

# Improved Approximation for Longest Common Subsequence over Small Alphabets

Shyan Akmal  

MIT, EECS and CSAIL, Cambridge, MA, USA

Virginia Vassilevska Williams 

MIT, EECS and CSAIL, Cambridge, MA, USA

---

## Abstract

---

This paper investigates the approximability of the Longest Common Subsequence (LCS) problem. The fastest algorithm for solving the LCS problem exactly runs in essentially quadratic time in the length of the input, and it is known that under the Strong Exponential Time Hypothesis the quadratic running time cannot be beaten. There are no such limitations for the approximate computation of the LCS however, except in some limited scenarios. There is also a scarcity of approximation algorithms. When the two given strings are over an alphabet of size  $k$ , returning the subsequence formed by the most frequent symbol occurring in both strings achieves a  $1/k$  approximation for the LCS. It is an open problem whether a better than  $1/k$  approximation can be achieved in truly subquadratic time ( $O(n^{2-\delta})$  time for constant  $\delta > 0$ ).

A recent result [Rubinstein and Song SODA'2020] showed that a  $1/2 + \epsilon$  approximation for the LCS over a binary alphabet is possible in truly subquadratic time, provided the input strings have the same length. In this paper we show that if a  $1/2 + \epsilon$  approximation (for  $\epsilon > 0$ ) is achievable for binary LCS in truly subquadratic time when the input strings can be unequal, then for every constant  $k$ , there is a truly subquadratic time algorithm that achieves a  $1/k + \delta$  approximation for  $k$ -ary alphabet LCS for some  $\delta > 0$ . Thus the binary case is the hardest. We also show that for every constant  $k$ , if one is given two strings of *equal* length over a  $k$ -ary alphabet, one can obtain a  $1/k + \epsilon$  approximation for some constant  $\epsilon > 0$  in truly subquadratic time, thus extending the Rubinstein and Song result to all alphabets of constant size.

**2012 ACM Subject Classification** Theory of computation  $\rightarrow$  Design and analysis of algorithms

**Keywords and phrases** approximation algorithms, longest common subsequence, subquadratic

**Digital Object Identifier** 10.4230/LIPIcs.ICALP.2021.13

**Category** Track A: Algorithms, Complexity and Games

**Funding** *Shyan Akmal*: Supported by NSF Grant CCF-1909429.

*Virginia Vassilevska Williams*: Supported by an NSF CAREER Award, NSF Grant CCF-1909429, a BSF Grant BSF:2012338, a Google Research Fellowship and a Sloan Research Fellowship.

**Acknowledgements** We thank Jenny Kaufmann for suggesting several helpful revisions.

## 1 Introduction

In a large variety of applications, from spell-checkers to DNA sequence alignment, one seeks to compute how similar two given sequences of letters are. Arguably the most popular measures of sequence similarity are the *longest common subsequence (LCS)* and the *edit distance*. The LCS of two given sequences  $A$  and  $B$  (as the name suggests) measures the maximum length of a sequence whose symbols appear in both  $A$  and  $B$  in the same order. The edit distance, on the other hand, measures how far apart two strings are by counting the minimum number of insertions, deletions and substitutions of characters that must be performed on one string to transform it into the other. These two measures are related: the complement of the LCS is the version of edit distance in which one minimizes only the



© Shyan Akmal and Virginia Vassilevska Williams;  
licensed under Creative Commons License CC-BY 4.0

48th International Colloquium on Automata, Languages, and Programming (ICALP 2021).

Editors: Nikhil Bansal, Emanuela Merelli, and James Worrell; Article No. 13; pp. 13:1–13:18

Leibniz International Proceedings in Informatics



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



number of insertions and deletions (in fact, both versions of edit distance are the same up to a factor of 2, so with respect to constant factor approximation algorithms they are equivalent). Both the LCS and the edit distance of two length  $n$  strings can be computed in  $O(n^2)$  time using a classic dynamic programming approach. The fastest algorithm for both problems is the  $O(n^2/(\log n)^2)$  time algorithm of Masek and Paterson [20].

Hardness results from fine-grained complexity have shown that truly subquadratic time algorithms (those running in time  $O(n^{2-\delta})$  for some constant  $\delta > 0$ ) for LCS and edit distance cannot exist under the Strong Exponential Time Hypothesis [2, 9, 14] and other even more believable hypotheses [3, 16]. Consequently, much of the recent interest around LCS and edit distance has concerned approximation algorithms for the problems. A long chain of progress on edit distance approximation (e.g. [19, 11, 19, 10, 8, 6, 12]) has culminated in the breakthrough constant-factor approximation algorithm of [15] running in truly subquadratic time. Several improvements followed this breakthrough including [5, 18, 13], with the most recent result being a constant factor approximation algorithm running in near linear time [7].

In contrast, much less is known about how well LCS can be approximated in truly subquadratic time. For inputs over non-constant size alphabets, some fine-grained hardness results [1, 4, 16] and nontrivial super-constant approximations [17] are known. When the input strings come from a fixed alphabet of constant size  $k$ , there is a trivial  $1/k$ -approximation algorithm that returns in linear time the longest common *unary* subsequence of two inputs. Despite the simplicity of this algorithm, until recently no better constant-factor approximation algorithm was known for *any* constant size alphabet. There are also no existing hardness results that rule out better approximation algorithms.

Recently, for the case of *binary* strings of *equal* length, Rubinstein and Song [21] were able to improve upon the simple  $1/2$ -approximation algorithm described above, obtaining for some constants  $\epsilon, \delta > 0$  an  $O(n^{2-\delta})$  time algorithm that returns a  $(1/2 + \epsilon)$ -approximation. It is still open, however, whether one can obtain such an algorithm for binary strings of unequal length, and whether one can extend the result to achieve a truly subquadratic time algorithm achieving a better than  $k$ -approximation for strings over an alphabet of size  $k$ , for every constant  $k$ .

## Our results

In this paper we present two results. The first result shows that if one can obtain a truly subquadratic time, better than  $2$ -approximation algorithm for the LCS of binary strings of possibly unequal lengths, then one can use this algorithm to obtain a truly subquadratic time, better than  $k$ -approximation algorithm for the LCS of strings over a  $k$ -ary alphabet, for any constant  $k$ .

► **Theorem 1.** *For any fixed integer  $k \geq 2$  there is an  $O(n)$  time algorithm that given an instance of LCS for strings of length at most  $n$  over an alphabet of size  $k$ , reduces it to  $O(k^2)$  instances of LCS over binary strings of length at most  $n$ , so that  $(1/2 + \epsilon)$ -approximate solutions (for  $\epsilon > 0$ ) for these LCS instances can be translated in  $O(n)$  time into a  $(1/k + \epsilon_k)$  approximation of the  $k$ -ary alphabet LCS instance, where  $\epsilon_k > 0$  is a constant only depending on  $\epsilon$  and  $k$ .*

In other words, in order to beat the longstanding  $1/k$ -approximation algorithm for LCS, one merely needs to obtain a better than  $1/2$  approximation for the binary case, i.e. to extend the Rubinstein and Song result to strings of possibly unequal length.

Our second result generalizes the Rubinstein-Song result by proving that one can beat the simple  $1/k$ -approximation algorithm for every constant  $k$ , as long as the input strings have equal length.

► **Theorem 2.** *Given two strings  $A$  and  $B$  of length  $n$  over an arbitrary alphabet of size  $k$ , there exist positive constants  $\epsilon$  and  $\delta$  such that we can compute a  $1/k + \epsilon$  approximation for the longest common subsequence of  $A$  and  $B$  in  $O(n^{2-\delta})$  time.*

In fact, our algorithm can actually  $(1/k + \epsilon)$ -approximate the LCS in near-linear time. Note that our result applies only to strings with equal input lengths. The relevance of this restriction is discussed in Section 2, which also contains the proof of Theorem 1. We present the proof of Theorem 2 in Section 3.

## Preliminaries

We write approximation ratios as constants less than 1, so that for example a  $1/2 + \epsilon$  approximation algorithm for the LCS of  $A$  and  $B$  is an algorithm that returns a common subsequence of  $A$  and  $B$  with length at least  $(1/2 + \epsilon) \cdot \text{LCS}(A, B)$ .

When we discuss edit distance in the rest of the paper we mean the version of edit distance that does not allow symbol substitutions but only measures the number of insertions and deletions. For constant factor approximation algorithms, this version of the problem is equivalent to the original.

## 2 Reduction to Binary Alphabets & Input Length Conditions

We begin by showing how to reduce nontrivial constant factor approximations of LCS over large alphabets to better than  $1/2$  approximations of LCS over binary alphabets. Although we do not directly apply this reduction in our proof of Theorem 2, the reduction is elegant and motivates the approach we end up using. Moreover, the reduction works even for strings of non-equal lengths, thus showing that one merely needs to extend the Rubinfeld-Song result to non-equal length strings in order to truly improve upon the trivial alphabet-size approximation algorithm. We will need the following definition.

► **Definition 3 (Restrictions).** *Given an alphabet  $\Sigma$ , we call a subset  $\Sigma' \subseteq \Sigma$  a subalphabet. Given a string  $A$  from alphabet  $\Sigma$  the restriction of  $A$  to a subalphabet  $\Sigma'$  is the maximum subsequence of  $A$  whose characters are all in  $\Sigma'$ .*

► **Theorem 4.** *Fix integers  $s$  and  $\ell$  with  $s > \ell \geq 2$ . Suppose that there is a  $T(n, \ell)$  time algorithm that achieves an  $1/(\ell - \epsilon)$ -approximation of the LCS of two strings of length at most  $n$  from an alphabet of size  $\ell$ . Then, there is also a  $O((n + T(n, \ell)) \binom{s}{\ell})$  time algorithm that achieves an  $1/(s(1 - \epsilon/\ell))$ -approximation of the LCS of two strings of length at most  $n$  from an alphabet of size  $s$ .*

**Proof.** We will show how to reduce the LCS for two strings of length at most  $n$  over a  $s$ -ary alphabet, to the LCS for two strings of length at most  $n$  over an  $\ell$ -ary alphabet for any  $\ell < s$ . The reduction runs in  $O(n \binom{s}{\ell})$  time and produces  $\binom{s}{\ell}$  instances of  $\ell$ -ary alphabet LCS.

Let  $A$  and  $B$  two strings of length at most  $n$  over an alphabet  $\Sigma$  of size  $s$ . Let  $C$  be the longest common subsequence of  $A$  and  $B$  (we do not know  $C$ ). For the sake of argument, sort the alphabet symbols according to their number of occurrences in  $C$ .

Let  $x$  be the collection of the  $\ell$  most frequent alphabet symbols in  $C$ . Let  $C_x$  be the subsequence of  $C$  obtained by restricting  $C$  to the subalphabet of  $\Sigma$  that contains the symbols of  $x$ . Since  $x$  has the  $\ell$  most frequent symbols in  $C$ ,  $C_x$  contains at least an  $\ell/s$  fraction of  $C$ .

Now, let us describe our algorithm. Given  $A$  and  $B$ , we consider all subsets of the alphabet consisting of precisely  $\ell$  symbols (one of these subsets will be  $x$ .) For each such collection  $y$ , consider the sub-instance of the LCS instance restricted to the symbols of  $y$ . Let  $\text{OPT}(y)$  be the optimal LCS for this instance.

## 13:4 Improved Approximation for Longest Common Subsequence over Small Alphabets

We know that  $|\text{OPT}(x)| \geq |C_x| \geq (\ell/s)|C|$ . So when we consider  $y = x$ , if we can efficiently obtain a  $1/(\ell - \epsilon)$  approximation for  $\text{OPT}(x)$ , we will get a common subsequence of  $A$  and  $B$  of length at least

$$\frac{|\text{OPT}(x)|}{\ell - \epsilon} \geq \frac{|C|}{s(1 - \epsilon/\ell)},$$

and thus yields the desired approximation for the LCS of  $A$  and  $B$ . The running time is multiplied by  $\binom{s}{\ell}$  which is a constant, as long as  $s$  is. ◀

By setting  $\epsilon = \ell\delta s/(1 + \delta s)$  in the above Theorem statement, we obtain a linear time reduction from obtaining a  $1/s + \delta$ -approximation for  $s$ -ary strings to a  $1/\ell + (\delta s)/\ell$ -approximation for  $\ell$ -ary strings. We obtain Theorem 1 as a corollary:

► **Corollary 5.** *Fix an integer  $s \geq 3$  and a constant  $\delta > 0$ . The problem of obtaining a  $1/s + \delta$  approximation for the LCS of two strings from an alphabet of size  $s$  can be reduced in linear time to the problem of obtaining a  $1/2 + (\delta s)/2$  approximation for the LCS of two strings binary strings (i.e. strings from an alphabet of size two).*

**Proof.** This follows from Theorem 4 by taking  $\ell = 2$ . The reduction has  $O(s^2)$  overhead. ◀

The reason we cannot prove Theorem 2 by combining Corollary 5 with the result of [21] is that the latter gives a better than  $1/2$  approximation for strings from an alphabet of size 2 only when the input strings have *equal length*. Note that in the reduction from the proof of Theorem 4, the subsequences obtained from restrictions to subalphabets may be of different lengths even if the original strings have equal lengths.

Although at first it may seem that extending the Rubinstein-Song result to strings of differing length should not be too hard, generalizing the result does not appear straightforward. In the following we will discuss why this not simple. The main hurdles come from the lemma and algorithm used in [21] stated below.

► **Lemma 6 (LCS and Edit Distance Connection).** *For any strings  $X$  and  $Y$  of length  $n$  and  $m$  respectively, we have*

$$2 \cdot \text{LCS}(X, Y) + \text{ED}(X, Y) = n + m.$$

For the sake of completeness, we include a proof of the above lemma.

**Proof.** Consider an optimal alignment between  $X$  and  $Y$ , which matches the maximum possible number of characters of  $X$  with identical characters of  $Y$  while respecting the order in which the characters appear in each string. The characters that are matched in  $X$  and  $Y$  correspond to a longest common subsequence. This is because if there were a longer common subsequence, we could get a larger alignment by matching the characters of the subsequence in  $X$  and  $Y$ , but this would contradict the optimality of  $X$  and  $Y$ .

Similarly, the unmatched characters correspond to a minimum set of symbols that need to be deleted from  $X$  and inserted from  $Y$  to turn  $X$  into  $Y$ . If there were a smaller edit distance computation, then all the characters which were not deleted or inserted could be paired up to form a larger alignment, again contradicting optimality.

Thus there are exactly  $2 \cdot \text{LCS}(X, Y)$  characters paired up in the alignment and  $\text{ED}(X, Y)$  unmatched characters. These encompass all the characters in  $X$  and  $Y$ , and thus account for  $n + m$  symbols. ◀

► **Definition 7** (Approximating LCS through Edit Distance). *Given strings  $A$  and  $B$  of length  $n$ , the algorithm  $\text{ApproxED}(A, B)$  approximates the edit-distance between  $A$  and  $B$  and then returns the lower bound on the LCS implied by this. More precisely, the algorithm computes an approximate edit distance  $\widetilde{\text{ED}}(A, B)$  and then returns*

$$n - \frac{1}{2} \cdot \widetilde{\text{ED}}(A, B). \quad (1)$$

As stated  $\text{ApproxED}$  can use any edit distance approximation  $\widetilde{\text{ED}}$  as a black box. For concreteness, we will take  $\widetilde{\text{ED}}$  to be the edit distance algorithm from [5] which can achieve an approximation ratio of  $c$  for any constant  $c > 3$  and runs in truly subquadratic time.

In [21], the authors use the  $\text{ApproxED}$  algorithm to handle the case of binary strings with large LCS. In this case, they notice that if the LCS is large then the edit distance must be small. Then constant factor approximations to edit distance will give good lower bounds on the LCS because in the calculation from Equation (1) we are subtracting off a small quantity from the maximum possible LCS value of  $n$ .

However, if we tried extending this algorithm to the case where the inputs  $X$  and  $Y$  have lengths  $n$  and  $m$  with  $n = 100m$  (for example) by using the identity from Lemma 6, then even when the LCS is large (say of length  $(1 - \epsilon)m$  for some small positive  $\epsilon$ ) the edit distance will still be very large compared to the length of the smaller string (at least  $(99 + \epsilon)m$ ). So even a 3-approximation to edit-distance would incur massive error when trying to approximate LCS by computing

$$\frac{1}{2} \cdot (n + m - \widetilde{\text{ED}}(X, Y))$$

and the result would not give any nontrivial lower bound for the LCS. In other words, when the input strings have very different lengths it is not clear how to use approximate edit distance in general to obtain good approximations for LCS. This is essentially why the algorithm from [21] and Theorem 2 both require equal length inputs.

### 3 Extending Alphabet Size

This section proves Theorem 8 which restates Theorem 2 from the introduction slightly.

► **Theorem 8.** *Given two strings  $A$  and  $B$  of length  $n$  over an arbitrary alphabet  $\Sigma$  of size  $s$ , there exist a positive constant  $\epsilon$  such that we can compute a  $1/s + \epsilon$  approximation for the longest common subsequence  $\text{LCS}(A, B)$  of  $A$  and  $B$  in truly subquadratic time.*

Throughout the rest of this section, we assume  $A$  and  $B$  refer to strings satisfying the conditions of Theorem 8 and that  $s \geq 3$ . We begin by establishing lemmas corresponding to easy instances of the problem. The following definition is useful for identifying these easy cases.

► **Definition 9** (Balanced Strings). *Given a string  $A$  of length  $n$  from an alphabet of size  $s$  and a parameter  $\rho > 0$ , we say a string is  $\rho$ -balanced if all its character frequencies are within  $\rho n$  of  $n/s$ .*

► **Lemma 10** (Balanced Inputs, adapted from Lemma 3.2 of [21]). *For all sufficiently small  $\rho > 0$ , if either  $A$  or  $B$  is  $\rho$ -balanced, we can  $(1/s + \gamma)$ -approximate  $\text{LCS}(A, B)$  in truly subquadratic time, where  $\gamma$  is some positive constant depending on  $\rho$ .*

## 13:6 Improved Approximation for Longest Common Subsequence over Small Alphabets

**Proof.** We reduce the problem to approximating the edit distance between  $A$  and  $B$ .

Without loss of generality assume  $A$  is  $\rho$ -balanced. This means that all of its character frequencies are at least  $(1/s - \rho)n$ . So there exists a unary common subsequence of  $A$  and  $B$  of at least this length. If this is a  $(1/s + \gamma)$  approximation we are done. Otherwise the LCS must be quite large:

$$\text{LCS}(A, B) > s(1/s - \rho)n - s\gamma n = (1 - s(\rho + \gamma))n.$$

Recall from Lemma 6 that

$$\text{ED}(A, B) + 2 \cdot \text{LCS}(A, B) = 2n,$$

where  $\text{ED}(A, B)$  denotes the (no substitutions) edit distance of  $A$  and  $B$ . Using the ED approximation of [5] with some approximation ratio  $c > 3$  we recover an LCS approximation of length at least

$$n - c(n - \text{LCS}(A, B)) > n(1 - cs(\rho + \gamma))$$

so as long as we have  $1 - cs(\rho + \gamma) \geq 1/s + \gamma$ , this approximation is strong enough. This inequality holds when

$$\gamma \leq \frac{s - 1 - cs^2\rho}{s(1 + cs)}.$$

We can ensure it does by picking  $\rho$  small enough that the numerator of the right hand side above is positive and then taking  $\gamma$  smaller than the right hand side. Note that smaller values of  $\rho$  correspond to larger values of  $\gamma$ . ◀

► **Lemma 11 (Balanced LCS).** *If the LCS of  $A$  and  $B$  is not  $\rho$ -balanced, then in linear time we can  $(1/s + \rho/(s - 1))$ -approximate  $\text{LCS}(A, B)$ .*

**Proof.** Returning the longest common unary subsequence gives the desired approximation. To see this, let  $\sigma_{\max}$  and  $\sigma_{\min}$  be the most frequent and least frequent characters in the LCS respectively. Since the LCS is not balanced, either  $\sigma_{\max}$  makes up more than a  $(1/s + \rho)$  fraction of all symbols in the LCS, or  $\sigma_{\min}$  makes up fewer than a  $(1/s - \rho)$  fraction of the symbols in the LCS.

In the latter case, the  $s - 1$  members of the alphabet besides  $\sigma_{\min}$  must account for at least an  $((s - 1)/s + \rho)$  fraction of characters in the LCS. Among these,  $\sigma_{\max}$  appears the most often, which means by averaging that  $\sigma_{\max}$  accounts for at least a

$$\frac{1}{s} + \frac{\rho}{s - 1}$$

fraction of all symbols in the LCS.

In either case,  $\sigma_{\max}$  makes up at least a  $(1/s + \rho/(s - 1))$  fraction of the symbols in the LCS. Then the string consisting of  $\sigma_{\max}$  repeated  $\min(\sigma_{\max}(A), \sigma_{\max}(B))$  times is a common subsequence of  $A$  and  $B$  which has at least as many instances of  $\sigma_{\max}$  as the LCS does, and thus yields the desired approximation. ◀

After handling these easy cases, our approach is to restrict  $A$  and  $B$  to binary strings and invoke the frequency arguments from previous work. As we noted before, we cannot directly use the reduction in Corollary 5 because the better than  $1/2$  approximation of [21] only applies when the inputs have the same length. It turns out however, that the arguments

of [21] do hold for strings of differing length as long as the inputs satisfy a nice frequency condition. The following lemma demonstrates how by careful choice of subalphabets we can find restrictions meeting this condition. To describe this condition, we introduce some new notation: given a string  $A$  from an alphabet  $\Sigma$ , for any symbol  $\sigma \in \Sigma$  we let  $\sigma(A)$  denote the number of times  $\sigma$  appears in  $A$ .

► **Lemma 12** (Binary Restriction). *Suppose neither  $A$  nor  $B$  are  $\rho$ -balanced. Then there exists  $\Sigma' \subseteq \Sigma$  with  $|\Sigma'| = 2$  such that the restrictions of  $A$  and  $B$  to  $\Sigma'$  are not  $\rho/s$ -balanced.*

**Proof.** Let  $\alpha_1 \leq \dots \leq \alpha_s$  be the number of times the distinct symbols  $\sigma_1, \dots, \sigma_s$  of  $\Sigma$  appear respectively in  $A$ , so that  $\sigma_s$  is the most frequent symbol of  $A$  and  $\sigma_1$  is the least common character in  $A$ . By averaging we know that  $\alpha_s \geq n/s \geq \alpha_1$ . Since  $A$  is not  $\rho$ -balanced we deduce that

$$\alpha_s - \alpha_1 > \rho n.$$

We can decompose the left hand side of the above equation to

$$\alpha_s - \alpha_1 = \sum_{i=2}^s (\alpha_i - \alpha_{i-1}).$$

Since all the summands on the right hand side are positive, there exists some index  $j$  with

$$\alpha_j - \alpha_{j-1} > (\rho/s)n. \quad (2)$$

Note that we can find such a  $j$  in linear time by scanning through  $A$  and  $B$  and keeping counts of all the characters that appear. Now, consider the  $(s-1)$  two-element sets

$$\{\sigma_s, \sigma_{j-1}\}, \{\sigma_{s-1}, \sigma_{j-1}\}, \dots, \{\sigma_j, \sigma_{j-1}\}, \{\sigma_j, \sigma_{j-2}\}, \dots, \{\sigma_j, \sigma_1\}. \quad (3)$$

We claim that one of these sets satisfies the properties of  $\Sigma'$  from the lemma statement. First, note that Equation (2) ensures that the restriction of  $A$  to any of the above sets is not  $\rho/s$ -balanced, so it suffices to verify that the restriction of  $B$  will not be balanced for one of these sets. Suppose to the contrary that none of the sets from Equation (3) satisfies the desired conditions. We will use a triangle-inequality argument on the character frequencies of  $B$  to derive a contradiction. Let  $M$  and  $m$  be indices such that  $\sigma_M$  is the most frequent character of  $B$  and  $\sigma_m$  is the least frequent. Since  $B$  is not  $\rho$ -balanced, we know that  $M \neq m$ . If  $M \geq j > m$  we may write

$$|\sigma_M(B) - \sigma_m(B)| \leq |\sigma_M(B) - \sigma_{j-1}(B)| + |\sigma_{j-1}(B) - \sigma_j(B)| + |\sigma_j(B) - \sigma_m(B)|.$$

By assumption, the restrictions of  $B$  to the sets  $\{\sigma_M, \sigma_{j-1}\}$ ,  $\{\sigma_j, \sigma_{j-1}\}$ , and  $\{\sigma_j, \sigma_m\}$  are  $(\rho/s)$ -balanced. Thus, each addend on the right hand side of the above inequality is bounded above by  $(\rho/s)n$ . It follows that

$$|\sigma_M(B) - \sigma_m(B)| \leq (3\rho/s)n.$$

By similar reasoning, if  $m = j$  and  $M \geq j$  then we have

$$|\sigma_M(B) - \sigma_m(B)| = |\sigma_M(B) - \sigma_j(B)| \leq |\sigma_M(B) - \sigma_{j-1}(B)| + |\sigma_{j-1}(B) - \sigma_j(B)| \leq (2\rho/s)n.$$

If instead  $M, m > j$  we know that

$$|\sigma_M(B) - \sigma_m(B)| \leq |\sigma_M(B) - \sigma_{j-1}(B)| + |\sigma_{j-1}(B) - \sigma_m(B)| \leq (2\rho/s)n$$



## 13:8 Improved Approximation for Longest Common Subsequence over Small Alphabets

since now the subalphabets  $\{\sigma_M, \sigma_{j-1}\}$  and  $\{\sigma_m, \sigma_{j-1}\}$  both occur in Equation (3). Similar reasoning on the remaining cases of the values of  $M$  and  $m$  relative to  $j$  to establishes the inequality

$$\sigma_M(B) - \sigma_m(B) \leq (3\rho/s)n.$$

We have dropped absolute value signs because the left hand side of the above equation is positive by definition of  $M$  and  $m$ . Since  $s \geq 3$  this contradicts the assumption that  $B$  is not  $\rho$ -balanced, and the desired result follows. Note that we can find which of subalphabets from Equation (3) satisfies the conditions of the lemma in  $O(n + s)$  time by scanning through  $B$  to get the counts of each of its characters and trying out all the restrictions. ◀

In our proof of Theorem 2, we will require a stronger form of the result from Section 4 of [21]. This variant of their theorem is useful because it applies to strings of different length.

► **Lemma 13.** *Let  $X$  and  $Y$  be binary strings of length  $n$  and  $m$  respectively, where  $m \leq n$ . Suppose the frequencies  $0(Y)$  and  $1(X)$  are both at most  $(1/2 - \rho)m$  for some positive constant  $\rho$ . Then there exist positive constants  $\delta$  and  $\epsilon$  such that if  $0(Y)$  and  $1(X)$  are within  $\delta m$  of each other, we can compute a  $(1/2 + \epsilon)$  approximation of  $\text{LCS}(X, Y)$  in subquadratic time.*

Although the realization that this type of result holds for strings of differing length is novel, the proof of Lemma 13 itself is conceptually identical to the frequency analysis used in [21], requiring only minor changes to make the argument go through. For completeness we include the detailed casework proof of this result in Section 4.

Finally we apply a lemma that provides some frequency information about the strings.

► **Lemma 14** (Lemma 3.1 from [21]). *For any  $\delta > 0$  if*

$$\min(0(X), 0(Y)) > (1 + \delta) \min(1(X), 1(Y))$$

or

$$\min(1(X), 1(Y)) > (1 + \delta) \min(0(X), 0(Y))$$

then there is a unary  $(1 + \delta)/(2 + \delta)$  approximation for  $\text{LCS}(X, Y)$ .

**Proof.** Observe that for any binary strings  $X$  and  $Y$  we have

$$\text{LCS}(X, Y) \leq \min(0(X), 0(Y)) + \min(1(X), 1(Y)). \quad (4)$$

This equation holds because the LCS is a subsequence of  $X$  and  $Y$ , and thus cannot contain more zeros or ones than either of the strings  $X$  or  $Y$  individually have.

Suppose by symmetry that  $\min(0(X), 0(Y))$  is the larger of the two addends on the right. Then the we can return the all zeros string of this length. By the first inequality we get

$$\min(0(X), 0(Y)) > (1 + \delta) (\text{LCS}(X, Y) - \min(0(X), 0(Y)))$$

and then rearranging proves the claim. ◀

We now combine these results to improve LCS approximation on all alphabets.

**Proof of Theorem 8.** Let  $\rho$  and  $\rho'$  be positive parameters whose values will be specified later. If either of the strings  $A$  or  $B$  are  $\rho s$ -balanced, we are done by Lemma 10 (by taking  $\rho$  to be small enough so that the lemma applies). If the LCS of  $A$  and  $B$  is not  $\rho'$ -balanced we are done by Lemma 11. So, we may assume that neither  $A$  nor  $B$  are  $\rho s$ -balanced and that their LCS is  $\rho'$ -balanced.



By Lemma 12 we can find binary alphabet restrictions  $X$  and  $Y$  of  $A$  and  $B$  respectively with the property that neither  $X$  nor  $Y$  are  $\rho$ -balanced. Informally, since the LCS of  $A$  and  $B$  is balanced, a better than  $1/2$  approximation for the LCS of  $X$  and  $Y$  acts as a better than  $1/s$  approximation for  $\text{LCS}(A, B)$ . More precisely, since  $\text{LCS}(A, B)$  is  $\rho'$ -balanced, each of its characters occurs at least  $(1/s - \rho')$   $\text{LCS}(A, B)$  times in the LCS. By construction, we also know that  $\text{LCS}(X, Y)$  is at least as large as a binary alphabet restriction of  $\text{LCS}(A, B)$ . Thus, it follows that given a positive constant  $\epsilon'$ , a  $1/2 + \epsilon'$  approximation of  $\text{LCS}(X, Y)$  has length at least

$$(1/2 + \epsilon') \cdot 2(1/s - \rho') \text{LCS}(A, B) = (1/s + 2\epsilon'/s - \rho' - 2\epsilon'\rho') \text{LCS}(A, B).$$

By setting  $\epsilon = \epsilon'/s$  (for example) and  $\rho'$  sufficiently small in terms of  $\epsilon'$ , the above calculation shows that a  $1/2 + \epsilon'$  approximation for the LCS of  $X$  and  $Y$  acts as a  $1/s + \epsilon$  approximation for  $\text{LCS}(A, B)$ . Hence to complete the proof, it suffices to get a better than  $1/2$  approximation for  $\text{LCS}(X, Y)$ . We may assume that

$$\min(0(X), 0(Y)) \leq (1 + \delta) \min(1(X), 1(Y)) \quad (5)$$

and

$$\min(1(X), 1(Y)) \leq (1 + \delta) \min(0(X), 0(Y)) \quad (6)$$

for some constant  $\delta$  since otherwise Lemma 14 yields a better than  $1/2$  approximation.

As mentioned previously, [21] gives a better than  $1/2$  approximation for the LCS when  $X$  and  $Y$  have equal length. If  $X$  and  $Y$  have different lengths, without loss of generality we assume that  $|X| > |Y|$  and  $0(Y) \leq 1(Y)$ . We do casework on frequencies in  $X$ , relative to the frequencies in  $Y$ .

Since  $X$  is longer than  $Y$ , we cannot have both  $1(X) \leq 1(Y)$  and  $0(X) \leq 0(Y)$ .

If  $1(X) > 1(Y)$  and  $0(X) > 0(Y)$  simultaneously, then Equation (6) implies that

$$1(Y) \leq (1 + \delta) \cdot 0(Y)$$

which contradicts the fact that  $Y$  is not  $\rho$ -balanced as long as we take  $\delta \leq 2\rho$ .

If instead  $1(X) \leq 1(Y)$  and  $0(X) \leq 0(Y)$ , by Equation (6) we similarly have

$$0(X) \leq 0(Y) \leq 1(Y) \leq (1 + \delta) \cdot 0(X)$$

which again contradicts the fact that  $Y$  is not  $\rho$ -balanced for the same choice of  $\delta$ .

Thus, the only possibility is that  $1(X) \leq 1(Y)$  and  $0(X) > 0(Y)$ . Let  $m = |Y|$  be the length of string  $Y$ . Then Equation (5) implies that

$$0(Y) \leq (1 + \delta) \cdot 1(X) \leq 1(X) + \delta m$$

while Equation (6) implies that

$$1(X) \leq (1 + \delta) \cdot 0(Y) \leq 0(Y) + \delta m.$$

Consequently,  $0(Y)$  and  $1(X)$  are within  $\delta m$  of each other. Then by Lemma 13, as long as we take  $\delta$  small enough in terms of  $\rho$  we get a subquadratic  $1/2 + \epsilon'$  approximation for  $\text{LCS}(X, Y)$ . As noted earlier, this then yields the desired  $1/s + \epsilon$  approximation for  $\text{LCS}(A, B)$  in truly subquadratic time. In fact, because we only ever use subroutines that run in linear time or constant factor approximations to edit distance which take near-linear time, the overall algorithm takes near-linear time. ◀

#### 4 Proof of Lemma 13

This section is devoted to proving Lemma 13. We do this working through the individual arguments in the case analysis of [21] and verifying that the arguments still hold in the case where the strings have different lengths, as long as they satisfy the frequency requirements included in the hypotheses of the lemma. Throughout this section we fix the binary alphabet  $\Sigma = \{0, 1\}$  and assume that all strings come from this alphabet.

We carry over the following subroutines from [21].

► **Definition 15.** *Given strings  $A$  and  $B$  and a symbol  $\sigma$ , the algorithm  $\text{Match}(A, B, \sigma)$  returns the largest subsequence of  $A$  and  $B$  consisting entirely of copies of  $\sigma$ .*

► **Definition 16.** *Given strings  $A$  and  $B$ , the algorithm  $\text{BestMatch}(A, B)$  returns the longer of the strings*

$$\text{Match}(A, B, 0) \quad \text{and} \quad \text{Match}(A, B, 1).$$

*In other words, the algorithm returns the largest common unary subsequence.*

Note that  $\text{BestMatch}$  is a  $1/2$  approximation algorithm for LCS.

► **Definition 17.** *Given strings  $A_1, A_2$ , and  $B$ , the algorithm  $\text{Greedy}(A_1, A_2, B)$  returns*

$$\max_{B=B_1 \sqcup B_2} \text{BestMatch}(A_1, B_1) + \text{BestMatch}(A_2, B_2)$$

*taken over all possible splits of  $B$  into two contiguous left and right substrings  $B_1$  and  $B_2$ .*

These algorithms can all be implemented to run in linear time by scanning through the counting how many 0s and 1s appear in each string. The final procedure we will invoke utilizes an approximation algorithm for computing the edit distance of two strings as a black-box. Although any fast enough constant factor approximation for edit distance will work, for concreteness we assume that it leverages the algorithm  $\widetilde{\text{ED}}$  from [5], which runs in subquadratic time and can approximate the edit distance to a  $c = 3 + \epsilon'$  factor for any fixed positive constant  $\epsilon'$ . Note that this algorithm can also return a common subsequence achieving the given length.

**Proof of Lemma 13.** Let  $1(X) = \alpha m$  for some  $\alpha \in [0, 1]$ . By assumption, we know that  $\alpha < 1/2$  is bounded away from  $1/2$  by some constant amount. We can also assume that  $0(Y)$  is within  $\delta m$  of  $1(X)$  for some positive parameter  $\delta < 0$  to be picked later on. We will end up setting  $\delta$  to be some sufficiently small constant depending on  $\alpha$ . We will use the notation  $\approx$  to denote quantities that are within  $\delta m$  of each other. For example  $1(X) \approx 0(Y)$ . This lets us avoid having to stick in  $\pm \delta m$  symbols in all the inequalities and helps make the arguments cleaner without affecting the correctness (since we just care about getting a  $1/2 + \epsilon$  approximation for *some* constant  $\epsilon > 0$ ).

Note that by Equation (4) the LCS of the input strings

$$\text{LCS}(X, Y) \leq (2\alpha + \delta)m \tag{7}$$

cannot be too large, because the LCS contain at most  $\alpha m$  1s from  $X$  and  $\approx \alpha m$  0s from  $Y$ .

Let  $R_X$  and  $R_Y$  denote the substrings of  $X$  and  $Y$  consisting of their rightmost  $\alpha m$  characters respectively. Similarly define  $L_X$  and  $L_Y$  as the substrings of  $X$  and  $Y$  consisting of the leftmost  $\alpha m$  characters. Set  $M_X = X \setminus (L_X \cup R_X)$  and  $M_Y = Y \setminus (L_Y \cup R_Y)$  to be the middle substrings of  $X$  and  $Y$  that remain when these left and right ends are chopped off.

We now follow the casework of [21], explaining at each step why the arguments still hold in our more general setting. These cases are based off the frequencies of 0s and 1s in the left and right ends of the inputs, and consider separately the situation where these substrings are pseudorandom (balanced) or structured (imbalanced). In the former case we can appeal to edit distance as in the proof of Lemma 10, and in the latter situation we can exploit the imbalance in the strings to use the simpler unary algorithms described in definitions 15, 16, and 17. Intuitively, we succeed in using edit distance approximation arguments (and overcome the barrier described the end of Section 2) in this particular case because even though  $X$  and  $Y$  may have different size, we identify right and left substrings which all have equal length and only employ edit distance approximation around these areas.

Recall that we assumed that  $\alpha < 1/2 - \rho$  for some constant  $\rho$ . Take a parameter  $\beta < \rho/20$  to represent deviation from the balanced case. By choosing  $\delta$  sufficiently small in terms of  $\beta$  we may additionally assume that

$$1(X), 0(Y) < (1/2 - 10\beta)m \quad (8)$$

since in the hypothesis of the lemma we supposed that  $1(X) < (1/2 - \rho)m$  and that  $0(Y) \approx 1(X)$ . Finally, as one last piece of notation, we write  $|A|$  to denote length of an arbitrary string  $A$ . We now begin the casework, maintaining consistency with [21].

**Case 1(a):**  $1(R_Y), 0(R_X) \in [(\alpha/2 - 4\beta)m, (\alpha/2 + 4\beta)m]$

In this case both right ends of the strings are balanced. We will give a good approximation for the LCS by splitting these strings at the right ends and using the aforementioned algorithms.

Consider an optimal alignment between  $X$  and  $Y$  (i.e. a maximum partial matching of identical characters in  $X$  and  $Y$ , corresponding to the LCS of  $X$  and  $Y$ ). Let  $\widehat{R}_Y$  be the minimal suffix (from the right end) of the string  $Y$  with the property that every character from  $R_X$  which is matched in the alignment is paired up with some character in  $\widehat{R}_Y$ .

Without loss of generality, we may assume that  $\widehat{R}_Y$  is a substring of  $R_Y$ . This is because if  $\widehat{R}_Y$  was not contained in  $R_Y$ , we could define an analogous substring of  $\widehat{R}_X$  of  $X$  satisfying  $\widehat{R}_X \subseteq R_X$  and then use a symmetric argument to get the desired approximation. Let  $\widehat{L}_Y = Y \setminus \widehat{R}_Y$  be the left substring of  $Y$  that remains after chopping off  $\widehat{R}_Y$ .

Optimality of the alignment implies that

$$\text{LCS}(X, Y) = \text{LCS}(X \setminus R_X, \widehat{L}_Y) + \text{LCS}(R_X, \widehat{R}_Y). \quad (9)$$

Define the quantities

$$f_L = \min(1(X \setminus R_X), 1(\widehat{L}_Y)) + \min(0(X \setminus R_X), 0(\widehat{L}_Y))$$

and

$$f_R = \min(1(R_X), 1(\widehat{R}_Y)) + \min(0(R_X), 0(\widehat{R}_Y))$$

which represent frequency-based upper bounds for the LCS terms from the right hand side of Equation (9). They are useful because Equation (4) together with Equation (9) implies that

$$\text{LCS}(X, Y) \leq f_L + f_R. \quad (10)$$

We also introduce the quantity

$$Z = \max\left(\min(1(X \setminus R_X), 1(\widehat{L}_Y)), \min(0(X \setminus R_X), 0(\widehat{L}_Y))\right)$$

which is the larger of the two addends defining  $f_L$ , and equal to the length of the string returned by

$$\text{BestMatch}(X \setminus R_X, \widehat{L}_Y).$$

We further subdivide into cases based off the size of  $Z$ .

## 13:12 Improved Approximation for Longest Common Subsequence over Small Alphabets

**Case 1(a)(i):**  $Z > (\alpha/2 + 10\beta) m$

When  $Z$  is large we can combine two unary subsequences to get a good enough LCS approximation. We first show that  $Z$  is bigger than  $f_L/2$  by a constant fraction of  $m$ .

By definition we have

$$f_L - Z = \min\left(\min(1(X \setminus R_X), 1(\widehat{L}_Y)), \min(0(X \setminus R_X), 0(\widehat{L}_Y))\right) \leq 1(X \setminus R_X).$$

Since  $X$  has  $\alpha m$  ones and  $R_X$  has length  $\alpha m$  we get that

$$1(X \setminus R_X) = \alpha m - 1(R_X) = \alpha m - (\alpha m - 0(R_X)) = 0(R_X).$$

Then using the case assumptions we have

$$0(R_X) \leq (\alpha/2 + 4\beta) m < Z - 6\beta m.$$

Chaining these inequalities together and rearranging we deduce that

$$Z > f_L/2 + 3\beta m.$$

Now if we make a single call to the Greedy routine we obtain a string of length

$$\text{Greedy}(X \setminus R_X, R_X, Y) \geq \text{BestMatch}(X \setminus R_X, \widehat{L}_Y) + \text{BestMatch}(R_X, \widehat{R}_Y). \quad (11)$$

From our earlier discussion we have

$$\text{BestMatch}(X \setminus R_X, \widehat{L}_Y) \geq Z > f_L/2 + 3\beta m.$$

Moreover

$$\text{BestMatch}(R_X, \widehat{R}_Y) = \max(\min(1(R_X), 1(\widehat{R}_Y)), \min(0(R_X), 0(\widehat{R}_Y))) \geq f_R/2$$

since we are taking the maximum over two addends that sum to  $f_R$ .

Finally, if we substitute the above inequalities into Equation (11) and apply Equation (10) we get that

$$\begin{aligned} \text{Greedy}(X \setminus R_X, R_X, Y) &\geq (f_L + f_R)/2 + 3\beta m \geq \text{LCS}(X, Y)/2 + 3\beta m \\ &\geq (1/2 + 3\beta) \text{LCS}(X, Y). \end{aligned}$$

Thus running Greedy gives us a better than  $1/2$  approximation in this case.

**Case 1(a)(ii):**  $Z \leq (\alpha/2 + 10\beta) m$

Intuitively, in this case  $Z$  is too small for us to ensure a good approximation using frequency guarantees alone. However, because  $Z$  is so small, the LCS also cannot be too large. Because of this, we will be able to get a good approximation by combining a common subsequence of the (balanced) right ends of the strings with a common subsequence of the strings with the right ends removed.

More concretely, we leverage the ApproxED algorithm from Definition 7. From our previous observation we have  $f_L \leq 2Z \leq (\alpha + 20\beta) m$ . So via Equation (9) we can bound the LCS by

$$\text{LCS}(X, Y) \leq f_L + \text{LCS}(R_X, \widehat{R}_Y) \leq (\alpha + 20\beta) m + \text{LCS}(R_X, R_Y) \quad (12)$$

where in the last step we also used the fact that  $\widehat{R}_Y \subseteq R_Y$ .

We will get an approximation by returning a unary subsequence from the left parts of the strings, and using edit distance approximation on the right ends. First, we can make a call to `BestMatch` and get a common subsequence of length at least

$$\text{BestMatch}(X \setminus R_X, Y \setminus R_Y) \geq \text{Match}(X \setminus R_X, Y \setminus R_Y, 0) \geq (\alpha/2 - 4\beta) m.$$

This last inequality above follows from combining the case 1 assumptions about the frequencies of characters in  $0(R_X)$  and  $1(R_Y)$ , together with the facts that  $0(Y) \approx 1(X) = \alpha m$  and  $|R_X| = |R_Y| = \alpha m$ .

Now, since  $R_X$  and  $R_Y$  are  $(4\beta/\alpha)$ -balanced we can apply Lemma 10 with  $n = \alpha m$  as long as we take  $\beta$  sufficiently small in terms of  $\alpha$ . Note that for fixed alphabet size and  $\rho$ , the parameter  $\beta$  remains  $\Omega(1)$ . This ensures that in subquadratic time we can compute a common subsequence of  $R_X$  and  $R_Y$  with length at least

$$(1/2 + \gamma) \text{LCS}(R_X, R_Y)$$

where  $\gamma < 1$  is the constant from Lemma 10, which is larger for smaller values of  $\beta$ . By the frequency assumptions on  $R_X$  and  $R_Y$ , we know that

$$\text{LCS}(R_X, R_Y) \geq (\alpha/2 - 4\beta) m$$

so in fact the subsequence of  $R_X$  and  $R_Y$  returned by Lemma 10 has length at least

$$(1/2 + \gamma) \text{LCS}(R_X, R_Y) \geq \text{LCS}(R_X, R_Y)/2 + \gamma(\alpha/2 - 4\beta)m.$$

Now if we combine these two subsequences together, we get a common subsequence of  $X$  and  $Y$  of length at least

$$\text{LCS}(R_X, R_Y)/2 + (\alpha/2 - 4\beta) m + \gamma(\alpha/2 - 4\beta)m$$

which can be written in the form

$$((\alpha + 20\beta)m + \text{LCS}(R_X, R_Y)) / 2 + (\gamma\alpha/2 - 14\beta - 4\gamma\beta) m.$$

By applying Equation (12) and the fact that  $\gamma < 1$  we see that this is at least

$$\text{LCS}(X, Y)/2 + (\gamma\alpha/2 - 18\beta) m.$$

Finally, by picking  $\beta$  small enough this expression is at least

$$\text{LCS}(X, Y)/2 + (\gamma\alpha/3) m \geq (1/2 + \gamma/6) \text{LCS}(X, Y)$$

where in the last step we have used Equation (4) together with  $1(X) = \alpha m$  and  $0(Y) \approx \alpha m$ . This proves that we can attain a better than  $1/2$  approximation for the LCS as claimed.

**Case 1(b):  $1(R_Y) < (\alpha/2 - 4\beta) m$  and  $0(R_X) \leq (\alpha/2 + 2\beta) m$**

In this case the right ends of  $X$  and  $Y$  each do not contain too much their respective strings' most common characters. We show that this implies the LCS of both strings must be so small that simply returning a unary string yields a better than  $1/2$  approximation.

As in case 1(a), consider an optimal alignment between  $X$  and  $Y$ . Now define the substring  $\widehat{R}_Y$  to be the minimal suffix of  $Y$  which contains all characters of  $Y$  that  $R_X$  is aligned to. Let  $\widehat{L}_Y = Y \setminus \widehat{R}_Y$  be what remains of  $Y$  after  $R_Y$  is removed. Since the alignment is optimal, Equation (9) holds.

We now subdivide into further case based off how the right ends of strings are aligned.

## 13:14 Improved Approximation for Longest Common Subsequence over Small Alphabets

**Case 1(b)(i): every character of  $R_X$  is matched to some character of  $R_Y$  in the alignment**

In this case  $\widehat{R}_Y$  is a substring of  $R_Y$ . From Equation (4) and the case assumptions we get that

$$\text{LCS}(X \setminus R_X, \widehat{L}_Y) \leq 1(X \setminus R_X) + 0(\widehat{L}_Y) \leq (\alpha/2 + 2\beta)m + 0(\widehat{L}_Y)$$

and

$$\text{LCS}(R_X, \widehat{R}_Y) \leq 1(\widehat{R}_Y) + 0(\widehat{R}_Y) \leq (\alpha/2 - 4\beta)m + 0(\widehat{R}_Y).$$

Note that in this second inequality we are using the inequality  $1(\widehat{R}_Y) \leq 1(R_Y)$  which follows from the earlier observation that  $\widehat{R}_Y \subseteq R_Y$ . By adding these inequalities together and substituting the result into Equation (9) we deduce that

$$\text{LCS}(X, Y) \leq (\alpha - 2\beta)m + \left(0(\widehat{L}_Y) + 0(\widehat{R}_Y)\right) \leq (\alpha - 2\beta)m + \alpha m = (2\alpha - 2\beta)m$$

where we have used the fact that 0 occurs in  $Y \approx \alpha m$  times.

**Case 1(b)(ii): some character of  $R_X$  is matched outside  $R_Y$  in the alignment**

In this case  $R_Y$  is a substring of  $\widehat{R}_Y$ .

Using Equation (4) again we find that

$$\text{LCS}(X \setminus R_X, \widehat{L}_Y) \leq 1(X \setminus R_X) + 0(\widehat{L}_Y) \tag{13}$$

From  $0(R_X) \leq (\alpha/2 + 2\beta)m$  we know that  $1(R_X) \geq (\alpha/2 - 2\beta)m$  since  $R_X$  has length  $\alpha m$ . It follows that

$$1(X \setminus R_X) = 1(X) - 1(R_X) \leq (\alpha/2 + 2\beta)m$$

since  $1(X) = \alpha m$ .

Similarly, since  $1(R_Y) \leq (\alpha/2 - 4\beta)m$  and  $R_Y$  has length  $\alpha m$  we know that  $0(R_Y) > (\alpha/2 + 4\beta)m$ . Since  $R_Y \subseteq \widehat{R}_Y$  it must be the case that  $0(\widehat{R}_Y) > (\alpha/2 + 4\beta)m$  which means that

$$0(\widehat{L}_Y) = 0(Y) - 0(\widehat{R}_Y) \approx \alpha m - 0(\widehat{R}_Y) < (\alpha/2 - 4\beta)m.$$

Adding these two inequalities and substituting into Equation (13) proves that

$$\text{LCS}(X \setminus R_X, \widehat{L}_Y) \leq (\alpha - 2\beta)m.$$

Then applying Equation (9) and using the fact that  $R_X$  has length  $\alpha m$  proves that

$$\text{LCS}(X, Y) = \text{LCS}(X \setminus R_X, \widehat{L}_Y) + \text{LCS}(R_X, \widehat{R}_Y) \leq (\alpha - 2\beta)m + \alpha m = (2\alpha - 2\beta)m.$$

Thus in both this subcase and the previous one, the LCS is at most  $(2\alpha - 2\beta)m$ . Hence, returning the string of  $\approx \alpha m$  zeros obtained by calling  $\text{Match}(X, Y, 0)$  gives a better than  $1/2$  approximation as desired.

**Case 1(c):  $1(R_Y) \leq (\alpha/2 + 2\beta)m$  and  $0(R_X) < (\alpha/2 - 4\beta)m$**

This case is symmetric to case 1(b) and similar reasoning handles it.

**Case 2:**  $1(L_Y), 0(L_X) \leq (\alpha/2 + 2\beta)m$

Combining cases 1(a), 1(b), and 1(c) resolves the situation where

$$1(R_Y), 0(R_X) \leq \alpha/2 + 2\beta.$$

Consequently, case 2 is symmetric to case 1. In particular, we can flip the strings (by replacing  $X$  with the string  $X'$  which consists of the symbols of  $X$  read in reverse from right to left, and replacing  $Y$  with its analogous reverse string  $Y'$ ) and apply the arguments from case 1 to handle this case.

**Case 3:**  $1(L_Y), 1(R_Y) \leq (\alpha/2 + \beta)m$  and  $0(L_X), 0(R_X) > (\alpha/2 + 2\beta)m$

In this case both ends of the inputs have many 0s, so repeated calls to Match will be enough to guarantee a better than  $1/2$  approximation.

Since  $L_Y$  has length  $\alpha m$  and at most  $(\alpha/2 + \beta)m$  instances of 1, it must have at least

$$0(L_Y) \geq \alpha m - (\alpha/2 + \beta)m = (\alpha/2 - \beta)m$$

occurrences of 0. Since  $1(R_Y) \leq (\alpha/2 + \beta)m$  the same reasoning shows that

$$0(R_Y) \geq (\alpha/2 - \beta)m.$$

Consequently we can get common subsequences of the left and right ends of the strings consisting entirely of 0s with lengths

$$\text{Match}(L_X, L_Y, 0) = \min(0(L_X), 0(L_Y)) \geq (\alpha/2 - \beta)m$$

and

$$\text{Match}(R_X, R_Y, 0) = \min(0(R_X), 0(R_Y)) \geq (\alpha/2 - \beta)m.$$

Since the ends have many 0s, we expect the middle substrings to have many 1s. Indeed, since the left end of  $X$

$$1(L_X) = |L_X| - 0(L_X) < \alpha m - (\alpha/2 + 2\beta)m = (\alpha/2 - 2\beta)m$$

does not have many 1s and similar reasoning shows that

$$1(R_X) < (\alpha/2 - 2\beta)m$$

this holds for the right end as well, the middle of  $X$  has only

$$1(M_X) = 1(X) - 1(L_X) - 1(R_X) > \alpha m - 2(\alpha/2 - 2\beta)m = 4\beta m$$

appearances of 1.

To argue that  $M_Y$  has enough 1s we will need to appeal to the fact that  $Y$  is not balanced. From Equation (8) we know that

$$1(Y) = m - 0(Y) > m - (1/2 - 10\beta)m = (1/2 + 10\beta)m.$$

We can use the case assumptions together  $\alpha < 1/2$  to then deduce

$$1(M_Y) = 1(Y) - 1(L_Y) - 1(R_Y) > (\alpha + 10\beta)m - 2(\alpha/2 + \beta)m = 8\beta m.$$



## 13:16 Improved Approximation for Longest Common Subsequence over Small Alphabets

Consequently the largest all 1s subsequence between the middle substrings has length at least

$$\text{Match}(M_X, M_Y, 1) = \min(1(M_X), 1(M_Y)) \geq 4\beta m.$$

Thus, by combining these three subsequences from three calls to match we can recover in linear time a common subsequence of size

$$\text{Match}(L_X, L_Y, 0) + \text{Match}(M_X, M_Y, 1) + \text{Match}(R_X, R_Y, 0) > 2(\alpha/2 - \beta)m + 4\beta m = (\alpha + 2\beta)m.$$

By Equation (7) the true LCS has length at most  $\approx 2\alpha m$ , so a string of length  $(\alpha + 2\beta)m$  is a  $1/2 + \epsilon$  approximation for some constant  $\epsilon < 2\beta/\alpha$  as desired.

**Case 4:**  $0(L_X), 0(R_X) \leq (\alpha/2 + \beta)m$  and  $1(L_Y), 1(R_Y) > (\alpha/2 + 2\beta)m$

This case is symmetric to case 3 and similar reasoning handles it.

**Case 5:**  $0(L_X), 1(R_Y) > (\alpha/2 + \beta)m$

In this case, the ends of the strings have unusually large instances of either 0 or 1. This will enable us to combine two calls to Match to get the desired approximation.

We can check that the right end of  $Y$  does not have many 0s

$$0(R_Y) = |R_Y| - 1(R_Y) < \alpha m - (\alpha/2 + \beta)m = (\alpha/2 - \beta)m.$$

It follows that the remainder of the string  $Y \setminus R_Y = L_Y \cup M_Y$  has many zeros

$$0(Y \setminus R_Y) = 0(Y) - 0(R_Y) \approx \alpha m - 0(R_Y) > (\alpha/2 + \beta)m.$$

Similar reasoning shows that

$$1(L_X) = \alpha m - 0(L_X) < (\alpha/2 - \beta)m$$

so that the remaining portion of  $X$  has many ones

$$1(X \setminus L_X) = 1(X) - 1(L_X) = \alpha m - 1(L_X) > (\alpha/2 + \beta)m.$$

It follows that

$$\text{Match}(L_X, Y \setminus R_Y, 0) + \text{Match}(X \setminus L_X, R_Y, 1) > (\alpha + 2\beta)m$$

yields a better than  $1/2$  approximation since by Equation (7) the LCS is at most  $\approx 2\alpha m$ .

**Case 6:**  $1(L_Y), 0(R_X) > (\alpha/2 + \beta)m$

This is symmetric to case 5 and a similar argument proves the result holds in this situation.

By inspection or by referring to Table 1 of [21], we can verify that these cases handle all possible input strings satisfying the conditions of the lemma. Since in every case we obtain a better than  $1/2$  approximation for the LCS in subquadratic time, we have proven the claim.  $\blacktriangleleft$

## 5 Open Problems

Theorem 2 of our work shows how to obtain better than  $1/|\Sigma|$  approximations for the longest common subsequence of *equal-length strings* over an alphabet  $\Sigma$ . It remains an open problem to get such approximations in the setting where the input strings have different length (and by Corollary 5, to solve this problem it suffices to obtain an improvement in the setting of binary alphabets, where  $|\Sigma| = 2$ ).

Additionally, although our algorithm beats the longstanding trivial approximation ratio for LCS, it does so only by a modest amount. It would be interesting to get better approximation ratios, both in terms of their concrete value for small  $|\Sigma|$  and in terms of their growth as a function of the alphabet size.

---

### References

- 1 Amir Abboud and Arturs Backurs. Towards hardness of approximation for polynomial time problems. In *8th Innovations in Theoretical Computer Science Conference, ITCS 2017, January 9-11, 2017, Berkeley, CA, USA*, volume 67 of *LIPICs*, pages 11:1–11:26, 2017.
- 2 Amir Abboud, Arturs Backurs, and Virginia Vassilevska Williams. Tight hardness results for LCS and other sequence similarity measures. In *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*, pages 59–78. IEEE Computer Society, 2015.
- 3 Amir Abboud, Thomas Dueholm Hansen, Virginia Vassilevska Williams, and Ryan Williams. *Simulating Branching Programs with Edit Distance and Friends: Or: A Polylog Shaved is a Lower Bound Made*, page 375–388. Association for Computing Machinery, New York, NY, USA, 2016. doi:10.1145/2897518.2897653.
- 4 Amir Abboud and Aviad Rubinfeld. Fast and deterministic constant factor approximation algorithms for LCS imply new circuit lower bounds. In Anna R. Karlin, editor, *9th Innovations in Theoretical Computer Science Conference, ITCS 2018, January 11-14, 2018, Cambridge, MA, USA*, volume 94 of *LIPICs*, pages 35:1–35:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018.
- 5 Alex Andoni. Simpler constant-factor approximation to edit distance problems, 2019.
- 6 Alexandr Andoni, Robert Krauthgamer, and Krzysztof Onak. Polylogarithmic approximation for edit distance and the asymmetric query complexity. In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23-26, 2010, Las Vegas, Nevada, USA*, pages 377–386. IEEE Computer Society, 2010.
- 7 Alexandr Andoni and Negev Shekel Nosatzki. Edit distance in near-linear time: it’s a constant factor. *CoRR*, abs/2005.07678, 2020. arXiv:2005.07678.
- 8 Alexandr Andoni and Krzysztof Onak. Approximating edit distance in near-linear time. *SIAM J. Comput.*, 41(6):1635–1648, 2012.
- 9 Arturs Backurs and Piotr Indyk. Edit distance cannot be computed in strongly subquadratic time (unless SETH is false). *SIAM J. Comput.*, 47(3):1087–1097, 2018.
- 10 Ziv Bar-Yossef, T. S. Jayram, Robert Krauthgamer, and Ravi Kumar. Approximating edit distance efficiently. In *45th Symposium on Foundations of Computer Science (FOCS 2004), 17-19 October 2004, Rome, Italy, Proceedings*, pages 550–559. IEEE Computer Society, 2004.
- 11 Tugkan Batu, Funda Ergün, and Süleyman Cenk Sahinalp. Oblivious string embeddings and edit distance approximations. In *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2006, Miami, Florida, USA, January 22-26, 2006*, pages 792–801. ACM Press, 2006.
- 12 Mahdi Boroujeni, Soheil Ehsani, Mohammad Ghodsi, Mohammad Taghi Hajiaghayi, and Saeed Seddighin. Approximating edit distance in truly subquadratic time: Quantum and mapreduce. In Artur Czumaj, editor, *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, pages 1170–1189. SIAM, 2018.

- 13 Joshua Brakensiek and Aviad Rubinfeld. Constant-factor approximation of near-linear edit distance in near-linear time. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2020, page 685–698, New York, NY, USA, 2020. Association for Computing Machinery. doi:10.1145/3357713.3384282.
- 14 Karl Bringmann and Marvin Künnemann. Multivariate fine-grained complexity of longest common subsequence. In *SODA*, 2018.
- 15 Diptarka Chakraborty, Debarati Das, Elazar Goldenberg, Michal Koucký, and Michael E. Saks. Approximating edit distance within constant factor in truly sub-quadratic time. In Mikkel Thorup, editor, *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018*, pages 979–990. IEEE Computer Society, 2018.
- 16 Lijie Chen, Shafi Goldwasser, Kaifeng Lyu, Guy N. Rothblum, and Aviad Rubinfeld. Fine-grained complexity meets IP = PSPACE. In Timothy M. Chan, editor, *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 1–20. SIAM, 2019.
- 17 MohammadTaghi Hajiaghayi, Masoud Seddighin, Saeed Seddighin, and Xiaorui Sun. Approximating LCS in linear time: Beating the  $\sqrt{n}$  barrier. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 1181–1200. SIAM, 2019.
- 18 Michal Koucký and Michael Saks. Constant factor approximations to edit distance on far input pairs in nearly linear time. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2020, pages 699–712, New York, NY, USA, 2020. Association for Computing Machinery. doi:10.1145/3357713.3384307.
- 19 Gad M. Landau, Eugene W. Myers, and Jeanette P. Schmidt. Incremental string comparison. *SIAM Journal on Computing*, 27(2):557–582, 1998. doi:10.1137/s0097539794264810.
- 20 William J. Masek and Mike Paterson. A faster algorithm computing string edit distances. *J. Comput. Syst. Sci.*, 20(1):18–31, 1980.
- 21 Aviad Rubinfeld and Zhao Song. Reducing approximate longest common subsequence to approximate edit distance. In *Proceedings of the Thirty-First Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '20, page 1591–1600, USA, 2020. Society for Industrial and Applied Mathematics.