germany
Email: karl.bringmann@mpi-inf.mpg.de

***Abstract*—The Fréchet distance is a well-studied and very popular measure of similarity of two curves. Many variants and extensions have been studied since Alt and Godau introduced this measure to computational geometry in 1991. Their original algorithm to compute the Fréchet distance of two polygonal curves with n vertices has a runtime of $\mathcal{O}(n^2 \log n)$. More than 20 years later, the state of the art algorithms for most variants still take time more than $\mathcal{O}(n^2/\log n)$, but no matching lower bounds are known, not even under reasonable complexity theoretic assumptions.**

To obtain a conditional lower bound, in this paper we assume the Strong Exponential Time Hypothesis or, more precisely, that there is no $\mathcal{O}^*((2-\delta)^N)$ algorithm for CNF-SAT for any $\delta > 0$. Under this assumption we show that the Fréchet distance cannot be computed in strongly subquadratic time, i.e., in time $\mathcal{O}(n^{2-\delta})$ for any $\delta > 0$. This means that finding faster algorithms for the Fréchet distance is as hard as finding faster CNF-SAT algorithms, and the existence of a strongly subquadratic algorithm can be considered unlikely.

Our result holds for both the continuous and the discrete Fréchet distance. We extend the main result in various directions. Based on the same assumption we (1) show non-existence of a strongly subquadratic 1.001-approximation, (2) present tight lower bounds in case the numbers of vertices of the two curves are imbalanced, and (3) examine realistic input assumptions (c-packed curves).

***Keywords*-lower bounds; computational geometry; curves; inapproximability;**

I. Introduction

Intuitively, the (continuous) Fréchet distance of two curves P, Q is the minimal length of a leash required to connect a dog to its owner, as they walk along P or Q, respectively, without backtracking. The Fréchet distance is a very popular measure of similarity of two given curves. In contrast to distance notions such as the Hausdorff distance, it takes into account the order of the points along the curve, and thus better captures the similarity as perceived by human observers [3].

Alt and Godau introduced the Fréchet distance to computational geometry in 1991 [5, 24]. For polygonal curves P and Q with n and m vertices[1], respectively, they presented an $\mathcal{O}(nm\log(nm))$ algorithm. Since Alt and Godau's seminal paper, Fréchet distance has become a rich field of research, with various directions such as generalizations to surfaces (see, e.g., [4]), approximation algorithms for realistic input curves ([6, 7, 21]), the geodesic and homotopic Fréchet distance (see, e.g., [15, 17]), and many more variants (see, e.g., [11, 20, 29, 31]). Being a natural measure for curve similarity, the Fréchet distance has found applications in various areas such as signature verification (see, e.g., [32]), map-matching tracking data (see, e.g., [9]), and moving objects analysis (see, e.g., [12]).

A particular variant that we will also discuss in this paper is the *discrete* Fréchet distance. Here, intuitively the dog and its owner are replaced by two frogs, and in each time step each frog can jump to the next vertex along its curve or stay at its current vertex. Defined in [22], the original algorithm for the discrete Fréchet distance has runtime $\mathcal{O}(nm)$.

Recently, improved algorithms have been found for some variants. Agarwal et al. [2] showed how to compute the discrete Fréchet distance in (mildly) subquadratic time $\mathcal{O}\big(nm\frac{\log\log n}{\log n}\big)$. Buchin et al. [13] gave algorithms for the continuous Fréchet distance that runs in time $\mathcal{O}(n^2\sqrt{\log n}(\log\log n)^{3/2})$ on the Real RAM and $\mathcal{O}(n^2(\log\log n)^2)$ on the Word RAM. However, the problem remains open whether there is a *strongly subquadratic*[2] algorithm for the Fréchet distance, i.e., an algorithm with runtime $\mathcal{O}(n^{2-\delta})$ for any $\delta > 0$. For a particular variant, the discrete Fréchet distance with shortcuts, strongly subquadratic algorithms have been found recently [8], however, this seems to have no implications for the classical continuous or discrete Fréchet distance.

The only known lower bound shows that the Fréchet distance takes time $\Omega(n \log n)$ (in the algebraic decision tree model) [10]. The typical way of proving (conditional) quadratic lower bounds for geometric problems is via 3SUM [23], in fact, Alt conjectured that the Fréchet distance is 3SUM-hard. Buchin et al. [13] argued that the Fréchet distance is unlikely to be 3SUM-hard, because it has strongly subquadratic decision trees. However, their argument breaks down in light of a recent result showing strongly subquadratic decision trees also for 3SUM [25].

[1] We always assume that $m \leqslant n$.

[2] We use the term *strongly subquadratic* to differentiate between this runtime and the *(mildly) subquadratic* $\mathcal{O}(n^2\log\log n/\log n)$ algorithm from [2].

0272-5428/14 $31.00 © 2014 IEEE
DOI 10.1109/FOCS.2014.76

Hence, it is completely open whether the Fréchet distance is 3SUM-hard.

A. Strong Exponential Time Hypothesis

The Exponential Time Hypothesis (ETH) and the Strong Exponential Time Hypothesis (SETH), both introduced by Impagliazzo, Paturi, and Zane [27, 28], provide alternative ways of proving conditional lower bounds. ETH asserts that 3-SAT has no $2^{o(N)}$ algorithm, where N is the number of variables, and can be used to prove matching lower bounds for a wealth of problems, see [30] for a survey. However, since this hypothesis does not specify the exact exponent, it is not suited for proving polynomial time lower bounds, where the exponent is important.

The stronger hypothesis SETH asserts that there is no $\delta > 0$ such that k-SAT has an $\mathcal{O}((2-\delta)^N)$ algorithm for all k. In this paper, we will use the following slightly weaker variant, which has also been used in [33, 34].

Hypothesis SETH'**:** *There is no $\mathcal{O}^*((2-\delta)^N)$ algorithm for* CNF-SAT *for any $\delta > 0$. Here, $\mathcal{O}^*$ hides polynomial factors in the number of variables N and the number of clauses M.*

While SETH deals with formulas of width k, SETH' deals with CNF-SAT, i.e., unbounded width clauses. Thus, it is a weaker assumption and more likely to be true. Note that exhaustive search takes time $\mathcal{O}^*(2^N)$, and the fastest known algorithms for CNF-SAT are only slighly faster than that, namely of the form $\mathcal{O}^*(2^{N(1-C/\log(M/N))})$ for some positive constant C [14, 19]. Thus, SETH' is a reasonable assumption that can be considered unlikely to fail. It has been observed that one can use SETH and SETH' to prove lower bounds for polynomial time problems such as k-Dominating Set and others [33], the diameter of sparse graphs [34], and dynamic connectivity problems [1]. However, it seems to be applicable only for few problems, e.g., it seems to be a wide open problem to prove that 3SUM has no strongly subquadratic algorithms unless SETH fails, similarly for matching, maximum flow, edit distance, and other classic problems.

B. Main result

Our main theorem gives strong evidence that the Fréchet distance may have no strongly subquadratic algorithms by relating it to the Strong Exponential Time Hypothesis.

Theorem I.1. *There is no $\mathcal{O}(n^{2-\delta})$ algorithm for the (continuous or discrete) Fréchet distance for any $\delta > 0$, unless* SETH' *fails.*

Since SETH and its weaker variant SETH' are reasonable hypotheses, by this theorem one can consider it unlikely that the Fréchet distance has strongly subquadratic algorithms. In particular, any strongly subquadratic algorithm for the Fréchet distance would not only give improved algorithms for CNF-SAT that are much faster than exhaustive search, but also for various other problems such as Hitting Set, Set Splitting, and NAE-SAT via the reductions in [18]. Alternatively, in the spirit of [33], one can view the above theorem as a possible attack on CNF-SAT, as algorithms for the Fréchet distance now could provide a route to faster CNF-SAT algorithms. In any case, anyone trying to find strongly subquadratic algorithms for the Fréchet distance should be aware that this is as hard as finding improved CNF-SAT algorithms, which might be impossible.

We remark that all our lower bounds (unless stated otherwise) hold in the Euclidean plane, and thus also in $\mathbb{R}^d$ for any $d \geqslant 2$.

C. Extensions

We extend our main result in two important directions: We show approximation hardness and we prove that the lower bound still holds for restricted classes of curves.

First, it would be desirable to have good approximation algorithms in strongly subquadratic time, say a near-linear time approximation scheme. We exclude such algorithms by proving that there is no 1.001-approximation for the Fréchet distance in strongly subquadratic time unless SETH' fails. Hence, within $n^{o(1)}$-factors any 1.001-approximation takes as much time as an exact algorithm. We did not try to optimize the constant 1.001, but only to find the asymptotically largest possible approximation ratio, which seems to be a constant. We leave it as an open problem whether there is a strongly subquadratic $\mathcal{O}(1)$-approximation. The literature so far contains no strongly subquadratic approximation algorithms for general curves at all.

Second, it might be conceivable that if one curve has much fewer vertices than the other, i.e., $m \ll n$, then after some polynomial preprocessing on the smaller curve we can compute the Fréchet distance of the two curves quickly, e.g., in total time $\mathcal{O}((n+m^3)\log n)$. Note that such a runtime is not ruled out by the trivial argument that any algorithm needs time $\Omega(n+m)$ for reading the input, and is also not ruled out by Theorem I.1, since the runtime is not subquadratic for $n = m$. We rule out such runtimes by proving that there is no $\mathcal{O}((nm)^{1-\delta})$ algorithm "for any m", unless SETH' fails. More precisely, we prove this lower bound for the "special case" $m \approx n^\gamma$ for any constant $0 \leqslant \gamma \leqslant 1$. To make this formal, for any input parameter α and constants $\gamma_0 < \gamma_1$ in $\mathbb{R} \cup \{-\infty, \infty\}$, we say that a statement holds *for any polynomial restriction of* $n^{\gamma_0} \leqslant \alpha \leqslant n^{\gamma_1}$ if it holds restricted to instances with $n^{\gamma-\delta} \leqslant \alpha \leqslant n^{\gamma+\delta}$ for any constants $\delta > 0$ and $\gamma_0 + \delta \leqslant \gamma \leqslant \gamma_1 - \delta$. We obtain the following extension of the main result Theorem I.1, which yields tight lower bounds for any behaviour of m and any $(1+\varepsilon)$-approximation with $0 \leqslant \varepsilon \leqslant 0.001$.

Theorem I.2. *There is no 1.001-approximation with runtime $\mathcal{O}((nm)^{1-\delta})$ for the (continuous or discrete) Fréchet distance for any $\delta > 0$, unless* SETH' *fails. This holds for any polynomial restriction of $1 \leqslant m \leqslant n$.*

D. Realistic input curves

In attempts to capture the properties of realistic input curves, strongly subquadratic algorithms have been devised for restricted classes of inputs such as backbone curves [7], κ-bounded and κ-straight [6], and ϕ-low density curves [21]. The most popular model are *c-packed curves*, which have been used for various generalizations of the Fréchet distance [16, 20, 26]. Driemel et al. [21] introduced this model and presented a $(1+\varepsilon)$-approximation for the continuous Fréchet distance that runs in time $\mathcal{O}(cn/\varepsilon + cn\log n)$, which works in any $\mathbb{R}^d$, $d \geqslant 2$.

While the algorithm of [21] is near-linear for small c and $1/\varepsilon$, is is not clear whether its dependence on c and $1/\varepsilon$ is optimal for c and $1/\varepsilon$ that grow with n. We give strong evidence that the algorithm of [21] has optimal dependence on c for any constant $0 < \varepsilon \leqslant 0.001$.

Theorem I.3. *There is no* 1.001*-approximation with runtime* $\mathcal{O}((cn)^{1-\delta})$ *for the (continuous or discrete) Fréchet distance on c-packed curves for any* $\delta > 0$*, unless* SETH′ *fails. This holds for any polynomial restriction of* $1 \leqslant c \leqslant n$*.*

Since we prove this claim for any polynomial restriction $c \approx n^\gamma$, the above result excludes 1.001-approximations with runtime, say, $\mathcal{O}(c^2+n)$.

Regarding the dependence on ε, in any dimension $d \geqslant 5$ we can prove a conditional lower bound that matches the dependency on ε of [21] up to a polynomial. Due to space limitations, we omit its proof in this extended abstract.

Theorem I.4. *There is no* $(1+\varepsilon)$*-approximation for the (continuous or discrete) Fréchet distance on c-packed curves in* $\mathbb{R}^d$*,* $d \geqslant 5$*, with runtime* $\mathcal{O}(\min\{cn/\sqrt{\varepsilon}, n^2\}^{1-\delta})$ *for any* $\delta > 0$*, unless* SETH′ *fails. This holds for sufficiently small* $\varepsilon > 0$ *and any polynomial restriction of* $1 \leqslant c \leqslant n$ *and* $\varepsilon \leqslant 1$*.*

E. Outline of the main result

To prove the main result we present a reduction from CNF-SAT to the Fréchet distance. Given a CNF-SAT instance φ, we partition its variables into sets V_1, V_2 of equal size. In order to find a satisfying assignment of φ we have to choose (partial) assignments a_1 of V_1 and a_2 of V_2. We will construct curves P_1, P_2 where P_k is responsible for choosing a_k. To this end, P_k consists of *assignment gadgets*, one for each assignment of V_k. Assignment gadgets are built of *clause gadgets*, one for each clause. The assignment gadgets of assignments a_1 of V_1 and a_2 of V_2 are constructed such that they have Fréchet distance at most 1 if and only if (a_1, a_2) forms a satisfying assignment of φ. In P_1 and P_2 we connect these assignment gadgets with some additional curves to implement an OR-gadget, which forces any traversal of (P_1, P_2) to walk along two assignment gadgets in parallel. If φ is not satisfiable, then any pair of assignment gadgets has Fréchet distance larger than 1, so that P_1, P_2 have Fréchet distance larger than 1. If, on the other hand, a satisfying assignment (a_1, a_2) of φ exists, then we ensure that there is a traversal of P_1, P_2 that essentially only traverses the assignment gadgets of a_1 and a_2 in parallel, so that it always stays in distance 1.

To argue about the runtime, since P_k contains an assignment gadget for every assignment of one half of the variables, and every assignment gadget has polynomial size in M, there are $n = \mathcal{O}^*(2^{N/2})$ vertices on each curve. Thus, any $\mathcal{O}(n^{2-\delta})$ algorithm for the Fréchet distance would yield an $\mathcal{O}^*(2^{(1-\delta/2)N})$ algorithm for CNF-SAT, contradicting SETH′.

F. Remark: Orthogonal Vectors

Let Orthog be the problem of 'finding a pair of orthogonal vectors": given two sets $S_1, S_2 \subseteq \{0,1\}^d$ of n vectors each, determine if there are $u \in S_1$ and $v \in S_2$ with $\langle u, v\rangle = \sum_{i=1}^d u_i v_i = 0$, where the sum is computed over the integers, see [35, 36]. Clearly, Orthog can be solved in time $\mathcal{O}(n^2 d)$. However, Orthog has no strongly subquadratic algorithms unless SETH′ fails. More precisely, in [35] it was shown that SETH′ implies the following statement.

OrthogHypothesis**:** *There is no algorithm for* Orthog *with runtime* $\mathcal{O}(n^{2-\delta} d^{\mathcal{O}(1)})$ *for any* $\delta > 0$*.*

All known conditional lower bounds based on SETH′ implicitly go through Orthog or some variant of this problem. In fact, this is also the case for our results, as is easily seen by going through the proof in [35] and noting that we use the same tricks.[3]

Hence, in Theorems I.1, I.3, and I.4 we could replace the assumption "unless SETH′ fails" by the weaker assumption "unless OrthogHypothesis fails". This is a stronger statement, since there is only more reason to believe that Orthog has no strongly subquadratic algorithms than that there is for believing that CNF-SAT takes time $2^{N-o(N)}$. Moreover, it shows a relation between two polynomial time problems, Orthog and the Fréchet distance.

For Theorem I.2 we would need an imbalanced version of the OrthogHypothesis, where the two sets S_1, S_2 have different sizes n_1, n_2. Then unless SETH′ fails there is no $\mathcal{O}((n_1 n_2)^{1-\delta} d^{\mathcal{O}(1)})$ algorithm for any $\delta > 0$, and this holds for any polynomial restriction of $1 \leqslant n_1 \leqslant n_2$, which follows from a slight generalization of [35]. If we state this implication of SETH′ as a hypothesis OrthogHypothesis*,

[3]Specifically, given a CNF-SAT instance ϕ on variables $x_1, \ldots, x_N$ and clauses $C_1, \ldots, C_M$ we split the variables into two halves V_1, V_2 of equal size and enumerate all assignments A_k of true and false to V_k. Then every clause C_i specifies sets $B_k^i \subseteq A_k$ of partial assignments that do not make C_i become true. Clearly, a satisfying assignment $(a_1, a_2) \in A_1 \times A_2$ has to evade $B_1^i \times B_2^i$ for all i. This problem is equivalent to an instance of Orthog with $d = M$ and $n = 2^{N/2}$, where S_k contains a vector for every partial assignment $a_k \in A_k$ and the i-th position of this vector is 1 or 0, depending on whether $a_k \in B_k^i$ or not. In our proof, we could replace this instance by an arbitrary instance of Orthog, yielding a reduction from Orthog to the Fréchet distance.

then in Theorem I.2 we could replace "unless SETH' fails" by the weaker assumption "unless $\mathsf{OrthogHypothesis}^*$ fails".

G. Organization

We start by defining the variants of the Fréchet distance, c-packedness, and other basic notions in Section II. Section III deals with general curves. We prove the main result for the discrete Fréchet distance in less than 2 pages in Section III-A. This construction also already proves inapproximability. We generalize the proof to the continuous Fréchet distance in Section III-B (which is more tedious than in the discrete case) and to $m \ll n$ in Section III-C (which is an easy trick). Section IV deals with c-packed curves. We present a new OR-gadget that generates less packed curves; plugging in the curves constructed in the main result proves Theorem I.3. We omit the proof of Theorem I.4 in this extended abstract. We close with open problems in Section V.

II. Preliminaries

For $N \in \mathbb{N}$ we let $[N] := \{1, \ldots, N\}$. A (polygonal) curve P is defined by its vertices $p_1, \ldots, p_n$. We view P as a continuous function $P\colon [0, n] \to \mathbb{R}^d$ with $P(i + \lambda) = (1 - \lambda)p_i + \lambda p_{i+1}$ for $i \in [n - 1]$, $\lambda \in [0, 1]$. We write $|P| = n$ for the number of vertices of P. For two curves P_1, P_2 we let $P_1 \circ P_2$ be the curve on $|P_1| + |P_2|$ vertices that first follows P_1, then walks along the segment from $P_1(|P_1|)$ to $P_2(0)$, and then follows P_2. In particular, for two points $p, q \in \mathbb{R}^d$ the curve $p \circ q$ is the segment from p to q, and any curve P on vertices $p_1, \ldots, p_n$ can be written as $P = p_1 \circ \ldots \circ p_n$.

Consider a curve P and two points $p_1 = P(\lambda_1)$, $p_2 = P(\lambda_2)$ with $\lambda_1, \lambda_2 \in [0, n]$. We say that p_1 *is within distance* D *of* p_2 *along* P if the length of the subcurve of P between $P(\lambda_1)$ and $P(\lambda_2)$ is at most D.

Variants of the Fréchet distance: Let Φ_n be the set of all continuous and non-decreasing functions ϕ from $[0, 1]$ onto $[0, n]$. The *continuous Fréchet distance* between two curves P_1, P_2 with $|P_1| = n$, $|P_2| = m$ is defined as

$$d_{\mathrm{F}}(P_1, P_2) := \inf_{\substack{\phi_1 \in \Phi_n \\ \phi_2 \in \Phi_m}} \max_{t \in [0,1]} \|P_1(\phi_1(t)) - P_2(\phi_2(t))\|,$$

where $\|.\|$ denotes the Euclidean distance. We call (ϕ_1, ϕ_2) a (continuous) *traversal* of (P_1, P_2), and say that it has *width* D if $\max_{t \in [0,1]} \|P_1(\phi_1(t)) - P_2(\phi_2(t))\| \leqslant D$.

In the discrete case, we let Δ_n be the set of all non-decreasing functions ϕ from $[0, 1]$ onto $[n]$. The *discrete Fréchet distance* between two curves P_1, P_2 with $|P_1| = n$, $|P_2| = m$ is then defined as

$$d_{\mathrm{dF}}(P_1, P_2) := \inf_{\substack{\phi_1 \in \Delta_n \\ \phi_2 \in \Delta_m}} \max_{t \in [0,1]} \|P_1(\phi_1(t)) - P_2(\phi_2(t))\|.$$

We obtain an analogous notion of a (discrete) *traversal* and its *width*. Note that any $\phi \in \Delta_n$ is a staircase function attaining all values in $[n]$. Hence, $(\phi_1(t), \phi_2(t))$ changes only at finitely many points in time t. At any such *time step* we jump to the next vertex in P_1 or P_2 or both.

It is known that for any curves P_1, P_2 we have $d_{\mathrm{F}}(P_1, P_2) \leqslant d_{\mathrm{dF}}(P_1, P_2)$ [22].

Realistic input curves: As an example of input restrictions that resemble practical input curves we consider the model of [21]. A curve P is *c-packed* if for any point $q \in \mathbb{R}^d$ and any radius $r > 0$ the total length of P inside the ball $B(q, r)$ is at most cr. Here, $B(q, r)$ is the ball of radius r around q. In this paper, we say that a curve P is $\Theta(c)$-*packed*, if there are constants $\alpha > \beta > 0$ such that P is αc-packed but not βc-packed.

This model is well motivated from a practical point of view. Examples of classes of c-packed curves are boundaries of convex polygons and γ-fat shapes as well as algebraic curves of bounded maximal degree (see [21]).

Satisfiability: In $\mathsf{CNF\text{-}SAT}$ we are given a formula φ on variables $x_1, \ldots, x_N$ and clauses $C_1, \ldots, C_M$ in conjunctive normal form with unbounded clause width. Let V be any subset of the variables of φ. Let a be any assignment of T (true) or F (false) to the variables of V. We call a a *partial assignment* and say that a *satisfies* a clause $C = \bigvee_{i \in I} x_i \vee \bigvee_{i \in J} \neg x_i$ if for some $i \in I \cap V$ we have $a(x_i) = \mathsf{T}$ or for some $i \in J \cap V$ we have $a(x_i) = \mathsf{F}$. We denote by $\mathsf{sat}(a, C)$ whether partial assignment a satisfies clause C. Note that assignments a of V and a' of the remaining variables V' form a satisfying assignment (a, a') of φ if and only if we have $\mathsf{sat}(a, C_i) \vee \mathsf{sat}(a', C_i) = \mathsf{T}$ for all $i \in \{1, \ldots, M\}$.

All bounds that we prove in this paper assume the hypothesis SETH' (see Section I), which asserts that $\mathsf{CNF\text{-}SAT}$ has no $\mathcal{O}^*((2 - \delta)^N)$ algorithm for any $\delta > 0$. Here, $\mathcal{O}^*$ hides polynomials factors in N and M. The following is an easy corollary of SETH'.

Lemma II.1. *There is no $\mathcal{O}^*((2 - \delta)^N)$ algorithm for* $\mathsf{CNF\text{-}SAT}$ *restricted to formulas with N variables and $M \leqslant 2^{\delta' N}$ clauses for any $\delta, \delta' > 0$, unless* SETH' *fails.*

Proof: Any such algorithm would imply an $\mathcal{O}^*((2 - \delta)^N)$ algorithm for $\mathsf{CNF\text{-}SAT}$ (with no restrictions on the input), since for $M \leqslant 2^{\delta' N}$ we can run the given algorithm, while for $M > 2^{\delta' N}$ we can decide satisfiability in time $\mathcal{O}(M 2^N) = \mathcal{O}(M^{1+1/\delta'}) = \mathcal{O}^*(1)$. ■

III. General curves

We first present a reduction from $\mathsf{CNF\text{-}SAT}$ to the Fréchet distance and show that it proves Theorem I.1 for the discrete Fréchet distance. In Section III-B we then show that the same construction also works for the continuous Fréchet distance. Finally, in Section III-C we generalize these results to curves with imbalanced numbers of vertices n, m to show Theorem I.2.

A. The basic reduction, discrete case

Let φ be a given CNF-SAT instance with variables $x_1, \ldots, x_N$ and clauses $C_1, \ldots, C_M$. We split the variables into two halves $V_1 := \{x_1, \ldots, x_{N/2}\}$ and $V_2 := \{x_{N/2+1}, \ldots, x_N\}$. For $k \in \{1,2\}$ let A_k be all assignments[4] of T or F to the variables in V_k, so that $|A_k| = 2^{N/2}$. In the whole section we let $\varepsilon := 1/1000$.

We will construct two curves P_1, P_2 such that $d_{\mathrm{dF}}(P_1, P_2) \leqslant 1$ if and only if φ is satisfiable. In the construction we will use gadgets as follows.

Clause gadgets: This gadget encodes whether a partial assignment satisfies a clause. We set for $i \in \{0,1\}$

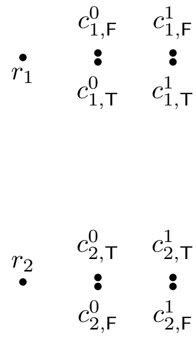

$$c^i_{1,\mathsf{T}} := \left(i/3, \tfrac{1}{2} - \varepsilon\right),$$
$$c^i_{1,\mathsf{F}} := \left(i/3, \tfrac{1}{2} + \varepsilon\right),$$
$$c^i_{2,\mathsf{T}} := \left(i/3, -\tfrac{1}{2} + \varepsilon\right),$$
$$c^i_{2,\mathsf{F}} := \left(i/3, -\tfrac{1}{2} - \varepsilon\right).$$

Let $k \in \{1,2\}$. For any partial assignment $a_k \in A_k$ and clause C_i, $i \in [M]$, we construct a clause gadget consisting of a single point,

$$CG(a_k, i) := c^{i \bmod 2}_{k, \mathsf{sat}(a_k, C_i)}.$$

Thus, if assignment a_k satisfies clause C_i then the corresponding clause gadget is nearer to the clause gadgets associated with A_{3-k}. Explicitly calculating all pairwise distances of these points, we obtain the following lemma.

Lemma III.1. *Let $a_k \in A_k$, $k \in \{1,2\}$, and $i, j \in [M]$. If $i \equiv j \pmod 2$ and $\mathsf{sat}(a_1, C_i) \vee \mathsf{sat}(a_2, C_j) = \mathsf{T}$ then $\|CG(a_1, i) - CG(a_2, j)\| \leqslant 1$. Otherwise $\|CG(a_1, i) - CG(a_2, j)\| \geqslant 1 + 2\varepsilon$.*

Assignment gadgets: This gadget consists of clause gadgets and encodes the set of satisfied clauses for an assignment. We set

$$r_1 := (-\tfrac{1}{3}, \tfrac{1}{2}), \quad r_2 := (-\tfrac{1}{3}, -\tfrac{1}{2}).$$

The assignment gadget for any $a_k \in A_k$ consists the starting point r_k followed by all clause gadgets of a_k,

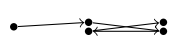

$$AG(a_k) := r_k \circ \bigcirc_{i \in [M]} CG(a_k, i),$$

(recall the definition of $\circ$ in Section II). The above figure shows an assignment gadget on $M = 2$ clauses at the top and an assignment gadget on $M = 4$ clauses at the bottom. The arrows indicate the order in which the segments are traversed.

[4]In later sections we will replace V_1, V_2 by different partitionings and A_1, A_2 by subsets of all assignments. The lemmas in this section are proven in a generality that allows this extension.

Lemma III.2. *Let $a_k \in A_k$, $k \in \{1,2\}$. If (a_1, a_2) is a satisfying assignment of φ then $d_{\mathrm{dF}}(AG(a_1), AG(a_2)) \leqslant 1$. If (a_1, a_2) is not satisfying then $d_{\mathrm{dF}}(AG(a_1), AG(a_2)) > 1 + \varepsilon$, and we even have $d_{\mathrm{dF}}(AG(a_1) \circ \pi_1, AG(a_2) \circ \pi_2) > 1 + \varepsilon$ for any curves π_1, π_2.*

Proof: If (a_1, a_2) is satisfying then the parallel traversal $(r_1, r_2), (CG(a_1, 1), CG(a_2, 1)), \ldots, (CG(a_1, M), CG(a_2, M))$ has width 1 by Lemma III.1.

Assume for the sake of contradiction that (a_1, a_2) is not satisfying but there is a traversal of $(AG(a_1) \circ \pi_1, AG(a_2) \circ \pi_2)$ with width $1 + \varepsilon$. Observe that $\|r_1 - r_2\| = 1$ and $\|r_k - c^i_{3-k,x}\| \geqslant 1 + 2\varepsilon$ for any $k \in \{1,2\}, i \in \{0,1\}, x \in \{\mathsf{T}, \mathsf{F}\}$. Thus, the traversal has to start at positions (r_1, r_2) and then step to positions $(CG(a_1, 1), CG(a_2, 1))$, as advancing in only one of the curves leaves us in distance larger than $1 + \varepsilon$. Inductively and using Lemma III.1, the same argument shows that in the i-th step we are at positions $(CG(a_1, i), CG(a_2, i))$ for any $i \in [M]$. Since there is an unsatisfied clause C_i, so that $\|CG(a_1, i) - CG(a_2, i)\| \geqslant 1 + 2\varepsilon$ by Lemma III.1, we obtain a contradiction. ∎

Construction of the curves: The curve P_k will consist of all assignment gadgets for assignments A_k, $k \in \{1,2\}$, plus some additional points. The additional points implement an OR-gadget over the assignment gadgets, by enforcing that any traversal of (P_1, P_2) with width $1 + \varepsilon$ has to traverse two assignment gadgets in parallel, and traversing one pair of assignment gadgets in parallel suffices. We define the following control points,

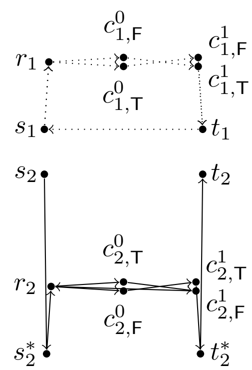

$$s_1 := (-\tfrac{1}{3}, \tfrac{1}{5}), \quad s_2 := (-\tfrac{1}{3}, 0), \quad s^*_2 := (-\tfrac{1}{3}, -\tfrac{4}{5}),$$
$$t_1 := (\tfrac{1}{3}, \tfrac{1}{5}), \quad t_2 := (\tfrac{1}{3}, 0), \quad t^*_2 := (\tfrac{1}{3}, -\tfrac{4}{5}).$$

Finally, we set

$$P_1 := \bigcirc_{a_1 \in A_1} \big(s_1 \circ AG(a_1) \circ t_1\big),$$
$$P_2 := s_2 \circ s^*_2 \circ \Big(\bigcirc_{a_2 \in A_2} AG(a_2)\Big) \circ t^*_2 \circ t_2.$$

The above figure shows P_1 (dotted) and P_2 (solid) in an example with $M = 2$ clauses and (unrealistically) only two assignments.

Let Q_k be the vertices that may appear in P_k, i.e., $Q_1 = \{s_1, t_1, r_1, c^0_{1,\mathsf{F}}, c^0_{1,\mathsf{T}}, c^1_{1,\mathsf{F}}, c^1_{1,\mathsf{T}}\}$ and $Q_2 = \{s_2, t_2, r_2, s^*_2, t^*_2, c^0_{2,\mathsf{F}}, c^0_{2,\mathsf{T}}, c^1_{2,\mathsf{F}}, c^1_{2,\mathsf{T}}\}$. Explicitly calculating all pairwise distances of all points, we obtain the following lemma.

Lemma III.3. *No pair $(q_1, q_2) \in Q_1 \times Q_2$ has $\|q_1 - q_2\| \in (1, 1+\varepsilon]$. Moreover, the set $\{(q_1, q_2) \in Q_1 \times Q_2 \mid \|q_1 - q_2\| \leqslant$*

$1\}$ *consists of the following pairs:*

$(q, s_2), (q, t_2)$ *for any* $q \in Q_1$,
(s_1, q) *for any* $q \in Q_2 \setminus \{t_2^*\}$,
(t_1, q) *for any* $q \in Q_2 \setminus \{s_2^*\}$,
(r_1, r_2),
$(c_{1,x}^i, c_{2,y}^i)$ *for* $x \vee y = \mathsf{T}$ *where* $i \in \{0, 1\}, x, y \in \{\mathsf{T}, \mathsf{F}\}$.

Correctness: We show that if φ is satisfiable then $d_{\mathrm{dF}}(P_1, P_2) \leqslant 1$, while otherwise $d_{\mathrm{dF}}(P_1, P_2) > 1 + \varepsilon$.

Lemma III.4. *If* $d_{\mathrm{dF}}(P_1, P_2) \leqslant 1 + \varepsilon$ *then* $A_1 \times A_2$ *contains a satisfying assignment.*

Proof: By Lemma III.3 any traversal with width $1 + \varepsilon$ also has width 1. Consider any traversal of (P_1, P_2) with width 1. Consider any time step T at which we are at position s_2^* in P_2. The only point in P_1 that is within distance 1 of s_2^* is s_1, say we are at the copy of s_1 that comes right before assignment gadget $AG(a_1)$, $a_1 \in A_1$. Following time step T, we have to start traversing $AG(a_1)$, so consider the first time step T' where we are at the point r_1 in $AG(a_1)$. The only points in P_2 within distance 1 of r_1 are s_2, t_2, and r_2. Note that we already passed s_2^* in P_2 by time T, so we cannot be in s_2 at time T'. Moreover, in between T and T' we are only at s_1 and r_1 in P_1, which have distance larger than 1 to t_2^*. Thus, we cannot pass t_2^*, and we cannot be at t_2 at time T'. Hence, we are at r_2, say at the copy of r_2 in assignment gadget $AG(a_2)$ for some $a_2 \in A_2$. The yet untraversed remainder of P_k is of the form $AG(a_k) \circ \pi_k$ for $k \in \{1, 2\}$. Since our traversal of (P_1, P_2) has width 1, we obtain $d_{\mathrm{dF}}(AG(a_1) \circ \pi_1, AG(a_2) \circ \pi_2) \leqslant 1$. By Lemma III.2, (a_1, a_2) forms a satisfying assignment of φ. ∎

Lemma III.5. *If* $A_1 \times A_2$ *contains a satisfying assignment then* $d_{\mathrm{dF}}(P_1, P_2) \leqslant 1$.

Proof: Let $(a_1, a_2) \in A_1 \times A_2$ be a satisfying assignment of φ. We describe a traversal through P_1, P_2 with width 1. We start at $s_2 \in P_2$ and the first point of P_1. We stay at s_2 and follow P_1 until we arrive at the copy of s_1 that comes right before $AG(a_1)$ (note that s_2 has distance 1 to any point in P_1). Then we stay at s_1 and follow P_2 until we arrive at the copy of r_2 in $AG(a_2)$ (note that the only point that is too far away from s_1 is t_2^*, but this point comes after all assignment gadgets in P_2). In the next step we go to positions (r_1, r_2) (in $AG(a_1), AG(a_2)$). Then we follow the clause gadgets $(CG(a_1, i), CG(a_2, i))$ in parallel, always staying within distance 1 by Lemma III.1. In the next step we stay at $CG(a_2, M)$ and go to t_1 in P_1 (which has distance 1 to any point in P_2 except for s_2^*, which we will never encounter again). We stay at t_1 in P_1 and follow P_2 completely until we arrive at its endpoint t_2. Since t_2 has distance 1 to any point in P_1, we can now stay at t_2 in P_2 and follow P_1 to its end. ∎

Proof of Theorem I.1, discrete case: Note that we have $n = \max\{|P_1|, |P_2|\} = \mathcal{O}(M) \cdot \max\{|A_1|, |A_2|\} = \mathcal{O}(M \cdot 2^{N/2})$. Moreover, the instance (P_1, P_2) can be constructed in time $\mathcal{O}(NM2^{N/2})$. Any $(1 + \varepsilon)$-approximation can decide whether $d_{\mathrm{dF}}(P_1, P_2) \leqslant 1$ or $d_{\mathrm{dF}}(P_1, P_2) > 1 + \varepsilon$, which by Lemmas III.4 and III.5 yields an algorithm that decides whether φ is satisfiable. If such an algorithm runs in time $\mathcal{O}(n^{2-\delta})$ for any small $\delta > 0$, then the resulting $\mathsf{CNF\text{-}SAT}$ algorithm runs in time $\mathcal{O}(M^2 2^{(1-\delta/2)N})$, contradicting SETH'.

B. Continuous case

The construction from the last section also works for the continuous Fréchet distance. However, for unsatisfiable formulas it becomes tedious to argue that continuous traversals are not much better than discrete traversals. For instance, we have to argue that we cannot stay at a fixed point between the clause gadgets $c_{1,\mathsf{T}}^0$ and $c_{1,\mathsf{T}}^1$ while traversing more than one clause gadget in P_2.

To adapt the proof from the last section, we have to reprove Lemmas III.4 and III.5. We will make use of the following property. Here, we set $\mathrm{sym}(CG(a_1, i)) := CG(a_2, i)$ and $\mathrm{sym}(r_1) := r_2$ and interpolate linearly between them to obtain a symmetric point in $AG(a_2)$ for every point in $AG(a_1)$ (for any fixed $a_1 \in A_1$, $a_2 \in A_2$). We also set $\mathrm{sym}(\mathrm{sym}(p_1)) := p_1$, to obtain a symmetric point in $AG(a_1)$ for every point in $AG(a_2)$.

Lemma III.6. *Consider any points* p_k *in* $AG(a_k)$, $k \in \{1, 2\}$, *with* $\|p_1 - p_2\| \leqslant 1 + \varepsilon$. *Then we have* $\|p_2 - \mathrm{sym}(p_1)\| \leqslant \frac{1}{9}$ *and* $\|\mathrm{sym}(p_2) - p_1\| \leqslant \frac{1}{9}$.

Proof: Let $p_k = (x_k, y_k)$ and note that we have $|y_1 - y_2| \geqslant 1 - 2\varepsilon$. Thus, if $|x_1 - x_2| > \frac{1}{9} - 2\varepsilon$ then we have (recall that $\varepsilon = 1/1000$)

$$\|p_1 - p_2\| > \sqrt{(\tfrac{1}{9} - 2\varepsilon)^2 + (1 - 2\varepsilon)^2} > 1 + \varepsilon,$$

a contradiction. Since $\mathrm{sym}(p_1) = (x_1, y_1')$ with $|y_1' - y_2| \leqslant 2\varepsilon$, we obtain

$$\|p_2 - \mathrm{sym}(p_1)\| \leqslant \sqrt{(\tfrac{1}{9} - 2\varepsilon)^2 + (2\varepsilon)^2} \leqslant \tfrac{1}{9}.$$

and the same bound holds for $\|\mathrm{sym}(p_2) - p_1\|$. ∎

Lemma III.7. (Analogue of Lemma III.4) *If* $d_{\mathrm{F}}(P_1, P_2) \leqslant 1 + \varepsilon = 1.001$ *then* $A_1 \times A_2$ *contains a satisfying assignment.*

Proof: In this proof, we say that two points $p_1 = (x_1, y_1)$, $p_2 = (x_2, y_2)$ have *y-distance* D if $|y_1 - y_2| \leqslant D$.

Consider any traversal of (P_1, P_2) with width $1 + \varepsilon$. Consider any time step T where we are at position s_2^* in P_2. The only points in P_1 that are within distance $1 + \varepsilon$ of s_2^* are within distance $1/20$ and y-distance ε of s_1 (since no point in P_1 has lower y-value than s_1 and $\sqrt{1 + (1/20)^2} > 1 + \varepsilon$). Say we are near the copy of s_1 that comes right before assignment gadget $AG(a_1)$, $a_1 \in A_1$. Following time step

T, we have to start traversing $AG(a_1)$, so consider the first time step T' where we are at the point r_1 in $AG(a_1)$. The only points in P_2 within distance $1+\varepsilon$ of r_1 are near s_2, t_2, or r_2. Note that we already passed s_2^* in P_2 by time T, so we cannot be near s_2 at time T'. Moreover, in between T and T' we are always near s_1 or between s_1 and r_1 in P_1, so we are always above and to the left of $s_1 + (1/20, 0)$, which has distance larger than $1+\varepsilon$ to t_2^*. Thus, we cannot pass t_2^*, and we cannot be near t_2 at time T'. Hence, we are near r_2, more precisely, we are in distance 1/20 and y-distance ε of r_2 (this is the same situation as for s_1 and s_2^*). After that, the traversal has to further traverse $AG(a_1)$ and/or $AG(a_2)$. Consider the first time step at which we are at $CG(a_1, 1)$ or $CG(a_2, 1)$, say we reach $CG(a_1, 1)$ first. By Lemma III.6, we are within distance 1/9 of $CG(a_2, 1)$. Since we were near r_2 at time T', we now passed r_2, and since we did not pass $CG(a_2, 1)$ yet, we are even within distance 1/9 of $CG(a_2, 1)$ *along the curve* P_2. This proves the induction base of the following inductive claim.

Claim III.8. *Let T_i be the first step in time at which the traversal is at $CG(a_1, i)$ or $CG(a_2, i)$, $i \in [M]$. At time T_i the traversal is within distance 1/9 of $CG(a_k, i)$ along the curve P_k for both $k \in \{1, 2\}$.*

Proof: Note that at all times T_i (and in between) Lemma III.6 is applicable, so we clearly are within distance 1/9 of $CG(a_k, i+1)$ at time T_{i+1} for any $i \in [M]$, $k \in \{1, 2\}$. Since $\|CG(a_k, i) - CG(a_k, i+1)\| \geqslant 1/3$, points within distance 1/9 of $CG(a_k, i)$ are not within distance 1/9 of $CG(a_k, i+1)$. Hence, if we are within distance 1/9 of $CG(a_k, i)$ along P_k for both $k \in \{1, 2\}$ at time T_i, then at time T_{i+1} we passed $CG(a_k, i)$ and did not pass $CG(a_k, i+1)$ yet (by definition of T_{i+1}), so that we are within distance 1/9 of $CG(a_k, i+1)$ along P_k for both $k \in \{1, 2\}$. ∎

Finally, we show that the above claim implies that (a_1, a_2) is a satisfying assignment. Assume for the sake of contradiction that some clause C_i is not satisfied by both a_1 and a_2. Say at time T_i we are at $CG(a_1, i)$ (if we are at $CG(a_2, i)$ instead, then a symmetric argument works). At the same time we are at some point p in $AG(a_2)$. By the above claim, p is within distance 1/9 of $CG(a_2, i)$ along P_2. Note that p lies on any of the line segments $c^0_{2,\mathsf{T}} \circ c^1_{2,\mathsf{F}}$, $c^0_{2,\mathsf{F}} \circ c^1_{2,\mathsf{T}}$, $c^0_{2,\mathsf{F}} \circ c^1_{2,\mathsf{F}}$, or $r_2 \circ c^0_{2,\mathsf{F}}$, since $\mathsf{sat}(a_2, C_i) = \mathsf{F}$. In any case, the current distance $\|p - CG(a_1, i)\|$ is at least the distance from the point $c^0_{1,\mathsf{F}}$ to the line through $c^0_{2,\mathsf{F}}$ and $c^1_{2,\mathsf{T}}$. We compute this distance as

$$\left(\tfrac{1}{3}(1+2\varepsilon)\right) / \sqrt{(\tfrac{1}{3})^2 + (2\varepsilon)^2} > 1 + \varepsilon,$$

which contradicts the traversal having width $1+\varepsilon$. ∎

Lemma III.9. (Analogue of Lemma III.5) *If $A_1 \times A_2$ contains a satisfying assignment then $d_{\mathrm{F}}(P_1, P_2) \leqslant 1$.*

Proof: Follows from Lemma III.5 and the general inequality $d_{\mathrm{F}}(P_1, P_2) \leqslant d_{\mathrm{dF}}(P_1, P_2)$. ∎

C. Generalization to imbalanced numbers of vertices

Assume that the input curves P_1, P_2 have different numbers of vertices $n = |P_1|$, $m = |P_2|$ with $n \geqslant m$. We show that there is no $\mathcal{O}((nm)^{1-\delta})$ algorithm for the Fréchet distance for any $\delta > 0$, even for any polynomial restriction of $1 \leqslant m \leqslant n$. More precisely, for any $\delta \leqslant \gamma \leqslant 1 - \delta$ we show that there is no $\mathcal{O}((nm)^{1-\delta})$ algorithm for the Fréchet distance restricted to instances with $n^{\gamma-\delta} \leqslant m \leqslant n^{\gamma+\delta}$.

To this end, given a $\mathsf{CNF\text{-}SAT}$ instance φ we partition its variables $x_1, \ldots, x_N$ into[5] $V'_1 := \{x_1, \ldots, x_\ell\}$ and $V'_2 := \{x_{\ell+1}, \ldots, x_N\}$ and let A'_k be all assignments of V'_k, $k \in \{1, 2\}$. Note that $|A'_1| = 2^{|V'_1|} = 2^\ell$ and $|A'_2| = 2^{N-\ell}$. Now we use the same construction as in Section III-A but replace V_k by V'_k and A_k by A'_k. Again we obtain that any 1.001-approximation for the Fréchet distance of the constructed curves P_1, P_2 decides satisfiability of φ. Observe that the constructed curves contain a number of points of

$$n = |P_1| = \Theta(M \cdot |A'_1|), \quad m = |P_2| = \Theta(M \cdot |A'_2|).$$

Hence, any 1.001-approximation of the Fréchet distance with runtime $\mathcal{O}((nm)^{1-\delta})$ for any small $\delta > 0$ yields an algorithm for $\mathsf{CNF\text{-}SAT}$ with runtime $\mathcal{O}(M^2(2^\ell 2^{N-\ell})^{1-\delta}) = \mathcal{O}(M^2 2^{(1-\delta)N})$, contradicting SETH'.

Finally, we set $\ell := N/(\gamma+1)$ (rounded in any way) so that $|A'_1| = \Theta(2^{N/(\gamma+1)})$ and $|A'_2| = \Theta(2^{N\gamma/(\gamma+1)})$. Using Lemma II.1 we can assume that $1 \leqslant M \leqslant 2^{\delta N/4}$. Hence, we have

$$\Omega(2^{N/(\gamma+1)}) \leqslant n \leqslant \mathcal{O}(2^{N/(\gamma+1)+\delta N/4}),$$
$$\Omega(2^{N\gamma/(\gamma+1)}) \leqslant m \leqslant \mathcal{O}(2^{N\gamma/(\gamma+1)+\delta N/4}),$$

which implies $\Omega(n^{\gamma-\delta/2}) \leqslant m \leqslant \mathcal{O}(n^{\gamma+\delta/2})$. For sufficiently large n, we obtain the desired polynomial restriction $n^{\gamma-\delta} \leqslant m \leqslant n^{\gamma+\delta}$. This proves Theorem I.2.

IV. Realistic inputs: c-packed curves

The curves constructed in Section III-A are highly packed, since all assignment gadgets lie roughly in the same area. Specifically they are not $o(n)$-packed. In this section we want to construct c-packed instances and show that there is no 1.001-approximation with runtime $\mathcal{O}((cn)^{1-\delta})$ for any $\delta > 0$ for the Fréchet distance unless SETH' fails, not even restricted to instances with $n^{\gamma-\delta} \leqslant c \leqslant n^{\gamma+\delta}$ for any $\delta \leqslant \gamma \leqslant 1 - \delta$. This proves Theorem I.3.

To this end, we again consider a $\mathsf{CNF\text{-}SAT}$ instance φ, partition its variables $x_1, \ldots, x_N$ into two sets V_1, V_2 of size $N/2$, and consider the set A_k of all assignments of T and F to the variables in V_k. Now we partition A_k into sets $A_k^1, \ldots, A_k^\ell$ of size $\Theta(2^{N/2}/\ell)$, where we fix $1 \leqslant \ell \leqslant 2^{N/2}$ later. Formula φ is satisfiable if and only

[5]For the impatient reader: we will set $\ell := N/(\gamma+1)$.

if for some pair $(j_1, j_2) \in [\ell]^2$ the set $A_1^{j_1} \times A_2^{j_2}$ contains a satisfying assignment. This suggests to use the construction of Section III-A after replacing A_1 by $A_1^{j_1}$ and A_2 by $A_2^{j_2}$, yielding a pair of curves $(P_1^{j_1 j_2}, P_2^{j_1 j_2})$. Now, φ is satisfiable if and only if $d_{\mathrm{F}}(P_1^{j_1 j_2}, P_2^{j_1 j_2}) \leqslant 1$ for some $(j_1, j_2) \in [\ell]^2$. For the sake of readability, we rename the constructed curves slightly so that we have curves (P_1^j, P_2^j) for $j \in [\ell^2]$.

OR-gadget: In the whole section we let $\rho := 1/\sqrt{2}$. We present an OR-construction over the gadgets (P_1^j, P_2^j) that is not too packed, in contrast to the OR-construction over assignment gadgets that we used in Section III-A. We start with two building blocks, where for any $j \in \mathbb{N}$ we set

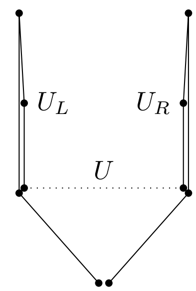

$$U_L(j) := (j\rho, 0) \circ ((j-1)\rho, \rho) \circ ((j-1)\rho, 3\rho) \circ ((j-1)\rho, 2\rho) \circ ((j-1)\rho, \rho),$$
$$U_R(j) := ((j+1)\rho, \rho) \circ ((j+1)\rho, 2\rho) \circ ((j+1)\rho, 3\rho) \circ ((j+1)\rho, \rho) \circ (j\rho, 0).$$

Moreover, we set $U(j) := U_L(j) \circ U_R(j)$. For a curve π and $z \in \mathbb{R}$ we let $\mathrm{tr}_z(\pi)$ be the curve π translated by z in x-direction. The OR-gadget now consists of the following two curves,

$$R_1 := \bigcirc_{j=1}^{\ell^2} \big(U_L(2j) \circ \mathrm{tr}_{2j\rho}(P_1^j) \circ U_R(2j)\big),$$
$$R_2 := U(1) \circ \bigcirc_{j=1}^{\ell^2} \big(\mathrm{tr}_{2j\rho}(P_2^j) \circ U(2j+1)\big).$$

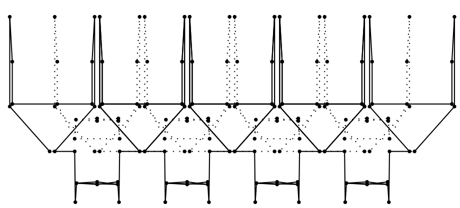

The above figure shows R_1 (dotted) and R_2 (solid) for $\ell^2 = 4$, see below for a figure showing $\ell^2 = 1$ with more details visible.

We denote by R_1^j the j-th "summand" of R_1, i.e., $R_1^j = U_L(2j) \circ \mathrm{tr}_{2j\rho}(P_1^j) \circ U_R(2j)$. Informally, we will use the term *U-shape* for the subcurves R_1^j and $U(2j+1)$, since they resemble the letter U. Moreover, we consider "summands" of R_2, namely $R_2^j := U(2j-1) \circ \mathrm{tr}_{2j\rho}(P_2^j) \circ ((2j+1)\rho, 0)$ and $\tilde{R}_2^j := ((2j-1)\rho, 0) \circ \mathrm{tr}_{2j\rho}(P_2^j) \circ U(2j+1)$.

Intuition: Considering traversals that stay within distance 1, we can traverse one U-shape in R_1 and one neighboring U-shape in R_2 together. Such traversals can be stitched together to a traversal of any number j of neighboring U-shapes in both curves. So far we can only traverse the same number of U-shapes in both curves, but R_2 has one more U-shape than R_1. We will show that we can traverse two U-shapes in R_2 while traversing only one U-shape in R_1, if these parts contain a satisfying assignment.

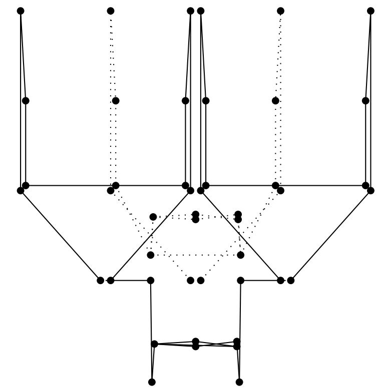

In the unsatisfiable case, essentially we show that we cannot traverse two U-shapes in R_2 while traversing only one U-shape in R_1, which implies a contradiction since the number of U-shapes in R_2 is larger than in R_1. We make this intuition formal in the remainder of this section.

Analysis: In order to be able to replace the curves P_1^j, P_2^j constructed above by other curves in the next section, we analyse the OR-gadget in a rather general way. To this end, we first specify a set of properties and show that the curves P_1^j, P_2^j constructed above satisfy these properties. Then we analyse the OR-gadget using only these properties of P_1^j, P_2^j.

Property IV.1. *(i) If φ is satisfiable then for some $j \in [\ell^2]$ we have $d_{\mathrm{dF}}(P_1^j, P_2^j) \leqslant 1$.*
(ii) If φ is not satisfiable then for all $j \in [\ell^2]$ and curves $\sigma_1, \sigma_2, \pi_1, \pi_2$ such that σ_1 stays to the left and above $(-\rho, \rho)$ and π_1 stays to the right and above (ρ, ρ), we have $d_{\mathrm{F}}(\sigma_1 \circ P_1^j \circ \pi_1, \sigma_2 \circ P_2^j \circ \pi_2) > \beta$, for some $\beta > 1$.
(iii) P_k^j is $\Theta(c)$-packed for some $c \geqslant 1$ for all $j \in [\ell^2]$, $k \in \{1, 2\}$.
(iv) $(0, \rho)$ is in distance 1 of any point in P_1^j for all $j \in [\ell^2]$.
(v) $(0, 0)$ is in distance 1 of any point in P_2^j for all $j \in [\ell^2]$.

Lemma IV.2. *The curves (P_1^j, P_2^j) constructed above satisfy Property IV.1 with $\beta = 1.001$ and $c = \Theta(M \cdot 2^{N/2}/\ell)$. Moreover, we have $|P_k^j| = \Theta(M \cdot 2^{N/2}/\ell)$ for all $j \in [\ell^2]$, $k \in \{1, 2\}$.*

Due to space limitations, the proof of this lemma is omitted and the proof of the next lemma is partially omitted.

In the following lemma we analyse the OR-gadget.

Lemma IV.3. *For any curves (P_1^j, P_2^j) that satisfy Property IV.1, the OR-gadget (R_1, R_2) satisfies:*
(i) $|R_k| = \Theta\big(\sum_{j=1}^{\ell^2} |P_k^j|\big)$ for $k \in \{1, 2\}$.
(ii) R_1 and R_2 are $\Theta(c)$-packed,
(iii) If φ is satisfiable then $d_{\mathrm{F}}(R_1, R_2) \leqslant d_{\mathrm{dF}}(R_1, R_2) \leqslant 1$,
(iv) If φ is not satisfiable then $d_{\mathrm{dF}}(R_1, R_2) \geqslant d_{\mathrm{F}}(R_1, R_2) > \min\{\beta, 1.2\}$.

Proof: (i) Precisely, we have $|R_k| = \sum_{j=1}^{\ell^2} (|P_k^j| + 10) + 10(k-1)$ for $k \in \{1, 2\}$.

(ii) Let $k \in \{1, 2\}$ and consider any ball $B = B(q, r)$. If $r \leqslant 1$ then B hits $\mathcal{O}(1)$ of the curves P_k^j. Since these curves are c-packed, their contribution to the total length of R_k in B is at most $\mathcal{O}(cr)$. Moreover, B hits $\mathcal{O}(1)$ segments of U or U_L, U_R, and the connecting segments to P_k^j. Each of these segments has length at most $2r$ inside B. This yields a total length of R_k in B of $\mathcal{O}((c+1)r)$.

Similarly, if $r > 1$ then B hits $\mathcal{O}(r)$ of the curves P_k^j. Note that the total length of P_k^j is at most c, since the curve is c-packed and contained in a ball of radius 1 around $(0,0)$ or $(0,\rho)$ by Property IV.1. Hence, the total length of of the curves P_k^j in B is $\mathcal{O}(cr)$. Moreover, B hits $\mathcal{O}(r)$ segments of U, U_L, U_R, and the connectors to P_k^j, each of constant length. This yields a total length of R_k in B of $\mathcal{O}((c+1)r)$.

In total, the curve R_k is $\mathcal{O}(c+1)$-packed. As $c \geqslant 1$, it is also $\mathcal{O}(c)$-packed. Since for some $\alpha > 0$ the curve P_k^j is not αc-packed, also R_k is not αc-packed, so R_k is even $\Theta(c)$-packed.

(iii) Note that $d_{\mathrm{F}}(R_1, R_2) \leqslant d_{\mathrm{dF}}(R_1, R_2)$ holds in general, so we only have to show that if φ is satisfiable then $d_{\mathrm{dF}}(R_1, R_2) \leqslant 1$. First we show that we can traverse one U-shape in R_1 and one neighboring U-shape in R_2 together.

Claim IV.4. *For any $j \in [\ell^2]$, we have $d_{\mathrm{dF}}(R_1^j, U(2j-1)) \leqslant 1$ and $d_{\mathrm{dF}}(R_1^j, U(2j+1)) \leqslant 1$.*

We can stitch these traversals together so that we traverse any number j of neighboring U-shapes in both curves together, because the parts in between the U-shapes are near to a single point, as shown by the following claim. Note that $(2j\rho, 0) \circ ((2j+2)\rho, 0)$ is the connecting segment in R_1 between $U_R(2j)$ and $U_L(2j+2)$, while $((2j-1)\rho, 0) \circ \mathrm{tr}_{2j\rho}(P_2^j) \circ ((2j+1)\rho, 0)$ is the part in R_2 between $U(2j-1)$ and $U(2j+1)$.

Claim IV.5. *For any $j \in [\ell^2]$,*

$$d_{\mathrm{dF}}((2j\rho, 0) \circ ((2j+2)\rho, 0), ((2j+1)\rho, 0)) \leqslant 1,$$
$$d_{\mathrm{dF}}((2j\rho, 0), ((2j-1)\rho, 0) \circ \mathrm{tr}_{2j\rho}(P_2^j) \circ ((2j+1)\rho, 0)) \leqslant 1.$$

Thus, we can stitch together traversals of U-shapes in both curves. However, so far we can only traverse the same number of U-shapes in both curves, but R_2 has one more U-shape than R_1. Consider $J \in [\ell^2]$ with $d_{\mathrm{dF}}(P_1^J, P_2^J) \leqslant 1$, which exists since φ is satisfiable, see Property IV.1.(i). Consider the two subcurves (also see the above figure)

$$R_1' := R_1^J = U_L(2J) \circ \mathrm{tr}_{2J\rho}(P_1^J) \circ U_R(2J),$$
$$R_2' := U(2J-1) \circ \mathrm{tr}_{2J\rho}(P_2^J) \circ U(2J+1).$$

We show that $d_{\mathrm{dF}}(R_1', R_2') \leqslant 1$, i.e., we can traverse two U-shapes in R_2 while traversing only one U-shape in R_1, using $d_{\mathrm{dF}}(P_1^J, P_2^J) \leqslant 1$. Adding simple traversals of U-shapes before and after (R_1', R_2'), we obtain a traversal of (R_1, R_2) with width 1, proving $d_{\mathrm{dF}}(R_1, R_2) \leqslant 1$. It is left to show the following claim.

Claim IV.6. $d_{\mathrm{dF}}(R_1', R_2') \leqslant 1$.

(iv) Note that the inequality $d_{\mathrm{dF}}(R_1, R_2) \geqslant d_{\mathrm{F}}(R_1, R_2)$ holds in general, so we only have to show that if φ is not satisfiable then $d_{\mathrm{F}}(R_1, R_2) > \min\{\beta, 1.2\}$. Assume for the sake of contradiction that there is a traversal of (R_1, R_2) with width $\min\{\beta, 1.2\}$. Essentially we show that it cannot traverse 2 U-shapes in R_2 while traversing only one U-shape in R_1, which implies a contradiction since the number of U-shapes in R_2 is larger than in R_1.

Let Y_ρ be the line $\{(x, y) \in \mathbb{R}^2 \mid y = \rho\}$. We inductively prove the following claims.

Claim IV.7. *(i) For any $0 \leqslant j \leqslant \ell^2$, when the traversal is in R_2 at the left highest point $(2j\rho, 3\rho)$ of $U(2j+1)$, then in R_1 we fully traversed R_1^j and are above the line Y_ρ.*

(ii) For any $1 \leqslant j \leqslant \ell^2$, when the traversal is in R_1 at the right highest point $((2j+1)\rho, 3\rho)$ of R_1^j, then in R_2 it is in $U(2j-1)$.

Note that claim (i) for $j = \ell^2$ yields the desired contradiction, since after traversing $R_1^{\ell^2}$ the curve R_1 has ended (at the point $(2\ell^2\rho, 0)$), so that we cannot go above the line Y_ρ anymore. ■

Proof of Theorem I.3: Finally, we use the OR-gadget (Lemma IV.3) together with the curves P_1^j, P_2^j we obtained from Section III-A (Lemma IV.2) to prove a runtime bound for c-packed curves: Any 1.001-approximation for the (discrete or continuous) Fréchet distance of (R_1, R_2) decides satisfiability of φ. Note that R_1 and R_2 are c-packed with $c = \Theta(M \cdot 2^{N/2}/\ell)$ and $n = \max\{|R_1|, |R_2|\} = \Theta(\ell^2 M \cdot 2^{N/2}/\ell)$. Thus, any $\mathcal{O}((cn)^{1-\delta})$ algorithm for the Fréchet distance implies a $\mathcal{O}(M^2 2^{(1-\delta)N})$ algorithm for $\mathsf{CNF\text{-}SAT}$, contradicting SETH'. Moreover, using Lemma II.1 we can assume that $1 \leqslant M \leqslant 2^{\delta N/4}$. Setting $\ell := \Theta(2^{\frac{1-\gamma}{1+\gamma}N/2})$ for any $0 \leqslant \gamma \leqslant 1$ we obtain

$$\Omega(2^{\frac{2}{1+\gamma}N/2}) \leqslant n \leqslant \mathcal{O}(2^{(\frac{2}{1+\gamma}+\delta/2)N/2}),$$
$$\Omega(2^{\frac{2\gamma}{1+\gamma}N/2}) \leqslant c \leqslant \mathcal{O}(2^{(\frac{2\gamma}{1+\gamma}+\delta/2)N/2}).$$

From this it follows that $\Omega(n^{\gamma-\delta/2}) \leqslant c \leqslant \mathcal{O}(n^{\gamma+\delta/2})$, which implies the desired polynomial restriction $n^{\gamma-\delta} \leqslant c \leqslant n^{\gamma+\delta}$ for sufficiently large n.

V. Open Problems

Our results leave three particularly interesting open questions for the Fréchet distance, asking for new algorithms or improved lower bounds. Here, $\tilde{\mathcal{O}}$ ignores any polylogarithmic factors in n, c, and $1/\varepsilon$. (1) Is there a strongly subquadratic $\mathcal{O}(1)$-approximation on general curves? (2) In any dimension $d \in \{2,3,4\}$, is there a $(1+\varepsilon)$-approximation with runtime $\tilde{\mathcal{O}}(cn)$ on c-packed curves? (3) In any dimension $d \geqslant 5$, is there a $(1+\varepsilon)$-approximation with runtime $\tilde{\mathcal{O}}(cn/\sqrt{\varepsilon})$ on c-packed curves?

Acknowledgments: The author wants to thank Wolfgang Mulzer, Marvin Künnemann, and Nabil H. Mustafa. K.B. is a recipient of the *Google Europe Fellowship in Randomized Algorithms*, and this research is supported in part by this Google Fellowship.

References

[1] A. Abboud and V. Vassilevska Williams. Popular conjectures imply strong lower bounds for dynamic problems. *Preprint arXiv:1402.0054*, 2014.

[2] P. Agarwal, R. B. Avraham, H. Kaplan, and M. Sharir. Computing the discrete Fréchet distance in subquadratic time. In *Proc. 24th Annu. ACM-SIAM Sympos. Discrete Algorithms (SODA'13)*, pages 156–167. SIAM, 2013.

[3] H. Alt. The computational geometry of comparing shapes. In *Efficient Algorithms*, volume 5760 of *LNCS*, pages 235–248. Springer, 2009.

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

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

[6] H. Alt, C. Knauer, and C. Wenk. Comparison of distance measures for planar curves. *Algorithmica*, 38(1):45–58, 2004.

[7] B. Aronov, S. Har-Peled, C. Knauer, Y. Wang, and C. Wenk. Fréchet distance for curves, revisited. In *Proc. 14th Annu. European Symp. Algorithms (ESA'06)*, pages 52–63. Springer, 2006.

[8] R. B. Avraham, O. Filtser, H. Kaplan, M. J. Katz, and M. Sharir. The discrete Fréchet distance with shortcuts via approximate distance counting and selection techniques. *Preprint arXiv:1310.5245*, 2013.

[9] S. Brakatsoulas, D. Pfoser, R. Salas, and C. Wenk. On map-matching vehicle tracking data. In *Proc. 31st International Conf. Very Large Data Bases (VLDB'05)*, pages 853–864, 2005.

[10] K. Buchin, M. Buchin, C. Knauer, G. Rote, and C. Wenk. How difficult is it to walk the dog? In *Proc. 23rd Euro. Workshop on Comput. Geom. (EuroCG'07)*, pages 170–173, 2007.

[11] K. Buchin, M. Buchin, and Y. Wang. Exact algorithms for partial curve matching via the Fréchet distance. In *Proc. 20th Annu. ACM-SIAM Symp. Discrete Algorithms (SODA'09)*, pages 645–654. SIAM, 2009.

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

[13] K. Buchin, M. Buchin, W. Meulemans, and W. Mulzer. Four soviets walk the dog - with an application to Alt's conjecture. In *Proc. 25th Annu. ACM-SIAM Sympos. Discrete Algorithms (SODA'14)*, pages 1399–1413. SIAM, 2014.

[14] C. Calabro, R. Impaglaizzo, and R. Paturi. A duality between clause width and clause density for SAT. In *Proc. 21st Annu. IEEE Conf. Computational Complexity (CCC'06)*, pages 252–260. IEEE, 2006.

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

[16] D. Chen, A. Driemel, L. J. Guibas, A. Nguyen, and C. Wenk. Approximate map matching with respect to the Fréchet distance. In *Proc. 13th Workshop on Algorithm Engineering and Experiments (ALENEX'11)*, pages 75–83. SIAM, 2011.

[17] A. F. Cook and C. Wenk. Geodesic Fréchet distance inside a simple polygon. *ACM Transactions on Algorithms*, 7(1): 193–204, 2010.

[18] M. Cygan, H. Dell, D. Lokshtanov, D. Marx, J. Nederlof, Y. Okamoto, R. Paturi, S. Saurabh, and M. Wahlström. On problems as hard as CNF-SAT. In *Proc. 2012 IEEE Conf. Computational Complexity (CCC'12)*, pages 74–84, 2012.

[19] E. Dantsin and E. A. Hirsch. Worst-case upper bounds. *Handbook of Satisfiability*, 185:403–424, 2009.

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

[21] A. Driemel, S. Har-Peled, and C. Wenk. Approximating the Fréchet distance for realistic curves in near linear time. *Discrete & Computational Geometry*, 48(1):94–127, 2012.

[22] T. Eiter and H. Mannila. Computing discrete Fréchet distance. Technical Report CD-TR 94/64, Christian Doppler Laboratory for Expert Systems, TU Vienna, Austria, 1994.

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

[24] M. Godau. A natural metric for curves - computing the distance for polygonal chains and approximation algorithms. In *Proc. 8th Sympos. Theoret. Aspects Comput. Sci. (STACS'91)*, volume 480 of *LNCS*, pages 127–136. Springer, 1991.

[25] A. Grønlund and S. Pettie. Threesomes, degenerates, and love triangles. In *Proc. 55th Annu. IEEE Sympos. Foundations of Computer Science (FOCS'14)*, 2014. To appear.

[26] S. Har-Peled and B. Raichel. The Fréchet distance revisited and extended. In *Proc. 27th Annu. Symp. Comp. Geometry (SoCG'11)*, pages 448–457, New York, NY, USA, 2011. ACM.

[27] R. Impagliazzo and R. Paturi. On the complexity of k-SAT. *Journal of Computer and System Sciences*, 62:367–375, 2001.

[28] R. Impagliazzo, R. Paturi, and F. Zane. Which problems have strongly exponential complexity? *J. Comput. Syst. Sci.*, 63(4): 512–530, 2001.

[29] P. Indyk. Approximate nearest neighbor algorithms for Fréchet distance via product metrics. In *Proc. 18th Annu. Symp. Comp. Geometry (SoCG'02)*, pages 102–106. ACM, 2002.

[30] D. Lokshtanov, D. Marx, and S. Saurabh. Lower bounds based on the exponential time hypothesis. *Bulletin of the EATCS*, 105:41–72, 2011.

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

[32] M. E. Munich and P. Perona. Continuous dynamic time warping for translation-invariant curve alignment with applications to signature verification. In *Proc. 7th Intl. Conf. Comp. Vision*, volume 1, pages 108–115. IEEE, 1999.

[33] M. Pătraşcu and R. Williams. On the possibility of faster SAT algorithms. In *Proc. 21nd Annu. ACM-SIAM Symp. Discrete Algorithms (SODA'10)*, pages 1065–1075. SIAM, 2010.

[34] L. Roditty and V. Vassilevska Williams. Fast approximation algorithms for the diameter and radius of sparse graphs. In *Proc. 45th Annu. ACM Symp. Theory of Computing (STOC'13)*, pages 515–524. ACM, 2013.

[35] R. Williams. A new algorithm for optimal constraint satisfaction and its implications. In *Proc. 31th Int. Colloq. Automata, Languages, and Programming (ICALP'04)*, volume 3142 of *LNCS*, pages 1227–1237. Springer, 2004.

[36] R. Williams and H. Yu. Finding orthogonal vectors in discrete structures. In *Proc. 25th Annu. ACM-SIAM Sympos. Discrete Algorithms (SODA'14)*, pages 1867–1877. SIAM, 2014.