

# Tighter Connections Between Formula-SAT and Shaving Logs\*

Amir Abboud<sup>†</sup>

Karl Bringmann<sup>‡</sup>

## Abstract

A noticeable fraction of Algorithms papers in the last few decades improve the running time of well-known algorithms for fundamental problems by logarithmic factors. For example, the  $O(n^2)$  dynamic programming solution to the Longest Common Subsequence problem (LCS) was improved to  $O(n^2/\log^2 n)$  in several ways and using a variety of ingenious tricks. This line of research, also known as *the art of shaving log factors*, lacks a tool for proving negative results. Specifically, how can we show that it is unlikely that LCS can be solved in time  $O(n^2/\log^3 n)$ ?

Perhaps the only approach for such results was suggested in a recent paper of Abboud, Hansen, Vassilevska W. and Williams (STOC'16). The authors blame the hardness of shaving logs on the hardness of solving satisfiability on boolean formulas (Formula-SAT) faster than exhaustive search. They show that an  $O(n^2/\log^{1000} n)$  algorithm for LCS would imply a major advance in circuit lower bounds. Whether this approach can lead to tighter barriers was unclear.

In this paper, we push this approach to its limit and, in particular, prove that a well-known barrier from complexity theory stands in the way for shaving five additional log factors for fundamental combinatorial problems. For LCS, regular expression pattern matching, as well as the Fréchet distance problem from Computational Geometry, we show that an  $O(n^2/\log^{7+\epsilon} n)$  runtime would imply new Formula-SAT algorithms.

Our main result is a reduction from SAT on formulas of size  $s$  over  $n$  variables to LCS on sequences of length  $N = 2^{n/2} \cdot s^{1+o(1)}$ . Our reduction is essentially as efficient as possible, and it greatly improves the previously known reduction for LCS with  $N = 2^{n/2} \cdot s^c$ , for some  $c \geq 100$ .

## 1 Introduction

Since the early days of Algorithms research, a noticeable fraction of papers each year *shave log factors* for fundamental problems: they reduce the best known upper bound on the time complexity from  $T(n)$  to  $T(n)/\log^c n$ , for some  $c > 0$ . While in some cases a cynic would call such results “hacks” and “bit tricks”, there is no doubt that they often involve ingenious algorithmic ideas and suggest fundamental new ways to look at the problem at hand. In his survey, Timothy Chan calls this kind of research “The Art of Shaving Logs” [37]. In many cases, we witness a race of shaving logs for some problem, in which a new upper bound is found every few months, without giving any hints on when this race is going to halt. For example, in the last few years, the upper bound for combinatorial Boolean Matrix Multiplication dropped from  $O(n^3/\log^2 n)$  [16], to  $O(n^3/\log^{2.25} n)$

---

\*Part of the work was performed while visiting the Simons Institute for the Theory of Computing, Berkeley, CA. The work was completed when A.A. was at Stanford University and was supported by Virginia Vassilevska Williams’ NSF Grants CCF-1417238 and CCF-1514339, and BSF Grant BSF:2012338.

<sup>†</sup>IBM Almaden Research Center, [abboud@cs.stanford.edu](mailto:abboud@cs.stanford.edu)

<sup>‡</sup>Max Planck Institute for Informatics, Saarland Informatics Campus, Germany, [kbringma@mpi-inf.mpg.de](mailto:kbringma@mpi-inf.mpg.de)

[20], to  $O(n^3/\log^3 n)$  [38], and most recently to  $O(n^3/\log^4 n)$  [102]. Perhaps the single most important missing technology for this kind of research is a tool for proving *lower bounds*.

Consider the problem of computing the Longest Common Subsequence (LCS) of two strings of length  $n$ . LCS has a simple  $O(n^2)$  time dynamic programming algorithm [95, 47]. Several approaches have been utilized in order to shave log factors such as the “Four Russians” technique [16, 63, 76, 23, 60], utilizing bit-parallelism [10, 48, 64], and working with compressed strings [49, 56]. The best known upper bounds are  $O(n^2/\log^2 n)$  for constant size alphabets [76], and  $O(n^2 \log \log n/\log^2 n)$  for large alphabets [60]. But can we do better? Can we solve LCS in  $O(n^2/\log^3 n)$  time? While the mathematical intrigue is obvious, we remark that even such mild speedups for LCS could be significant in practice. Besides its use as the *diff* operation in unix, LCS is at the core of highly impactful similarity measures between biological data. A heuristic algorithm called BLAST for a generalized version of LCS (namely, the Local Alignment problem [89]) has been cited more than sixty thousand times [14]. While such heuristics are much faster than the near-quadratic time algorithms above, they are not guaranteed to return an optimal solution and are thus useless in many applications, and biologists often fall back to (highly optimized implementations of) the quadratic solutions, see, e.g. [73, 74].

How would one show that it is hard to shave logs for some problem? A successful line of work, inspired by NP-hardness, utilizes “fine-grained reductions” to prove statements of the form: a small improvement over the known runtime for problem A implies a breakthrough algorithm for problem B, refuting a plausible hypothesis about the complexity of B. For example, it has been shown that if LCS can be solved in  $O(n^{2-\varepsilon})$  time, where  $\varepsilon > 0$ , then there is a breakthrough  $(2-\delta)^n$  algorithm for CNF-SAT, and the Strong Exponential Time Hypothesis (SETH, defined below) is refuted [2, 29]. Another conjecture that has been used to derive interesting lower bounds states that the 3-SUM problem<sup>1</sup> cannot be solved in  $O(n^{2-\varepsilon})$  time. It is natural to ask: can we use these conjectures to rule out log-factor improvements for problems like LCS? And even more optimistically, one might hope to base the hardness of LCS on a more standard assumption like  $P \neq NP$ . Unfortunately, we can formally prove that these assumptions are not sharp enough to lead to any consequences for log-factor improvements, if only Turing reductions are used. In Section 3 we prove the following theorem which also shows that an  $O(f(n)/\log^c(f(n)))$  time algorithm for problem A cannot imply, via a fine-grained reduction, an  $O(g(n)^{1-\varepsilon})$  algorithm for problem B, unless B is (unconditionally) solvable in  $O(g(n)^{1-\delta})$  time.

**Theorem 1.1** (Informally). *If for some  $c > 0$  there is a fine-grained reduction proving that LCS is not in  $O(n^2/\log^c n)$  time unless SETH fails, then SETH is false.*

Note that it also does not suffice to simply make SETH stronger by postulating a higher running time lower bound for CNF-SAT, since *superpolynomial improvements are known* for this problem [83, 34, 50, 8]. Similarly, we cannot base a study of log-factor improvements on the APSP conjecture, since *superlogarithmic improvements are known* for APSP [99]. (However, 3SUM could be a candidate to base higher lower bounds on, since only log-factor improvements are known [61, 55, 58, 21], see Section A for a discussion.)

Thus, in a time when super-linear lower bounds for problems like LCS are far out of reach, and our only viable approach to obtaining such negative results is reductions-based, we are left with two options. We could either leave the study of log-factor improvements in limbo, without a

---

<sup>1</sup>3-SUM asks, given a list of  $n$  numbers, to find three that sum to zero. The best known upper bound is  $O((n^2/\log^2 n)(\log \log n)^{O(1)})$  for real numbers [61, 55, 58, 39] and  $O(n^2(\log \log n/\log n)^2)$  for integers [21].

technology for proving negative results, or we could search for natural and convincing assumptions that are more fine-grained than SETH that could serve as the basis for the negative results we desire. Such assumptions were recently proposed by Abboud, Hansen, Vassilevska Williams and Williams [3]. The authors blame the hardness of shaving logs on the hardness of solving satisfiability on boolean formulas (Formula-SAT) faster than exhaustive search<sup>2</sup>, by polynomial factors (which are log-factors in the runtime), a task for which there are well known “circuit lower bound” barriers in complexity theory. They show that an  $O(n^2/\log^{1000} n)$  algorithm for LCS would imply a major advance in circuit lower bounds. In the final section of this paper, we give a more detailed argument in favor of this approach. Whether one should expect it to lead to *tight* barriers, i.e. explaining the lack of  $O(n^2/\log^3 n)$  algorithms for LCS or any other natural problem, was completely unclear.

**The Machine Model** We use the *Word-RAM* model on words of size  $\Theta(\log n)$ , where there is a set of operations on words that can be performed in time  $O(1)$ . Most papers do not fix the concrete set of allowed operations, and instead refer to “typical Boolean and arithmetic operations”. In this paper, we choose a set of operations  $\mathcal{P}$  that is *robust* with respect to changing the word size: For any operation  $\circ \in \mathcal{P}$ , given two words  $a, b$  (of size  $\Theta(\log n)$ ) we can compute  $a \circ b$  in time  $(\log n)^{1+o(1)}$  on a Word RAM with word size  $\Theta(\log \log n)$  and operation set  $\mathcal{P}$ . In other words, if we split  $a, b$  into  $\Theta(\log n/\log \log n)$  words of size  $\Theta(\log \log n)$  then  $a \circ b$  can still be computed very efficiently.

This robustness in particular holds for the following standard set of operations: initializing a cell with a constant, bitwise AND, OR, NOT, shift, addition, subtraction, multiplication, and division with remainder (since multiplication and division have near-linear time algorithms).

The results in this paper will get gradually weaker as we relax the restriction on near-linear time per operation to higher runtimes, however, even with this restriction, to the best of our knowledge this model captures all log shaving results in the literature (on the “standard” Word RAM model without fancy word operations).

**Formula-SAT** A boolean formula over  $n$  input variables can be viewed as a tree in which every leaf is marked by an input variable or its negation and every internal node or *gate* represents some basic boolean operation. Throughout this introduction we will only talk about *deMorgan* formulas, in which every gate is from the set  $\{\wedge, \vee\}$ . The *size* of the formula is defined to be the number of leaves in the tree.

In the Formula-SAT problem we are given a formula  $F$  of size  $s$  over  $n$  inputs, and we have to decide whether there is an input  $\{0, 1\}^n$  that makes it output 1. A naive algorithm takes  $O(2^n \cdot s)$  time, since evaluating the formula on some input takes  $O(s)$  time. Can we do better? We will call a SAT algorithm *non-trivial*<sup>3</sup> if it has a runtime at most  $O(\frac{2^n}{n^\varepsilon})$ , for some  $\varepsilon > 0$ .

It seems like a clever algorithm must look at the given formula  $F$  and try to gain a speedup by analyzing it. The more complicated  $F$  can be, the harder the problem becomes. Indeed, Dantsin and Hirsch [50] survey dozens of algorithms for SAT on CNF formulas which exploit their structure. For  $k$ -CNF formulas of size  $s$  there are  $2^n s/2^{\Omega(n/k)}$  time algorithms (e.g. [83]), and for general CNF

---

<sup>2</sup>In [3] the authors focus on SAT on Branching Programs (BPs) rather than formulas, but due to standard transformations between BPs and formulas, the two problems are equivalent up to polynomial factors. Focusing on Formula-SAT will be crucial to the progress we make in this paper.

<sup>3</sup>Some works on SAT algorithms used this term for runtimes of the form  $2^n \text{poly}(s)/n^{\omega(1)}$ . In our context, we need to be a bit more fine-grained.

formulas the bound is  $2^n s / 2^{\Omega(n/\log \Delta)}$  where  $\Delta = s/n$  is the clause-to-variable ratio [34, 50, 8]. The popular SETH [68, 35] essentially says that this is close to optimal, and that there is no  $2^n s / 2^{\Omega(n)}$  algorithm for CNF-SAT. For arbitrary deMorgan formulas, the upper bounds are much worse. A FOCS'10 paper by Santhanam [86] and several recent improvements [42, 44, 43, 72, 93] solve Formula-SAT on formulas of size  $s = n^{3-16\epsilon}$  in time  $2^n s^{O(1)} / 2^{n^\epsilon}$ , which is non-trivial only for  $s = o(n^3)$ , and going beyond cubic seems extremely difficult. This leads us to the first barrier which we will transform into a barrier for shaving logs.

**Hypothesis 1.2.** *There is no algorithm that can solve SAT on deMorgan formulas of size  $s = n^{3+\Omega(1)}$  in  $O(\frac{2^n}{n^\epsilon})$  time, for some  $\epsilon > 0$ , in the Word-RAM model.*

Perhaps the main reason to believe this hypothesis is that despite extensive algorithmic attacks on variants of SAT (perhaps the most extensively studied problem in computer science) over decades, none of the ideas that anyone has ever come up with seem sufficient to refute it. Recent years have been particularly productive in non-trivial algorithms designed for special cases of Circuit-SAT [86, 88, 66, 35, 101, 22, 41, 69, 65, 45, 85, 59] (in addition to the algorithms for deMorgan formulas above) and this hypothesis still stands.

A well-known “circuit lower bounds” barrier seems to be in the way for refuting Hypothesis 1.2: can we find an explicit boolean function that cannot be computed by deMorgan formulas of cubic size? Functions that require formulas of size  $\Omega(n^{1.5})$  [91] and  $\Omega(n^2)$  [71] have been known since the 60’s and 70’s, respectively. In the late 80’s, Andreev [15] proved an  $\Omega(n^{2.5})$  which was later gradually improved to  $\Omega(n^{2.55})$  by Nisan and Wigderson [67] and to  $\Omega(n^{2.63})$  by Paterson and Zwick [81] until Håstad proved his  $n^{3-o(1)}$  lower bound in FOCS’93 [62] (a recent result by Tal improves the  $n^{o(1)}$  term [92]). All these lower bound results use the “random restrictions” technique, first introduced in this context by Subbotovskaya in 1961 [91], and it is known that a substantially different approach must be taken in order to go beyond the cubic barrier. What does this have to do with Formula-SAT algorithms? Interestingly, this same “random restrictions” technique was crucial to *all* the non-trivial Formula-SAT algorithms mentioned above. This is not a coincidence, but only one out of the many examples of the intimate connection between the task of designing non-trivial algorithms for SAT on a certain class  $\mathcal{F}$  of formulas or circuits and the task of proving lower bounds against  $\mathcal{F}$ . This connection is highlighted in many recent works and in several surveys [87, 80, 98]. The intuition is that both of these tasks seem to require identifying a strong structural property of functions in  $\mathcal{F}$ . There is even a formal connection shown by Williams [97], which in our context implies that solving Formula-SAT on formulas of size  $O(n^{3.1})$  in  $O(2^n/n^{10})$  time (which is only slightly stronger than refuting Hypothesis 1.2) is sufficient in order to prove that there is a function in the class  $\mathbf{E}^{\text{NP}}$  that cannot be computed by formulas of size  $O(n^{3.1})$  (see [3] for more details). This consequence would be the first polynomial progress on the fundamental question of worst case formula lower bounds since Håstad’s result.

## 1.1 Our Results: New Reductions

Many recent papers have reduced CNF-SAT to fundamental problems in  $\mathbf{P}$  to prove SETH-based lower bounds (e.g. [82, 84, 6, 4, 27, 18, 7, 1, 33, 5, 19, 77, 40]). Abboud et al. [3] show that even SAT on formulas, circuits, and more, can be efficiently reduced to combinatorial problems in  $\mathbf{P}$ . In particular, they show that Formula-SAT on formulas of size  $s$  over  $n$  inputs can be reduced to an instance of LCS on sequences of length  $N = O(2^{n/2} \cdot s^{1000})$ . This acts as a barrier for shaving logs

as follows. A hypothetical  $O(N^2/\log^c N)$  time algorithm for LCS can be turned into an

$$n^{1+o(1)} \cdot (2^{n/2} \cdot s^{1000})^2 / (\log 2^{\Omega(n)})^c = O(2^n \cdot s^{2000} / n^{c-1})$$

time algorithm for Formula-SAT, which for a large enough  $c \geq 2001$  would refute Hypothesis 1.2. The first  $n^{1+o(1)}$  factor in the runtime comes from the jump from  $n$  to  $N = 2^n$  and our Word-RAM machine model: whenever the LCS algorithm wants to perform a unit-cost operation on words of size  $\Theta(\log N)$  (this is much more than the word size of our SAT algorithm which is only  $\Theta(\log n) = \Theta(\log \log N)$ ), the SAT algorithm can simulate it in  $(\log N)^{1+o(1)} = n^{1+o(1)}$  time in the Word-RAM model with words of size  $\Theta(\log n)$ .

Our main result is a *much* more efficient reduction to LCS. For large but constant size alphabets, we get a near-linear dependence on the formula size, reducing the  $s^{1000}$  factor to just  $s^{1+o(1)}$ .

**Theorem 1.3.** *Formula-SAT on formulas of size  $s$  on  $n$  inputs can be reduced to an instance of LCS on two sequences over an alphabet of size  $\sigma$  of length  $N = 2^{n/2} \cdot s^{1+O(1/\log \log \sigma)}$ , in  $O(N)$  time.*

Thus, if LCS on sequences of length  $N$  and alphabet of size  $\omega(1)$  can be solved in  $O(N^2/\log^c N)$  time, then Formula-SAT can be solved in  $2^n \cdot \frac{s^{2+o(1)}}{n^c} \cdot n^{1+o(1)}$  time. Recall that the known upper bound for LCS is  $O(n^2/\log^c n)$  for any constant alphabet size, with  $c = 2$ , and we can now report that the barrier of cubic formulas stands in the way of improving it to  $c > 7$  (see Corollary 1.6 below).

The novelty in the proof of Theorem 1.3 over [3] is discussed in Section 2. As an alternative to Theorem 1.3, in Section D we present another reduction to LCS which is much simpler than all previously known reductions, but uses a larger alphabet.

**Fréchet Distance** An important primitive in computational geometry is to judge how similar are two basic geometric objects, such as polygonal curves, represented as sequences of points in  $d$ -dimensional Euclidean space. Such curves are ubiquitous, since they arise naturally as trajectory data of moving objects, or as time-series data of stock prices and other measures. The most popular similarity measure for curves in computational geometry is the Fréchet distance, also known as dog-leash-distance. For formal definitions see Section F. The Fréchet distance has found many applications (see, e.g., [78, 26, 30]) and developed to a rich field of research with many generalizations and variants (see, e.g., [11, 17, 13, 53, 36, 46, 32, 52, 75, 70]).

This distance measure comes in two variants: the continuous and the discrete. A classic algorithm by Alt and Godau [12, 57] computes the continuous Fréchet distance in time  $O(n^2 \log n)$  for two given curves with  $n$  vertices. The fastest known algorithm runs in time  $O(n^2(\log \log n)^2)$  (on the Word RAM) [31]. If we only want to decide whether the Fréchet distance is at most a given value  $\delta$ , this algorithm runs in time  $O(n^2(\log \log n)^2/\log n)$ . For the discrete Fréchet distance, the original algorithm has running time  $O(n^2)$  [54], which was improved to  $O(n^2 \log \log n/\log n)$  by Agarwal et al. [9]. Their algorithm runs in time  $O(n^2 \log \log n/\log^2 n)$  for the decision version. It is known that both versions of the Fréchet distance are SETH-hard [27]. However, this does not rule out log factor improvements. In particular, no reduction from versions of SETH on formulas or branching programs is known.

In this paper we focus on the *decision version of the discrete Fréchet distance* (which we simply call “Fréchet distance” from now on). We show that Fréchet distance suffers from the same barriers for shaving logs like LCS. In particular, this reduction allows us to base the usual  $\Omega(n^{2-\epsilon})$  lower bound on a weaker assumption than SETH, such as NC-SETH (see the discussion in [3]). This is the first NC-SETH hardness for a problem that does not admit alignment gadgets (as in [29]).

**Theorem 1.4.** *Formula-SAT on formulas of size  $s$  on  $n$  inputs can be reduced to an instance of the Fréchet distance on two curves of length  $N = O(2^{n/2} \cdot s)$ , in  $O(N)$  time.*

**Regular Expression Pattern Matching** Our final example is the fundamental *Regular Expression Pattern Matching* problem: Decide whether a given regular expression of length  $m$  matches a substring of a text of length  $n$ . Again, there is a classical  $O(nm)$  algorithm [94], and the applicability and interest in this problem resulted in algorithms shaving log factors; the first one by Myers [79] was improved by Bille and Thorup [24] to time  $O(mn/\log^{1.5} n)$ . Recently, Backurs and Indyk proved an  $n^{2-o(1)}$  SETH lower bound [19], and performed an impressive study of the exact time complexity of the problem with respect to the complexity of the regular expression. This study was essentially completed by Bringmann, Grønlund, and Larsen [28], up to  $n^{o(1)}$  factors. In Section E we show that this problem is also capable of efficiently simulating formulas and thus has the same barriers as LCS and Fréchet distance.

**Theorem 1.5.** *Formula-SAT on formulas of size  $s$  on  $n$  inputs can be reduced to an instance of Regular Expression Pattern Matching on text and pattern of length  $N = O(2^{n/2} \cdot s \log s)$  over a constant size alphabet, in  $O(N)$  time.*

**Consequences of the Cubic Formula Barrier** We believe that SAT on formulas can be tightly connected to many other natural problems in P. As we discuss in the next section, such reductions seem to require problem-specific engineering and are left for future work. The main point of this paper is to demonstrate the possibility of basing such ultra fine-grained lower bounds on one common barrier. Our conditional lower bounds are summarized in the following corollary, which shows that current log-shaving algorithms are very close to the well-known barrier from complexity theory of cubic formula lower bounds.

**Corollary 1.6.** *For all  $\varepsilon > 0$ , solving any of the following problems in  $O(n^2/\log^{7+\varepsilon} n)$  time refutes Hypothesis 1.2, and solving them in  $O(n^2/\log^{17+\varepsilon} n)$  time implies that  $\mathbf{E}^{\text{NP}}$  cannot be computed by non-uniform formulas of cubic size:*

- *LCS over alphabets of size  $\omega(1)$*
- *The Fréchet distance on two curves in the plane*
- *Regular Expression Pattern Matching over constant size alphabets.*

The main reason that our lower bounds above are not tight (the gap between 2 and 7) is that we need to start from SAT on *cubic* size formulas rather than linear size ones, due to the fact that clever algorithms do exist for smaller formulas. We remark that throughout the paper we will work with a class of formulas we call  $\mathcal{F}_1$  (see Section B), also known as bipartite formulas, that are more powerful than deMorgan formulas yet our reduction to LCS can support them as well. This makes our results stronger, since  $\mathcal{F}_1$ -Formula-SAT could be a harder problem than SAT on deMorgan formulas. In fact, in an earlier version of the paper we had suggested the hypothesis that  $\mathcal{F}_1$ -Formula-SAT does not have non-trivial algorithms even on *linear size* formulas. This stronger hypothesis would give higher lower bounds. However, Avishay Tal (personal communication) told us about such a non-trivial algorithm for formulas of size up to  $n^{2-\Omega(1)}$  using tools from quantum query complexity. We are optimistic that one could borrow such ideas or the “random restrictions” technique from SAT algorithms in order to shave more logs for combinatorial problems such as LCS. This is an intriguing direction for future work.

## 2 Technical Overview and the Reduction to LCS

All the reductions from SAT to problems in P mentioned above start with a split-and-list reduction to some “pair finding” problem. In the SETH lower bounds, CNF-SAT is reduced to the Orthogonal-Vectors problem of finding a pair  $a \in A, b \in B, A, B \subseteq \{0, 1\}^d$  that are orthogonal [96]. When starting from Formula-SAT, we get a more complex pair-finding problem. In Section B we show a simple reduction from SAT on formulas from the class  $\mathcal{F}_1$  (which contains deMorgan formulas) to the following problem.

**Definition 2.1** (FORMULA-PAIR Problem). *Given a deMorgan formula over  $2m$  variables  $F = F(x_1, \dots, x_m, y_1, \dots, y_m)$  (each appearing once in  $F$ ), and two sets of vectors  $A, B \subseteq \{0, 1\}^m$  of size  $n$ , decide if there is a pair  $a \in A, b \in B$  such that  $F(a_1, \dots, a_m, b_1, \dots, b_m) = \text{true}$ .*

In Section B we show a Four-Russians type algorithm that solves FORMULA-PAIR in  $O(n^2m/\log^2 n)$  time, and even when  $m = |F| = (\log n)^{1+o(1)}$  no  $O(n^2/\log^{1+\varepsilon} n)$  upper bound is known. By our reduction, such an upper bound would imply a non-trivial algorithm for SAT on formulas from  $\mathcal{F}_1$ . Moreover, Hypothesis 1.2 implies that we cannot solve FORMULA-PAIR in  $O(n^2/\log^\varepsilon n)$  time, for  $m = (\log n)^{3+\Omega(1)}$ . In the next sections, we reduce FORMULA-PAIR to LCS, from which Theorem 1.3 follows. A simpler reduction using much larger alphabet size can be found in Section D.

**Theorem 2.2.** *FORMULA-PAIR on formulas of size  $s$  and lists of size  $n$  can be reduced to an instance of LCS on two strings over alphabet of size  $\sigma \geq 2$  of length  $O(n \cdot s^{1+O(1/\log \log \sigma)})$ , in linear time.*

The reduction constructs strings  $x, y$  and a number  $\rho$  such that  $\text{LCS}(x, y) \geq \rho$  holds if and only if the given Formula-Pair instance  $(F, A, B)$  is satisfiable. The approach is similar to the reductions from Orthogonal-Vectors to sequence alignment problems (e.g. [6, 27, 18, 2, 29]). The big difference is that our formula  $F$  can be much more complicated than a CNF, and so we will need more powerful gadgets. Sequence gadgets that are able to simulate the evaluation of deMorgan formulas were (implicitly) constructed in [3] with a recursive approach. Our main contribution is an *extremely efficient* implementation of such gadgets with LCS.

The main part of the reduction is to construct *gate gadgets*: for any vectors  $a, b \in \{0, 1\}^m$  and any gate  $g$  of  $F$ , we construct strings  $x(g, a)$  and  $y(g, b)$  whose LCS determines whether gate  $g$  evaluates to true for input  $(a, b)$  to  $F$  (see Section 2.1). Once we have this, to find a pair of vectors  $a \in A, b \in B$  satisfying  $F$ , we combine the strings  $x(r, a), y(r, b)$ , constructed for the root  $r$  of  $F$ , using a known construction of so-called alignment gadgets [2, 29] from previous work (see Section C.1).

Let us quickly explain how [3] constructed gate gadgets and the main ideas that go into our new construction. There are two kinds of gadgets, corresponding to the two types of gates in  $F$ : AND and OR gates. Since the AND gadgets will be relatively simple, let us consider the OR gadgets. Fix two inputs  $a, b$ , and let  $g = (g_1 \vee g_2)$  be an OR gate, and assume that we already constructed gate gadgets for  $g_1, g_2$ , namely  $x_1 = x(g_1, a), y_1 = y(g_1, b), x_2 = x(g_2, a), y_2 = y(g_2, b)$  so that for  $i \in \{1, 2\}$  we have that  $\text{LCS}(x_i, y_i)$  is large if the gate  $g_i$  outputs true on input  $(a, b)$ , and it is smaller otherwise. In [3], these gadgets were combined as follows. Let  $\beta$  be an upper bound on the total length of the gadgets  $x_i, y_i$ . We add a carefully chosen padding of 0’s and 1’s, so that any optimal matching of the two strings will have to match *either*  $x_1, y_1$  or  $x_2, y_2$  but not both.

$$\begin{aligned}
x &:= 0^{4\beta} x_1 1^\beta x_2 0^{4\beta} \\
y &:= y_1 1^\beta 0^{4\beta} 1^\beta y_2
\end{aligned}$$

One then argues that, in any optimal LCS matching of  $x, y$ , the  $0^{4\beta}$  block of  $y$  must be matched either left or right. If it's matched left, then the total score will be equal to  $4\beta + \beta + LCS(x_2, y_2)$  while if it's matched right, we will get  $4\beta + \beta + LCS(x_1, y_1)$ . Thus,  $LCS(x, y)$  is determined by the OR of  $g_1, g_2$ . The blowup of this construction is a multiplicative factor of 11 with every level of the formula, and the length of the gadget of the root will end up roughly  $11^{\text{depth}(F)}$ . To obtain our tight lower bounds, we will need to decrease this blowup to  $1 + \varepsilon_\sigma$  at every level, where  $\varepsilon_\sigma$  goes to 0 when the alphabet size  $\sigma$  tends to infinity. With the above construction, decreasing the length of the padding will allow the optimal LCS matching to cheat, e.g. by matching  $y_1$  to both  $x_1$  and  $x_2$ , and no longer corresponding to the OR of  $g_1, g_2$ .

Our first trick is an ultra-efficient OR gadget in case we are allowed unbounded alphabet size. We take  $x_1, y_1$  and transform all their letters into a new alphabet  $\Sigma^{g_1}$ , and we take  $x_2, y_2$  and transform their letters into a disjoint alphabet  $\Sigma^{g_2}$ . Then our OR gadget does not require any padding at all:

$$\begin{aligned}
x &:= x_1 x_2 \\
y &:= y_2 y_1
\end{aligned}$$

The crossing structure of this construction means that any LCS matching that matches letters from  $x_1, y_1$  cannot also match letters from  $x_2, y_2$ , and vice versa, while the disjoint alphabets make sure that there can be no matches between  $x_1, y_2$  or  $x_2, y_1$ . With such gadgets we can encode a formula of size  $s$  with  $O(s)$  letters, for details see Section D.

But how would such an idea work for *constant size alphabets*? Once we allow  $x_1$  and  $y_2$  to share even a single letter, this argument breaks. Natural attempts to simulate this construction with smaller alphabets, e.g. by replacing each letter with a random sequence, do not seem to work, and we do not know how to construct such an OR gadget with a smaller alphabet *in a black box way*. The major part of our proof will be a careful examination of the formula and the sub-gadgets  $g_1, g_2$  in order to reduce the alphabet size to a large enough constant, while using padding that is only  $1 + \varepsilon_\sigma$  times the length of the sub-gadgets. We achieve this by combining this crossing gadget with a small padding that will reuse letters from alphabets that were used much deeper in the formula, and we will argue that the noise we get from recycling letters is dominated by our paddings, in any optimal matching.

We remark that the reduction of [3] can be implemented in a generic way with any problem that admits *alignment gadgets* as defined in [29], giving formula-gadgets of size  $s^{O(1)}$ . The list of such problems includes LCS and Edit-Distance on binary strings. However, to get gadgets of length  $s^{1+o(1)}$  it seems that problem-specific reductions are necessary. A big open question left by our work is to find the most efficient reduction from Formula-SAT to Edit-Distance. A very efficient OR gadget, even if the alphabet is unbounded, might be (provably) impossible. Can we use this intuition to shave more log factors for Edit-Distance?

Fréchet Distance falls outside the alignment gadgets framework of [29] and no reduction from Formula-SAT was known before. In Section F we prove such a reduction by a significant boosting of the SETH-lower bound construction of [27]. In order to implement recursive AND/OR gadgets,

our new proof utilizes the geometry of the curves, in contrast to [27] which only used ten different points in the plane.

In the remainder of this section we present the details of the reduction to LCS. Some missing proofs can be found in Section C.

## 2.1 Implementing Gates

Fix vectors  $a, b \in \{0, 1\}^m$  (where  $2m$  is the number of inputs to  $F$ ). In this section we prove the following lemma which demonstrates our main construction.

**Lemma 2.3.** *For any sufficiently large  $\sigma > 0$  let  $\tau = (\log \sigma)^{1/4}$ . We can inductively construct, for each gate  $g$  of  $F$ , strings  $x(g) = x(g, a)$  and  $y(g) = y(g, b)$  over alphabet size  $5\sigma^2$  and a number  $\rho(g)$  such that for  $L(g) := \text{LCS}(x(g), y(g))$  we have (1)  $L(g) \leq \rho(g)$  and (2)  $L(g) = \rho(g)$  if and only if gate  $g$  evaluates to true on input  $(a, b)$  to  $F$ . Moreover, we have  $|x(g)| = |y(g)| = n(g) \leq 6\tau \cdot |F_g|(1 + 7/\tau)^{\text{depth}(F_g)}$ , where  $F_g$  is the subformula of  $F$  below  $g$ .*

In this construction, we use disjoint size-5 alphabets  $\Sigma_1, \dots, \Sigma_{\sigma^2}$ , determining the total alphabet size as  $5\sigma^2$ . Each gate  $g$  is assigned an alphabet  $\Sigma_{f(g)}$ . We fix the function  $f$  later.

In the following, consider any gate  $g$  of  $F$ , and write the gate alphabet as  $\Sigma_{f(g)} = \{0, 1, 2, 3, 4\}$ . For readability, we write  $x = x(g)$  and similarly define  $y, n, L, \rho$ . If  $g$  has fanin 2, write  $g_1, g_2$  for the children of  $g$ . Moreover, let  $x_1 = x(g_1)$  and similarly define  $y_1, n_1, L_1, \rho_1$  and  $x_2, y_2, n_2, L_2, \rho_2$ .

**Input Gate** The base case is an input bit  $a_i$  to  $F$  (input bits  $b_j$  are symmetric). Interpreting  $a_i$  as a string of length 1 over alphabet  $\{0, 1\}$ , note that  $\text{LCS}(a_i, 1) = a_i$ . Hence, the strings  $x = a_i$  and  $y = 1$ , with  $n = \rho = 1$ , trivially simulate the input bit  $a_i$ .

**AND Gates** Consider an AND gate  $g$  and let  $\beta := \lceil (n_1 + n_2)/\tau^2 \rceil$ . We construct strings  $x, y$  as

$$\begin{aligned} x &:= x_1 0^\beta 1^\beta x_2 \\ y &:= y_1 0^\beta 1^\beta y_2 \end{aligned}$$

**Lemma 2.4.** *If  $\text{LCS}(x_2, y_1), \text{LCS}(x_1, y_2) \leq \beta/4$  and the symbols 0, 1 appear at most  $\beta/16$  times in each of  $x_1, x_2, y_1$ , and  $y_2$ , then we have  $L = \text{LCS}(x, y) = 2\beta + L_1 + L_2$ .*

Later we will choose the gate alphabets  $\Sigma_{f(g)}$  such that the precondition of the above lemma is satisfied. Setting  $\rho := 2\beta + \rho_1 + \rho_2$  we thus inductively obtain (1)  $L \leq \rho$  and (2)  $L = \rho$  if and only if  $g_1$  and  $g_2$  both evaluate to true. Thus, we correctly simulated the AND gate  $g$ . It remains to prove the lemma.

*Proof.* Clearly, we have  $L \geq \text{LCS}(x_1, y_1) + \text{LCS}(0^\beta, 0^\beta) + \text{LCS}(1^\beta, 1^\beta) + \text{LCS}(x_2, y_2) = 2\beta + L_1 + L_2$ . For the other direction, consider any LCS  $z$  of  $x, y$ . If  $z$  does not match any symbol of the left half of  $x$ ,  $x_1 0^\beta$ , with any symbol of the right half of  $y$ ,  $1^\beta y_2$ , and it does not match any symbol of the right half of  $x$ ,  $1^\beta x_2$ , with any symbol of the left half of  $y$ ,  $y_1 0^\beta$ , then we can split both strings in the middle and obtain

$$L = |z| \leq \text{LCS}(x_1 0^\beta, y_1 0^\beta) + \text{LCS}(1^\beta x_2, 1^\beta y_2).$$

Greedy suffix/prefix matching now yields

$$L \leq (\text{LCS}(x_1, y_1) + \beta) + (\beta + \text{LCS}(x_2, y_2)) = 2\beta + L_1 + L_2.$$

In the remaining case, there is a matching from some left half to some right half. By symmetry, we can assume that there is a matching from the left half of  $x$  to the right half of  $y$ . We can moreover assume that  $z$  matches a symbol of  $x_1$  with a symbol of  $1^\beta y_2$ , since the case that  $z$  matches a symbol of  $y_2$  with a symbol of  $x_1 0^\beta$  is symmetric. Now no symbol in  $0^\beta$  in  $x$  can be matched with a symbol in  $0^\beta$  in  $y$ . We obtain a rough upper bound on  $L = |z|$  by summing up the LCS length of all remaining  $4 \cdot 4 - 1 = 15$  pairs of a part  $x' \in \{x_1, 0^\beta, 1^\beta, x_2\}$  in  $x$  and a part  $y' \in \{y_1, 0^\beta, 1^\beta, y_2\}$  in  $y$ . This yields  $L \leq L_1 + L_2 + \beta + 2 \cdot \beta/4 + 8 \cdot \beta/16 = 2\beta + L_1 + L_2$ , finishing the proof.  $\square$

**OR Gates** Consider an OR gate  $g$  and again let  $\beta := \lceil (n_1 + n_2) / \tau^2 \rceil$ . We first make the LCS target values equal by adding  $4^{|\rho_1 - \rho_2|}$  to the shorter of  $x_2/y_2$  and  $x_1/y_1$ , i.e., we set  $x'_1 := 4^{\max\{0, \rho_2 - \rho_1\}} x_1$  and similarly  $y'_1 := 4^{\max\{0, \rho_2 - \rho_1\}} y_1$ ,  $x'_2 := 4^{\max\{0, \rho_1 - \rho_2\}} x_2$ ,  $y'_2 := 4^{\max\{0, \rho_1 - \rho_2\}} y_2$ . Note that the resulting strings satisfy  $L'_1 := \text{LCS}(x'_1, y'_1) \leq \rho' := \max\{\rho_1, \rho_2\}$  and  $L'_1 = \rho'$  if and only if  $g_1$  evaluates to true, and similarly  $L'_2 := \text{LCS}(x'_2, y'_2) \leq \rho'$  and  $L'_2 = \rho'$  if and only if  $g_2$  evaluates to true. We construct the strings  $x, y$  as

$$\begin{aligned} x &:= 0^\beta 1^\beta x'_1 2^\beta 3^\beta x'_2 0^\beta 1^\beta \\ y &:= 2^\beta 3^\beta y'_2 0^\beta 1^\beta y'_1 2^\beta 3^\beta \end{aligned}$$

**Lemma 2.5.** *If  $\text{LCS}(x_2, y_1), \text{LCS}(x_1, y_2) \leq \beta/8$  and the symbols  $0, 1, 2, 3$  appear at most  $\beta/48$  times in each of  $x_1, x_2, y_1$ , and  $y_2$ , then  $L = \text{LCS}(x, y) = 4\beta + \max\{L'_1, L'_2\}$ .*

Later we will choose the gate alphabets  $\Sigma_{f(g)}$  such that the precondition of the above lemma is satisfied. Setting  $\rho := 4\beta + \rho' = 4\beta + \max\{\rho_1, \rho_2\}$  we thus inductively obtain (1)  $L \leq \rho$  and (2)  $L = \rho$  if and only if at least one of  $g_1$  and  $g_2$  evaluates to true, so we correctly simulated the OR gate  $g$ . The proof of the Lemma is in Section C.

**Analyzing the Length** Note that the above constructions inductively yields strings  $x(g), y(g)$  simulating each gate  $g$ . We inductively prove bounds for  $n(g)$  and  $\rho(g)$ . See Section C.

**Lemma 2.6.** *We have  $n(g) \leq 6\tau \cdot |F_g|(1 + 7/\tau)^{\text{depth}(F_g)}$  and  $\rho(g) \leq 6|F_g|(1 + 7/\tau)^{\text{depth}(F_g)}$  for any gate  $g$ , where  $F_g$  is the subformula of  $F$  below  $g$ .*

**Fixing the Gate Alphabets** Now we fix the gate alphabet  $\Sigma_{f(g)}$  for any gate  $g$ . Again let  $\Sigma_{(i,j)}$ ,  $i, j \in [\sigma]$ , be disjoint alphabets of size 5, and let  $\Sigma := \bigcup_{i,j} \Sigma_{(i,j)}$ . For any gate  $g$  of  $F$ , we call its distance to the root the *height*  $h(g)$ . For any  $h$ , order the gates with height  $h$  from left to right, and let  $\iota(g)$  be the index of gate  $g$  in this order, for any gate  $g$  with height  $h$ . Note that  $(h(g), \iota(g))$  is a unique identifier of gate  $g$ . We define  $f(g) := (h(g) \bmod \sigma, \iota(g) \bmod \sigma)$ , i.e., we set the gate alphabet of  $g$  to  $\Sigma_{f(g)} = \Sigma_{(h(g) \bmod \sigma, \iota(g) \bmod \sigma)}$ . Note that the overall alphabet  $\Sigma$  has size  $5\sigma^2$ . Recall that we set  $\tau := (\log \sigma)^{1/4}$ .

It remains to show that the preconditions of Lemmas 2.4 and 2.5 are satisfied. Specifically, consider a gate  $g$  with children  $g_1, g_2$ . As before, let  $x, y, n$  be the strings and string length constructed for gate  $g$ , and let  $x_i, y_i, n_i$  be the corresponding objects for  $g_i$ ,  $i \in \{1, 2\}$ . We need to show:

- (1)  $\text{LCS}(x_2, y_1), \text{LCS}(x_1, y_2) \leq (n_1 + n_2)/(8\tau^2)$ , and  
(2) each symbol  $c \in \Sigma_{f(g)}$  appears at most  $(n_1 + n_2)/(48\tau^2)$  times in each of  $x_1, x_2, y_1$ , and  $y_2$ .

We call a gate  $g'$  in the subformula  $F_g$  *d-deep* if  $h(g') \geq h(g) + d$ , and *d-shallow* otherwise. For each symbol  $c$  in  $x$  or  $y$  we can trace our construction to find the gate  $g'$  in  $F_g$  at which we introduced  $c$  to  $x$  or  $y$ . In other words, each symbol in  $x, y$  stems from some gate  $g'$  below  $g$ .

First consider (2). Observe that all symbols in  $x, y$  stemming from  $\sigma$ -shallow gates do not belong to the gate alphabet  $\Sigma_{f(g)}$ , since the function  $f(g')$  has  $(h(g') \bmod \sigma)$  as the first component, which repeats only every  $\sigma$  levels. Thus, if a symbol  $c \in \Sigma_{f(g)}$  occurs in  $x_i$  or  $y_i$ , then this occurrence stems from a  $\sigma$ -deep gate. We now argue that only few symbols in  $x, y$  stem from deep gates. For any  $d > 0$ , let  $N_d$  be the number of symbols in  $x$  (or, equivalently,  $y$ ) stemming from  $d$ -deep gates. Note that  $N_d$  is equal to the total string length  $\sum n(g')$ , summed over all gates  $g'$  in  $F_g$  with height  $h(g') = h(g) + d$ . Observe that our construction increases the string lengths in each step by at least a factor  $1 + 1/\tau^2$ , i.e.,  $N_d \geq (1 + 1/\tau^2)N_{d+1}$  holds for any  $d$ . It follows that  $N_\sigma \leq N_1/(1 + 1/\tau^2)^{\sigma-1} = (n_1 + n_2)/(1 + 1/\tau^2)^{\sigma-1}$ . Hence, each symbol in  $\Sigma_{f(g)}$  appears at most  $(n_1 + n_2)/(1 + 1/\tau^2)^{\sigma-1}$  times in each of  $x_1, x_2, y_1, y_2$ . Since  $\tau = (\log \sigma)^{1/4}$ , we have  $(1 + 1/\tau^2)^{\sigma-1} = 2^{\Omega(\sigma/\sqrt{\log \sigma})} \geq 48\sqrt{\log \sigma} = 48\tau^2$  for sufficiently large  $\sigma$ . This proves (2).

For (1), remove all  $\log(\sigma)$ -deep symbols from  $x_1$  and  $y_2$  to obtain strings  $x'_1, y'_2$ . Note that we removed exactly  $N_{\log \sigma}$  symbols from each of  $x_1, y_2$ . This yields  $\text{LCS}(x_1, y_2) \leq 2N_{\log \sigma} + \text{LCS}(x'_1, y'_2)$ . For  $x'_1, y'_2$ , we claim that any  $\log(\sigma)$ -shallow gates  $g'_1 \neq g'_2$  in  $F_g$  have disjoint alphabets  $\Sigma_{f(g'_1)}, \Sigma_{f(g'_2)}$ . Indeed, if  $h(g'_1) \neq h(g'_2)$  then since the first component  $(h(g') \bmod \sigma)$  of  $f(g')$  repeats only every  $\sigma$  levels we have  $f(g'_1) \neq f(g'_2)$ . If  $h(g'_1) = h(g'_2) =: h$ , then note that each gate  $g'$  in height  $h$  has a unique label  $\iota(g') \bmod \sigma$ , since there are  $\sigma$  such labels and there are at most  $2^{h-h(g)} < \sigma$  gates with height  $h$  in  $F_g$ . Hence,  $x'_1$  and  $y'_2$  use disjoint alphabets, and we obtain  $\text{LCS}(x'_1, y'_2) = 0$ . Thus,  $\text{LCS}(x_1, y_2) \leq 2N_{\log \sigma}$ . As above, we bound  $N_{\log \sigma} \leq (n_1 + n_2)/(1 + 1/\tau^2)^{\log \sigma - 1}$ , so that  $\text{LCS}(x_1, y_2) \leq 2(n_1 + n_2)/(1 + 1/\tau^2)^{\log \sigma - 1}$ . Since  $\tau = (\log \sigma)^{1/4}$ , we have  $(1 + 1/\tau^2)^{\log \sigma - 1}/2 = 2^{\Omega(\sqrt{\log \sigma})} \geq 8\sqrt{\log \sigma} = 8\tau^2$  for sufficiently large  $\sigma$ . This yields (1), since the strings  $x_2, y_1$  are symmetric. This finishes the proof of Lemma 2.3.

**Finalizing the Proof** Let us sketch how we complete the proof of Theorem 2.2. The full details are in Section C.1. First, for all vectors  $a \in A, b \in B$  we construct gate gadgets for the output gate of the formula, i.e. *formula gadgets*, by invoking Lemma 2.3. Then we combine all these gadgets by applying a standard alignment gadget [2, 29] to get our final sequences of length  $O(n\tau|F|(1 + 7/\tau)^{\text{depth}(F)})$  and with alphabet of size  $O(\sigma^2)$ . The LCS of the final sequence will be determined by the existence of a satisfying pair. Since a priori the depth of  $F$  could be as large as  $|F|$ , the factor  $(1 + 7/\tau)^{\text{depth}(F)}$  in our length bound is not yet satisfactory. Thus, as a preprocessing before the above construction, we decrease the depth of  $F$  using a depth-reduction result of Bonet and Buss [90, 25]: for all  $k \geq 2$  there is an equivalent formula  $F'$  with depth at most  $(3k \ln 2) \log |F|$  and size  $|F'| \leq |F|^{1+1/(1+\log(k-1))}$ . Choosing the parameters correctly, we get final sequences of length  $O(n|F|^{1+O(1/\log \log \sigma)})$ .

### 3 On the Limitations of Fine-Grained Reductions

With the increasingly complex web of reductions and conjectures used in the ‘‘Hardness in P’’ research, one might oppose to our use of nonstandard assumptions. Why can’t we base the hardness

of shaving logs on one of the more established assumptions such as SETH, or even better, on  $P \neq NP$ ? We conclude the paper with a proof that such results are not possible if one is restricted to fine-grained reductions, which is essentially the only tool we have in this line of research.

Let  $A$  be a problem with best known upper bound of  $T_A(n)$  on inputs of size  $n$ , and let  $B$  be a problem with best known upper bound of  $T_B(n)$  on inputs of size  $n$ . Throughout this section we assume that these runtime are non-decreasing functions, such as  $2^n$  or  $n^2$ . A fine-grained reduction from “solving  $A$  in time  $T_A(n)/g(n)$ ” to “solving  $B$  in time  $T_B(n)/f(n)$ ” proves that improving  $T_B(n)$  to  $T_B(n)/f(n)$  improves  $T_A$  to  $T_A(n)/g(n)$ . Formally, it is an algorithm  $X$  that solves  $A$  and it is allowed to call an oracle for problem  $B$ , as long as the following bound holds. Let  $n_i$  be the size of the instance in the  $i^{\text{th}}$  call to problem  $B$  that our algorithm performs, where  $i \leq t$  for some value  $t$ , and let  $T_X(n)$  be the runtime of  $X$  excluding the time it takes to answer all the instances of problem  $B$ . It must be that  $T_X(n) + \sum_{i=1}^t T_B(n_i)/f(n_i) \leq T_A(n)/g(n)$ . This is a natural adaptation of the definition of fine-grained reductions from previous works, where the improvements were restricted to be by polynomial factors.

We can now give a formal version of Theorem 1.1 from the introduction. Note that  $k$ -SAT on  $n$  variables and  $m$  clauses can be solved in time  $\text{poly}(n, m)2^n$ .

**Theorem 3.1.** *If for some  $c, \varepsilon > 0$  and all  $k \geq 2$  there is a fine-grained reduction from solving  $k$ -SAT in time  $\text{poly}(n, m)2^n/2^{\varepsilon n}$  to solving LCS in time  $O(n^2/\log^c n)$ , then SETH is false.*

*Proof.* Assume there was a fine-grained reduction from  $k$ -SAT to LCS as above. This means that there is an algorithm  $X$  for  $k$ -SAT that makes  $t$  calls to LCS with instances of size  $n_1, \dots, n_t$  such that:

$$T_X(n) + \sum_{i=1}^t n_i^2 / \log^c n_i = O(\text{poly}(n, m)2^n/2^{\varepsilon n})$$

But then consider algorithm  $X'$  which simulates  $X$  and whenever  $X$  makes a call to the LCS oracle with an instance of size  $n_i$ , our algorithm will execute the known quadratic time solution for LCS that takes  $O(n_i^2)$  time. Let  $n_{\max}$  be the size of the largest instance we call, and note that  $n_{\max} < 2^n$ . Simple calculations show that  $X'$  solves  $k$ -SAT and has a running time of

$$T_X(n) + \sum_{i=1}^t n_i^2 = O(\text{poly}(n, m)2^n/2^{\varepsilon n}) \cdot \log^c n_{\max} = O(\text{poly}(n, m)2^n/2^{\varepsilon n})$$

for all  $k$ , refuting SETH. □

**Acknowledgements** We are grateful to Avishay Tal for telling us about his algorithm for SAT on bipartite formulas, and for very helpful discussions. We also thank Mohan Paturi, Rahul Santhanam, Srikanth Srinivasan, and Ryan Williams for answering our questions about the state of the art of Formula-SAT algorithms, and Arturs Backurs, Piotr Indyk, Mikkel Thorup, and Virginia Vassilevska Williams for helpful discussions regarding regular expressions. We also thank an anonymous reviewer for ideas leading to shaving off a second log-factor for FORMULA-PAIR, and other reviewers for helpful suggestions.

## References

- [1] A. Abboud, A. Backurs, T. D. Hansen, V. Vassilevska Williams, and O. Zamir. Subtree isomorphism revisited. In *Proc. of 27th SODA*, pages 1256–1271, 2016.
- [2] A. Abboud, A. Backurs, and V. Vassilevska Williams. Tight Hardness Results for LCS and other Sequence Similarity Measures. In *Proc. of 56th FOCS*, pages 59–78, 2015.
- [3] A. Abboud, T. D. Hansen, V. V. Williams, and R. Williams. Simulating branching programs with edit distance and friends: or: a polylog shaved is a lower bound made. In *Proc. of the 48th STOC*, pages 375–388, 2016.
- [4] A. Abboud and V. Vassilevska Williams. Popular conjectures imply strong lower bounds for dynamic problems. In *Proc. of 55th FOCS*, pages 434–443, 2014.
- [5] A. Abboud, V. Vassilevska Williams, and J. R. Wang. Approximation and fixed parameter subquadratic algorithms for radius and diameter in sparse graphs. In *Proc. of 27th SODA*, pages 377–391, 2016.
- [6] A. Abboud, V. Vassilevska Williams, and O. Weimann. Consequences of faster sequence alignment. In *Proc. of 41st ICALP*, pages 39–51, 2014.
- [7] A. Abboud, V. Vassilevska Williams, and H. Yu. Matching triangles and basing hardness on an extremely popular conjecture. In *Proc. of 47th STOC*, pages 41–50, 2015.
- [8] A. Abboud, R. Williams, and H. Yu. More applications of the polynomial method to algorithm design. In *Proc. of 26th SODA*, pages 218–230, 2015.
- [9] P. Agarwal, R. B. Avraham, H. Kaplan, and M. Sharir. Computing the discrete Fréchet distance in subquadratic time. In *Proc. 24th ACM-SIAM Symposium on Discrete Algorithms (SODA '13)*, pages 156–167, 2013.
- [10] L. Allison and T. I. Dix. A bit-string longest-common-subsequence algorithm. *Information Processing Letters*, 23(5):305–310, 1986.
- [11] H. Alt and M. Buchin. Can we compute the similarity between surfaces? *Discrete & Computational Geometry*, 43(1):78–99, 2010.
- [12] H. Alt and M. Godau. Computing the Fréchet distance between two polygonal curves. *International Journal of Computational Geometry & Applications*, 5(1-2):78–99, 1995.
- [13] H. Alt, C. Knauer, and C. Wenk. Comparison of distance measures for planar curves. *Algorithmica*, 38(1):45–58, 2004.
- [14] S. F. Altschul, T. L. Madden, A. A. Schäffer, J. Zhang, Z. Zhang, W. Miller, and D. J. Lipman. Gapped blast and psi-blast: a new generation of protein database search programs. *Nucleic acids research*, 25(17):3389–3402, 1997.
- [15] A. E. Andreev. About one method of obtaining more than quadratic effective lower bounds of complexity of pi-schemes, 1987.

- [16] V. Arlazarov, E. Dinic, I. Faradzev, and M. Kronrod. On economic construction of the transitive closure of a direct graph. In *Sov. Math (Doklady)*, volume 11, pages 1209–1210, 1970.
- [17] B. Aronov, S. Har-Peled, C. Knauer, Y. Wang, and C. Wenk. Fréchet distance for curves, revisited. In *Proc. 14th Annual European Symposium on Algorithms (ESA'06)*, volume 4168 of *LNCS*, pages 52–63. Springer, 2006.
- [18] A. Backurs and P. Indyk. Edit Distance Cannot Be Computed in Strongly Subquadratic Time (unless SETH is false). In *Proc. of 47th STOC*, pages 51–58, 2015.
- [19] A. Backurs and P. Indyk. Which regular expression patterns are hard to match? In *FOCS*, 2016.
- [20] N. Bansal and R. Williams. Regularity lemmas and combinatorial algorithms. In *Proc. of 50th FOCS*, pages 745–754, 2009.
- [21] I. Baran, E. D. Demaine, and M. Patrascu. Subquadratic algorithms for 3sum. *Algorithmica*, 50(4):584–596, 2008.
- [22] P. Beame, R. Impagliazzo, and S. Srinivasan. Approximating  $ac^0$  by small height decision trees and a deterministic algorithm for  $\#ac^0$  sat. In *Computational Complexity (CCC), 2012 IEEE 27th Annual Conference on*, pages 117–125. IEEE, 2012.
- [23] P. Bille and M. Farach-Colton. Fast and compact regular expression matching. *Theoretical Computer Science*, 409(3):486 – 496, 2008.
- [24] P. Bille and M. Thorup. Faster regular expression matching. In *International Colloquium on Automata, Languages, and Programming*, pages 171–182. Springer, 2009.
- [25] M. L. Bonet and S. R. Buss. Size-depth tradeoffs for boolean fomulae. *Inf. Process. Lett.*, 49(3):151–155, 1994.
- [26] S. Brakatsoulas, D. Pfoser, R. Salas, and C. Wenk. On map-matching vehicle tracking data. In *Proc. 31st International Conference on Very Large Data Bases (VLDB'05)*, pages 853–864, 2005.
- [27] K. Bringmann. Why walking the dog takes time: Frechet distance has no strongly subquadratic algorithms unless seth fails. In *Proc. of 55th FOCS*, pages 661–670, 2014.
- [28] K. Bringmann, A. Grønlund, and K. G. Larsen. A dichotomy for regular expression membership testing. *CoRR*, abs/1611.00918, 2016.
- [29] K. Bringmann and M. Künnemann. Quadratic Conditional Lower Bounds for String Problems and Dynamic Time Warping. In *Proc. of 56th FOCS*, pages 79–97, 2015.
- [30] K. Buchin, M. Buchin, J. Gudmundsson, M. Löffler, and J. Luo. Detecting commuting patterns by clustering subtrajectories. *International Journal of Computational Geometry & Applications*, 21(3):253–282, 2011.

- [31] K. Buchin, M. Buchin, W. Meulemans, and W. Mulzer. Four soviets walk the dog - with an application to Alt’s conjecture. In *Proc. 25th ACM-SIAM Symposium on Discrete Algorithms (SODA’14)*, pages 1399–1413, 2014.
- [32] K. Buchin, M. Buchin, and Y. Wang. Exact algorithms for partial curve matching via the Fréchet distance. In *Proc. 20th ACM-SIAM Symposium on Discrete Algorithms (SODA’09)*, pages 645–654, 2009.
- [33] M. Cairo, R. Grossi, and R. Rizzi. New bounds for approximating extremal distances in undirected graphs. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 363–376, 2016.
- [34] C. Calabro, R. Impagliazzo, and R. Paturi. A duality between clause width and clause density for SAT. In *Proc. of 21st CCC*, pages 252–260, 2006.
- [35] C. Calabro, R. Impagliazzo, and R. Paturi. The complexity of satisfiability of small depth circuits. In *Proc. of 4th IWPEC*, pages 75–85, 2009.
- [36] 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.
- [37] T. M. Chan. The art of shaving logs. In *Proc. of the 13th WADS*, page 231, 2013.
- [38] T. M. Chan. Speeding up the four russians algorithm by about one more logarithmic factor. In *Proc. of 26th SODA*, pages 212–217, 2015.
- [39] T. M. Chan. More logarithmic-factor speedups for 3sum, (median,+)-convolution, and some geometric 3sum-hard problems. In *29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA’18)*, pages 881–897, 2018.
- [40] K. Chatterjee, W. Dvorák, M. Henzinger, and V. Loitzenbauer. Model and objective separation with conditional lower bounds: Disjunction is harder than conjunction. *CoRR*, abs/1602.02670, 2016.
- [41] R. Chen. Satisfiability algorithms and lower bounds for boolean formulas over finite bases. In *International Symposium on Mathematical Foundations of Computer Science*, pages 223–234. Springer, 2015.
- [42] R. Chen and V. Kabanets. Correlation bounds and #sat algorithms for small linear-size circuits. In *Computing and Combinatorics - 21st International Conference, COCOON 2015, Beijing, China, August 4-6, 2015, Proceedings*, pages 211–222, 2015.
- [43] R. Chen, V. Kabanets, A. Kolokolova, R. Shaltiel, and D. Zuckerman. Mining circuit lower bound proofs for meta-algorithms. *computational complexity*, 24(2):333–392, 2015.
- [44] R. Chen, V. Kabanets, and N. Saurabh. An improved deterministic # sat algorithm for small de morgan formulas. In *International Symposium on Mathematical Foundations of Computer Science*, pages 165–176. Springer, 2014.

- [45] R. Chen and R. Santhanam. Satisfiability on mixed instances. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science, Cambridge, MA, USA, January 14-16, 2016*, pages 393–402, 2016.
- [46] A. F. Cook and C. Wenk. Geodesic Fréchet distance inside a simple polygon. *ACM Transactions on Algorithms*, 7(1):193–204, 2010.
- [47] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to algorithms*, volume 6. MIT press Cambridge, 2001.
- [48] M. Crochemore, C. S. Iliopoulos, Y. J. Pinzon, and J. F. Reid. A fast and practical bit-vector algorithm for the longest common subsequence problem. *Information Processing Letters*, 80(6):279–285, 2001.
- [49] M. Crochemore, G. M. Landau, and M. Ziv-Ukelson. A subquadratic sequence alignment algorithm for unrestricted scoring matrices. *SIAM journal on computing*, 32(6):1654–1673, 2003.
- [50] E. Dantsin and E. A. Hirsch. Worst-case upper bounds. In *Handbook of Satisfiability*, pages 403–424. 2009.
- [51] E. Dantsin and A. Wolpert. Exponential complexity of satisfiability testing for linear-size boolean formulas. In *International Conference on Algorithms and Complexity*, pages 110–121. Springer, 2013.
- [52] A. Driemel and S. Har-Peled. Jaywalking your dog: computing the Fréchet distance with shortcuts. *SIAM Journal on Computing*, 42(5):1830–1866, 2013.
- [53] 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.
- [54] 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.
- [55] A. Freund. Improved subquadratic 3sum. *Algorithmica*, 77(2):440–458, 2017.
- [56] P. Gawrychowski. Faster algorithm for computing the edit distance between slp-compressed strings. In *International Symposium on String Processing and Information Retrieval*, pages 229–236. Springer, 2012.
- [57] M. Godau. A natural metric for curves - computing the distance for polygonal chains and approximation algorithms. In *Proc. 8th Symposium on Theoretical Aspects of Computer Science (STACS'91)*, volume 480 of *LNCS*, pages 127–136. Springer, 1991.
- [58] O. Gold and M. Sharir. Improved Bounds for 3SUM, k-SUM, and Linear Degeneracy. In *25th Annual European Symposium on Algorithms (ESA 2017)*, volume 87, pages 42:1–42:13, 2017.
- [59] A. Golovnev, A. S. Kulikov, A. Smal, and S. Tamaki. Circuit size lower bounds and # sat upper bounds through a general framework. In *Electronic Colloquium on Computational Complexity (ECCC)*, volume 23, page 22, 2016.

- [60] S. Grabowski. New tabulation and sparse dynamic programming based techniques for sequence similarity problems. In *Stringology*, pages 202–211, 2014.
- [61] A. Grønlund and S. Pettie. Threesomes, degenerates, and love triangles. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 621–630, 2014.
- [62] J. Håstad. The shrinkage exponent of de morgan formulas is 2. *SIAM J. Comput.*, 27(1):48–64, 1998.
- [63] J. Hopcroft, W. Paul, and L. Valiant. On time versus space. *Journal of the ACM (JACM)*, 24(2):332–337, 1977.
- [64] H. Hyvrö. Bit-parallel lcs-length computation revisited. In *Proc. 15th Australasian Workshop on Combinatorial Algorithms (AWOCA 2004)*, pages 16–27. Citeseer, 2004.
- [65] R. Impagliazzo, S. Lovett, R. Paturi, and S. Schneider. 0-1 integer linear programming with a linear number of constraints. *Electronic Colloquium on Computational Complexity (ECCC)*, 21:24, 2014.
- [66] R. Impagliazzo, W. Matthews, and R. Paturi. A satisfiability algorithm for ac 0. In *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms*, pages 961–972. SIAM, 2012.
- [67] R. Impagliazzo and N. Nisan. The effect of random restrictions on formula size. *Random Structures & Algorithms*, 4(2):121–133, 1993.
- [68] R. Impagliazzo and R. Paturi. On the complexity of k-sat. *Journal of Computer and System Sciences*, 62(2):367–375, 2001.
- [69] R. Impagliazzo, R. Paturi, and S. Schneider. A satisfiability algorithm for sparse depth two threshold circuits. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26-29 October, 2013, Berkeley, CA, USA*, pages 479–488, 2013.
- [70] P. Indyk. Approximate nearest neighbor algorithms for Fréchet distance via product metrics. In *Proc. 18th Annual Symposium on Computational Geometry (SoCG'02)*, pages 102–106, 2002.
- [71] V. M. Khrapchenko. Method of determining lower bounds for the complexity of p-schemes. *Mathematical Notes*, 10(1):474–479, 1971.
- [72] I. Komargodski, R. Raz, and A. Tal. Improved average-case lower bounds for demorgan formula size. In *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*, pages 588–597. IEEE, 2013.
- [73] I. T. Li, W. Shum, and K. Truong. 160-fold acceleration of the smith-waterman algorithm using a field programmable gate array (fpga). *BMC bioinformatics*, 8(1):1, 2007.
- [74] Y. Liu, A. Wirawan, and B. Schmidt. Cudasw++ 3.0: accelerating smith-waterman protein database search by coupling cpu and gpu simd instructions. *BMC bioinformatics*, 14(1):1, 2013.

- [75] A. Maheshwari, J.-R. Sack, K. Shahbaz, and H. Zarrabi-Zadeh. Fréchet distance with speed limits. *Computational Geometry*, 44(2):110–120, 2011.
- [76] W. J. Masek and M. S. Paterson. A faster algorithm computing string edit distances. *Journal of Computer and System sciences*, 20(1):18–31, 1980.
- [77] D. Moeller, R. Paturi, and S. Schneider. Subquadratic algorithms for succinct stable matching. In *Computer Science - Theory and Applications - 11th International Computer Science Symposium in Russia, CSR 2016, St. Petersburg, Russia, June 9-13, 2016, Proceedings*, pages 294–308, 2016.
- [78] M. E. Munich and P. Perona. Continuous dynamic time warping for translation-invariant curve alignment with applications to signature verification. In *Proc. 7th IEEE International Conference on Computer Vision*, volume 1, pages 108–115, 1999.
- [79] G. Myers. A four russians algorithm for regular expression pattern matching. *Journal of the ACM (JACM)*, 39(2):432–448, 1992.
- [80] I. C. Oliveira. Algorithms versus circuit lower bounds. *arXiv preprint arXiv:1309.0249*, 2013.
- [81] M. S. Paterson and U. Zwick. Shrinkage of de morgan formulae under restriction. *Random Structures & Algorithms*, 4(2):135–150, 1993.
- [82] M. Patrascu and R. Williams. On the possibility of faster SAT algorithms. In *Proc. of 21st SODA*, pages 1065–1075, 2010.
- [83] R. Paturi, P. Pudlák, M. E. Saks, and F. Zane. An improved exponential-time algorithm for  $k$ -sat. *J. ACM*, 52(3):337–364, 2005.
- [84] L. Roditty and V. Vassilevska Williams. Fast approximation algorithms for the diameter and radius of sparse graphs. In *Proc. of 45th STOC*, pages 515–524, 2013.
- [85] T. Sakai, K. Seto, S. Tamaki, and J. Teruyama. A satisfiability algorithm for depth-2 circuits with a symmetric gate at the top and and gates at the bottom. In *Electronic Colloquium on Computational Complexity (ECCC)*, 2015.
- [86] R. Santhanam. Fighting perebor: New and improved algorithms for formula and QBF satisfiability. In *Proc. of the 51th FOCS*, pages 183–192, 2010.
- [87] R. Santhanam et al. Ironic complicity: Satisfiability algorithms and circuit lower bounds. *Bulletin of EATCS*, 1(106), 2013.
- [88] K. Seto and S. Tamaki. A satisfiability algorithm and average-case hardness for formulas over the full binary basis. *computational complexity*, 22(2):245–274, 2013.
- [89] T. F. Smith and M. S. Waterman. Identification of common molecular subsequences. *Journal of molecular biology*, 147(1):195–197, 1981.
- [90] P. M. Spira. On time-hardware complexity tradeoffs for boolean functions. In *Proceedings of the 4th Hawaii Symposium on System Sciences*, pages 525–527, 1971.

- [91] B. A. Subbotovskaya. Realizations of linear functions by formulas using+. *Doklady Akademii Nauk SSSR*, 136(3):553–555, 1961.
- [92] A. Tal. Shrinkage of de morgan formulae by spectral techniques. In *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*, pages 551–560. IEEE, 2014.
- [93] A. Tal. #sat algorithms from shrinkage. *Electronic Colloquium on Computational Complexity (ECCC)*, 22:114, 2015.
- [94] K. Thompson. Programming techniques: Regular expression search algorithm. *Communications of the ACM*, 11(6):419–422, 1968.
- [95] R. A. Wagner and M. J. Fischer. The string-to-string correction problem. *Journal of the ACM (JACM)*, 21(1):168–173, 1974.
- [96] R. Williams. A new algorithm for optimal 2-constraint satisfaction and its implications. *Theoretical Computer Science*, 348(2):357–365, 2005.
- [97] R. Williams. Improving exhaustive search implies superpolynomial lower bounds. *SIAM Journal on Computing*, 42(3):1218–1244, 2013.
- [98] R. Williams. Algorithms for Circuits and Circuits for Algorithms: Connecting the Tractable and Intractable. In *Proceedings of the International Congress of Mathematicians*, 2014.
- [99] R. Williams. Faster all-pairs shortest paths via circuit complexity. In *Proc. of 46th STOC*, pages 664–673, 2014.
- [100] R. Williams. New algorithms and lower bounds for circuits with linear threshold gates. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pages 194–202. ACM, 2014.
- [101] R. Williams. Nonuniform ACC circuit lower bounds. *J. ACM*, 61(1):2:1–2:32, 2014.
- [102] H. Yu. An improved combinatorial algorithm for boolean matrix multiplication. In *Proc. of 42nd ICALP*, pages 1094–1105, 2015.

## A Discussion

As shown above, the popular conjectures are not fine-grained enough for our purposes and our only viable option is to start from assumptions about the hardness of shaving logs for some problem. The approach taken in this paper and in [3] is to start with variants of SAT. Another option would have been to conjecture that 3-SUM cannot be solved in  $O(n^2/\log^{2+\epsilon} n)$  time, but SAT has several advantages. First, SAT is deeply connected to fundamental topics in complexity theory, which allows us to borrow barriers that complexity theorists have faced for decades. Moreover, there is a vast number of combinatorial problems that we can reduce SAT to, whereas 3-SUM seems more useful in geometric contexts, e.g. 3-SUM-hardness for LCS and Frechet might be impossible [31]. Thus, for the task of proving barriers for shaving logs, our approach seems as good as any.

Conditional lower bounds can even lead to better algorithms, by suggesting regimes of possible improvements. Phrased this way, our results leave the open problem of finding a  $O(n^2/(\log n)^{2+\epsilon})$

time algorithms for LCS, and perhaps more interestingly, shaving many more logs for the related-but-different Edit-Distance problem. The longstanding upper bound for Edit-Distance is  $O(n^2/\log^2 n)$  [76] and our approach does not give barriers higher than  $\Omega(n^2/\log^{20} n)$ .

Finally, regardless of the consequences of our reductions, we think that the statements themselves are intrinsically interesting as they reveal a surprisingly close connection between Formula-SAT (a problem typically studied by complexity theorists) and combinatorial problems that are typically studied by stringologists, computational biologists, and computational geometers, which are *a priori* completely different creatures. The runtime of the standard algorithm for SAT can be recovered almost exactly by encoding the formula into an LCS, Fréchet, or Pattern Matching instance and running the standard dynamic programming algorithms!

## B From Formula-SAT to Formula-Pair

In this section we show a chain of simple reductions starting from variants of Formula-SAT, which have  $2^n$  time complexity, and ending at  $n^2$  time variants of a problem we call FORMULA-PAIR.

A formula  $F$  of size  $s$  over  $n$  variables  $x_1, \dots, x_n$  is in the class  $\mathcal{F}_1$  iff it has the following properties. The gates in the first layer (nodes in the tree whose children are all leaves) compute arbitrary functions  $C : \{0, 1\}^{n/2} \rightarrow \{0, 1\}$ , as long as  $C$  can be computed in  $2^{o(n)}$  time and all children of a gate are marked with variables in  $\{x_1, \dots, x_{n/2}\}$  or with variables in  $\{x_{n/2+1}, \dots, x_n\}$  but not with both. W.l.o.g. we can assume that the inputs are only connected to nodes in the first layer. The gates in the other layers compute deMorgan gates, i.e., OR and AND gates. The size of  $F$  is considered to be the *number of gates in the first layer*. Since  $F$  is a formula and thus has fanout 1, our size measure is up to constant factors equal to the total number of all gates except the inputs. Note that the complexity of the functions in the first layer and their number of incoming wires, i.e. the number of leaves in the tree, do not count towards the size of  $F$ .

### $\mathcal{F}_1$ -FORMULA-SAT

*Input:* Formula  $F = F(x_1, \dots, x_n)$  of size  $s$  with  $n$  inputs from the class  $\mathcal{F}_1$

*Question:* Exist  $x_1, \dots, x_n \in \{0, 1\}$  such that  $F(x_1, \dots, x_n) = \text{true}$ ?

*Complexity:*  $O(2^n(s/n + 1))$ , even restricted to  $s \leq n^{1+o(1)}$

Note that many techniques developed for deMorgan formulas (and generalizations) are not applicable to  $\mathcal{F}_1$ -Formula-SAT, since the first layer is so general. In particular, techniques based on collapses by random restrictions [44, 43, 93] do not seem to work, as the first layer can be resistant to such collapses.

A simple algorithm achieves  $O(2^n \cdot s)$  runtime: Preprocess  $F$  to create a table that allows one to quickly lookup the value of each one of the first-layer gates on a given input. Constructing the table takes time  $s \cdot 2^{n/2} \cdot 2^{o(n)}$ , and after we have it, a brute-force SAT algorithm takes  $O(2^n \cdot s)$  time.

However, one can improve upon this simple algorithm and obtain time  $O(2^n s/n)$ . This is the best time complexity we are aware of, up to log-factors, unless we restrict the formula size to  $s \leq n^{2-\Omega(1)}$ . Our tightest barriers for shaving logs for LCS can be based on the assumption that  $\mathcal{F}_1$ -Formula-SAT cannot be solved in time  $O(2^n/n^{1+\varepsilon})$ , for any  $\varepsilon > 0$ , even when  $s = n^{2+o(1)}$ . Such upper bounds seem out of reach of current techniques.

**Theorem B.1.**  $\mathcal{F}_1$ -Formula-SAT can be solved in time  $O(2^n(s/n + 1))$ .

*Proof.* This follows by combining Lemma B.2 and Theorem B.4 below, noting that due to the change in the machine model we lose one log-factor. It is also easy to directly design an algorithm with the claimed running time by following the ideas in Theorem B.4, noting that reading a packed word now takes time  $O(n)$  instead of  $O(1)$ , and all accesses in the (exponential size) precomputed table now take time  $O(n)$ .  $\square$

Next, we move from  $2^n$  to  $n^2$  with a simple split-and-list reduction (similar to [96]) to the following “pair finding” problem.

**FORMULA-PAIR**

*Input:* A deMorgan Formula  $F = F(x_1, \dots, x_m, y_1, \dots, y_m)$  of size  $2m$  where each input is used exactly once, and  $A, B \subseteq \{0, 1\}^m$  of size  $n$

*Question:* Exist  $a \in A, b \in B$  such that  $F(a, b) = F(a_1, \dots, a_m, b_1, \dots, b_m) = \text{true}$ ?

*Complexity:*  $O(n^2m/\log^2 n + n^2/\log n)$  (Theorem B.4), even restricted to  $m \leq (\log n)^{1+o(1)}$

We remark that the assumption that  $F$  reads any input exactly once is w.l.o.g., since the sets  $A, B$  allow us to “copy” an input  $x_i$  to  $x_j$  by simply ensuring that  $a_i = a_j$  for all  $a \in A$ .

**Lemma B.2.** An instance of  $\mathcal{F}_1$ -Formula-SAT on a formula of size  $s$  over  $n$  inputs can be reduced to an instance of FORMULA-PAIR on two sets of size  $O(2^{n/2})$  and a formula of size  $m = O(s)$ , in linear time.

*Proof.* Let  $C_1, \dots, C_m$  be the gates of the first layer of  $F$  that compute functions of the first half of the variables, and let  $C'_1, \dots, C'_m$  be the rest of the gates of the first layer, for some  $k \leq s$ . For each partial assignment to the first  $n/2$  variables of the formula we compute a bit-string  $a \in \{0, 1\}^m$  such that  $a_i$  is 1 iff  $C_i$  outputs 1 on the corresponding partial assignment. Similarly, we compute a bit-string  $b \in \{0, 1\}^m$  for each partial assignment to the second half of the variables. We define our deMorgan formula  $F'$  to be equivalent to the layers of  $F$  that are above the first one. Now, to conclude the reduction, we observe that for any two partial assignments  $\alpha, \beta$  to  $F$ , the corresponding bit strings  $a, b$  contain the values of the first layer of  $F$  on the assignment  $(\alpha\beta)$  and so the value of  $F'$  on  $(ab)$  is exactly the value of  $F$  on  $(\alpha\beta)$ . Thus,  $F$  is satisfiable iff there is a pair  $a \in A, b \in B$  that satisfies  $F'$ . Since  $F$  has fanout 1, every input of  $F'$  is used exactly once.  $\square$

A trivial algorithm for FORMULA-PAIR takes  $O(n^2 \cdot m)$  time, and in Section B.1 we show how to shave two log factors to  $O(n^2m/\log^2 n)$ , for  $m = \Omega(\log n)$ . Via the above reduction, this yields an  $O(2^n s/n)$  algorithm for  $\mathcal{F}_1$ -Formula-SAT for  $s = \Omega(n)$ , due to the loss of a  $\log(n)$  factor in the jump from  $2^n$  to  $n$  and the machine model.

In the other sections we reduce FORMULA-PAIR to LCS and show that any improvement by more than another log-factor on its running time directly yields improved runtimes for all problems above, including  $\mathcal{F}_1$ -Formula-SAT. While we chose to present the barriers for shaving logs for LCS in terms of barriers for solving SAT faster, it is only safer to conjecture that the possibly harder FORMULA-PAIR problem cannot be solved faster.

Next, we consider different variants that we will be able to reduce to the Fréchet distance. A formula  $F$  of size  $s$  over  $n$  variables  $x_1, \dots, x_n$  is in the class  $\mathcal{F}_2$  iff it has the following properties.

The gates in the first layer (nodes in the tree whose children are all leaves) are as in  $\mathcal{F}_1$ . The gates in the second layer compute threshold functions  $\tau : \{0,1\}^n \rightarrow \{0,1\}$  of their inputs, i.e.  $\tau(x_1, \dots, x_n) = \text{true}$  iff  $\sum_{i=1}^n c_i x_i \leq T$  where  $c_1, \dots, c_n, T \in \{-M, -M+1, \dots, M\}$ . W.l.o.g. we can assume that the inputs are only connected to nodes in the first layer, and first layer nodes are only connected to second layer nodes. The gates in the other layers compute deMorgan gates, i.e., OR and AND gates. The size of  $F$  is considered to be the *number of gates in the first plus second layer*. Note that the complexity of the functions in the first layer and their number of incoming wires, i.e. the number of leaves in the tree, do not count towards the size of  $F$ .

**$\mathcal{F}_2$ -FORMULA-SAT**

*Input:* Formula  $F = F(x_1, \dots, x_n)$  of size  $s$  with  $n$  inputs from the class  $\mathcal{F}_2$  with coordinates bounded by  $M \leq 2^{O(n)}$   
*Question:* Exist  $x_1, \dots, x_n \in \{0,1\}$  such that  $F(x_1, \dots, x_n) = \text{true}$ ?  
*Complexity:*  $O(2^n s \cdot n^{o(1)})$ , even restricted to  $s \leq n^{1+o(1)}$

The  $O(2^n s \cdot n^{o(1)})$  algorithm follows from Lemma B.3 and the algorithm for INEQ-FORMULA-PAIR below, again noting that we lose one log-factor due to the change in the machine model. Our tight barriers for shaving logs for Fréchet will be based on the assumption that  $\mathcal{F}_2$ -Formula-SAT cannot be solved in  $2^n/n^\varepsilon$ , for some  $\varepsilon > 0$ , even when  $s = n^{1+o(1)}$ . Since  $\mathcal{F}_2$  is even more expressive than  $\mathcal{F}_1$ , such upper bounds seem out of reach.

*Remark:* Any formula in  $\mathcal{F}_2$  (of size  $s$  over  $n$  variables) can be transformed into a formula in  $\mathcal{F}_1$  over the same variables of size  $\text{poly}(s)$ , since arithmetic has efficient circuits. However, since we care about polynomial factors in  $s$  in this paper, the difference between  $\mathcal{F}_1$  and  $\mathcal{F}_2$  is non-negligible for our purposes. In fact, the fastest running times that we know for  $\mathcal{F}_1$ -Formula-SAT and  $\mathcal{F}_2$ -Formula-SAT differ by a factor  $n^{1+o(1)}$ .

We will work with an intermediate “pair finding” problem:

**INEQ-FORMULA-PAIR**

*Input:* Formula  $F = F(x_1, \dots, x_m)$  where each input is used exactly once, and  $A, B \subseteq \{-M, \dots, M\}^m$  of size  $n$ , where  $M \leq n^{O(1)}$   
*Question:* Exist  $a \in A, b \in B$  such that  $F([a_1 \leq b_1], \dots, [a_m \leq b_m]) = \text{true}$ ?  
*Complexity:*  $O(n^2 m \frac{\log \log n}{\log n})$ , even restricted to  $m \leq (\log n)^{1+o(1)}$

Similarly to the reductions above, we can prove the following. Similar reductions have been used in [100, 65].

**Lemma B.3.** *An instance of  $\mathcal{F}_2$ -Formula-SAT on a formula of size  $s$  over  $n$  inputs can be reduced to an instance of INEQ-FORMULA-PAIR on two sets of size  $O(2^{n/2})$  and a formula of size  $m = O(s)$ , in linear time.*

The claimed running time  $O(n^2 m \frac{\log \log n}{\log n})$  follows by combining our reduction from INEQ-FORMULA-PAIR to the Fréchet distance with the fastest known algorithm for deciding the Fréchet distance [9]. We think that one can also obtain this running time more directly by “Four Russians” tricks, as have been used for LCS for large alphabet size. The further tricks that shaved a second log-factor off the running time of FORMULA-PAIR do not seem to work for INEQ-FORMULA-PAIR.

## B.1 Algorithm for Formula-Pair

**Theorem B.4.** FORMULA-PAIR can be solved in time  $O(n^2m/\log^2 n + n^2/\log n)$ .

The main trick used to prove the above theorem is the following decomposition of formulas.

**Lemma B.5.** Given a deMorgan formula  $F$  of size  $m$  and a number  $L \geq 2$ , in polynomial time we can compute a decomposition  $\mathcal{D} = \{F_1, \dots, F_k\}$  such that

1. each  $F_i$  is a subformula of  $F$ , i.e., a set of gates of  $F$  forming a subtree,
2. for each wire  $e$  of  $F$ , exactly one  $F_i$  contains both endpoints of  $e$ ,
3. each  $F_i$  contains less than  $3L$  gates, and
4.  $|\mathcal{D}| \leq 4m/L + 1$ .

Furthermore, call a gate in  $F_i$  special if it is the root of another subformula  $F_j$ ,  $j \neq i$ . In our decomposition each  $F_i$  has at most 2 special gates.

Similar decompositions were known before, see, e.g., [51], however, we are not aware of a decomposition bounding the number of special gates by a constant.

*Proof.* Note that we can assume  $F$  to have indegrees bounded by 2. Initialize  $\mathcal{D} = \emptyset$ . We assign to each gate  $g$  of  $F$  a weight  $w(g)$  in  $\{1, L\}$ ; initially all weights are 1, and inner nodes will always have weight 1. Now repeatedly perform the following procedure. Start at the root of  $F$  and repeatedly go to the child whose subtree has larger total weight, until reaching a gate  $g$  whose subtree has total weight in  $[\frac{5}{4}L, 3L)$ . Add the subformula computed by  $g$  to  $\mathcal{D}$ . Finally, remove all gates below  $g$  from  $F$  and set the weight of  $g$  to  $L$ . Repeat this procedure until the total weight of  $F$  is less than  $\frac{5}{4}L$ . Add the remaining formula  $F$  to  $\mathcal{D}$ . This defines a decomposition  $\mathcal{D}$ .

Note that we always find a gate whose subtree has total weight in  $[\frac{5}{4}L, 3L)$ . Indeed, if the current gate  $g$  has total subtree weight  $w \geq 3L$ , then  $g$  is an inner node and thus has weight 1, and  $g$ 's children together have total subtree weight at least  $w - 1$ . Since we go to the heavier child, we obtain total subtree weight at least  $(w - 1)/2 \geq (3L - 1)/2 \geq \frac{5}{4}L$ , where we used  $L \geq 2$ . Hence, we will land in the interval  $[\frac{5}{4}L, 3L)$ .

Also note that  $|\mathcal{D}| \leq 4m/L + 1$ , since in each call of the procedure we remove weight at least  $\frac{5}{4}L$  (by deleting the gates below  $g$ ) and then increase the weight of the chosen gate  $g$  by at most  $L$ , so we lose weight at least  $L/4$  in every call (except possibly in the very last step).

Let  $\mathcal{D} = \{F_1, \dots, F_k\}$ . Each subformula  $F_i$  of  $F$  has total weight less than  $3L$ , and thus it consists of less than  $3L$  gates. Moreover, the gates  $g$  in  $F_i$  with weight  $L$  exactly correspond to the special gates of  $F_i$ . Since each such gate  $g$  has weight  $L$ , there are less than  $3L/L$  special gates, and thus each  $F_i$  has at most 2 special gates.  $\square$

*Proof of Theorem B.4.* To solve FORMULA-PAIR in time  $O(n^2m/\log^2 n + n^2/\log n)$ , we use the ‘‘Four Russians’’ trick as well as packing  $\log(n)$  bits in a word. Let  $\mathcal{D} = \{F_1, \dots, F_k\}$  be a decomposition of  $F$  as in Lemma B.5, where  $L := \varepsilon \log n$ , for a sufficiently small constant  $\varepsilon > 0$ . Note that each subformula  $F_i$  has size  $O(\varepsilon \log n)$  and thus contains  $O(\varepsilon \log n)$  input gates of  $F$ .

For each subformula  $F_i$  and each assignment  $\beta$  to the  $O(\varepsilon \log n)$   $b$ -variables appearing in  $F_i$ , let  $F_i^\beta$  be the simplified subformula after fixing the  $b$ -variables. Compute the value of  $F_i^\beta$  on all vectors in  $A$  and all possible assignments to the special gates in  $F_i$ . We store these values in memory as

follows. For each  $1 \leq \ell \leq n$  divisible by  $L$  and all length- $L$  bitstrings  $x, x'$ , store the values of  $F_i^\beta$  on the  $\ell + j$ -th vector in  $A$ , with the input of the special gates set to  $x_j, x'_j$  (recall that there are at most 2 special gates). Store these  $L$  bits, for  $0 \leq j < L$ , in one word. We thus stored for each of  $O(m/\log n)$  subformulas, each of  $n^{O(\epsilon)}$  assignments to the  $b$ -variables, each of  $O(n/L)$  bundles of vectors in  $A$  (indexed by  $\ell$ ), and each of  $n^{O(\epsilon)}$  assignments to the special gates, one word in memory. Computing these values takes time  $mn^{1+O(\epsilon)}$  and is thus negligible.

Now we can evaluate  $F$  on all pairs  $a \in A, b \in B$  as follows. Iterate over all  $b \in B$ , and over all bundles of  $L$  consecutive vectors  $A' \subseteq A$  (indexed by  $\ell$ , divisible by  $L$ ). Iterate over the subformulas  $F_i$  in a topological order (from bottom to top). For each subformula  $F_i$ , if  $F_i$  has no special gates then access  $F_i^\beta$ , the simplification of  $F_i$  under  $b$ , and determine by one table lookup the output of  $F_i^\beta$  on all vectors  $a \in A'$ . If  $F_i$  does have special gates, then it depends on the output of some subformulas  $F_j$ , for which we have already computed the output on  $b$  and any  $a \in A'$ . Plugging in these outputs as the vectors  $x, x'$  above, we can again determine by one table lookup the output of  $F_i^\beta$  on all vectors  $a \in A'$ . Repeating this for all subformulas  $F_i$  eventually yields the output of  $F$  on  $b$  and any  $a \in A'$ .

Note that for each vector  $b$  and bundle  $A'$  this procedure takes time  $O(|\mathcal{D}|) = O(m/\log n + 1)$ . Since there are  $n$  vectors  $b$  and  $O(n/\log n)$  bundles  $A'$ , the total running time is  $O(n^2m/\log^2 n + n^2/\log n)$ . In the end, we check whether the output on any pair  $a \in A, b \in B$  is 1 to decide the given FORMULA-PAIR instance.  $\square$

## C Missing Details in the LCS Proof

This section contains the missing details for completing the reduction from FORMULA-PAIR to LCS that we presented in Section 2.

We will make use of the fact that for any strings  $x, y$ , symbol  $c$ , and  $k \geq 1$  we have  $\text{LCS}(1^k x, 1^k y) = \text{LCS}(x1^k, y1^k) = k + \text{LCS}(x, y)$ , which allows to greedily match prefixes and suffixes (see, e.g., [29, Fact 7.1.(1)]).

*Proof of Lemma 2.5.* Clearly, we have  $L \geq \text{LCS}(0^\beta 1^\beta, 0^\beta 1^\beta) + \text{LCS}(x'_1, y'_1) + \text{LCS}(2^\beta 3^\beta, 2^\beta 3^\beta) = 4\beta + L'_1$  and  $L \geq \text{LCS}(2^\beta 3^\beta, 2^\beta 3^\beta) + \text{LCS}(x'_2, y'_2) + \text{LCS}(0^\beta 1^\beta, 0^\beta 1^\beta) = 4\beta + L'_2$ . For the other direction, note that no LCS of  $x, y$  can match symbols both in the  $0^\beta 1^\beta$ -prefix of  $x$  and in the  $2^\beta 3^\beta$ -prefix of  $y$ , since these prefixes use disjoint alphabets and thus any such matchings would cross. Hence,

$$L \leq \max \{ \text{LCS}(x'_1 2^\beta 3^\beta x'_2 0^\beta 1^\beta, y), \text{LCS}(x, y'_2 0^\beta 1^\beta y'_1 2^\beta 3^\beta) \}. \quad (1)$$

Consider the first term,  $\text{LCS}(\tilde{x}, y)$  for  $\tilde{x} = x'_1 2^\beta 3^\beta x'_2 0^\beta 1^\beta$ . By the same argument as above, no LCS of  $\tilde{x}, y$  can match symbols both in the  $0^\beta 1^\beta$ -suffix of  $\tilde{x}$  and in the  $2^\beta 3^\beta$ -suffix of  $y$ . Thus,

$$\text{LCS}(\tilde{x}, y) = \max \{ \text{LCS}(x'_1 2^\beta 3^\beta x'_2, y), \text{LCS}(\tilde{x}, 2^\beta 3^\beta y'_2 0^\beta 1^\beta y'_1) \}. \quad (2)$$

Note that  $\text{LCS}(x'_1 2^\beta 3^\beta x'_2, y) \leq \text{LCS}(2^\beta 3^\beta, y) + \text{LCS}(x'_1 x'_2, y) = 2\beta + \text{LCS}(x'_1 x'_2, y)$ . We bound  $\text{LCS}(x'_1 x'_2, y)$  by summing up the LCS length of all pairs of a part  $x' \in \{x'_1, x'_2\}$  in the first string and a part  $y' \in \{2^\beta, 3^\beta, y'_2, 0^\beta, 1^\beta, y'_1, 2^\beta, 3^\beta\}$  in  $y$ , noting that we can either match  $x'_1$  and  $y'_1$  or  $x'_2$  and  $y'_2$ , so that the two pairs  $(x'_1, y'_1), (x'_2, y'_2)$  together contribute at most  $\max\{L'_1, L'_2\}$ . This yields

$$\text{LCS}(x'_1 x'_2, y) \leq \max\{L'_1, L'_2\} + 2 \cdot \beta/8 + 12 \cdot \beta/48 \leq \beta/2 + \max\{L'_1, L'_2\}.$$

Together we obtain, as desired,

$$\text{LCS}(x'_1 2^\beta 3^\beta x'_2, y) \leq \frac{5}{2}\beta + \max\{L'_1, L'_2\}. \quad (3)$$

For the other term  $\text{LCS}(\tilde{x}, \tilde{y}) = \text{LCS}(\tilde{x}, 2^\beta 3^\beta y'_2 0^\beta 1^\beta y'_1)$  of equation (2), let  $z$  be an LCS and consider the number  $k$  of symbols in the  $2^\beta 3^\beta$ -prefix of  $\tilde{y}$  matched by  $z$ . Note that we can bound  $\text{LCS}(\tilde{x}, \tilde{y}) \leq k + \text{LCS}(\tilde{x}, y'_1 0^\beta 1^\beta y'_2) \leq k + \text{LCS}(x, y'_1 0^\beta 1^\beta y'_2)$ . By symmetry, we can bound  $\text{LCS}(x, y'_1 0^\beta 1^\beta y'_2)$  in the same way as equation (3) by  $\frac{5}{2}\beta + \max\{L'_1, L'_2\}$ . Thus, for  $k \leq \frac{3}{2}\beta$  we obtain  $\text{LCS}(\tilde{x}, \tilde{y}) \leq 4\beta + \max\{L'_1, L'_2\}$ , as desired.

In the remaining case  $k > \frac{3}{2}\beta$ , we argue that we can split  $\text{LCS}(\tilde{x}, \tilde{y}) = \text{LCS}(x'_1 2^\beta 3^\beta, 2^\beta 3^\beta) + \text{LCS}(x'_2 0^\beta 1^\beta, y'_2 0^\beta 1^\beta y'_1)$ . Note that at least  $k - \beta > \beta/2$  symbols 2 of the  $2^\beta 3^\beta$ -prefix of  $\tilde{y}$  are matched. Since  $x'_1$  and  $x'_2$  contain at most  $\beta/48$  symbols 2 by assumption, at least one symbol of the  $2^\beta$ -block in  $\tilde{y}$  is matched to the  $2^\beta$ -block in  $\tilde{x}$ . Similarly, it follows that at least one symbol of the  $3^\beta$ -block in  $\tilde{y}$  is matched to the  $3^\beta$ -block in  $\tilde{x}$ . Thus, the LCS  $z$  splits between  $2^\beta$  and  $3^\beta$ , i.e., we have

$$\text{LCS}(\tilde{x}, \tilde{y}) = \text{LCS}(x'_1 2^\beta, 2^\beta) + \text{LCS}(3^\beta x'_2 0^\beta 1^\beta, 3^\beta y'_2 0^\beta 1^\beta y'_1).$$

Using greedy prefix matching, we can remove the prefix  $3^\beta$  of the second term to get

$$\begin{aligned} \text{LCS}(\tilde{x}, \tilde{y}) &= \text{LCS}(x'_1 2^\beta, 2^\beta) + \text{LCS}(3^\beta, 3^\beta) + \text{LCS}(x'_2 0^\beta 1^\beta, y'_2 0^\beta 1^\beta y'_1) \\ &= \text{LCS}(x'_1 2^\beta 3^\beta, 2^\beta 3^\beta) + \text{LCS}(x'_2 0^\beta 1^\beta, y'_2 0^\beta 1^\beta y'_1). \end{aligned}$$

Hence, there exists an LCS of  $\tilde{x}, \tilde{y}$  matching only symbols in the prefixes  $x'_1 2^\beta 3^\beta$  and  $2^\beta 3^\beta$  and symbols in the suffixes  $x'_2 0^\beta 1^\beta$  and  $y'_2 0^\beta 1^\beta y'_1$ . By symmetry, we can also split off the suffixes  $0^\beta 1^\beta$  and  $0^\beta 1^\beta y'_2$ . It follows that

$$\text{LCS}(\tilde{x}, \tilde{y}) = \text{LCS}(x'_1 2^\beta 3^\beta, 2^\beta 3^\beta) + \text{LCS}(x'_2, y'_2) + \text{LCS}(0^\beta 1^\beta, 0^\beta 1^\beta y'_2) = 4\beta + L'_2,$$

as desired. This finishes the analysis of the first term in equation (1). The second term gives the symmetric bound  $4\beta + L'_1$ , so in total we obtain the desired bound  $L \leq 4\beta + \max\{L'_1, L'_2\}$ .  $\square$

*Proof of Lemma 2.6.* In the base case we construct strings of length 1, so that  $n(g), \rho(g) = 1$ .

For an inner gate  $g$ , we recursively construct strings of length  $n_1, n_2$  and LCS bounds  $\rho_1, \rho_2$  for the children on subformulas  $F_1, F_2$ . If  $g$  is an AND gate we have

$$\begin{aligned} n(g) &= n_1 + n_2 + 2\beta = n_1 + n_2 + 2\lceil (n_1 + n_2)/\tau^2 \rceil \leq (n_1 + n_2)(1 + 2/\tau^2) + 2, \\ \rho(g) &= \rho_1 + \rho_2 + 2\beta \leq \rho_1 + \rho_2 + 2(n_1 + n_2)/\tau^2 + 2. \end{aligned}$$

If  $g$  is an OR gate we pad the strings to the same LCS length, thus increasing  $n(g)$  by  $|\rho_1 - \rho_2|$ ,

$$\begin{aligned} n(g) &= n_1 + n_2 + |\rho_1 - \rho_2| + 6\beta \leq (n_1 + n_2)(1 + 6/\tau^2) + 6 + \rho_1 + \rho_2, \\ \rho(g) &= 4\beta + \max\{\rho_1, \rho_2\} \leq \rho_1 + \rho_2 + 4(n_1 + n_2)/\tau^2 + 4. \end{aligned}$$

In both cases, we have

$$\begin{aligned} n(g) &\leq (n_1 + n_2)(1 + 6/\tau^2) + 6 + \rho_1 + \rho_2, \\ \rho(g) &\leq \rho_1 + \rho_2 + 4(n_1 + n_2)/\tau^2 + 4. \end{aligned}$$

Plugging in the inductive hypothesis yields, since  $F_1, F_2$  have depth at most  $\text{depth}(F_g) - 1$ ,

$$n(g) \leq 6\tau(|F_1| + |F_2|)(1 + 7/\tau)^{\text{depth}(F_g)-1}(1 + 6/\tau^2) + 6 + 6(|F_1| + |F_2|)(1 + 7/\tau)^{\text{depth}(F_g)-1}.$$

We use the fact  $|F_g| = |F_1| + |F_2| + 1$  to cancel the additive term  $+6$  at the cost of increasing the first  $|F_1| + |F_2|$  to  $|F_g|$ . We obtain

$$\begin{aligned} n(g) &\leq 6\tau|F_g|(1 + 7/\tau)^{\text{depth}(F_g)-1}(1 + 6/\tau^2) + 6|F_g|(1 + 7/\tau)^{\text{depth}(F_g)-1} \\ &= \frac{1 + 6/\tau^2 + 1/\tau}{1 + 7/\tau} \cdot 6\tau|F_g|(1 + 7/\tau)^{\text{depth}(F_g)}. \end{aligned}$$

Since  $\tau = (\log \sigma)^{1/4} \geq 1$ , as  $\sigma$  is sufficiently large, we have  $1 + 6/\tau^2 + 1/\tau \leq 1 + 7/\tau$ , which yields the desired bound. For  $\rho(g)$  we similarly obtain

$$\rho(g) \leq 6(|F_1| + |F_2|)(1 + 7/\tau)^{\text{depth}(F_g)-1} + 4 \cdot 6\tau(|F_1| + |F_2|)(1 + 7/\tau)^{\text{depth}(F_g)-1}/\tau^2 + 4$$

Again we cancel the additive  $+4$  by increasing  $|F_1| + |F_2|$  to  $|F_g|$ . We obtain

$$\begin{aligned} \rho(g) &\leq 6|F_g|(1 + 7/\tau)^{\text{depth}(F_g)-1} + 4 \cdot 6\tau|F_g|(1 + 7/\tau)^{\text{depth}(F_g)-1}/\tau^2 \\ &= \frac{1 + 4/\tau}{1 + 7/\tau} \cdot 6|F_g|(1 + 7/\tau)^{\text{depth}(F_g)}, \end{aligned}$$

which yields the desired bound and finishes the proof.  $\square$

## C.1 The Outer OR - Finalizing the Proof

We now finish the reduction from Formula-Pair to LCS using alignment gadgets [2, 29]. The following result is implicit in [2, Lemmas 3 and 4], and a similar (independently found) result is contained in [29] under the name of ‘‘alignment gadget’’.

**Lemma C.1.** *Let  $x_1, \dots, x_n$  and  $y_1, \dots, y_n$  be strings of length  $\lambda$ . Let  $\rho \geq \text{LCS}(x_i, y_j)$  for all  $i, j$ . Then one can construct strings  $x, y$  of length  $O(n\lambda)$  and an integer  $\rho'$  such that (1)  $\text{LCS}(x, y) \leq \rho'$  and (2)  $\text{LCS}(x, y) = \rho'$  if and only if there exist  $i, j$  with  $\text{LCS}(x_i, y_j) = \rho$ . If the strings  $x_i, y_j$  use alphabet  $\Sigma$ , then  $x, y$  use alphabet size at most  $|\Sigma| + 4$ .*

For any vectors  $a \in A, b \in B$ , invoking Lemma 2.3 for the output gate of  $F$ , we obtain strings  $x(a), y(b)$  and a number  $\rho$  such that (1)  $\text{LCS}(x(a), y(b)) \leq \rho$  and (2)  $\text{LCS}(x(a), y(b)) = \rho$  if and only if  $F(a, b)$  evaluates to true. Tracing the proof of Lemma 2.3 one can see that  $\rho$  does not depend on the choice of  $a, b$ , and similarly  $x(a)$  does not depend on  $b$  and  $y(b)$  does not depend on  $a$ . Hence, we obtain  $n$  strings  $x(a), a \in A$  and  $n$  strings  $y(b), b \in B$  and a threshold  $\rho$  on which we can invoke Lemma C.1. This yields strings  $x, y$  and a threshold  $\rho'$  such that (1)  $\text{LCS}(x, y) \leq \rho'$  and (2)  $\text{LCS}(x, y) = \rho'$  if and only if there exist  $a \in A, b \in B$  such that  $F(a, b)$  evaluates to true, so indeed from the LCS of the constructed strings  $x, y$  one can solve the Formula-Pair instance.

The strings  $x, y$  use alphabet size  $5\sigma^2 + 4$ . Using the length bound of Lemma 2.3,  $x, y$  have length  $O(n\tau|F|(1 + 7/\tau)^{\text{depth}(F)})$ . Here, as in Lemma 2.3, we may choose any sufficiently large  $\sigma > 0$  and set  $\tau = (\log \sigma)^{1/4}$ .

Since a priori the depth of  $F$  could be as large as  $|F|$ , the factor  $(1 + 7/\tau)^{\text{depth}(F)}$  in our length bound is not yet satisfactory. Thus, as a preprocessing before the above construction, we decrease the depth of  $F$  using the following result of Bonet and Buss [90, 25].

**Theorem C.2** ([25]). *Let  $F$  be a formula. Then for all  $k \geq 2$  there is an equivalent formula  $F'$  with depth at most  $(3k \ln 2) \log |F|$  and size  $|F'| \leq |F|^{1+1/(1+\log(k-1))}$ . This formula can be constructed in time  $|F|^{O(1)}$ .*

Given an instance  $(F, A, B)$  of Formula-Pair, we first run Bonet-Buss with  $k := \sqrt{\tau}$  to obtain a formula  $F'$ . Then we use the reduction to LCS on  $F'$ . This yields an equivalent LCS instance  $(x, y)$  of length  $O(n\tau|F'|^{1+1/(1+7/\tau)^{\text{depth}(F')}})$ , as shown above. Plugging in the bounds on  $|F'|$  from Bonet-Buss, we obtain length

$$O\left(n|F|^{1+1/(1+\log(k-1))}(1+7/\tau)^{(3k \ln 2) \log |F|}\right) = O\left(n|F|^{1+1/(1+\log(k-1))+O(k/\tau)}\right),$$

where we used  $1+x \leq \exp(x) = 2^{O(x)}$ . Since we set  $k = \sqrt{\tau}$ , this evaluates to  $O(n|F|^{1+O(1/\log \tau)})$ . Since furthermore  $\tau = (\log \sigma)^{1/4}$ , we obtain length bound  $O(n|F|^{1+O(1/\log \log \sigma)})$ . This finishes the proof of Theorem 2.2 for any sufficiently large alphabet size  $5\sigma^2 + 4$ . Note that it suffices to prove Theorem 2.2 for sufficiently large alphabet size, since for any constant  $C \geq 2$  and any alphabet size  $|\Sigma| \leq C$  the statement follows from the reduction in [3], which has length bound  $O(n|F|^{O(1)})$  (with an explicit constant  $O(1)$  in the exponent).

Now assume that LCS on strings of length  $n$  over sufficiently large alphabet size  $\sigma$  has an algorithm in time  $T(n) = O(n^2/(\log n)^{2+C/\log \log \sigma + \varepsilon})$ , for some sufficiently large  $C > 0$  and any (small)  $\varepsilon > 0$ . Then from the above reduction we obtain an algorithm for Formula-Pair on formulas of size  $|F| \leq (\log n)^{1+o(1)}$  in time

$$T\left(n(\log n)^{1+O(1/\log \log \sigma)+o(1)}\right) = O\left(n^2(\log n)^{2+O(1/\log \log \sigma)+o(1)}/(\log n)^{2+C/\log \log \sigma + \varepsilon}\right),$$

which is bounded by  $O(n^2/(\log n)^{\varepsilon/2})$ . This would break the Formula-Pair barrier, so we obtain a conditional lower bound for LCS of  $n^2/(\log n)^{2+O(1/\log \log \sigma)+o(1)}$ . This implies Corollary 1.6.

## D Simple Reduction to LCS using Large Alphabet

In this section we present a simpler alternative to the reduction from FORMULA-PAIR to LCS given in Sections 2 and C, using however a much larger alphabet size.

**Theorem D.1.** *FORMULA-PAIR on formulas of size  $s$  and lists of size  $n$  can be reduced to an instance of LCS on two strings over alphabet of size  $\Theta(s)$  of length  $O(n \cdot s \cdot 2^{O(\sqrt{\log s})})$ , in linear time.*

Similarly as before, we construct strings  $x, y$  and a number  $\rho$  such that  $\text{LCS}(x, y) \geq \rho$  holds if and only if the given Formula-Pair instance  $(F, A, B)$  is satisfiable. Fix vectors  $a, b \in \{0, 1\}^m$  (where  $2m$  is the number of inputs to  $F$  and thus  $m = \Theta(s)$ ). For any gate  $g$  of  $F$ , we construct strings as follows.

**Lemma D.2.** *We can inductively construct, for each gate  $g$  of  $F$ , strings  $x(g) = x(g, a)$  and  $y(g) = y(g, b)$  over alphabet size  $|F_g|$  (where  $F_g$  is the subformula of  $F$  below  $g$ ) and a number  $\rho(g)$  such that for  $L(g) := \text{LCS}(x(g), y(g))$  we have  $L(g) \leq \rho(g)$ , with equality if and only if gate  $g$  evaluates to true on input  $(a, b)$  to  $F$ . Moreover, we have  $|x(g)| = |y(g)| = n(g) \leq (\text{depth}(F_g) + 1) \cdot |F_g|$ .*

*Proof.* Consider any gate  $g$  of  $F$ , and write for readability  $x = x(g)$  and similarly define  $y, n, L, \rho$ . If  $g$  has fanin 2, write  $g_1, g_2$  for the children of  $g$ . Moreover, let  $x_1 = x(g_1)$  and similarly define  $y_1, n_1, L_1, \rho_1$  and  $x_2, y_2, n_2, L_2, \rho_2$ .

**Input Gate** The base case is an input bit  $a_i$  to  $F$  (input bits  $b_j$  are symmetric). Interpreting  $a_i$  as a string of length 1 over alphabet  $\{0, 1\}$ , note that  $\text{LCS}(a_i, 1) = a_i$ . Hence, the strings  $x = a_i$  and  $y = 1$ , with  $n = \rho = 1$ , trivially simulate the input bit  $a_i$ .

**AND Gates** For a gate  $g = (g_1 \wedge g_2)$  we first replace all symbols in  $x_2, y_2$  by fresh symbols, so that  $x_1 y_1$  and  $x_2 y_2$  use disjoint alphabets. We then construct

$$\begin{aligned} x &:= x_1 x_2 \\ y &:= y_1 y_2. \end{aligned}$$

Since these strings use disjoint alphabets, it is easy to see that  $\text{LCS}(x, y) = \text{LCS}(x_1, y_1) + \text{LCS}(x_2, y_2)$ . Hence, inductively we obtain for  $\rho := \rho_1 + \rho_2$  that  $L = \text{LCS}(x, y) \leq \rho$ , with equality if and only if both gates  $g_1$  and  $g_2$  evaluate to true. Thus, we correctly implemented an AND gate. Note that  $n = n_1 + n_2$ .

**OR Gates** For a gate  $g = (g_1 \vee g_2)$  we first replace all symbols in  $x_2, y_2$  by fresh symbols, so that  $x_1 y_1$  and  $x_2 y_2$  use disjoint alphabets. Without loss of generality assume that  $\rho_1 \geq \rho_2$  (the other case is symmetric). Let  $\delta := \rho_1 - \rho_2$  and let  $\$$  be a fresh symbol. We construct the strings

$$\begin{aligned} x &:= x_1 \quad x_2 \$^\delta \\ y &:= y_2 \$^\delta y_1. \end{aligned}$$

Since  $x_1, y_1$  and  $x_2 \$^\delta, y_2 \$^\delta$  use disjoint alphabets, we can either match symbols in the first pair or in the second. It follows that  $\text{LCS}(x, y) = \max\{\text{LCS}(x_1, y_1), \text{LCS}(x_2 \$^\delta, y_2 \$^\delta)\}$ . Since  $\$$  is a fresh symbol, we moreover have  $\text{LCS}(x_2 \$^\delta, y_2 \$^\delta) = \text{LCS}(x_2, y_2) + \delta$ . Together, we obtain  $L = \text{LCS}(x, y) \leq \max\{\rho_1, \rho_2 + \delta\} = \rho_1$ , with equality if and only if  $g_1$  or  $g_2$  evaluate to true. Thus, we correctly implemented an OR gate. Note that  $n = n_1 + n_2 + |\rho_1 - \rho_2|$  and  $\rho = \max\{\rho_1, \rho_2\}$ .

**Analyzing the Length** Note that both for AND and OR gates, the inequality  $\rho \leq \rho_1 + \rho_2$  holds, with the base case  $\rho = 1$  at input gates. This yields  $\rho(g) \leq |F_g|$ . Furthermore, for both AND and OR gates the inequality  $n \leq n_1 + n_2 + |\rho_1 - \rho_2|$  holds. Since  $|\rho_1 - \rho_2| \leq \rho_1 + \rho_2 \leq |F_{g_1}| + |F_{g_2}| \leq |F_g|$ , we obtain  $n \leq n_1 + n_2 + |F_g|$ . It now follows inductively that  $n \leq (\text{depth}(F_g) + 1) \cdot |F_g|$ , since

$$\begin{aligned} n &\leq n_1 + n_2 + |F_g| \leq (\text{depth}(F_{g_1}) + 1) \cdot |F_{g_1}| + (\text{depth}(F_{g_2}) + 1) \cdot |F_{g_2}| + |F_g| \\ &\leq \text{depth}(F_g) \cdot (|F_{g_1}| + |F_{g_2}|) + |F_g| \leq (\text{depth}(F_g) + 1) \cdot |F_g|. \end{aligned}$$

Also note that since each gate can introduce one new symbol the alphabet size is at most  $|F_g|$ . This proves the lemma.  $\square$

**Outer OR** We now finish the reduction from Formula-Pair to LCS using alignment gadgets, see Lemma D.2. For any vectors  $a \in A, b \in B$ , invoking Lemma D.2 for the output gate of  $F$ , we obtain strings  $x(a), y(b)$  and a number  $\rho$  such that  $\text{LCS}(x(a), y(b)) \leq \rho$ , with equality if and only if  $F(a, b)$  evaluates to true. Tracing the proof of Lemma D.2 one can see that  $\rho$  does not depend on the choice of  $a, b$ , and similarly  $x(a)$  does not depend on  $b$  and  $y(b)$  does not depend on  $a$ . Hence, we obtain  $n$  strings  $x(a), a \in A$  and  $n$  strings  $y(b), b \in B$  and a threshold  $\rho$  on which we can invoke Lemma C.1. This yields strings  $x, y$  and a threshold  $\rho'$  such that  $\text{LCS}(x, y) \leq \rho'$ , with equality if

and only if there exist  $a \in A, b \in B$  such that  $F(a, b)$  evaluates to true, so indeed from the LCS of the constructed strings  $x, y$  one can solve the Formula-Pair instance. Note that so far the resulting strings have length  $O(n \cdot \text{depth}(F) \cdot |F|)$ .

**Finalizing the Proof** Since a priori the depth of  $F$  could be as large as  $|F|$ , the factor  $\text{depth}(F)$  in our length bound is not yet satisfactory. Thus, as a preprocessing before the above construction, we use the depth reduction of Bonet and Buss, see Theorem C.2. Given an instance  $(F, A, B)$  of Formula-Pair, we first run Bonet-Buss with  $k := 2^{\Theta(\sqrt{\log(|F|)})}$  to obtain a formula  $F'$ . Then we use the reduction to LCS on  $F'$ . This yields an equivalent LCS instance  $(x, y)$  of length  $O(n \cdot (\text{depth}(F') + 1) \cdot |F'|)$ , as shown above. Plugging in the bounds on  $|F'|$  from Bonet-Buss, we obtain length  $O(n \cdot |F| \cdot 2^{\Theta(\sqrt{\log(|F|)})}) = O(n \cdot |F|^{1+o(1)})$ . This finishes the proof of Theorem D.1.

## E The Reduction to Pattern Matching

In this section we reduce the FORMULA-PAIR problem to regular expression pattern matching. By the reductions in Section B, this gives us barriers based on  $\mathcal{F}_1$ -Formula-SAT and proves Theorem 1.5. Let us start with a formal definition of the problem.

**Regular Expression Pattern Matching** Regular expressions over alphabet  $\Sigma$  are inductively defined as follows. For any  $c \in \Sigma$ , the regular expression  $c$  matches only itself, i.e.,  $L(c) = \{c\}$ . The regular expression  $r = r_1 \circ r_2$  matches all concatenations of strings matched by  $r_1$  and  $r_2$ , i.e.,  $L(r) = \{s_1 s_2 \mid s_1 \in L(r_1), s_2 \in L(r_2)\}$ . We often abbreviate  $r_1 \circ r_2$  by  $r_1 r_2$ . The regular expression  $r = r_1 | r_2$  matches the union of all strings matched by  $r_1$  and  $r_2$ , i.e.,  $L(r) = L(r_1) \cup L(r_2)$ . Finally, the regular expression  $r = r_1^*$  matches all sequences of strings matched by  $r_1$ , i.e.,  $L(r) = \{s_1 \dots s_k \mid k \geq 0, s_1, \dots, s_k \in L(r_1)\}$ . The *regular expression pattern matching problem* is, given a string  $t$  (the text) and a regular expression  $p$  (the pattern), to decide whether  $p$  matches some substring of  $t$ .

We will prove the following theorem. Together with the reductions in Section B, this implies Theorem 1.5.

**Theorem E.1.** FORMULA-PAIR *on formulas of size  $s$  and lists of size  $n$  can be reduced to an instance of regular expression pattern matching with text length and pattern size  $O(n \cdot s \log s)$ , in linear time. The alphabet is  $\Sigma = \{0, 1\}$ .*

In our construction, we will use the following encoding of numbers. For any integer  $m \geq 0$ , let  $m_1 \dots m_k$  be its representation in binary, where  $k = \max\{1, \lfloor \log_2 m \rfloor + 1\}$ . We define the string

$$\#(m) := 110m_10m_20\dots 0m_{k-1}0m_k011.$$

We will make use of the following simple observation: For any  $x, y_1, \dots, y_\ell \geq 0$  the string  $\#(x)$  is a substring of the concatenation  $\#(y_1) \dots \#(y_k)$  if and only if we have  $x = y_i$  for some  $i$ . (In particular,  $\#(x)$  cannot appear as a suffix of  $\#(y_1)$  followed by a prefix of  $\#(y_2)$ .)

For each gate  $g$  in the formula  $F$  and vectors  $a, b \in \{0, 1\}^k$  we design a text  $t(g, a)$  and a pattern  $p(g, b)$  such that  $p(g, b)$  matches  $t(g, a)$  if and only if on input  $a, b$  in the formula  $F$  gate  $g$  evaluates to true. In our construction we ensure that for any gate  $g$  of height  $h$  the text  $t(g, a)$  can be written as  $\#(y_1) \dots \#(y_k)$  for some  $k \geq 1$  and  $0 \leq y_1, \dots, y_k \leq h + 1$ .

In the following, fix a gate  $g$  of  $F$  with children  $g_1$  and  $g_2$ . For simplicity, we write  $t = t(g, a), p = p(g, b)$  and similarly  $t_1, p_1, t_2, p_2$  for  $g_1$  and  $g_2$ . We denote by  $h$  the height of  $g$ , i.e., the length of the longest path from  $g$  to any of its descendants.

**Input Gates** The base case is an input bit  $a_i$  to  $F$  (input bits  $b_j$  are symmetric). Since  $a_i$  is a number in  $\{0, 1\}$ , we can set  $t := \#(a_i)$  and  $p := \#(1)$ . Then the pattern matches the text if and only if  $a_i = 1$ . Note that the height is 0 and we indeed only used  $\#(0)$  and  $\#(1)$ .

**AND Gates** Consider a gate  $g = (g_1 \wedge g_2)$  of height  $h \geq 1$  and let  $h' := h + 1$ . We set

$$\begin{aligned} t &:= t_1 \#(h') t_2, \\ p &:= p_1 \#(h') p_2. \end{aligned}$$

Since  $t_1$  and  $t_2$  only contain  $\#(y)$  with  $y < h'$ , the only way for  $p$  to match  $t$  is to match the occurrences of  $\#(h')$  in  $p$  and  $t$  in the natural way, so that  $p_1$  must match  $t_1$  and  $p_2$  must match  $t_2$ . Thus, we correctly implemented an AND gate.

**OR Gates** Consider a gate  $g = (g_1 \vee g_2)$  of height  $h \geq 1$  and let  $h' := h + 1$ . We set

$$\begin{aligned} t &:= t_1 \#(h') t_2, \\ p &:= p_1 \#(h') (0|1)^* \mid (0|1)^* \#(h') p_2. \end{aligned}$$

This time we can choose which of the two sub-expressions of  $p$  we want to match. In any case, there is an occurrence of  $\#(h')$  in the pattern, and since  $t_1$  and  $t_2$  only contain  $\#(y)$  with  $y < h'$ , the only way for the pattern to match  $t$  is to match the occurrences of  $\#(h')$  in  $p$  and  $t$  in the natural way. Since  $(0|1)^*$  matches any string, in particular  $t_1$  and  $t_2$ , we obtain that  $p_1$  matches  $t_1$  or  $p_2$  matches  $t_2$ , depending on which sub-expression of  $p$  we choose. Thus, we correctly implemented an OR gate.

This finishes the construction of the texts  $t(g, a)$  and patterns  $p(g, a)$ .

**Outer OR** Invoking this construction for the root  $r$  of  $F$  yields a text  $t(a) := t(r, a)$  and a pattern  $p(b) := p(r, b)$  such that  $p(b)$  matches  $t(a)$  if and only if on input  $a, b$  the formula  $F$  evaluates to true. Let  $H$  be two plus the height of  $r$ , and let  $A = \{a^1, \dots, a^n\}$  and  $B = \{b^1, \dots, b^n\}$ . We construct the final text and pattern as

$$\begin{aligned} t &= \#(H) t(a^1) \#(H) \dots \#(H) t(a^n) \#(H), \\ p &= \#(H) p(b^1) \#(H) \mid \dots \mid \#(H) p(b^n) \#(H). \end{aligned}$$

To match the pattern  $p$ , we have to choose one of the sub-expressions  $p^j := \#(H) p(b^j) \#(H)$ . Since  $\#(H)$  does not appear in any  $t(a^i)$ , in order to match  $p^j$  to a substring of  $t$  we must match the two occurrences of  $\#(H)$  in  $p^j$  to two top-level occurrences of  $\#(H)$  in  $t$ , and thus we must match  $p(b^j)$  to some  $t(a^i)$ . As this is possible if and only if  $F$  is satisfied by  $a^i, b^j$ , we have constructed an instance of regular expression pattern matching that is equivalent to the given instance of FORMULA-PAIR.

Note that the height of any gate is bounded by  $O(s)$ , where  $s$  is the size of the formula  $F$ , and thus any encoding  $\#(y)$  appearing in our construction has length  $O(\log s)$ . The text length and pattern length both satisfy a recursion of the type

$$L \leq L_1 + L_2 + O(\log s),$$

which yields the bound  $O(s \log s)$  on the length of  $t(a)$  and the size of  $p(b)$ . In total, we obtain length and size  $O(n \cdot s \log s)$ . This proves Theorem 1.5.

## F The Reduction to Fréchet Distance

In this section we reduce the INEQ-FORMULA-PAIR problem to the Fréchet distance. By the reductions in Section B, this gives us barriers based on  $\mathcal{F}_2$ -FORMULA-SAT and proves Theorem 1.4. Familiarity with the SETH-hardness proof of [27] will help to understand our proof. Let us start with a formal definition of the problem.

**Fréchet distance** A curve  $P$  is a sequence  $(p_1, \dots, p_n)$  of points in the Euclidean space  $\mathbb{R}^d$ . We call  $p_1, \dots, p_n$  the *vertices* of  $P$ . In the construction in this paper, we will always work in dimension  $d = 2$ , and the coordinates of vertices will be rationals with small bitlength. For two curves  $P = (p_1, \dots, p_n)$  and  $Q = (q_1, \dots, q_m)$  and any monotone non-decreasing and onto functions  $\phi: \{1, \dots, n+m\} \rightarrow \{1, \dots, n\}$  and  $\psi: \{1, \dots, n+m\} \rightarrow \{1, \dots, m\}$  we call  $\tau = (\phi, \psi)$  a *traversal*. We say that at time step  $t$  the traversal  $\tau$  is at vertex  $p_{\phi(t)}$  in  $P$  and at vertex  $q_{\psi(t)}$  in  $Q$ . Observe that  $\tau$  describes one way of a man and its dog walking along  $P$  and  $Q$ , respectively, from their starting vertices to their ending vertices, where in each time step man and dog may step to the next vertex on their curve or stay at their current vertex. We say that the traversal  $\tau$  *stays in distance*  $\delta$  if at any time step  $1 \leq t \leq n+m$  the current vertices  $p_{\phi(t)}$  and  $q_{\psi(t)}$  are within distance  $\delta$  of each other. Finally, the *Fréchet distance*  $d_F(P, Q)$  is the minimal number  $\delta$  such that there exists a traversal  $\tau$  of  $P, Q$  staying in distance  $\delta$ .

We will prove the following theorem. Together with the reductions in Section B, this implies Theorem 1.4.

**Theorem F.1.** *INEQ-FORMULA-PAIR on formulas of size  $s$  and lists of size  $n$  can be reduced to an instance of Fréchet on two curves of length  $O(n \cdot s)$ , in linear time.*

### F.1 Outline of the Reduction

We say that curves  $P, Q$  are  $\delta$ -placed if

- the  $x$ -coordinates of all vertices of  $P$  and  $Q$  are in  $[-\delta, \delta]$ ,
- the  $y$ -coordinates of all vertices of  $P$  are in  $[1 - \delta^2, 1 + \delta^2]$ , and
- the  $y$ -coordinates of all vertices of  $Q$  are in  $[-\delta^2, \delta^2]$ .

This ensures that  $P$  is contained in a small region around the point  $(1, 0)$ , and  $Q$  is close to  $(0, 0)$ .

**Lemma F.2.** *For each gate  $g$  in the formula  $F$  and vectors  $a, b \in \{-M, \dots, M\}^k$ , and any  $\delta \in (0, 1)$ , we design curves  $P = P_\delta(g, a)$  and  $Q = Q_\delta(g, b)$  such that*

- (P1)  $d_F(P, Q) \leq 1$  if and only if on input  $a, b$  in the formula  $F$  gate  $g$  evaluates to true, and
- (P2)  $P, Q$  are  $\delta$ -placed.

The curves  $P$  and  $Q$  have  $O(|F|)$  vertices. Coordinates of the vertices of  $P, Q$  use  $O(\log(1/\delta) + h + \log M)$  bits, where  $h$  is the height of gate  $g$ .

Property (P1) makes sure that we correctly simulate formula  $F$  by the Fréchet distance. We use property (P2) to gain control over the Fréchet distance of recursively defined subcurves. Before proving the above lemma, in Section F.2 we first design the “outer OR” of the reduction from INEQ-FORMULA-PAIR, which chooses a pair of vectors  $a \in A, b \in B$ . We then show how to implement the different types of gates in Section F.3, proving Lemma F.2.

## F.2 Outer OR

Let  $r$  be the root of formula  $F$  and consider the curves  $P_a := P_{1/16}(r, a)$  and  $Q_b := Q_{1/16}(r, b)$  for vectors  $a \in A, b \in B$  given by Lemma F.2. We want to construct curves  $P, Q$  such that  $d_F(P, Q) \leq 1$  if and only if there exist  $a \in A, b \in B$  with  $d_F(P_a, Q_b) \leq 1$ , which holds if and only if  $F(a, b)$  evaluates to true.

We study this situation slightly more generally as follows. Let  $\delta \in (0, 1/2]$  and let  $P_1, Q_1, \dots, P_k, Q_k$  be  $\delta/8$ -placed curves. Define the following auxiliary points:

$$\begin{aligned} s_P &:= (-\delta/2, 1 - \delta^2), & s_Q &:= (-\delta/2, \delta^2), & s_Q^* &:= (-\delta/2, -\delta^2), \\ t_P &:= (\delta/2, 1 - \delta^2), & t_Q &:= (\delta/2, \delta^2), & t_Q^* &:= (\delta/2, -\delta^2), \\ b_P &:= (-\delta/2, 1), & b_Q &:= (-\delta/2, 0), \\ e_P &:= (\delta/2, 1), & e_Q &:= (\delta/2, 0). \end{aligned}$$

We define the curve  $P = P_\delta^{OR}(P_1, \dots, P_k)$  as

$$P := \bigcirc_{\ell=1}^k s_P \circ b_P \circ P_\ell \circ e_P \circ t_P,$$

i.e. we start with vertices in  $s_P$  and  $b_P$ , then follow  $P_1$ , then add vertices at  $e_P, t_P, s_P$ , and  $b_P$ , then follow  $P_2$  and so on. We define  $Q = Q_\delta^{OR}(Q_1, \dots, Q_k)$  by

$$Q := s_Q \circ s_Q^* \circ \left( \bigcirc_{\ell=1}^k b_Q \circ Q_\ell \circ e_Q \right) \circ t_Q^* \circ t_Q.$$

**Lemma F.3** (Outer OR). *For any  $\delta \in (0, 1/2]$  and any  $\delta/8$ -placed curves  $P_1, Q_1, \dots, P_k, Q_k$ , the curves  $P = P_\delta^{OR}(P_1, \dots, P_k)$  and  $Q = Q_\delta^{OR}(Q_1, \dots, Q_k)$  constructed above are  $\delta$ -placed and satisfy  $d_F(P, Q) \leq 1$  if and only if there are  $1 \leq i, j \leq k$  with  $d_F(P_i, Q_j) \leq 1$ .*

Note that for  $\delta := 1/2$  the curves  $P_a = P_{1/16}(r, a)$  and  $Q_b = Q_{1/16}(r, b)$  satisfy the requirements of the above construction, in particular they are  $(\delta/8 = 1/16)$ -placed, and thus we obtain curves  $P, Q$  such that  $d_F(P, Q) \leq 1$  if and only if there exist  $a \in A, b \in B$  with  $d_F(P_a, Q_b) \leq 1$ . This yields a reduction from INEQ-FORMULA-PAIR on  $n$  vectors and formula  $F$  to the Fréchet distance on curves of length  $O(n|F|)$ , as  $P_a, Q_b$  have length  $O(|F|)$  by Lemma F.2. The hypothesis that for  $|F| = (\log n)^{1+o(1)}$  INEQ-FORMULA-PAIR has no  $n^2/(\log n)^\epsilon$  algorithm now implies that (the decision variant of) the Fréchet distance has no  $n^2/(\log n)^{2+\epsilon'}$  algorithm, as desired. Also note that the coordinates of the vertices of the resulting curves use  $O(\text{depth}(F) + \log M)$  bits. Since by Bonet and Buss’ depth reduction (Theorem C.2) we can assume that  $\text{depth}(F) = O_\epsilon(\log |F|)$  and  $|F| = (\log n)^{1+o(1)}$ , these coordinates fit into a memory cell for any  $M = n^{O(1)}$ . Thus, the reduction not only works on the Real RAM, but also on the Word RAM.

*Proof of Lemma F.3.* From the definition of the auxiliary points and since  $P_\ell, Q_\ell$  are  $\delta/8$ -placed for all  $\ell$ , it follows that  $P, Q$  are  $\delta$ -placed.

If there are  $i, j$  such that  $d_F(P_i, Q_j) \leq 1$  then we find a traversal of  $P, Q$  staying in distance 1 as follows. Denote by  $s_P^{(\ell)}$  the  $\ell$ -th occurrence of  $s_P$  on curve  $P$ , i.e. the copy of  $s_P$  that comes before  $P_\ell$  in  $P$ , and similarly define  $t_P^{(\ell)}, b_P^{(\ell)}, e_P^{(\ell)}, b_Q^{(\ell)}, e_Q^{(\ell)}$ . Note that  $s_Q, s_Q^*, t_Q, t_Q^*$  appear only once in  $Q$ , so they do not get a superscript. Any traversal of  $P, Q$  starts in  $s_P^{(1)}$  and  $s_Q$ . We stay at  $s_Q$  and walk along  $P$  until we reach  $s_P^{(i)}$ , staying in distance 1 by claim (U1) below. Then we stay in  $s_P^{(i)}$  and walk along  $Q$  until we reach  $b_Q^{(j)}$ , using (U2). Now we step to  $b_P^{(j)}$  in  $P$ , using (U4). After a simultaneous step in  $P, Q$  we are at the starting vertices of  $P_i, Q_j$  and we proceed by optimally traversing them, staying in distance 1 since  $d_F(P_i, Q_j) \leq 1$  by assumption. We then make a simultaneous step to  $e_P^{(i)}$  and  $e_Q^{(j)}$ , using (U4). Now we make a single step in  $P$  to  $t_P^{(i)}$ , using (U3). We stay in this vertex of  $P$  and walk along  $Q$  until its final vertex  $t_Q$ , using (U3). Finally, we stay in  $t_Q$  and walk along  $P$  until its final vertex  $t_P^{(k)}$ , using (U1). It remains to show the following.

**Claim F.4.** *We have the following upper bounds on vertex distances:*

- (U1)  $s_Q$  and  $t_Q$  are in distance 1 of any vertex of  $P$ ,
- (U2)  $s_P$  is in distance 1 of any vertex of  $Q$  except for  $t_Q^*$ ,
- (U3)  $t_P$  is in distance 1 of any vertex of  $Q$  except for  $s_Q^*$ ,
- (U4)  $b_P$  and  $b_Q$  as well as  $e_P$  and  $e_Q$  are in distance 1.

*Proof.* For (U1) consider  $s_Q$ , since  $t_Q$  is symmetric. Observe that every vertex on  $P$  has  $y$ -coordinate at most  $1 + (\delta/8)^2 \leq 1 + \delta^2/2$  (and at least  $1 - \delta^2$ ) and  $x$ -coordinate at most  $\delta/2$  (and at least  $-\delta/2$ ). Thus, their squared distance to  $s_Q = (-\delta/2, \delta^2)$  is at most

$$(-\delta/2 - \delta/2)^2 + ((1 + \delta^2/2) - \delta^2)^2 = \delta^2/16 + (1 - \delta^2 + \delta^4/4) \leq 1,$$

proving (U1). For (U2) we similarly use that all vertices on  $Q$ , except for  $s_Q^*$  and  $t_Q^*$  have  $y$ -coordinate at least  $-(\delta/8)^2$  to obtain that  $s_P$  is within distance 1 of these vertices. The distance between  $s_P$  and  $s_Q^*$  is exactly 1 by definition. The only exception is  $t_Q^*$ , and indeed the distance between  $s_P$  and  $t_Q^*$  is larger than 1. Thus,  $s_P$  is in distance 1 of any point on  $Q$  except for  $t_Q^*$ . Claim (U3) is symmetric, and (U4) follows directly from the definitions.  $\square$

For the other direction, assume that there is a traversal  $\tau$  of  $P, Q$  staying in distance 1. We want to show that  $d_F(P_i, Q_j) \leq 1$  holds for some  $i, j$ . Consider any time step where  $\tau$  is in  $s_Q^*$  on  $Q$  and some vertex  $x$  on  $P$ . By claim (L1) below,  $x$  has to be a vertex  $s_P^{(i)}$ , for some  $1 \leq i \leq k$ . We may stay in  $s_P^{(i)}$  for some steps, until we at some point in time make a step from  $s_P^{(i)}$  to  $b_P^{(i)}$ . We show that at this time we have to be in  $b_Q^{(j)}$  for some  $j$ . Consider the possible vertices of  $Q$  where we can be. By (L1) and (L2) the points  $s_Q^*$  and  $t_Q^*$  are too far from  $b_P^{(i)}$ . The vertices  $s_Q$  and  $t_Q$  are both in distance 1 of  $b_P^{(i)}$ , however, we already visited  $s_Q^*$ , and since  $s_Q$  comes before  $s_Q^*$  on  $Q$  we cannot be at  $s_Q$ . Moreover, since the point in time at which we were at  $(s_P^{(i)}, s_Q^*)$  we only visited the vertices  $s_P^{(i)}$  and  $b_P^{(i)}$  on  $P$ , and both have distance more than 1 to the bottleneck vertex  $t_Q^*$ , so we cannot walk to  $t_Q$ , which comes after  $t_Q^*$  on  $Q$ . Hence, while being at  $b_Q^{(i)}$  we cannot be at  $s_Q$  or  $t_Q$ . The vertices of  $Q_\ell$ , for any  $\ell$ , are also not in distance 1, by (L3). Finally, by (L5) we cannot

be at  $e_Q$ . Hence, while vertex at  $b_P^{(i)}$  we have to be at  $b_Q^{(j)}$  for some  $j$ . By (L3) and (L4), the next step is simultaneous in  $P$  and  $Q$ , since the distance from  $b_P$  to  $Q_j$  as well as from  $b_Q$  to  $P_i$  is larger than 1. After this simultaneous step we are at the starting vertices of  $P_i$  and  $Q_j$ . By (L3) and (L4), we cannot make a step to  $e_P$  while staying in  $Q_j$  (or to  $e_Q$  while staying in  $P_i$ ), since their distance is too large. Hence,  $\tau$  contains a traversal  $\tau'$  of  $P_i, Q_j$ , after which it makes a simultaneous step to  $e_P^{(i)}, e_Q^{(j)}$ . Since  $\tau$  stays in distance 1, also  $\tau'$  stays in distance 1, and we obtain  $d_F(P_i, Q_j) \leq 1$  (for some  $i, j$ ). It remains to show the following.

**Claim F.5.** *We have the following lower bounds on vertex distances:*

- (L1)  $s_Q^*$  has distance larger than 1 to any vertex of  $P$  except  $s_P$ ,
- (L2)  $t_Q^*$  has distance larger than 1 to any vertex of  $P$  except  $t_P$ ,
- (L3)  $b_P$  and  $e_P$  have distance larger than 1 to any vertex of  $Q_\ell$  for all  $\ell$ ,
- (L4)  $b_Q$  and  $e_Q$  have distance larger than 1 to any vertex of  $P_\ell$  for all  $\ell$ ,
- (L5)  $b_P$  and  $e_Q$  as well as  $e_P$  and  $b_Q$  have distance larger than 1,

*Proof.* For (L1), since  $s_Q^* = (-\delta/2, -\delta^2)$  and all vertices of  $P$  have  $y$ -coordinate at least  $1 - \delta^2$ , vertices in distance 1 to  $s_Q^*$  on  $P$  are of the form  $(-\delta/2, 1 - \delta^2)$ , which is  $s_P$ . (L2) is symmetric.

For (L3), since  $P_\ell, Q_\ell$  are  $\delta/8$ -placed, all vertices of  $Q_\ell$  have  $x$ -coordinate at least  $-\delta/8$  and  $y$ -coordinate at most  $(\delta/8)^2$ . Thus, the squared distance from any point on  $Q_\ell$  to  $b_P = (-\delta/2, 1)$  is

$$\geq (\delta/2 - \delta/8)^2 + (1 - \delta^2/64)^2 \geq \delta^2/16 + (1 - \delta^2/32 + \delta^4/4096) > 1.$$

The distance to  $e_P$  is symmetric, as is (L4). Finally, (L5) follows directly from the definitions.  $\square$

$\square$

### F.3 Implementing Gates

It remains to prove Lemma F.2, i.e. to implement the different types of gates. Let  $\delta \in (0, 1)$  and fix vectors  $a \in A, b \in B$ .

**Comparison Gate** First consider an input gate  $g$  to formula  $F$ , i.e., a comparison  $[a_i \leq b_i]$  for some  $i$ . We define  $P = P_\delta(g, a)$  and  $Q = Q_\delta(g, b)$  as degenerate curves consisting only of a single vertex. Specifically, we set  $P_\delta(g, a) := ((0, 1 + \delta^2 a_i/M))$  and  $Q_\delta(g, b) := ((0, \delta^2 b_i/M))$ .

**Lemma F.6** (Comparison Gate). *The curves  $P_\delta(g, a), Q_\delta(g, b)$  satisfy (P1) and (P2).*

*Proof.* Note that the Fréchet distance of degenerate curves consisting of one vertex is simply the Euclidean distance of the vertices. Thus,  $d_F(P, Q) = |1 + \delta^2 a_i/M - \delta^2 b_i/M|$ . Since  $\delta \in (0, 1)$  and  $a_i, b_i \in \{-M, \dots, M\}$ , this equals  $1 + \delta^2(a_i - b_i)/M$ , which is at most 1 if and only if  $a_i \leq b_i$ . This shows property (P1), and property (P2) holds by definition.  $\square$

**AND Gate** Now consider an AND gate  $g$  in  $F$ , and let  $g_1, \dots, g_k$  be the children of  $g$  in  $F$ . We (inductively) construct the curves  $P_\ell := P_{\delta'}(g_\ell, a)$  and  $Q_\ell := Q_{\delta'}(g_\ell, b)$  for  $1 \leq \ell \leq k$  and some  $\delta' \leq \delta/4$ . The only facts about  $P_\ell, Q_\ell$  that we will use are property (P1) and that  $P_\ell, Q_\ell$  are  $\delta/4$ -placed (by (P2) and  $\delta' \leq \delta/4$ ). We translate the curves  $P_\ell$  and  $Q_\ell$  along the  $x$ -axis by  $+\delta/2$ , if  $\ell$  is odd, and by  $-\delta/2$ , if  $\ell$  is even, i.e. we add  $\delta/2$  or  $-\delta/2$  to the  $x$ -coordinate of every vertex of  $P_\ell$  and  $Q_\ell$ . This results in curves  $P'_\ell, Q'_\ell$ . Note that translating both curves does not change the Fréchet distance and thus we have  $d_F(P'_\ell, Q'_\ell) \leq 1$  if and only if  $g_\ell$  evaluates to true, by (P1). We form the final curve  $P = P_\delta(g, a)$  by concatenating  $P'_1, \dots, P'_k$ , and similarly we form  $Q = Q_\delta(g, b)$  by concatenating  $Q'_1, \dots, Q'_k$ .

**Lemma F.7** (AND Gate). *The curves  $P_\delta(g, a), Q_\delta(g, b)$  satisfy (P1) and (P2).*

*Proof.* The curves  $P_\ell, Q_\ell$  are  $\delta/4$ -placed since  $\delta' \leq \delta/4$ . The translation by  $\pm\delta/2$  is sufficiently small to ensure that  $P'_\ell, Q'_\ell$  and thus also  $P, Q$  are  $\delta$ -placed, so property (P2) holds.

For property (P1), first note that if all gates  $g_i$  evaluate to true then  $d_F(P'_i, Q'_i) \leq 1$  for all  $i$ , so we can traverse  $P'_i$  and  $Q'_i$  staying in distance 1, and concatenating these traversals for  $1 \leq i \leq k$  (with a simultaneous step in  $P$  and  $Q$  between any  $P'_i, Q'_i$  and  $P'_{i+1}, Q'_{i+1}$ ) yields a traversal of  $P, Q$  staying in distance 1. Thus, if  $g$  evaluates to true then  $d_F(P, Q) \leq 1$ .

For the other direction, we first show that any vertices in  $P'_i$  and  $Q'_j$  have distance larger than 1 for  $i \not\equiv j \pmod{2}$ . For simplicity, say that  $i$  is odd and  $j$  is even. Since we translated  $P_i$  by  $+\delta/2$  and  $Q_j$  by  $-\delta/2$ , we obtain that  $P'_i$  is contained in  $[\delta/2 - \delta/4, \delta/2 + \delta/4] \times [1 - (\delta/4)^2, 1 + (\delta/4)^2]$  and  $Q'_j$  is contained in  $[-\delta/2 - \delta/4, -\delta/2 + \delta/4] \times [-(\delta/4)^2, (\delta/4)^2]$ . Thus, a lower bound for the distance between any two vertices of  $P'_i$  and  $Q'_j$  is the distance from the lower left corner of the region containing  $P'_i$  and the upper right corner of the region containing  $Q'_j$ , which is

$$\left( ((\delta/2 - \delta/4) - (-\delta/2 + \delta/4))^2 + (1 - 2(\delta/4)^2)^2 \right)^{1/2} = (\delta^2/4 + 1 - 4(\delta/4)^2 + 4(\delta/4)^4)^{1/2} > 1.$$

Hence, no traversal staying in distance 1 can be simultaneously in  $P'_i$  and  $Q'_j$ , for  $i \not\equiv j \pmod{2}$ .

Now assume that  $d_F(P, Q) \leq 1$  and consider a traversal  $\tau$  realizing the Fréchet distance. Note that  $\tau$  starts in the first vertices of  $P'_1, Q'_1$ , and consider the first time where the traversal reaches a vertex in  $P'_2$  or  $Q'_2$ . Observe that the previous move has to be a simultaneous move from the last vertices in  $P'_1, Q'_1$  to the first vertices in  $P'_2, Q'_2$ , since we cannot be in  $P'_1$  and  $Q'_2$  at the same time, or symmetrically in  $P'_2$  and  $Q'_1$ , as  $1 \not\equiv 2 \pmod{2}$ . Thus, we can split the traversal  $\tau$  into a traversal of  $P'_1, Q'_1$  and a traversal of the remainder of  $P, Q$ . Proceeding inductively, we can split  $\tau$  into traversals  $\tau_i$  of  $P'_i, Q'_i$  for  $1 \leq i \leq k$ . Since  $\tau$  stays in distance 1, each  $\tau_i$  also stays in distance 1. This shows  $d_F(P'_i, Q'_i) \leq 1$ , and by  $d_F(P'_i, Q'_i) = d_F(P_i, Q_i)$  and property (P1) we obtain that all gates  $g_i$  evaluate to true. This proves correctness of the AND gate construction.  $\square$

**OR Gate** We simulate OR gates by a combination of all previous constructions in Section F. Consider any OR gate  $g$  in  $F$ , and let  $g_1, \dots, g_k$  be its children in  $F$ , corresponding to subformulas  $F_1, \dots, F_k$ . Note that the outer OR construction from Lemma F.3 does not suffice to simulate an OR gate: If we recursively build curves  $P_\ell, Q_\ell$  simulating subformula  $F_\ell$  and combine them with an outer OR then it is possible to pair up  $P_i$  and  $Q_j$  for  $i \neq j$ , thus not respecting the structure of the formula. To prevent these mismatches, we adapt the subformulas  $F_\ell$  as follows. We first present a sketch of this adaptation, for the details see Lemma F.8. We consider the auxiliary formulas  $F'_\ell = [\ell \leq \ell] \wedge F_\ell \wedge [-\ell \leq -\ell]$ , and construct the curves  $P'_\ell, Q'_\ell$  that our conversion yields

when run on  $F'_\ell$ . This might seem like an unnecessary step at first sight, since the comparisons  $[\ell \leq \ell], [-\ell \leq -\ell]$  are always true, so  $F'_\ell$  simplifies to  $F_\ell$ . However, consider  $i \neq j$  and the constructed curves  $P'_i$  and  $Q'_j$ . Then the comparison  $[\ell \leq \ell]$  becomes  $[i \leq j]$ , since the left operand of the comparison is encoded in  $P$  and the right operand is encoded in  $Q$ . Similarly, the comparison  $[-\ell \leq -\ell]$  becomes  $[-i \leq -j]$ . One of these comparisons is wrong for  $i \neq j$ , so that the AND evaluates to false. Thus, the Fréchet distance of  $P'_i, Q'_j$  is larger than 1 for any  $i \neq j$ .

This trick allows us to use the outer OR construction to finish our OR gate construction. Indeed, running Lemma F.3 on the curves  $P'_1, Q'_1, \dots, P'_k, Q'_k$  yields curves  $P, Q$  such that  $d_F(P, Q) \leq 1$  if and only if there are  $i, j$  with  $d_F(P'_i, Q'_j) \leq 1$ . Since by Lemma F.8.(1) we must have  $i = j$  for Fréchet distance at most 1, we have  $d_F(P, Q) \leq 1$  if and only if  $d_F(P'_\ell, Q'_\ell) \leq 1$  for some  $\ell$ . Finally, by Lemma F.8.(2) we have  $d_F(P, Q) \leq 1$  if and only if  $F_\ell$  evaluates to true for some  $\ell$ . Hence, we correctly simulated the OR gate  $g$ .

**Lemma F.8.** *We can construct  $\delta/8$ -placed curves  $P'_1, Q'_1, \dots, P'_k, Q'_k$  satisfying (1)  $d_F(P'_i, Q'_j) > 1$  for  $i \neq j$  and (2)  $d_F(P'_\ell, Q'_\ell) \leq 1$  if and only if subformula  $F_\ell$  evaluates to true on  $(a, b)$ .*

*Proof.* As above,  $P'_\ell, Q'_\ell$  are the curves produced by our conversion when run on  $F'_\ell = [\ell \leq \ell] \wedge F_\ell \wedge [-\ell \leq -\ell]$ . More precisely, we recursively construct the curves  $P_\ell := P_{\delta/32}(g_\ell, a)$  and  $Q_\ell := Q_{\delta/32}(g_\ell, b)$  for  $1 \leq \ell \leq k$ . For the comparison  $[\ell \leq \ell]$  we construct the degenerate curves  $C_\ell^P = ((0, 1 + \delta/32 \cdot \ell/n))$  and  $C_\ell^Q = ((0, \delta/32 \cdot \ell/n))$  (see Lemma F.6), and for the comparison  $[-\ell \leq -\ell]$  we construct the degenerate curves  $D_\ell^P = ((0, 1 - \delta/32 \cdot \ell/n))$  and  $D_\ell^Q = ((0, -\delta/32 \cdot \ell/n))$ . Note that  $C_\ell^P, C_\ell^Q, D_\ell^P, D_\ell^Q, P_\ell, Q_\ell$  are  $\delta/32$ -placed, so we can apply the AND gate construction resulting in  $\delta/8$ -placed curves, see Lemma F.7. Specifically, we translate  $C_\ell^P, C_\ell^Q, D_\ell^P, D_\ell^Q$  by  $+\delta/16$  along the  $x$ -axis, and  $P_\ell, Q_\ell$  by  $-\delta/16$ . Finally, we concatenate (the translated versions of)  $C_\ell^P, P_\ell, D_\ell^P$  to obtain  $P'_\ell$  and we concatenate (the translated versions of)  $C_\ell^Q, Q_\ell, D_\ell^Q$  to obtain  $Q'_\ell$ . Note that  $P'_\ell, Q'_\ell$  are  $\delta/8$ -placed, so we can indeed use Lemma F.3 to combine  $P'_1, Q'_1, \dots, P'_k, Q'_k$  to  $\delta$ -placed curves  $P, Q$  as in the paragraph before the lemma.

From the correctness of the AND gate and comparison constructions we obtain that  $P'_\ell, Q'_\ell$  simulate  $F'_\ell$ , i.e.  $d_F(P'_\ell, Q'_\ell) \leq 1$  if and only if  $F'_\ell(a, b)$  evaluates to true. Since  $F'_\ell$  simplifies to  $F_\ell$ , this yields (2). For (1), consider  $P'_i, Q'_j$  for  $i \neq j$ . One of the comparisons  $[i \leq j], [-i \leq -j]$  is wrong, so we obtain  $d_F(C_i^P, C_j^Q) > 1$  or  $d_F(D_i^P, D_j^Q) > 1$ . Since  $C_i^P, C_j^Q$  and  $D_i^P, D_j^Q$  and  $P_i, Q_j$  are  $\delta/32$ -placed (before we translate them), the correctness argument of the AND gate, Lemma F.7, yields  $d_F(P'_i, Q'_j) > 1$ , as desired.  $\square$

**Finishing the Proof** We obtain Lemma F.2 by putting together the gate constructions of this section. This yields a recursive construction of curves  $P_\delta(g, a), Q_\delta(g, a)$  simulating gate  $g$ , correctness follows from a simple inductive argument. Note that if  $g$  has fanin  $k$  then we add  $O(k)$  vertices to the recursively constructed curves corresponding to the  $k$  children, yielding a bound of  $O(|F|)$  on the length of the constructed curves. Finally, for the factor by which we decrease  $\delta$  from one gate to its children we note that  $1/32$  is sufficiently small. Thus, we may produce  $\delta/32^d$ -placed curves for each gate in depth  $d$  below  $g$ . The coordinates of our AND and OR gate constructions thus use  $O(\log(1/\delta) + d)$  bits. The only exception is the comparison gate construction, for which we need  $O(\log M) = O(\log n)$  additional bits to represent numbers in  $\{-M, \dots, M\}$ . Together this yields Lemma F.2.