



Tight Conditional Lower Bounds for Longest Common Increasing Subsequence

Lech Duraj¹ · Marvin Künnemann² · Adam Polak¹ 

Received: 28 December 2017 / Accepted: 16 July 2018
© The Author(s) 2018

Abstract

We consider the canonical generalization of the well-studied Longest Increasing Subsequence problem to multiple sequences, called k -LCIS: Given k integer sequences X_1, \dots, X_k of length at most n , the task is to determine the length of the longest common subsequence of X_1, \dots, X_k that is also strictly increasing. Especially for the case of $k = 2$ (called LCIS for short), several algorithms have been proposed that require quadratic time in the worst case. Assuming the Strong Exponential Time Hypothesis (SETH), we prove a tight lower bound, specifically, that no algorithm solves LCIS in (strongly) subquadratic time. Interestingly, the proof makes no use of normalization tricks common to hardness proofs for similar problems such as Longest Common Subsequence. We further strengthen this lower bound (1) to rule out $\mathcal{O}((nL)^{1-\varepsilon})$ time algorithms for LCIS, where L denotes the solution size, (2) to rule out $\mathcal{O}(n^{k-\varepsilon})$ time algorithms for k -LCIS, and (3) to follow already from weaker variants of SETH. We obtain the same conditional lower bounds for the related Longest Common Weakly Increasing Subsequence problem.

Keywords Fine-grained complexity · Combinatorial pattern matching · Sequence alignments · Parameterized complexity · SETH

An extended abstract of this paper has appeared in the proceedings of IPEC 2017. Lech Duraj was partially supported by Polish National Science Center Grant 2016/21/B/ST6/02165. Adam Polak was partially supported by Polish Ministry of Science and Higher Education program *Diamentowy Grant*.

✉ Adam Polak
polak@tcs.uj.edu.pl

Lech Duraj
duraj@tcs.uj.edu.pl

Marvin Künnemann
marvin@mpi-inf.mpg.de

¹ Theoretical Computer Science, Faculty of Mathematics and Computer Science, Jagiellonian University, Kraków, Poland

² Max Planck Institute for Informatics, Saarland Informatics Campus, Saarbrücken, Germany

1 Introduction

The longest common subsequence problem (LCS) and its variants are computational primitives with a variety of applications, which includes, e.g., uses as similarity measures for spelling correction [37,43] or DNA sequence comparison [5,39], as well as determining the differences of text files as in the UNIX DIFF utility [28]. LCS shares characteristics of both an easy and a hard problem: (*Easy*) A simple and elegant dynamic-programming algorithm computes an LCS of two length- n sequences in time $\mathcal{O}(n^2)$ [43], and in many practical settings, certain properties of typical input sequences can be exploited to obtain faster, “tailored” solutions (e.g., [7,27,29,38]; see also [14] for a survey). (*Hard*) At the same time, no polynomial improvements over the classical solution are known, thus exact computation may become infeasible for very long general input sequences. The research community has sought for a resolution of the question “*Do subquadratic algorithms for LCS exist?*” already shortly after the formalization of the problem [4,21].

Recently, an answer conditional on the Strong Exponential Time Hypothesis (SETH; see Sect. 2 for a definition) could be obtained: Based on a line of research relating the satisfiability problem to quadratic-time problems [3,15,41,44] and following a breakthrough result for Edit Distance [9], it has been shown that unless SETH fails, there is no (strongly) subquadratic-time algorithm for LCS [1,16]. Subsequent work [2] strengthens these lower bounds to hold already under weaker assumptions and even provides surprising consequences of sufficiently strong polylogarithmic improvements.

Due to its popularity and wide range of applications, several variants of LCS have been proposed. This includes the heaviest common subsequence (HCS) [32], which introduces weights to the problem, as well as notions that constrain the structure of the solution, such as the longest common increasing subsequence (LCIS) [46], LCS k [13], constrained LCS [8,20,42], restricted LCS [26], and many other variants (see, e.g., [6,19,33]). Most of these variants are (at least loosely) motivated by biological sequence comparison tasks. To the best of our knowledge, in the above list, LCIS is the only LCS variant for which (1) the best known algorithms run in quadratic time in the worst case and (2) its definition does not include LCS as a special case (for such generalizations of LCS, the quadratic-time SETH hardness of LCS [1,16] would transfer immediately). As such, it is open to determine whether there are (strongly) subquadratic algorithms for LCIS or whether such algorithms can be ruled out under SETH. The starting point of our work is to settle this question.

1.1 Longest Common Increasing Subsequence (LCIS)

The Longest Common Increasing Subsequence problem on k sequences (k -LCIS) is defined as follows: Given integer sequences X_1, \dots, X_k of length at most n , determine the length of the longest sequence Z such that Z is a strictly increasing sequence of integers and Z is a subsequence of each $X_i, i \in \{1, \dots, k\}$. For $k = 1$, we obtain the well-studied longest increasing subsequence problem (LIS; we refer to [22] for an overview), which has an $\mathcal{O}(n \log n)$ time solution and a matching lower bound in

the decision tree model [25]. The extension to $k = 2$, denoted simply as LCIS, has been proposed by Yang, Huang, and Chao [46], partially motivated as a generalization of LIS and by potential applications in bioinformatics. They obtained an $\mathcal{O}(n^2)$ time algorithm, leaving open the natural question whether there exists a way to extend the near-linear time solution for LIS to a near-linear time solution for multiple sequences.

Interestingly, already a classic connection between LCS and LIS combined with a recent conditional lower bound of Abboud et al. [1] yields a partial negative answer assuming SETH.

Observation 1 (Folklore reduction, implicit in [29], explicit in [32]) *After $\mathcal{O}(kn^2)$ time preprocessing, we can solve k -LCS by a single call to $(k - 1)$ -LCIS on sequences of length at most n^2 .*

Proof Let $L(\sigma)$ denote the decreasing sequence of positions j with $X_1[j] = \sigma$. We define sequences $X'_i = L(X_i[0]) \cdots L(X_i[|X_i| - 1])$ for all $i \in \{2, \dots, k\}$. We claim that for any ℓ , there exists a length- ℓ increasing common subsequence of X'_2, \dots, X'_k if and only if there is a length- ℓ common subsequence of X_1, \dots, X_k . Thus, the length of the LCIS of X'_2, \dots, X'_k is equal to the length of the LCS of X_1, \dots, X_k , and the claim follows since $|L(\sigma)| \leq n$ for all σ .

To prove this claim, let $(\sigma_0, \dots, \sigma_{\ell-1})$ be any common subsequence of X_1, \dots, X_k . In particular, we have $\sigma_j = X_1[p_j]$ for some strictly increasing sequence of positions $(p_0, \dots, p_{\ell-1})$. We claim that $(p_0, \dots, p_{\ell-1})$ is a common increasing sequence of X'_2, \dots, X'_k . Indeed, for any $j \in \{0, \dots, \ell - 1\}$, p_j belongs to $L(\sigma_j)$ by definition, so $(p_0, \dots, p_{\ell-1})$ is a subsequence of $L(\sigma_0) \dots L(\sigma_{\ell-1})$, which in turn is a subsequence of X'_i for any $i \geq 2$.

Conversely, let $(p_0, \dots, p_{\ell-1})$ be any common increasing subsequence of X'_2, \dots, X'_k . Let $\sigma_j = X_1[p_j]$ for $j = 0, \dots, \ell - 1$. The sequence $(\sigma_0, \dots, \sigma_{\ell-1})$ is trivially a subsequence of X_1 . For $i \in \{2, \dots, k\}$, observe that every p_j must belong to $L(X_i[r_j])$ for some $0 \leq r_j < |X_i|$, and that for any $j_1 < j_2$, we must have $r_{j_1} < r_{j_2}$, as $L(\sigma)$ is always a strictly decreasing sequence. Thus, $(X_i[r_0], \dots, X_i[r_{\ell-1}])$ is a subsequence of X_i , where $X_i[r_j] = \sigma_j$ must hold, since p_j appears in $L(X_i[r_j])$. Thus $(\sigma_0, \dots, \sigma_{\ell-1})$ is a length- ℓ subsequence of all the X_i sequences. \square

Corollary 1 *Unless SETH fails, there is no $\mathcal{O}(n^{\frac{3}{2}-\varepsilon})$ time algorithm for LCIS for any constant $\varepsilon > 0$.*

Proof Note that by the above reduction, an $\mathcal{O}(n^{\frac{3}{2}-\varepsilon})$ time LCIS algorithm would give an $\mathcal{O}(n^{3-2\varepsilon})$ time algorithm for 3-LCS. Such an algorithm would refute SETH by a result of Abboud et al. [1]. \square

While this rules out near-linear time algorithms, still an unsatisfying large polynomial gap between best upper and conditional lower bounds persists.

1.2 Our Results

Our first result is a tight SETH-based lower bound for LCIS.

Theorem 1 *Unless SETH fails, there is no $\mathcal{O}(n^{2-\varepsilon})$ time algorithm for LCIS for any constant $\varepsilon > 0$.*

We extend our main result in several directions.

1.2.1 Parameterized Complexity I: Solution Size

Subsequent work [18,35] improved over Yang et al.’s algorithm when certain input parameters are small. Here, we focus particularly on the solution size, i.e., the length L of the LCIS. Kutz et al. [35] provided an algorithm running in time $\mathcal{O}(nL \log \log n + n \log n)$. Clearly, L can be as large as n . However, when L is significantly smaller, say $L = n^{\frac{1}{2} \pm o(1)}$, this algorithm runs in strongly subquadratic time. Interestingly, exactly for this case, the reduction from 3-LCS to LCIS of Observation 1 already yields a matching SETH-based lower bound of $(Ln)^{1-o(1)} = n^{\frac{3}{2}-o(1)}$. However, for smaller L , this reduction yields no lower bound at all and only a non-matching lower bound for larger L . We remedy this situation by the following result.¹

Theorem 2 *Unless SETH fails, there is no $\mathcal{O}((nL)^{1-\varepsilon})$ time algorithm for LCIS for any constant $\varepsilon > 0$. This even holds restricted to instances with $L = n^{\gamma \pm o(1)}$, for arbitrarily chosen $0 < \gamma \leq 1$.*

1.2.2 Parameterized Complexity II: k -LCIS

For constant $k \geq 2$, $\mathcal{O}(n^k \text{polylog}(n))$ time algorithms for k -LCIS follow from [18,35], and a folklore DP approach yields an $\mathcal{O}(n^k)$ solution (see the appendix). While it is known that k -LCS cannot be computed in time $\mathcal{O}(n^{k-\varepsilon})$ for any constant $\varepsilon > 0$, $k \geq 2$ unless SETH fails [1], this does not directly transfer to k -LCIS, since the reduction in Observation 1 is not tight. However, by extending our main construction, we can prove the analogous result.

Theorem 3 *Unless SETH fails, there is no $\mathcal{O}(n^{k-\varepsilon})$ time algorithm for k -LCIS for any constant $k \geq 2$ and $\varepsilon > 0$.*

1.2.3 Longest Common Weakly Increasing Subsequence (LCWIS)

We consider a closely related variant of LCIS called the Longest Common Weakly Increasing Subsequence (k -LCWIS): Here, given integer sequences X_1, \dots, X_k of length at most n , the task is to determine the longest *weakly increasing* (i.e. non-decreasing) integer sequence Z that is a common subsequence of X_1, \dots, X_k . Again, we write LCWIS as a shorthand for 2-LCWIS. Note that the seemingly small change in the notion of increasing sequence has a major impact on algorithmic and hardness results: Any instance of LCIS in which the input sequences are defined over a small-sized alphabet $\Sigma \subseteq \mathbb{Z}$, say $|\Sigma| = \mathcal{O}(n^{1/2})$, can be solved in strongly subquadratic

¹ We mention in passing that a systematic study of the complexity of LCS in terms of such input parameters has been performed recently in [17].

time $\mathcal{O}(nL \log n) = \mathcal{O}(n^{3/2} \log n)$ [35], by using the fact that $L \leq |\Sigma|$. In contrast, LCWIS is quadratic-time SETH hard already over slightly superlogarithmic-sized alphabets [40]. We give a substantially different proof for this fact and generalize it to k -LCWIS.

Theorem 4 *Unless SETH fails, there is no $\mathcal{O}(n^{k-\varepsilon})$ time algorithm for k -LCWIS for any constant $k \geq 2$ and $\varepsilon > 0$. This even holds restricted to instances defined over an alphabet of size $|\Sigma| \leq f(n) \log n$ for any function $f(n) = \omega(1)$ growing arbitrarily slowly.*

1.2.4 Strengthening the Hardness

In an attempt to strengthen the conditional lower bounds for Edit Distance and LCS [1,9,16], particularly, to obtain barriers even for subpolynomial improvements, Abboud et al. [2] gave the first fine-grained reductions from the satisfiability problem on branching programs. Using this approach, the quadratic-time hardness of a problem can be explained by considerably weaker variants of SETH, making the conditional lower bound stronger. We show that our lower bounds also hold under these weaker variants. In particular, we prove the following.

Theorem 5 *There is no strongly subquadratic time algorithm for LCIS, unless there is, for some $\varepsilon > 0$, an $\mathcal{O}((2 - \varepsilon)^N)$ algorithm for the satisfiability problem on branching programs of width W and length T on N variables with $(\log W)(\log T) = o(N)$.*

1.3 Discussion, Outline and Technical Contributions

Apart from an interest in LCIS and its close connection to LCS, our work is also motivated by an interest in the *optimality of dynamic programming (DP) algorithms*.² Notably, many conditional lower bounds in P target problems with natural DP algorithms that are proven to be near-optimal under some plausible assumption (see, e.g., [1,3,9–11,15,16,23,34,45] for an introduction to the field). Even if we restrict our attention to problems that find optimal sequence alignments under some restrictions, such as LCS, Edit Distance and LCIS, the currently known hardness proofs differ significantly, despite seemingly small differences between the problem definitions. Ideally, we would like to classify the properties of a DP formulation which allow for matching conditional lower bounds.

One step in this direction is given by the *alignment gadget framework* [16]. Exploiting normalization tricks, this framework gives an abstract property of sequence similarity measures to allow for SETH-based quadratic lower bounds. Unfortunately, as it turns out, we cannot directly transfer the alignment gadget hardness proof for LCS to LCIS – some indication for this difficulty is already given by the fact that LCIS can be solved in strongly subquadratic time over sublinear-sized alphabets [35], while the LCS hardness proof already applies to binary alphabets. By collecting gadgetry needed to overcome such difficulties (that we elaborate on below), we hope to pro-

² We refer to [47] for a simple quadratic-time DP formulation for LCIS.

vide further tools to generalize more and more quadratic-time lower bounds based on SETH.

1.3.1 Technical Challenges

The known conditional lower bounds for global alignment problems such as LCS and Edit Distance work as follows. The reductions start from the quadratic-time SETH-hard Orthogonal Vectors problem (OV), that asks to determine, given two sets of $(0, 1)$ -vectors $\mathcal{U} = \{u_0, \dots, u_{n-1}\}, \mathcal{V} = \{v_0, \dots, v_{n-1}\} \subseteq \{0, 1\}^d$ over $d = n^{o(1)}$ dimensions, whether there is a pair i, j such that u_i and v_j are orthogonal, i.e., whose inner product $(u_i \cdot v_j) := \sum_{k=0}^{d-1} u_i[k] \cdot v_j[k]$ is 0 (over the integers). Each vector u_i and v_j is represented by a (normalized) vector gadget $\text{VG}_X(u_i)$ and $\text{VG}_Y(v_j)$, respectively. Roughly speaking, these gadgets are combined to sequences X and Y such that each candidate for an optimal alignment of X and Y involves locally optimal alignments between n pairs $\text{VG}_X(u_i), \text{VG}_Y(v_j)$ —the optimal alignment exceeds a certain threshold if and only if there is an orthogonal pair u_i, v_j .

An analogous approach does not work for LCIS: Let $\text{VG}_X(u_i)$ be defined over an alphabet Σ and $\text{VG}_X(u_{i'})$ over an alphabet Σ' . If Σ and Σ' overlap, then $\text{VG}_X(u_i)$ and $\text{VG}_X(u_{i'})$ cannot both be aligned in an optimal alignment without interference with each other. On the other hand, if Σ and Σ' are disjoint, then each vector v_j should have its corresponding vector gadget $\text{VG}_Y(v_j)$ defined over both Σ and Σ' in order to allow aligning $\text{VG}_X(u_i)$ with $\text{VG}_Y(v_j)$ as well as $\text{VG}_X(u_{i'})$ with $\text{VG}_Y(v_j)$. The latter option drastically increases the size of vector gadgets. Thus, we must define all vector gadgets over a common alphabet Σ and make sure that *only a single pair* $\text{VG}_X(u_i), \text{VG}_Y(v_j)$ is aligned in an optimal alignment (in contrast with n pairs aligned in the previous reductions for LCS and Edit Distance).

1.3.2 Technical Contributions and Proof Outline

Fortunately, a surprisingly simple approach works: As a key tool, we provide *separator sequences* $\alpha_0 \dots \alpha_{n-1}$ and $\beta_0 \dots \beta_{n-1}$ with the following properties: (1) for every $i, j \in \{0, \dots, n-1\}$ the LCIS of $\alpha_0 \dots \alpha_i$ and $\beta_0 \dots \beta_j$ has a length of $f(i+j)$, where f is a linear function, and (2) $\sum_i |\alpha_i|$ and $\sum_j |\beta_j|$ are bounded by $n^{1+o(1)}$. Note that existence of such a gadget is somewhat unintuitive: condition (1) for $i=0$ and $j=n-1$ requires $|\alpha_0| = \Omega(n)$, yet still the total length $\sum_i |\alpha_i|$ must not exceed the length of $|\alpha_0|$ significantly. Indeed, we achieve this by a careful inductive construction that generates such sequences with heavily varying block sizes $|\alpha_i|$ and $|\beta_j|$.

We apply these separator sequences as follows. We first define simple vector gadgets $\text{VG}_X(u_i), \text{VG}_Y(v_j)$ over an alphabet Σ such that the length of an LCIS of $\text{VG}_X(u_i)$ and $\text{VG}_Y(v_j)$ is $d - (u_i \cdot v_j)$. Then we construct the separator sequences as above over an alphabet $\Sigma_{<}$ whose elements are strictly smaller than all elements in Σ . Furthermore, we create analogous separator sequences $\alpha'_0 \dots \alpha'_{n-1}$ and $\beta'_0 \dots \beta'_{n-1}$ which satisfy a property like (1) for all suffixes instead of prefixes, using an alphabet $\Sigma_{>}$ whose elements are strictly larger than all elements in Σ . Now, we define

$$X = \alpha_0 \text{VG}_X(u_0) \alpha'_0 \dots \alpha_{n-1} \text{VG}_X(u_{n-1}) \alpha'_{n-1},$$

$$Y = \beta_0 \text{VG}_Y(v_0) \beta'_0 \dots \beta_{n-1} \text{VG}_Y(v_{n-1}) \beta'_{n-1}.$$

As we will show in Sect. 3, the length of an LCIS of X and Y is $C - \min_{i,j}(u_i \cdot v_j)$ for some constant C depending only on n and d .

In contrast to previous such OV-based lower bounds, we use heavily varying separators (padding) between vector gadgets.

2 Preliminaries

As a convention, we use capital or Greek letters to denote sequences over integers. Let X, Y be integer sequences. We write $|X|$ for the length of X , $X[k]$ for the k -th element in the sequence X ($k \in \{0, \dots, |X| - 1\}$), and $X \circ Y$ (or just XY , interchangeably) for the concatenation of X and Y . We say that Y is a subsequence of X if there exist indices $0 \leq i_0 < i_1 < \dots < i_{|Y|-1} \leq |X| - 1$ such that $X[i_k] = Y[k]$ for all $k \in \{0, \dots, |Y| - 1\}$. Given any number of sequences X_1, \dots, X_k , we say that Y is a common subsequence of X_1, \dots, X_k if Y is a subsequence of each $X_i, i \in \{1, \dots, k\}$. X is called strictly increasing (or weakly increasing) if $X[0] < X[1] < \dots < X[|X| - 1]$ (or $X[0] \leq X[1] \leq \dots \leq X[|X| - 1]$). For any k sequences X_1, \dots, X_k , we denote by $\text{lcs}(X_1, \dots, X_k)$ the length of their longest common subsequence that is strictly increasing.

2.1 Hardness Assumptions

All of our lower bounds hold assuming the Strong Exponential Time Hypothesis (SETH), introduced by Impagliazzo and Paturi [30,31]. It essentially states that no exponential speed-up over exhaustive search is possible for the CNF satisfiability problem.

Hypothesis 1 [Strong Exponential Time Hypothesis (SETH)] There is no $\varepsilon > 0$ such that for all $q \geq 3$ there is an $\mathcal{O}(2^{(1-\varepsilon)n})$ time algorithm for q -SAT.

This hypothesis implies tight hardness of the k -Orthogonal Vectors problem (k -OV), which will be the starting point of our reductions: Given k sets $\mathcal{U}_1, \dots, \mathcal{U}_k \subseteq \{0, 1\}^d$, each with $|\mathcal{U}_i| = n$ vectors over $d = n^{o(1)}$ dimensions, determine whether there is a k -tuple $(u_1, \dots, u_k) \in \mathcal{U}_1 \times \dots \times \mathcal{U}_k$ such that $\sum_{\ell=0}^{d-1} \prod_{i=1}^k u_i[\ell] = 0$. By exhaustive enumeration, it can be solved in time $\mathcal{O}(n^k d) = n^{k+o(1)}$. The following conjecture is implied by SETH by the well-known split-and-list technique of Williams [44] (and the sparsification lemma [31]).³

³ We sketch how the split-and-list technique reduces q -SAT to k -Orthogonal Vectors. Given a q -CNF formula on N variables and M clauses, the key idea is to split the variables into k sets V_1, \dots, V_k of roughly equal size. For each i , we then construct the vector set \mathcal{U}_i as follows: for each of the $\mathcal{O}(2^{N/k})$ possible assignments to the variables V_i , we include a vector in $\{0, 1\}^M$ that represents the clauses that are not already satisfied by this partial assignment. It is easy to see that the q -SAT instance is satisfiable if and only if there are vectors $(u_1, \dots, u_k) \in \mathcal{U}_1 \times \dots \times \mathcal{U}_k$ with $\sum_{\ell=0}^{M-1} \prod_{i=1}^k u_i[\ell] = 0$. This yields

Hypothesis 2 (*k-OV conjecture*) Let $k \geq 2$. There is no $\mathcal{O}(n^{k-\varepsilon})$ time algorithm for k -OV, with $d = \omega(\log n)$, for any constant $\varepsilon > 0$.

For the special case of $k = 2$, which we simply denote by OV, we obtain the following weaker conjecture.

Hypothesis 3 (*OV conjecture*) There is no $\mathcal{O}(n^{2-\varepsilon})$ time algorithm for OV, with $d = \omega(\log n)$, for any constant $\varepsilon > 0$. Equivalently, even restricted to instances with $|\mathcal{U}_1| = n$ and $|\mathcal{U}_2| = n^\gamma$, $0 < \gamma \leq 1$, there is no $\mathcal{O}(n^{1+\gamma-\varepsilon})$ time algorithm for OV, with $d = \omega(\log n)$, for any constant $\varepsilon > 0$.

A proof of the folklore equivalence of the statements for equal and unequal set sizes can be found, e.g., in [16].

3 Main Construction: Hardness of LCIS

In this section, we prove quadratic-time SETH hardness of LCIS, i.e., prove Theorem 1. We first introduce an *inflation* operation, which we then use to construct our separator sequences. After defining simple vector gadgets, we show how to embed an Orthogonal Vectors instance using our vector gadgets and separator sequences.

3.1 Inflation

We begin by introducing the inflation operation, which simulates weighing the sequences.

Definition 1 For a sequence $A = \langle a_0, a_1, \dots, a_{n-1} \rangle$ of integers we define:

$$\text{inflate}(A) = \langle 2a_0 - 1, 2a_0, 2a_1 - 1, 2a_1, \dots, 2a_{n-1} - 1, 2a_{n-1} \rangle.$$

Lemma 1 For any two sequences A and B ,

$$\text{lcs}(\text{inflate}(A), \text{inflate}(B)) = 2 \cdot \text{lcs}(A, B).$$

Proof Let C be the longest common increasing subsequence of A and B . Observe that $\text{inflate}(C)$ is a common increasing subsequence of $\text{inflate}(A)$ and $\text{inflate}(B)$ of length $2 \cdot |C|$, thus $\text{lcs}(\text{inflate}(A), \text{inflate}(B)) \geq 2 \cdot \text{lcs}(A, B)$.

Conversely, let \bar{A} denote $\text{inflate}(A)$ and \bar{B} denote $\text{inflate}(B)$. Let \bar{C} be the longest common increasing subsequence of \bar{A} and \bar{B} . If we divide all elements of \bar{C} by 2 and

Footnote 3 continued

a k -OV instance with sets of size $n = \mathcal{O}(2^{N/k})$ and vector dimension $d = M$. Roughly speaking, the sparsification lemma now allows us to assume that our q -CNF formula only has $M = \mathcal{O}(N) = \mathcal{O}(\log n)$ clauses, where the hidden constant depends on q . Thus, any $\mathcal{O}(n^{k-\varepsilon})$ algorithm for k -OV with vector dimension $d = \omega(\log n)$ would imply an $\mathcal{O}(2^{N/k \cdot (k-\varepsilon)}) = \mathcal{O}(2^{(1-\delta)N})$ algorithm for q -SAT for some $\delta > 0$ that is independent of q .

round up to the closest integer, we end up with a weakly increasing sequence. Now, if we remove duplicate elements to make this sequence strictly increasing, we obtain C , a common increasing subsequence of A and B . At most 2 distinct elements may become equal after division by 2 and rounding, therefore C contains at least $\lceil \text{lcis}(\bar{A}, \bar{B})/2 \rceil$ elements, so $2 \cdot \text{lcis}(A, B) \geq \text{lcis}(\bar{A}, \bar{B})$. This completes the proof. \square

3.2 Separator Sequences

Our goal is to construct two sequences A and B which can be split into n blocks, i.e. $A = \alpha_0\alpha_1 \dots \alpha_{n-1}$ and $B = \beta_0\beta_1 \dots \beta_{n-1}$, such that the length of the longest common increasing subsequence of the first i blocks of A and the first j blocks of B equals $i + j$, up to an additive constant. We call A and B *separator sequences*, and use them later to separate vector gadgets in order to make sure that only one pair of gadgets may interact with each other at the same time.

We construct the separator sequences inductively. For every $k \in \mathbb{N}$, the sequences A_k and B_k are concatenations of 2^k blocks (of varying sizes), $A_k = \alpha_k^0\alpha_k^1 \dots \alpha_k^{2^k-1}$ and $B_k = \beta_k^0\beta_k^1 \dots \beta_k^{2^k-1}$. Let s_k denote the largest element of both sequences. As we will soon observe, $s_k = 2^{k+2} - 3$.

The construction works as follows: for $k = 0$, we can simply set A_0 and B_0 as one-element sequences $\langle 1 \rangle$. We then construct A_{k+1} and B_{k+1} inductively from A_k and B_k in two steps. First, we inflate both A_k and B_k , then after each (now inflated) block we insert 3-element sequences, called *tail gadgets*, $\langle 2s_k + 2, 2s_k + 1, 2s_k + 3 \rangle$ for A_{k+1} and $\langle 2s_k + 1, 2s_k + 2, 2s_k + 3 \rangle$ for B_{k+1} . Formally, we describe the construction by defining blocks of the new sequences (see Figs. 1 and 2). For $i \in \{0, 1, \dots, 2^k - 1\}$,

$$\begin{aligned} \alpha_{k+1}^{2^i} &= \text{inflate}(\alpha_k^i) \circ \langle 2s_k + 2 \rangle, & \alpha_{k+1}^{2^i+1} &= \langle 2s_k + 1, 2s_k + 3 \rangle, \\ \beta_{k+1}^{2^i} &= \text{inflate}(\beta_k^i) \circ \langle 2s_k + 1 \rangle, & \beta_{k+1}^{2^i+1} &= \langle 2s_k + 2, 2s_k + 3 \rangle. \end{aligned}$$

Note that the symbols appearing in tail gadgets do not appear in the inflated sequences. The largest element of both new sequences s_{k+1} equals $2s_k + 3$, and solving the recurrence gives indeed $s_k = 2^{k+2} - 3$.

Now, let us prove two useful properties of the separator sequences.

Lemma 2 $|A_k| = |B_k| = (\frac{3}{2}k + 1) \cdot 2^k = \mathcal{O}(k2^k)$.

Proof Observe that $|A_{k+1}| = 2|A_k| + 3 \cdot 2^k$. Indeed, to obtain A_{k+1} we first double the size of A_k and then add 3 new elements for each of the 2^k blocks of A_k . Solving the recurrence completes the proof. The same reasoning applies to B_k . \square

Lemma 3 For every $i, j \in \{0, 1, \dots, 2^k - 1\}$, $\text{lcis}(\alpha_k^0 \dots \alpha_k^i, \beta_k^0 \dots \beta_k^j) = i + j + 2^k$.

Proof The proof is by induction on k . For $k = 0$, we have $\text{lcis}(\alpha_0^0, \beta_0^0) = \text{lcis}(\langle 1 \rangle, \langle 1 \rangle) = 1$, as desired. Assume the statement is true for k and let us prove it for $k + 1$.

The “ \geq ” direction. First, consider the case when both i and j are even. Observe that $\text{inflate}(\alpha_k^0 \dots \alpha_k^{i/2})$ and $\text{inflate}(\beta_k^0 \dots \beta_k^{j/2})$ are subsequences of $\alpha_{k+1}^0 \dots \alpha_{k+1}^i$

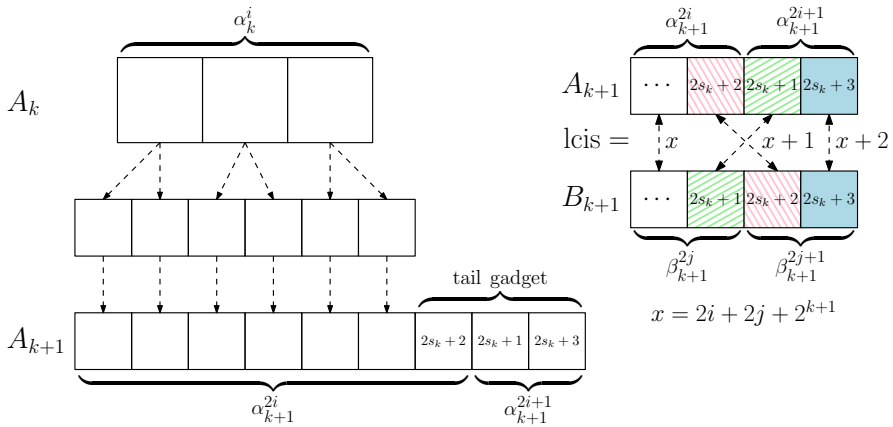


Fig. 1 Constructing A_{k+1} from A_k (left), and intuition behind tail gadgets (right)

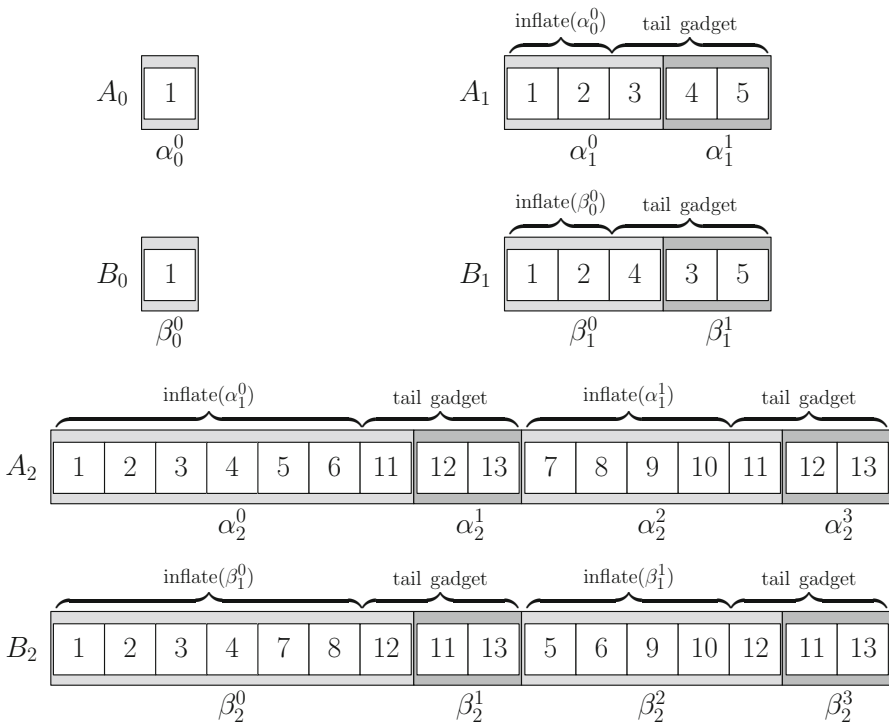


Fig. 2 Initial steps of the inductive construction of the separator sequences

and $\beta_{k+1}^0 \dots \beta_{k+1}^j$, respectively. Thus, using the induction hypothesis and inflation properties,

$$\begin{aligned} \text{lcis}(\alpha_{k+1}^0 \dots \alpha_{k+1}^i, \beta_{k+1}^0 \dots \beta_{k+1}^j) &\geq \\ &\geq \text{lcis}(\text{inflate}(\alpha_k^0 \dots \alpha_k^{i/2}), \text{inflate}(\beta_k^0 \dots \beta_k^{j/2})) = \\ &= 2 \cdot \text{lcis}(\alpha_k^0 \dots \alpha_k^{i/2}, \beta_k^0 \dots \beta_k^{j/2}) = 2 \cdot (i/2 + j/2 + 2^k) = i + j + 2^{k+1}. \end{aligned}$$

If i is odd and j is even, refer to the previous case to get a common increasing subsequence of $\alpha_{k+1}^0 \dots \alpha_{k+1}^{i-1}$ and $\beta_{k+1}^0 \dots \beta_{k+1}^j$ of length $i - 1 + j + 2^{k+1}$ consisting only of elements less than or equal to $2s_k$, and append the element $2s_k + 1$ to the end of it. Analogously, for i even and j odd, take such an LCIS of $\alpha_{k+1}^0 \dots \alpha_{k+1}^i$ and $\beta_{k+1}^0 \dots \beta_{k+1}^{j-1}$, and append $2s_k + 2$. Finally, for both i and j odd, take an LCIS of $\alpha_{k+1}^0 \dots \alpha_{k+1}^{i-1}$ and $\beta_{k+1}^0 \dots \beta_{k+1}^{j-1}$, and append $2s_k + 1$ and $2s_k + 3$.

The “ \leq ” direction. We proceed by induction on $i + j$. Fix i and j , and let L be a longest common increasing subsequence of $\alpha_{k+1}^0 \dots \alpha_{k+1}^i$ and $\beta_{k+1}^0 \dots \beta_{k+1}^j$.

If the last element of L is less than or equal to $2s_k$, L is in fact a common increasing subsequence of $\text{inflate}(\alpha_k^0 \dots \alpha_k^{\lfloor i/2 \rfloor})$ and $\text{inflate}(\beta_k^0 \dots \beta_k^{\lfloor j/2 \rfloor})$, thus, by the induction hypothesis and inflation properties, $|L| \leq 2 \cdot (\lfloor i/2 \rfloor + \lfloor j/2 \rfloor + 2^k) \leq i + j + 2^{k+1}$.

The remaining case is when the last element of L is greater than $2s_k$. In this case, consider the second-to-last element of L . It must belong to some blocks $\alpha_{k+1}^{i'}$ and $\beta_{k+1}^{j'}$ for $i' \leq i$ and $j' \leq j$, and we claim that $i = i'$ and $j = j'$ cannot hold simultaneously: by construction of separator sequences, if blocks $\alpha_{k+1}^{i'}$ and $\beta_{k+1}^{j'}$ have a common element larger than $2s_k$, then it is the only common element of these two blocks. Therefore, it cannot be the case that both $i = i'$ and $j = j'$, because the last two elements of L would then be located in $\alpha_{k+1}^{i'}$ and $\beta_{k+1}^{j'}$. As a consequence, $i' + j' < i + j$, which lets us apply the induction hypothesis to reason that the prefix of L omitting its last element is of length at most $i' + j' + 2^{k+1}$. Therefore, $|L| \leq 1 + i' + j' + 2^{k+1} \leq i + j + 2^{k+1}$, which completes the proof. \square

Observe that if we reverse the sequences A_k and B_k along with changing all elements to their negations, i.e. x to $-x$, we obtain sequences \hat{A}_k and \hat{B}_k such that \hat{A}_k splits into 2^k blocks $\hat{\alpha}_k^0 \dots \hat{\alpha}_k^{2^k-1}$, \hat{B}_k splits into 2^k blocks $\hat{\beta}_k^0 \dots \hat{\beta}_k^{2^k-1}$, and

$$\text{lcis}(\hat{\alpha}_k^i \dots \hat{\alpha}_k^{2^k-1}, \hat{\beta}_k^j \dots \hat{\beta}_k^{2^k-1}) = 2 \cdot (2^k - 1) - i - j + 2^k. \tag{1}$$

Finally, observe that we can add any constant to all elements of the sequences A_k and B_k (as well as \hat{A}_k and \hat{B}_k) without changing the property stated in Lemma 3 (and its analogue for \hat{A}_k and \hat{B}_k , i.e. Eq. (1)).

3.3 Vector Gadgets

Let $\mathcal{U} = \{u_0, \dots, u_{n-1}\}$ and $\mathcal{V} = \{v_0, \dots, v_{n-1}\}$ be two sets of d -dimensional $(0, 1)$ -vectors.

For $i \in \{0, 1, \dots, n-1\}$ let us construct the vector gadgets U_i and V_i as $2d$ -element sequences, by defining, for every $p \in \{0, 1, \dots, d-1\}$,

$$\begin{aligned} (U_i[2p], U_i[2p+1]) &= \begin{cases} (2p, 2p+1) & \text{if } u_i[p] = 0, \\ (2p, 2p) & \text{if } u_i[p] = 1, \end{cases} \\ (V_i[2p], V_i[2p+1]) &= \begin{cases} (2p+1, 2p) & \text{if } v_i[p] = 0, \\ (2p+1, 2p+1) & \text{if } v_i[p] = 1. \end{cases} \end{aligned}$$

Observe that at most one of the elements $2p$ and $2p+1$ may appear in the LCIS of U_i and V_j , and it happens if and only if $u_i[p]$ and $v_j[p]$ are not both equal to one. Therefore, $\text{lcis}(U_i, V_j) = d - (u_i \cdot v_j)$, and, in particular, $\text{lcis}(U_i, V_j) = d$ if and only if u_i and v_j are orthogonal.

3.4 Final Construction

To put all the pieces together, we plug vector gadgets U_i and V_j into the separator sequences from Sect. 3.2, obtaining two sequences whose LCIS depends on the minimal inner product of vectors u_i and v_j . We provide a general construction of such sequences, which will be useful in later sections.

Lemma 4 *Let $X_0, X_1, \dots, X_{n-1}, Y_0, Y_1, \dots, Y_{n-1}$ be integer sequences such that none of them has an increasing subsequence longer than δ . Then there exist sequences X and Y of length $\mathcal{O}(\delta \cdot n \log n) + \sum |X_i| + \sum |Y_j|$, constructible in linear time, such that:*

$$\text{lcis}(X, Y) = \max_{i,j} \text{lcis}(X_i, Y_j) + C$$

for a constant C that only depends on n and δ and satisfies $C = \mathcal{O}(n\delta)$.

Proof First, we can assume that $n = 2^k$ for some positive integer k . If not, we can add dummy one-element sequences as new X_i 's and Y_j 's such that they have no common element with any other sequences. This increases n at most twofold and $\sum |X_i|$ and $\sum |Y_j|$ by at most n .

Recall the sequences A_k, B_k, \hat{A}_k and \hat{B}_k constructed in Sect. 3.2. Let A, B, \hat{A}, \hat{B} be the sequences obtained from $A_k, B_k, \hat{A}_k, \hat{B}_k$ by applying inflation $\lceil \log_2 \delta \rceil$ times (thus increasing their length by a factor of $\ell = 2^{\lceil \log_2 \delta \rceil} \geq \delta$). Each of these four sequences splits into (now inflated) blocks, e.g. $A = \alpha_0 \alpha_1 \dots \alpha_{n-1}$, where $\alpha_i = \text{inflate}^{\lceil \log_2 \delta \rceil}(\alpha_k^i)$.

We subtract from A and B a constant large enough for all their elements to be smaller than all elements of every X_i and Y_j . Similarly, we add to A' and B' a constant large enough for all their elements to be larger than all elements of every X_i and Y_j . Now, we can construct the sequences X and Y as follows:

$$X = \alpha_0 X_0 \hat{\alpha}_0 \alpha_1 X_1 \hat{\alpha}_1 \dots \alpha_{n-1} X_{n-1} \hat{\alpha}_{n-1},$$

$$Y = \beta_0 Y_0 \hat{\beta}_0 \beta_1 Y_1 \hat{\beta}_1 \dots \beta_{n-1} Y_{n-1} \hat{\beta}_{n-1}.$$

We claim that

$$\text{lcis}(X, Y) = \ell \cdot (4n - 2) + M, \text{ where } M = \max_{i,j} \text{lcis}(X_i, Y_j).$$

Let X_i and Y_j be the pair of sequences achieving $\text{lcis}(X_i, Y_j) = M$. Recall that $\text{lcis}(\alpha_0 \dots \alpha_i, \beta_0 \dots \beta_j) = \ell \cdot (i + j + n)$, with all the elements of this common subsequence preceding the elements of X_i and Y_j in X and Y , respectively, and being smaller than them. In the same way $\text{lcis}(\hat{\alpha}_i \dots \hat{\alpha}_{n-1}, \hat{\beta}_j \dots \hat{\beta}_{n-1}) = \ell \cdot (2 \cdot (n - 1) - (i + j) + n)$ with all the elements of LCIS being greater and appearing later than those of X_i and Y_j . By concatenating these three sequences we obtain a common increasing subsequence of X and Y of length $\ell \cdot (4n - 2) + M$.

It remains to prove $\text{lcis}(X, Y) \leq \ell \cdot (4n - 2) + M$. Let L be any common increasing subsequence of X and Y . Observe that L must split into three (some of them possibly empty) parts $L = SG\hat{S}$ with S consisting only of elements of A and B , G – only elements of X_i and Y_j , and \hat{S} – elements of \hat{A} and \hat{B} .

Let x be the last element of S and \hat{x} the first element of \hat{S} . We know that x belongs to some blocks α_i of A and β_j of B , and \hat{x} belongs to some blocks $\hat{\alpha}_{\hat{i}}$ of \hat{A} and $\hat{\beta}_{\hat{j}}$ of \hat{B} . Obviously $i \leq \hat{i}$ and $j \leq \hat{j}$. By Lemma 3 and inflation properties we have $|S| \leq \ell \cdot (i + j + n)$ and $|\hat{S}| \leq \ell \cdot (2 \cdot (n - 1) - (\hat{i} + \hat{j}) + n)$. We consider two cases:

Case 1. If $i = \hat{i}$ and $j = \hat{j}$, then G may only contain elements of X_i and Y_j . Therefore

$$|L| \leq |S| + \text{lcis}(X_i, Y_j) + |\hat{S}| \leq \ell \cdot (4n - 2) + M.$$

Case 2. If $i < \hat{i}$ or $j < \hat{j}$, then G must be a strictly increasing subsequence of both $X_i \circ \dots \circ X_{\hat{i}}$ and $Y_j \circ \dots \circ Y_{\hat{j}}$ therefore its length can be bounded by

$$\begin{aligned} |G| &\leq \min(\delta \cdot (\hat{i} - i + 1), \delta \cdot (\hat{j} - j + 1)) \leq \ell \cdot (\min(\hat{i} - i, \hat{j} - j) + 1) \leq \\ &\leq \ell \cdot (\min(\hat{i} - i, \hat{j} - j) + \max(\hat{i} - i, \hat{j} - j)) = \ell \cdot (\hat{i} - i + \hat{j} - j). \end{aligned}$$

On the other hand, $|S| + |\hat{S}| \leq \ell \cdot (4n - 2 - (\hat{i} - i) - (\hat{j} - j))$. From that we obtain $|L| \leq \ell \cdot (4n - 2)$, as desired. □

We are ready to prove the main result of the paper.

Proof of Theorem 1 Let $\mathcal{U} = \{u_0, \dots, u_{n-1}\}$, $\mathcal{V} = \{v_0, \dots, v_{n-1}\}$ be two sets of binary vectors in d dimensions. In Sect. 3.3 we constructed vector gadgets U_i and V_j , for $i, j \in \{0, 1, \dots, n - 1\}$, such that $\text{lcis}(U_i, V_j) = d - (u_i \cdot v_j)$. To these sequences we apply Lemma 4, with $\delta = 2d$, obtaining sequences X and Y of length $\mathcal{O}(n \log n \text{poly}(d))$ such that $\text{lcis}(X, Y) = C + d - \min_{i,j} (u_i \cdot v_j)$ for a constant C . This reduction, combined with an $\mathcal{O}(n^{2-\epsilon})$ time algorithm for LCIS, would yield an

$\mathcal{O}(n^{2-\epsilon} \text{polylog}(n) \text{poly}(d))$ algorithm for OV, refuting Hypothesis 3 and, in particular, SETH. □

With the reduction above, one can not only determine whether there exist a pair of orthogonal vectors or not, but also, in the latter case, calculate the minimum inner product over all pairs of vectors. Formally, by the above construction, we can reduce even the Most Orthogonal Vectors problem, as defined in Abboud et al. [1] to LCIS. This bases hardness of LCIS already on the inability to improve over exhaustive search for the MAX-CNF-SAT problem, which is a slightly weaker conjecture than SETH.

4 Matching Lower Bound for Output-Dependent Algorithms

To prove our bivariate conditional lower bound of $(nL)^{1-o(1)}$, we provide a reduction from an OV instance with unequal vector set sizes.

Proof of Theorem 2 Let $0 < \gamma \leq 1$ be arbitrary and consider any OV instance with sets $\mathcal{U}, \mathcal{V} \subseteq \{0, 1\}^d$ with $|\mathcal{U}| = n, |\mathcal{V}| = m = n^\gamma$ and $d = n^{o(1)}$. We reduce this problem, in linear time in the output size, to an LCIS instance with sequences X and Y satisfying $|X| = |Y| = \mathcal{O}(nd \log n)$ and an LCIS of length $\mathcal{O}(n^\gamma d)$. Theorem 2 is an immediate consequence of the reduction: an $\mathcal{O}((nL)^{1-\epsilon})$ time LCIS algorithm would yield an OV algorithm running in time $\mathcal{O}(n^{1+\gamma-\epsilon'})$, which would refute Hypothesis 3 and, in particular, SETH.

It remains to show the reduction itself. Let $\mathcal{U} = \{u_0, \dots, u_{n-1}\}$ and $\mathcal{V} = \{v_0, \dots, v_{m-1}\}$ be two sets of d -dimensional $(0, 1)$ -vectors. By padding \mathcal{U} , if necessary, with some dummy 1^d vectors, we can assume without loss of generality that $n = q \cdot m$ for some integer q .

We start with the vector gadgets U_i and V_j from Sect. 3.4. This time, however, we group together every q consecutive gadgets, i.e., $(U_0, \dots, U_{q-1}), (U_q, \dots, U_{2q-1})$, and so on. Specifically, let $U_i^{[r]}$ be the i -th vector gadget shifted by an integer r (i.e. with r added to all its elements). We define, for each $l \in \{0, 1, \dots, m - 1\}$,

$$\bar{U}_l = U_{lq}^{[2qd]} U_{lq+1}^{[2qd-2d]} \dots U_{lq+q-1}^{[2d]}.$$

In a similar way, for $j \in \{0, 1, \dots, m - 1\}$, we replicate every V_j gadget q times with appropriate shifts, i.e.,

$$\bar{V}_j = V_j^{[2qd]} V_j^{[2qd-2d]} \dots V_j^{[2d]}.$$

Let us now determine $\text{lcs}(\bar{U}_l, \bar{V}_j)$. No two gadgets grouped in \bar{U}_l can contribute to an LCIS together, as the later one would have smaller elements. Therefore, only one U_i gadget can be used, paired with the one copy of V_j having the matching shift. This yields $\text{lcs}(\bar{U}_l, \bar{V}_j) = \max_{lq \leq i < lq+q} \text{lcs}(U_i, V_j)$, and in turn, also $\max_{l,j} \text{lcs}(\bar{U}_l, \bar{V}_j) = \max_{i,j} \text{lcs}(U_i, V_j) = d - \min_{i,j} (u_i \cdot v_j)$.

Observe that every \bar{U}_l is a concatenation of several U_i gadgets, each one shifted to make its elements smaller than previous ones. Therefore, any increasing subsequence

of \bar{U}_l must be contained in a single U_i , and thus cannot be longer than $2d$. The same argument applies to every \bar{V}_j . Therefore, we can apply Lemma 4, with $\delta = 2d$, to these sequences, obtaining \bar{X} and \bar{Y} satisfying:

$$\text{lcis}(\bar{X}, \bar{Y}) = C + d - \min_{i,j} (u_i \cdot v_j).$$

Recall that C is some constant dependent only on m and d , and $C = \mathcal{O}(md)$. The length of both \bar{X} and \bar{Y} is $\mathcal{O}(dm \log m + mqd) = \mathcal{O}(nd \log n)$, and the length of the output is $\mathcal{O}(md)$, as desired. \square

5 Hardness of k -LCIS

In this section we show that, assuming SETH, there is no $\mathcal{O}(n^{k-\epsilon})$ algorithm for the k -LCIS problem, i.e., we prove Theorem 3. To obtain this lower bound we show a reduction from the k -Orthogonal Vectors problem (for definition, see Sect. 2). There are two main ingredients of the reduction, i.e. separator sequences and vector gadgets, and both of them can be seen as natural generalizations of those introduced in Sect. 3.

5.1 Generalizing Separator Sequences

Please note that in this section we use a notation which is not consistent with the one from Sect. 3, because it has to accommodate indexing over k sequences.

The aim of this section is to show, for any N that is a power of two, how to construct k sequences A_1, A_2, \dots, A_k such that each of them can be split into N blocks, i.e. $A_i = \alpha_i^0 \alpha_i^1 \dots \alpha_i^{N-1}$, and for any choice of $j_1, j_2, \dots, j_k \in \{0, 1, \dots, N - 1\}$

$$\text{lcis}(\alpha_1^0 \dots \alpha_1^{j_1}, \alpha_2^0 \dots \alpha_2^{j_2}, \dots, \alpha_k^0 \dots \alpha_k^{j_k}) = j_1 + j_2 + \dots + j_k + N. \tag{2}$$

As before, we construct separator sequences inductively, doubling the number of blocks in each step. Again, for $N = 1$, we define the sequences by $A_i = \langle 1 \rangle, i \in \{1, \dots, k\}$.

Suppose we have N -block sequences $A_1, A_2, \dots, A_k, A_i = \alpha_i^0 \alpha_i^1 \dots \alpha_i^{N-1}$ as above. We show how to construct $2N$ -block sequences $B_1, B_2, \dots, B_k, B_i = \beta_i^0 \beta_i^1 \dots \beta_i^{2N-1}$. Note that inflation properties still hold for k sequences, as the proof of Lemma 1 works in exactly the same way, i.e. inflating all the sequences increases their LCIS by a factor of 2.

To obtain B_i , we first inflate A_i , and then append a tail gadget after each block α_i^j . However, tail gadgets are now more involved.

Let s denote the largest element appearing in A_1, A_2, \dots, A_k . Then the blocks of B_i are

$$\beta_i^{2j} = \text{inflate}(\alpha_i^j) \circ T_i^0, \quad \beta_i^{2j+1} = T_i^1,$$

where T_i^0 is the sorted sequence of numbers of the form $2s + x$ for $x \in \{1, \dots, 2^k - 1\}$ such that the i -th bit in the binary representation of x equals 0, while T_i^1 contains those with i -th bit set to 1. Note that for $k = 2$ this exactly leads to the construction from Sect. 3.

During one construction step, every block doubles its size, and constant number of elements (precisely, $2^k - 1$) is added for every original block. Therefore, the length $L(N)$ of N -block sequences satisfies the recursive equation:

$$L(2N) = 2 \cdot L(N) + (2^k - 1) \cdot N$$

which yields $L(N) = \mathcal{O}(N \log N)$. Note also that the size of the alphabet $S(N)$ used in N -block sequences gives the equation $S(2N) = 2S(N) + 2^k - 1$, as a constant number of elements is added in every step. Therefore $S(N) = \mathcal{O}(N)$.

Lemma 5 *The constructed sequences satisfy*

$$\text{lcis}(\beta_1^0 \dots \beta_1^{j_1}, \beta_2^0 \dots \beta_2^{j_2}, \dots, \beta_k^0 \dots \beta_k^{j_k}) = j_1 + j_2 + \dots + j_k + 2N,$$

for any $(j_1, j_2, \dots, j_k) \in \{0, 1, \dots, 2N - 1\}$.

Proof We prove the claim by induction on $j_1 + j_2 + \dots + j_k$. In fact, to make the induction work, we need to prove a stronger statement that there always exists a corresponding LCIS that ends on an element less than or equal to $2s + x(j_1, \dots, j_k)$, where $x(j_1, \dots, j_k)$ is the integer given by the binary representation $(j_1 \bmod 2, \dots, j_k \bmod 2)$.

By the inflation properties and the observation that T_1^0, \dots, T_k^0 have no common elements, we obtain $\text{lcis}(\beta_1^0, \dots, \beta_k^0) = 2 \cdot \text{lcis}(\alpha_1^0, \dots, \alpha_k^0) = 2N$, with a corresponding LCIS using only elements bounded by $2s$, which settles the base case for induction.

Let j_1, j_2, \dots, j_k be indices with $j_1 + \dots + j_k > 0$. Let us first construct a common increasing subsequence of length at least $j_1 + \dots + j_k + 2N$. If all indices j_1, \dots, j_k are even, then, for every $i \in \{1, \dots, k\}$, the prefix $\beta_i^0 \dots \beta_i^{j_i}$ contains $\text{inflate}(\alpha_i^0 \dots \alpha_i^{j_i/2})$ as a subsequence. Thus we can find, by inflation properties, a common increasing subsequence of length $2 \cdot (j_1/2 + \dots + j_k/2 + N) = j_1 + \dots + j_k + 2N$, as desired. Now, let j_i be any odd index, and let L be the LCIS of the prefixes corresponding to $j_1, \dots, j_{i-1}, j_i - 1, j_{i+1}, \dots, j_k$, which ends on an element bounded by $x(j_1, \dots, j_{i-1}, 0, j_{i+1}, \dots, j_k)$, of length $j_1 + \dots + j_k + 2N - 1$ (which exists by the induction hypothesis). Then $L \circ x(j_1, \dots, j_{i-1}, 1, j_{i+1}, \dots, j_k)$ is an LCIS for the prefixes corresponding to j_1, \dots, j_k : Indeed, $2s + x(j_1, \dots, j_k)$ is a common member of $T_1^{j_1 \bmod 2}, \dots, T_k^{j_k \bmod 2}$, the last parts of these prefixes, and this element is larger and appears later in the sequences than all elements in L (since all T_i^j 's are sorted in the increasing order).

For the converse, let L denote the LCIS of $\beta_1^0 \dots \beta_1^{j_1}, \beta_2^0 \dots \beta_2^{j_2}, \dots, \beta_k^0 \dots \beta_k^{j_k}$. Note that if the last symbol of L does not come from the last blocks, i.e. $\beta_1^{j_1}, \beta_2^{j_2}, \dots, \beta_k^{j_k}$, then L is an LCIS of prefixes corresponding to some j'_1, \dots, j'_k with $j'_1 + \dots + j'_k <$

$j_1 + \dots + j_k$ and the claim follows from the induction hypotheses. Thus, we may assume that L ends on a common symbol of the last blocks.

If all the indices are even, the last blocks share only elements less than or equal to $2s$ (since T_1^0, \dots, T_k^0 share no elements), thus L is the LCIS of inflate($\alpha_i^0, \dots, \alpha_i^{j_i/2}$), $i \in \{1, \dots, k\}$ and the claim follows from the inflation properties. Otherwise, the only element the last blocks have in common is $x(j_1, j_2, \dots, j_k)$, and thus $L = L' \circ x(j_1, \dots, j_k)$, where L' is the LCIS of prefixes corresponding to some j'_1, \dots, j'_k with $j'_1 + \dots + j'_k < j_1 + \dots + j_k$. Thus, $|L| \leq j'_1 + \dots + j'_k + 2N + 1 \leq j_1 + \dots + j_k + 2N$, as desired. \square

5.2 Generalizing Vector Gadgets

Each vector gadget is the concatenation of *coordinate gadgets*. Coordinate gadgets for j -th coordinate use elements from the range $\{kj + 1, \dots, kj + k\}$. If a coordinate is 0, the corresponding gadget contains all k elements sorted in decreasing order, otherwise the gadget for the i -th sequence skips the $kj + i$ element. Formally,

$$VG_i(u) = CG_i^0(u[0]) \circ CG_i^1(u[1]) \circ \dots \circ CG_i^{d-1}(u[d-1]),$$

where

$$CG_i^j(0) = \langle kj + k, kj + (k - 1), \dots, kj + 1 \rangle,$$

$$CG_i^j(1) = \langle kj + k, kj + (k - 1), \dots, kj + (i + 1), kj + (i - 1), \dots, kj + 1 \rangle.$$

Thus, if all k vectors have the j -th coordinate equal 1, there is no common element in the corresponding gadgets. Otherwise, if at least one, say i -th, vector has the j -th coordinate equal 0, the element $kj + i$ appears in all coordinate gadgets. Since the coordinate gadgets are sorted in decreasing order, their LCIS cannot exceed 1. Therefore,

$$lcis(CG_1^j(u_1), CG_2^j(u_2), \dots, CG_k^j(u_k)) = 1 - \prod_{i=1}^k u_i[j],$$

and ultimately

$$lcis(VG_1(u_1), VG_2(u_2), \dots, VG_k(u_k)) = d - \sum_{j=0}^{d-1} \prod_{i=1}^k u_i[j].$$

5.3 Putting Pieces Together

We can finally prove our lower bound for k -LCIS, i.e., Theorem 3.

Proof of Theorem 3 Let $\mathcal{U}_1, \dots, \mathcal{U}_k \subseteq \{0, 1\}^d$ be a k -OV instance with $|\mathcal{U}_i| = n$. By at most doubling the number of vectors in each set, we may assume without loss of generality that n is a power of two.

We construct separator sequences consisting of n blocks. Inflate the sequences $\lceil \log_2 kd \rceil$ times, thus increasing their length by a factor $\ell = 2^{\lceil \log_2 kd \rceil}$, and subtract from all their elements a constant large enough for them to become smaller than all elements of vector gadgets. Let $A_i = \alpha_i^0 \dots \alpha_i^{n-1}$ denote the thus constructed separator sequence corresponding to set \mathcal{U}_i .

Analogously (and as in the proof of Theorem 1), we construct, for each $i \in \{1, \dots, k\}$, the separator sequence $\hat{A}_i = \hat{\alpha}_i^0, \dots, \hat{\alpha}_i^{n-1}$ by reversing A_i , replacing each element by its additive inverse, and adding a constant large enough to make all the elements larger than vector gadgets (note that each $\hat{\alpha}_i^j$ equals the reverse of α_i^{n-j-1} , with negated elements, shifted by an additive constant). In this way, the analogous property to Equation (2) holds for suffixes instead of prefixes.

Finally, we construct sequences X_1, X_2, \dots, X_k by defining

$$X_i = \alpha_i^0 \text{VG}_i(u_i^0) \hat{\alpha}_i^0 \alpha_i^1 \text{VG}_i(u_i^1) \hat{\alpha}_i^1 \dots \alpha_i^{n-1} \text{VG}_i(u_i^{n-1}) \hat{\alpha}_i^{n-1},$$

where the VG_i are defined as in Sect. 5.2. It is straightforward to rework the proof of Theorem 1 to verify that these sequences fulfill

$$\text{lcis}(X_1, X_2, \dots, X_k) = \ell \cdot (k \cdot (n - 1) + 2n) + d - m,$$

where $m = \min_{u_1 \in \mathcal{U}_1, u_2 \in \mathcal{U}_2, \dots, u_k \in \mathcal{U}_k} \sum_{j=0}^{d-1} \prod_{i=1}^k u_i[j]$.

By this reduction, an $\mathcal{O}(n^{k-\varepsilon})$ time algorithm for k -LCIS would yield an $\mathcal{O}(n^{k-\varepsilon'})$ time k -OV algorithm (for any dimension $d = n^{o(1)}$), thus refuting Hypothesis 2 and, in particular, SETH. \square

6 Hardness of k -LCWIS

We shortly discuss the proof of Theorem 4.

Proof sketch of Theorem 4 Note that our lower bound for k -LCIS almost immediately yields a lower bound for k -LCWIS: Clearly, each common increasing subsequence of X_1, \dots, X_k is also a common *weakly* increasing subsequence. The claim then follows after carefully verifying that, in the constructed sequences, we cannot obtain longer common weakly increasing subsequences by reusing some symbols.

Our claim for k -LCWIS is slightly stronger, however. In particular, we aim to reduce the size of the alphabet over which all the sequences used in the reduction are defined. For this, the key insight is to replace the *inflation* operation $\text{inflate}(\langle a_0, \dots, a_{n-1} \rangle) = \langle 2a_0 - 1, 2a_0, \dots, 2a_{n-1} - 1, 2a_{n-1} \rangle$ by

$$\text{inflate}'(\langle a_0, \dots, a_{n-1} \rangle) = \langle a_0, a_0, \dots, a_{n-1}, a_{n-1} \rangle,$$

which does not increase the alphabet size, but still satisfies the desired property for k -LCWIS.

Replacing this notion in the proof of Theorem 3, we obtain final sequences X_1, \dots, X_k by combining separator gadgets over alphabets of size $\mathcal{O}(\log n)$ with vector gadgets over alphabets of size $\mathcal{O}(d)$, where d is the dimension of the vectors in the k -OV instance. Correctness of this construction under k -LCWIS can be verified by reworking the proof of Theorem 3. Thus, we construct hard k -LCWIS instances over an alphabet of size $\mathcal{O}(\log n + d)$, and the claim follows. \square

7 Strengthening the Hardness

In this section we show that a natural combination of constructions proposed in the previous sections with the idea of *reachability gadgets* introduced by Abboud et al. [2] lets us strengthen our lower bounds to be derived from considerably weaker assumptions than SETH. Before we do this, we first need to introduce the notion of branching programs.

A *branching program* of width W and length T on N Boolean input variables $x_1, x_2, \dots, x_N \in \{0, 1\}$ is a directed acyclic graph on $W \cdot T$ nodes, arranged into T layers of size W each. A node in the k -th layer may have outgoing edges only to the nodes in the $(k + 1)$ -th layer, and for every layer there is a variable x_i such that every edge leaving this layer is labeled with a constraint of the form $x_i = 0$ or $x_i = 1$. There is a single *start* node in the first layer and a single *accept* node in the last layer. We say that the branching program *accepts* an input $x \in \{0, 1\}^N$ if there is a path from the start node to the accept node which uses only edges that are labeled with constraints satisfied by the input x .

The expressive power of branching programs is best illustrated by the theorem of Barrington [12]. It states that any depth- d fan-in-2 Boolean circuit can be expressed as a branching program of width 5 and length 4^d . In particular, NC-circuits can be expressed as constant width quasipolynomial length branching programs.

Given a branching program P on N input variables, the Branching Program Satisfiability problem (BP-SAT) asks if there exists an assignment $x \in \{0, 1\}^N$ such that P accepts x . Abboud et al. [2] gave a reduction from BP-SAT to LCS (and some other related problems, such as Edit Distance) on two sequences of length $2^{N/2} \cdot T^{\mathcal{O}(\log W)}$. The reduction proves that a strongly subquadratic algorithm for LCS would imply, among others, exponential improvements over exhaustive search for satisfiability problems not only on CNF formulas (i.e. refuting SETH), but even NC-circuits and circuits representing $o(\sqrt{n})$ -space nondeterministic Turing machines. Moreover, even a sufficiently large polylogarithmic improvement would imply nontrivial results in circuit complexity. We refer to the original paper [2] for an in-depth discussion of these consequences.

In this section we prove Theorem 5 and thus show that a subquadratic algorithm for LCIS would have the same consequences. Our reduction from OV to LCIS (presented in Sect. 3) is built of two ingredients: (1) relatively straightforward vector gadgets, encoding vector inner product in the language of LCIS, and (2) more involved separator sequences, which let us combine many vector gadgets into a single sequence. In

order to obtain a reduction from BP-SAT we will need to replace vector gadgets with more complex reachability gadgets. Fortunately, reachability gadgets for LCIS can be constructed in a similar manner as reachability gadgets for LCS proposed in [2].

Proof sketch of Theorem 5 Given a branching program, as in [2], we follow the split-and-list technique of Williams [44]. Assuming for ease of presentation that N is even, we split the input variables into two halves: $x_1, \dots, x_{N/2}$ and $x_{N/2+1}, \dots, x_N$. Then, for each possible assignment $a \in \{0, 1\}^{N/2}$ of the first half we list a reachability gadget $RG_X(a)$, and similarly, for each possible assignment $b \in \{0, 1\}^{N/2}$ of the second half we list a reachability gadget $RG_Y(b)$. We shall define the gadgets such that there exists a constant C (depending only on the branching program size) such that $lcis(RG_X(a), RG_Y(b)) = C$ if and only if $a \circ b$ is an assignment accepted by the branching program, and otherwise $lcis(RG_X(a), RG_Y(b)) < C$. The reduction is finished by applying Lemma 4 to the constructed gadgets in order to obtain two sequences such that their LCIS lets us determine whether a satisfying assignment to the branching program exists. The rest of the proof is devoted to constructing suitable reachability gadgets.

We assume without loss of generality that $T = 2^t + 1$ for some integer t . For every $k \in \{0, 1, \dots, t\}$ and for every two nodes u, v being 2^k layers apart from each other we want to construct two reachability gadgets $RG_X^{u \rightarrow v}$ and $RG_Y^{u \rightarrow v}$ such that, for some constant C_k ,

$$lcis(RG_X^{u \rightarrow v}(a), RG_Y^{u \rightarrow v}(b)) \begin{cases} = C_k & \text{if there is a path } u \rightsquigarrow v \text{ satisfied by } a \circ b, \\ < C_k & \text{otherwise,} \end{cases}$$

holds for all $a, b \in \{0, 1\}^{N/2}$.

Consider $k = 0$, i.e., designing reachability gadgets for nodes in neighboring layers L_j and L_{j+1} . There is a variable x_i such that all edges between L_j and L_{j+1} are labeled with a constraint $x_i = 0$ or $x_i = 1$. We say *the left half is responsible for* x_i if x_i is among the first half $x_1, \dots, x_{N/2}$ of variables; otherwise, we say *the right half is responsible for* x_i . We set $RG_X^{u \rightarrow v}(a)$ to be an empty sequence if the left half is responsible for x_i and there is no edge from u to v labeled $x_i = a_i$; otherwise, we set $RG_X^{u \rightarrow v}(a) = \langle 0 \rangle$. Similarly, $RG_Y^{u \rightarrow v}(b)$ is an empty sequence if the right half is responsible and there is no edge from u to v labeled with $x_i = b_{i-N/2}$; otherwise $RG_Y^{u \rightarrow v}(b) = \langle 0 \rangle$. It is easy to verify that such reachability gadgets satisfy the desired property for $C_0 = 1$.

For $k > 0$, let w_1, w_2, \dots, w_W be the nodes in the layer exactly halfway between u and v . Observe that there exists a path from u to v if and only if there exists a path from u to w_i and from w_i to v for some $i \in \{1, 2, \dots, W\}$.

Let $\overline{RG}_X^{w_i \rightarrow v}$ and $\overline{RG}_Y^{w_i \rightarrow v}$ denote the sequences $RG_X^{w_i \rightarrow v}$ and $RG_Y^{w_i \rightarrow v}$ with every element increased by a constant large enough so that all elements are larger than all elements of $RG_X^{u \rightarrow w_i}$ and $RG_Y^{u \rightarrow w_i}$. Observe that $lcis(RG_X^{u \rightarrow w_i}(a) \circ \overline{RG}_X^{w_i \rightarrow v}(a), RG_Y^{u \rightarrow w_i}(b) \circ \overline{RG}_Y^{w_i \rightarrow v}(b))$ equals $2 \cdot C_{k-1}$ if there is a path $u \rightsquigarrow w_i \rightsquigarrow v$ satisfied by $a \circ b$, and otherwise it is less than $2 \cdot C_{k-1}$. Now, for every i take a different constant q_i and add it to both $RG_X^{u \rightarrow w_i} \circ \overline{RG}_X^{w_i \rightarrow v}$ and $RG_Y^{u \rightarrow w_i} \circ \overline{RG}_Y^{w_i \rightarrow v}$ so that their alphabets are disjoint, and therefore, for $i \neq j$, $lcis((RG_X^{u \rightarrow w_i}(a) \circ \overline{RG}_X^{w_i \rightarrow v}(a)) +$

$q_i, (RG_Y^{u \rightarrow w_j}(b) \circ \overline{RG_Y^{w_j \rightarrow v}}(b)) + q_j) = 0$ (where $+$ denotes element-wise addition). Finally, apply Lemma 4 to these W pairs of concatenated reachability gadgets (where we choose δ as the maximum length of these gadget) to obtain two reachability gadgets $RG_X^{u \rightarrow v}$ and $RG_Y^{u \rightarrow v}$ such that $lcis(RG_X^{u \rightarrow v}(a), RG_Y^{u \rightarrow v}(b))$ equals $C + 2 \cdot C_{k-1}$ (for a constant C resulting from the application of Lemma 4) if there exists (for some $i \in \{1, 2, \dots, W\}$) a path $u \rightsquigarrow w_i \rightsquigarrow v$ satisfied by $a \circ b$, and is strictly smaller otherwise, as desired.

Let u_{start} and u_{accept} denote the start node and the accept node of the branching program. Then, $RG_X = RG_X^{u_{start} \rightarrow u_{accept}}$ and $RG_Y = RG_Y^{u_{start} \rightarrow u_{accept}}$ satisfy the property that

$$lcis(RG_X(a), RG_Y(b)) \begin{cases} = C_t & \text{if the branching program accepts } a \circ b, \\ < C_t & \text{otherwise.} \end{cases}$$

Since $RG_X(a)$ and $RG_Y(b)$ are constructed in t steps of the inductive construction, and each step increases the length of gadgets by a factor of $\mathcal{O}(W \log W)$, their final length can be bounded by $\mathcal{O}((W \log W)^t)$, which is $T^{\mathcal{O}(\log W)}$. Combining the reachability gadgets $RG_X(a), a \in \{0, 1\}^{N/2}$ and $RG_Y(b), b \in \{0, 1\}^{N/2}$ using Lemma 4 (where we choose δ as the maximum length of the reachability gadgets) yields the desired strings X, Y of length $2^{N/2} \cdot N \cdot T^{\mathcal{O}(\log W)}$ whose LCIS lets us determine satisfiability of the given branching program, thus finishing the proof. \square

Similar techniques can be used to analogously strengthen other lower bounds in our paper.

8 Conclusion and Open Problems

We prove a tight quadratic lower bound for LCIS, ruling out strongly subquadratic time algorithms under SETH. It remains open whether LCIS admits mildly subquadratic algorithms, such as the Masek-Paterson algorithm for LCS [36]. Note, however, that our reduction from BP-SAT gives evidence that shaving many logarithmic factors is immensely difficult. Finally, we give tight SETH-based lower bounds for k -LCIS.

For the related variant LCWIS that considers weakly increasing sequences, strongly subquadratic-time algorithms are ruled out under SETH for slightly superlogarithmic alphabet sizes ([40] and Theorem 4). On the other hand, for binary and ternary alphabets, even linear time algorithms exist [24,35]. Can LCWIS be solved in time $\mathcal{O}(n^{2-f(|\Sigma|)})$ for some decreasing function f that yields strongly subquadratic-time algorithms for any constant alphabet size $|\Sigma|$?

Finally, by an easy observation (see the appendix), we can compute a $(1 + \varepsilon)$ -approximation of LCIS in $\mathcal{O}(n^{3/2} \varepsilon^{-1/2} \text{polylog}(n))$ time. Can we improve upon this running time or give a matching conditional lower bound? Note that a positive resolution seems difficult by the reduction in Observation 1: Any $n^\alpha, \alpha > 0$, improvement over this running time would yield a strongly subcubic $(1 + \varepsilon)$ -approximation for

3-LCS, which seems hard to achieve, given the difficulty to find strongly subquadratic $(1 + \varepsilon)$ -approximation algorithms for LCS.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

Appendix

Theorem 6 (folklore, generalization of [47]) *For any $k \geq 2$, the LCIS of k sequences of length n can be computed in $\mathcal{O}(n^k)$ time.*

Proof Let X_1, X_2, \dots, X_k be the input sequences. Let $X[0 : i]$ denote the prefix consisting of first i elements of X , with $X[0 : 0]$ being the empty prefix. Now, for every $i_1, \dots, i_k \in \{0, 1, \dots, n\}$ we define $R[i_1, \dots, i_k]$ to be the length of the LCIS of the prefixes $X_1[0 : i_1], X_2[0 : i_2], \dots, X_k[0 : i_k]$ with an additional assumption that this common subsequence must end with the element $X_k[i_k]$, i.e. the last element of the last prefix. Observe that it is enough to compute all $R[i_1, \dots, i_k]$, with the desired answer being simply $\max_{0 \leq i_k \leq n} R[n, \dots, n, i_k]$.

The algorithm is based on the fact that $R[i_1, \dots, i_k]$ satisfies the following two-case recurrence:

- **Case 1.** If $X_k[i_k] \neq X_s[i_s]$ for some $s < k$, the desired common subsequence ends with $X_k[i_k]$ and thus it cannot contain $X_s[i_s]$, so $R[i_1, \dots, i_s, \dots, i_k] = R[i_1, \dots, i_s - 1, \dots, i_k]$.
- **Case 2.** If $X_1[i_1] = \dots = X_k[i_k]$, let us call this common symbol σ , and observe that σ is the last element of the LCIS. Consider the next-to-last element: it must be certainly smaller than σ , and must appear in the X_k sequence at a position earlier than at $X_k[i_k]$. Therefore $R[i_1, \dots, i_k] = 1 + \max_{j < i_k, X_k[j] < \sigma} R[i_1 - 1, \dots, i_{k-1} - 1, j]$.

To obtain the values of R in $\mathcal{O}(n^k)$ time, the algorithm iterates through all possible i_1, \dots, i_k with the i_k loop being the innermost one. Obviously, $R[i_1, \dots, i_k] = 0$ if any of the indices is 0. Before every innermost i_k loop, with fixed i_1, i_2, \dots, i_{k-1} , the algorithm checks whether $X_1[i_1] = \dots = X_{k-1}[i_{k-1}]$. If so, it sets $\sigma = X_1[i_1] = \dots = X_{k-1}[i_{k-1}]$, otherwise $\sigma = \text{null}$.

If $\sigma \neq \text{null}$, for every $1 \leq i \leq n$ let $D[i] = \max_{j < i, X_k[j] < \sigma} R[i_1 - 1, \dots, i_{k-1} - 1, j]$. Observe that $D[i]$ can be obtained from $D[i - 1]$ and $R[i_1 - 1, \dots, i_{k-1} - 1, i - 1]$ in constant time. Therefore, before the start of the i_k loop, the algorithm can precompute all the $D[i]$ values in $\mathcal{O}(n)$ time, as all the needed $R[i_1 - 1, \dots, i_{k-1} - 1, i]$ values are already known from earlier iterations.

Throughout the i_k loop the algorithm checks if $X_k[i_k] = \sigma$, which corresponds to Case 2 above. If so, then $R[i_1, \dots, i_k] = 1 + \max_{j < i_k, X_k[j] < \sigma} R[i_1 - 1, \dots, i_{k-1} - 1, j] = 1 + D[i_k]$, which is already precomputed. If Case 1 holds, then $R[i_1, \dots, i_s, \dots, i_k] = R[i_1, \dots, i_s - 1, \dots, i_k]$ for some $s < k$. As the index s is

easy to find, and the necessary values in R have been computed earlier, this step also works in constant time (assuming k is fixed).

The above algorithm computes only the length of LCIS. However, it can be easily modified to reconstruct the sequence, using the common dynamic programming techniques (e.g. by storing with every value in R a link to the previous element of LCIS). \square

Theorem 7 *A $(1 + \varepsilon)$ -approximation of LCIS of sequences X, Y of length n can be computed in $\mathcal{O}(n^{3/2}\varepsilon^{-1/2}\text{polylog}(n))$ time.*

Proof First, delete all integers occurring more than $2\sqrt{n/\varepsilon}$ times in total in both of the sequences. Since there are at most $\sqrt{n\varepsilon}$ such integers, this operation decreases the length of the LCIS by at most $\sqrt{n\varepsilon}$. In the resulting instance, there are at most $n^{3/2}\varepsilon^{-1/2}$ matching pairs, i.e., indices i, j with $X[i] = Y[j]$. Thus, the exact LCIS in this instance can be computed in time $\mathcal{O}(n^{3/2}\varepsilon^{-1/2}\log n \log \log n)$ using an algorithm of Chan et al. [18] running in time $\mathcal{O}(M \log L \log n + n \log n)$, where L is the length of the LCIS of X and Y and M is the number of matching pairs. Now, consider two cases. If the algorithm returns a solution Z longer than $\sqrt{n/\varepsilon}$, then Z is a $(1 + \varepsilon)$ -approximation of the LCIS of the original instance, since the LCIS is bounded by $L \leq |Z| + \sqrt{n\varepsilon} \leq (1 + \varepsilon)|Z|$. In the remaining case, it is guaranteed that $L \leq |Z| + \sqrt{n\varepsilon} \leq (1 + \varepsilon)\sqrt{n/\varepsilon}$. Thus, we may compute the exact LCIS in $\mathcal{O}(n^{3/2}\varepsilon^{-1/2}\log n)$ time using the algorithm running in $\mathcal{O}(nL \log \log n + n \log n)$ time [35]. \square

References

1. Abboud, A., Backurs, A., Vassilevska Williams, V.: Quadratic-time hardness of LCS and other sequence similarity measures. In: Proceedings of the 56th Annual IEEE Symposium on Foundations of Computer Science (FOCS'15), pp. 59–78 (2015)
2. Abboud, A., Hansen, T.D., Vassilevska Williams, V., Williams, R.: Simulating branching programs with edit distance and friends or: a polylog shaved is a lower bound made. In: Proceedings of 48th Annual ACM Symposium on Symposium on Theory of Computing (STOC'16), pp. 375–388 (2016)
3. Abboud, A., Williams, V.V., Weimann, O.: Consequences of faster alignment of sequences. In: Proceedings of 41st International Colloquium on Automata, Languages, and Programming (ICALP'14), pp. 39–51 (2014)
4. Aho, A.V., Hirschberg, D.S., Ullman, J.D.: Bounds on the complexity of the longest common subsequence problem. *J. ACM* **23**(1), 1–12 (1976)
5. Altschul, S.F., Gish, W., Miller, W., Myers, E.W., Lipman, D.J.: Basic local alignment search tool. *J. Mol. Biol.* **215**(3), 403–410 (1990)
6. Ann, H., Yang, C., Tseng, C.: Efficient polynomial-time algorithms for the constrained LCS problem with strings exclusion. *J. Comb. Optim.* **28**(4), 800–813 (2014)
7. Apostolico, A., Guerra, C.: The longest common subsequence problem revisited. *Algorithmica* **2**(1), 316–336 (1987)
8. Arslan, A.N., Egecioglu, Ö.: Algorithms for the constrained longest common subsequence problems. *Int. J. Found. Comput. Sci.* **16**(6), 1099–1109 (2005)
9. Backurs, A., Indyk, P.: Edit distance cannot be computed in strongly subquadratic time (unless SETH is false). In: Proc. 47th Annual ACM Symposium on Theory of Computing (STOC'15), pp. 51–58 (2015)
10. Backurs, A., Indyk, P.: Which regular expression patterns are hard to match? In: Proceedings of 57th Annual Symposium on Foundations of Computer Science, (FOCS'16), pp. 457–466 (2016)
11. Backurs, A., Tzamos, C.: Improving viterbi is hard: Better runtimes imply faster clique algorithms. In: Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6–11 August 2017, pp. 311–321 (2017)

12. Barrington, D.A.: Bounded-width polynomial-size branching programs recognize exactly those languages in NC1. *J. Comput. Syst. Sci.* **38**(1), 150–164 (1989)
13. Benson, G., Levy, A., Maimoni, S., Noifeld, D., Shalom, B.R.: Lcsk: a refined similarity measure. *Theor. Comput. Sci.* **638**, 11–26 (2016)
14. Bergroth, L., Hakonen, H., Raita, T.: A survey of longest common subsequence algorithms. In: *Proc. 7th International Symposium on String Processing and Information Retrieval (SPIRE'00)*, pp. 39–48 (2000)
15. Bringmann, K.: Why walking the dog takes time: Frechet distance has no strongly subquadratic algorithms unless SETH fails. In: *Proc. 55th Annual IEEE Symposium on Foundations of Computer Science (FOCS'14)*, pp. 661–670 (2014)
16. Bringmann, K., Künnemann, M.: Quadratic conditional lower bounds for string problems and dynamic time warping. In: *Proceedings of the 56th Annual IEEE Symposium on Foundations of Computer Science (FOCS'15)*, pp. 79–97 (2015)
17. Bringmann, K., Künnemann, M.: Multivariate fine-grained complexity of longest common subsequence. In: *Proc. 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'18)*, pp. 1216–1235 (2018)
18. Chan, W., Zhang, Y., Fung, S.P.Y., Ye, D., Zhu, H.: Efficient algorithms for finding a longest common increasing subsequence. *J. Comb. Optim.* **13**(3), 277–288 (2007)
19. Chen, Y.C., Chao, K.M.: On the generalized constrained longest common subsequence problems. *J. Comb. Optim.* **21**(3), 383–392 (2011)
20. Chin, F.Y.L., Santis, A.D., Ferrara, A.L., Ho, N.L., Kim, S.K.: A simple algorithm for the constrained sequence problems. *Inf. Process. Lett.* **90**(4), 175–179 (2004)
21. Chvatal, V., Klarner, D.A., Knuth, D.E.: Selected combinatorial research problems. *Tech. Rep. CS-TR-72-292*, Stanford University, Department of Computer Science (1972)
22. Crochemore, M., Porat, E.: Fast computation of a longest increasing subsequence and application. *Inf. Comput.* **208**(9), 1054–1059 (2010)
23. Cygan, M., Mucha, M., Wegrzycki, K., Włodarczyk, M.: On problems equivalent to (min,+)-convolution. In: *Proceedings of 44th International Colloquium on Automata, Languages, and Programming (ICALP'17)*, pp. 22:1–22:15 (2017)
24. Duraj, L.: A linear algorithm for 3-letter longest common weakly increasing subsequence. *Inf. Process. Lett.* **113**(3), 94–99 (2013)
25. Fredman, M.L.: On computing the length of longest increasing subsequences. *Discret. Math.* **11**(1), 29–35 (1975)
26. Gotthilf, Z., Hermelin, D., Landau, G.M., Lewenstein, M.: Restricted LCS. In: *Proceedings of 17th International Symposium on String Processing and Information Retrieval (SPIRE'10)*, pp. 250–257 (2010)
27. Hirschberg, D.S.: Algorithms for the longest common subsequence problem. *J. ACM* **24**(4), 664–675 (1977)
28. Hunt, J.W., McIlroy, M.D.: An algorithm for differential file comparison. *Computing Science Technical Report 41*, Bell Laboratories (1975)
29. Hunt, J.W., Szymanski, T.G.: A fast algorithm for computing longest subsequences. *Commun. ACM* **20**(5), 350–353 (1977)
30. Impagliazzo, R., Paturi, R.: On the complexity of k-SAT. *J. Comput. Syst. Sci.* **62**(2), 367–375 (2001)
31. Impagliazzo, R., Paturi, R., Zane, F.: Which problems have strongly exponential complexity? *J. Comput. Syst. Sci.* **63**(4), 512–530 (2001)
32. Jacobson, G., Vo, K.: Heaviest increasing/common subsequence problems. In: *Proceedings of 3rd Annual Symposium Combinatorial Pattern Matching, CPM 92*, Tucson, Arizona, USA, April 29–May 1, 1992, pp. 52–66 (1992)
33. Jiang, T., Lin, G., Ma, B., Zhang, K.: The longest common subsequence problem for arc-annotated sequences. *J. Discret. Algorithms* **2**(2), 257–270 (2004)
34. Künnemann, M., Paturi, R., Schneider, S.: On the Fine-grained Complexity of One-Dimensional Dynamic Programming. In: *Proc. 44th International Colloquium on Automata, Languages, and Programming (ICALP'17)*, pp. 21:1–21:15 (2017)
35. Kutz, M., Brodal, G.S., Kaligosi, K., Katriel, I.: Faster algorithms for computing longest common increasing subsequences. *J. Discret. Algorithms* **9**(4), 314–325 (2011)
36. Masek, W.J., Paterson, M.: A faster algorithm computing string edit distances. *J. Comput. Syst. Sci.* **20**(1), 18–31 (1980)

37. Morgan, H.L.: Spelling correction in systems programs. *Commun. ACM* **13**(2), 90–94 (1970)
38. Myers, E.W.: An $O(ND)$ difference algorithm and its variations. *Algorithmica* **1**(2), 251–266 (1986)
39. Needleman, S.B., Wunsch, C.D.: A general method applicable to the search for similarities in the amino acid sequence of two proteins. *J. Mol. Biol.* **48**(3), 443–453 (1970)
40. Polak, A.: Why is it hard to beat $O(n^2)$ for longest common weakly increasing subsequence? *Inf. Process. Lett.* **132**, 1–5 (2018)
41. Roditty, L., Vassilevska Williams, V.: Fast approximation algorithms for the diameter and radius of sparse graphs. In: *Proceedings of the 45th Annual ACM Symposium on Symposium on Theory of Computing (STOC'13)*, pp. 515–524 (2013)
42. Tsai, Y.: The constrained longest common subsequence problem. *Inf. Process. Lett.* **88**(4), 173–176 (2003)
43. Wagner, R.A., Fischer, M.J.: The string-to-string correction problem. *J. ACM* **21**(1), 168–173 (1974)
44. Williams, R.: A new algorithm for optimal 2-constraint satisfaction and its implications. *Theor. Comput. Sci.* **348**(2), 357–365 (2005)
45. Williams, V.V.: Hardness of easy problems: Basing hardness on popular conjectures such as the strong exponential time hypothesis (invited talk). In: *Proceedings of 10th International Symposium on Parameterized and Exact Computation (IPEC'15)*, pp. 17–29 (2015)
46. Yang, I., Huang, C., Chao, K.: A fast algorithm for computing a longest common increasing subsequence. *Inf. Process. Lett.* **93**(5), 249–253 (2005)
47. Zhu, D., Wang, X.: A space efficient algorithm for lcis problem. In: *Proceedings of 10th International Conference Security, Privacy, and Anonymity in Computation, Communication, and Storage: SpaCCS 2017, Guangzhou, China, December 12–15, 2017*, pp. 70–77 (2017)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.