

Hardness of Dynamic Tree Edit Distance and Friends

Bingbing Hu   

University of California San Diego, La Jolla, CA, USA

Jakob Nogler¹   

Massachusetts Institute of Technology, Cambridge, MA, USA

Barna Saha   

University of California San Diego, La Jolla, CA, USA

Abstract

String Edit Distance is a more-than-classical problem whose behavior in the dynamic setting, where the strings are updated over time, is well studied. A single-character substitution, insertion, or deletion can be processed in time $\tilde{O}(nw)^2$ when operation costs are positive integers bounded by w [Charalampopoulos, Kociumaka, Mozes, CPM 2020][Gorbachev, Kociumaka, STOC 2025]. If the weights are further uniform (insertions and deletions have equal cost), also an $\tilde{O}(n\sqrt{n})$ -update time algorithm exists [Charalampopoulos, Kociumaka, Mozes, CPM 2020]. This is a substantial improvement over the static $\mathcal{O}(n^2)$ algorithm when $w \ll n$ or when we are dealing with uniform weights.

In contrast, for inherently related problems such as Tree Edit Distance, Dyck Edit Distance, and RNA Folding, it has remained unknown whether it is possible to devise dynamic algorithms with an advantage over the static algorithm. In this paper, we resolve this question by showing that (weighted) Tree Edit Distance, Dyck Edit Distance, and RNA Folding admit no dynamic speedup: under well-known fine-grained assumptions, we show that the best possible algorithm recomputes the solution from scratch after each update. Furthermore, we prove a quadratic per-update lower bound for unweighted Tree Edit Distance under the k -Clique Conjecture. This provides the first separation between dynamic unweighted String Edit Distance and unweighted Tree Edit Distance, problems whose relative difficulty in the static setting is still open.

2012 ACM Subject Classification Theory of computation → Dynamic graph algorithms

Keywords and phrases fine-grained complexity, dynamic lower bounds, pattern matching

Digital Object Identifier 10.4230/LIPIcs.ITCS.2026.78

Related Version *Full Version:* <https://arxiv.org/abs/2511.09842> [40]

Funding The authors are partially supported by NSF grants 1652303, 1909046, 2112533, and HDR TRIPODS Phase II grant 2217058.

1 Introduction

The Tree Edit Distance (TED) measures the minimum number of insertions, deletions, and substitutions required to transform one labeled, ordered tree into another³, and has found numerous applications in settings where tree structures arise: computational biology [35, 65, 39, 77], data analysis [16, 20, 33], image processing [11, 48, 47, 62], and more.

There is a long line of work spanning over thirty years [67, 78, 46, 28] on improving the running time of TED to arrive at a cubic algorithm [28]. It has only been shown recently that from a computational point of view, weighted TED (where there is a weight associated

¹ Work done while the author was visiting the EnCORE Institute at UC San Diego.

² $\tilde{O}(\cdot)$ hides polylogarithmic factors.

³ More specifically, upon deleting a node, the children of the deleted node become the children of the original parent of the deleted node. Insertion is the reverse process of deletion.



with every operation, and the goal is to minimize the total weight of the transformation) resembles the All-Pairs Shortest Paths (APSP) Problem. Bringmann *et al.* [13] proved that weighted TED cannot be solved in $\mathcal{O}(n^{3-\varepsilon})$ time under APSP, and subsequently Nogler *et al.* [57] showed that it is, in fact, *equivalent* to APSP. For unweighted TED, however, the situation is drastically different and truly subcubic algorithms are known [54, 32, 57], with the best running time to date being $\tilde{\mathcal{O}}(n^{(3+\omega)/2}) = \mathcal{O}(n^{2.6857})$ [57]⁴. On the other hand, from a lower bound perspective all we know is that unweighted TED is at least as hard as the String Edit Distance. It remains a major open problem whether unweighted TED is strictly harder than String Edit Distance, conditioned under some fine-grained hardness hypothesis.

In this paper, we study the complexity of TED (both weighted and unweighted) in the dynamic setting. While this question is well understood for the String Edit Distance problem, its complexity remains wide open for TED. Note that the String Edit Distance, also known as the Levenshtein distance, is defined as the minimum number of single-character insertions, deletions, and substitutions required to transform one string into the other. When each operation is assigned a weight and the goal is to minimize the total cost, we obtain the *weighted* edit distance. Computing the (weighted or unweighted) edit distance between two strings of length n is a standard introductory example of dynamic programming, solvable in $\mathcal{O}(n^2)$ time [75, 56, 76, 64]. Surprisingly, despite its simplicity, this algorithm is in fact optimal under the Strong Exponential Time Hypothesis (SETH) [9, 1, 4, 15].

The complexity of the Edit Distance Problem is well understood in the dynamic setting, where the strings undergo updates (insertions, deletions and substitutions) and we need to maintain the edit distance throughout the updates. After initial works that considered updates occurring only at the endpoints of the two strings [52, 43, 42, 70], Charalampopoulos, Kociumaka, and Mozes gave the first algorithm for the most general scenario [19], where updates can happen anywhere. In the unweighted setting, their algorithm supports updates in $\tilde{\mathcal{O}}(n)$ time, and this is conditionally optimal under SETH. In the uniform-cost setting (where insertions and deletions have equal positive integer cost) bounded by $w = n^{\mathcal{O}(1)}$, their algorithm supports updates in $\tilde{\mathcal{O}}(n \cdot \min(\sqrt{n}, w))$ time. Very recently, it has been shown that $\tilde{\mathcal{O}}(n\sqrt{n})$ is (conditionally) optimal when w is a large polynomial in n [17]. Moreover, an $\tilde{\mathcal{O}}(nw)$ update-time algorithm for general positive weights has been given [34]. Beyond this, dynamic string Edit Distance has been also studied in more specific settings: Gorbachev and Kociumaka showed in [34] that if the edit distance is *bounded* by k , then the update time improves to $\tilde{\mathcal{O}}(k)$. Moreover, for *approximate* (unweighted) Edit Distance, there is an algorithm achieving subpolynomial amortized update time with a corresponding approximation guarantee [49].

As stated earlier, it is a major open question whether unweighted TED is strictly harder than unweighted string edit distance, and this same question remains in the dynamic setting.

Question 1: Is unweighted TED strictly harder than String Edit Distance?

Very recently, TED has been extensively studied in settings such as the approximate case [12, 63] and the bounded case [25, 24, 50], where several breakthrough results have been achieved. However, almost nothing is known about its behavior in the dynamic setting. The only (very recent) work about TED that supports update is a dynamic algorithm with approximation factor $n^{1/2+o(1)}$ and subpolynomial update time [26]. However, this result is obtained through an embedding into String Edit Distance and employing known algorithms for strings, rather than an algorithm specific to trees.

⁴ $\omega < 2.372$ is the fast matrix multiplication exponent [8].

The lack of knowledge for dynamic (exact) TED stands out even more jarringly when contrasted with our good understanding of dynamic APSP [44, 45, 29, 30, 68, 69, 6, 36, 21, 55] (which, in the static setting, is equivalent to weighted TED) and dynamic String Edit Distance (which may still be no easier than unweighted TED). In fact, it is reasonable to assume from the recent work of [57] that dynamic improvements for TED might be possible: The authors introduce the notion of border-to-border distance on tree-based graphs. It is a generalization of the *string* alignment graph, which is used by [19] in their dynamic algorithm for String Edit Distance. When we represent the strings as two sequences of singleton trees, we recover exactly the string alignment graph. The seeming connection between the alignment graphs raises the question:

Question 2: Is there a dynamic advantage for dynamic Tree Edit Distance?

Dyck Edit Distance and RNA Folding. Another problem closely related to String Edit Distance is the Dyck Edit Distance Problem, which asks for the minimum number of edits needed to transform a given (not necessarily balanced) parentheses string into a well-balanced one. A further related problem is the RNA Folding Problem: given a string over the nucleotides A, U, C, and G, the goal is to find a maximum set of non-crossing base pairs (A,U and C,G) so that the string “folds” into a valid secondary structure. (See Section 3 for formal definitions and a discussion of why these problems are so closely connected.)

In the static setting, both problems allow for a simple $\mathcal{O}(n^3)$ dynamic programming solution, assuming constant grammar size for Dyck Edit Distance [7, 58]. A series of (inexhaustive) works [74, 14, 73, 22] improves this running time to the current best of $\tilde{\mathcal{O}}(n^{(3+\omega)/2})$ via bounded-monotone min-plus products [23]. A lot of attention has also been given to the approximation setting [60, 61, 27, 51, 12, 63]. In terms of hardness, Abboud *et al.* [2] show a *combinatorial* $\Omega(n^3)$ lower bound for both problems. However, this does not exclude the possibility of a faster dynamic (combinatorial) algorithm, as such an algorithm would be highly desirable in practice. However, in the dynamic setting, nothing is known of these two problems except for an approximation algorithm [26] for dynamic Dyck Edit Distance, which relies on an embedding into strings. This leads to our third question:

Question 3: Is there a dynamic advantage for dynamic Dyck Edit Distance and RNA Folding?

Our Results

In our paper, we make progress in answering the above questions. Technique-wise, we use static hardness conjectures to prove lower bounds for dynamic problems, a technique started by [5]. We base our hardness on the Min-Weight k -Clique Conjecture and the k -Clique Detection Conjecture in fine-grained complexity (see [71] for background). The two problems are defined as follows:

Min-Weight k -Clique

Input: A weighted graph $\mathbf{G} = (V, E, w)$, where $V = [1..n]$.

Output: $\min_{v_1, \dots, v_k \in V} \sum_{i < j} w(v_i, v_j)$ such that v_1, \dots, v_k is a k -clique.

k -Clique Detection

Input: A unweighted graph $\mathbf{G} = (V, E)$ on n nodes, where $V = [1..n]$.

Output: YES if there are $v_1, \dots, v_k \in V$ such that v_1, \dots, v_k is a k -clique, and NO otherwise.

► **Conjecture 1** (Min-Weight k -Clique Conjecture). *For any $\varepsilon > 0$, there exists a constant $c > 0$ such that for any $k \geq 3$, the Min-Weight k -Clique Problem with edge weights in $\{1, \dots, n^{ck}\}$ cannot be solved in $\mathcal{O}(n^{k(1-\varepsilon)})$ time.*

► **Conjecture 2** (Combinatorial k -Clique Detection Conjecture). *For any $\varepsilon > 0$, the k -Clique Detection Problem cannot be solved in $\mathcal{O}(n^{k-\varepsilon})$ time by any combinatorial algorithm.*

The term “combinatorial algorithms” has been used very frequently by the research community (including but not limited to [37, 59, 5, 38, 72, 3]), although it is not a well-defined term. In our paper, as well as many other papers, it is used to refer to algorithms that do not use “algebraic” techniques such as in Strassen’s algorithm [66], which underlie all known algorithms for fast matrix multiplication. (This definition has been first introduced by Ballard *et al.* [10].) Compared to algebraic algorithms, combinatorial algorithms are generally more feasible to implement in practice. Thus, a lower bound for combinatorial algorithms can be somehow also considered as a “practical” lower bound.

For the first question, we show that unweighted TED is strictly harder than String Edit Distance in the dynamic setting when restricted to combinatorial algorithms.

► **Theorem 19.** *Unless Conjecture 2 fails, for any $\varepsilon > 0$, there is no combinatorial algorithm for (unweighted) dynamic TED of size $\mathcal{O}(N)$ that satisfies $p(N) + N(u(N) + q(N)) = \mathcal{O}(N^{3-\varepsilon})$, where $p(N)$, $u(N)$, and $q(N)$ denote the preprocessing, update, and query times of the algorithm, respectively. This lower bound holds even when the alphabet size is constant.*

Thus, even though it still remains open whether unweighted TED can be solved in the same time complexity as String Edit Distance, we show that at least in the dynamic setting, these two problems are separate.

For the second question, we show that there is no dynamic advantage to weighted TED.

► **Theorem 25.** *Unless Conjecture 1 fails, for any $\varepsilon > 0$, there is no algorithm for dynamic TED of size $\mathcal{O}(N)$ and alphabet size $\mathcal{O}(N)$ that satisfies $p(N) + N(u(N) + q(N)) = \mathcal{O}(N^{4-\varepsilon})$, where $p(N)$, $u(N)$, and $q(N)$ denote the preprocessing, update, and query times of the algorithm, respectively.*

Thus, while both weighted String Edit Distance and APSP admit faster dynamic advantage, weighted TED does not: under the Min-Weight k -Clique Conjecture, it cannot be improved beyond recomputing from scratch after each update. This highlights a major difference between weighted TED and the two problems which it is most closely related to, both computationally and definition-wise. In fact, our lower bound holds even for incremental and decremental TED.

To answer the third question, we show that there is no dynamic advantage to RNA Folding and Dyck Edit Distance among combinatorial algorithms.

► **Theorem 29.** *Unless Conjecture 2 fails, for any $\varepsilon > 0$, there is no combinatorial algorithm for dynamic RNA Folding or dynamic Dyck Edit Distance of size $\mathcal{O}(N)$ that satisfy $p(N) + N(u(N) + q(N)) = \mathcal{O}(N^{4-\varepsilon})$, where $p(N)$, $u(N)$, and $q(N)$ are the preprocessing, update, and query times of the algorithm, respectively. This lower bound holds even when the alphabet size is constant.*

Finally, we remark that by combining existing results, one can obtain a dynamic advantage for RNA Folding and Dyck Edit Distance when the updates are only insertions at the end of the string (essentially placing us in an online setting). By combining [31] and [41], one can show an even stronger result: in this online setting, these problems achieve the same time

complexity as Online Matrix-Vector Multiplication (OMv) Problem [38]. The latter is one of the most recent studied problems in fine-grained complexity for dynamic lower bounds: given an $n \times n$ matrix M and a sequence of vectors v_1, \dots, v_n arriving one by one, the task is to output Mv_i before the arrival of v_{i+1} .

► **Lemma 3.** *There is a randomized algorithm that solves online RNA Folding and online Dyck Edit Distance of size $\mathcal{O}(N)$ in the same total time complexity as OMv (currently $N^3/2^{\Omega(\sqrt{\log(N)})}$ by [53]), and succeeds with high probability.*

2 Overview

Lower Bound for Dynamic Unweighted TED. We show how to reduce an instance of $3k$ -Clique Detection, G , to an instance of dynamic unweighted TED. The key ingredient is the construction of embeddings CLG and CNG, referred to as *clique gadgets*, which map two k -cliques X, Y in \mathbf{G} into strings $\text{CLG}(X)$ and $\text{CNG}(Y)$ of length $\mathcal{O}(n \log n)$. These gadgets are designed so that $\text{ed}(\text{CLG}(X), \text{CNG}(Y)) = C$ for some fixed constant C if and only if X and Y are fully connected; otherwise, the edit distance is strictly larger than C .

First, we enumerate the set \mathcal{S} of all k -cliques in \mathbf{G} . Then, for each fixed $Z \in \mathcal{S}$, we construct two trees $\mathbf{T}(Z)$ and $\mathbf{T}'(Z)$ that satisfy:

1. The tree edit distance $\text{ted}(\mathbf{T}(Z), \mathbf{T}'(Z))$ minimizes, over all $X, Y \in \mathcal{S}$, the sum

$$\text{ed}(\text{CLG}(X), \text{CNG}(Y)) + \text{ed}(\text{CNG}(X), \text{CLG}(Z)) + \text{ed}(\text{CLG}(Y), \text{CNG}(Z))$$

2. The sizes of $\mathbf{T}(Z)$ and $\mathbf{T}'(Z)$ are $\mathcal{O}(n^{k+1} \log n)$, while the part of the construction that depends on Z involves only $\mathcal{O}(n \log n)$ nodes in $\mathbf{T}(Z), \mathbf{T}'(Z)$.

Next, we run $|\mathcal{S}| = \mathcal{O}(n^k)$ rounds, one for each $Z \in \mathcal{S}$. In each round, we maintain $\mathbf{T}(Z), \mathbf{T}'(Z)$, update them to the next instance using only $\mathcal{O}(n \log n)$ updates, and then query for the tree edit distance. This allows us to solve the $3k$ -Clique Detection instance in $\mathcal{O}(n^{k+1} \log n)$ updates and $\mathcal{O}(n^k \log n)$ queries. A careful analysis of the factors, together with choosing k sufficiently large, yields Theorem 19.

Lower Bound for Dynamic Weighted TED. Our starting point is the TED lower bound construction of Bringmann *et al.* [13], which encodes a Minimum Weight Triangle instance (known to be equivalent to APSP) into a TED instance. We extend such a reduction into another one that simultaneously ensures the two following:

1. instead of yielding the minimum weight triangle, the instance should yield the minimum weight 4-clique with one of the four vertices fixed; and
2. the portion of the instance that depends on the chosen fixed vertex should involve only $\mathcal{O}(1)$ nodes.

This allows us to proceed similarly to the unweighted case. Given a 4-clique instance \mathbf{G} , the algorithm runs for n rounds, one for each vertex in \mathbf{G} . In each round, we maintain the instance corresponding to the fixed vertex, update it to the next instance using only $\mathcal{O}(1)$ updates, and then query for the minimum-weight 4-clique containing the current vertex. Taking the minimum over all queried values yields the minimum weight 4-clique.

We remark that although, at a high level, this reduction may appear similar to the one for the unweighted case, the unweighted instances $\mathbf{T}(Z)$ and $\mathbf{T}(Z')$ differ significantly from the weighted TED instances (compare Figure 1 with Figure 2b). The latter rely heavily on extremely large weights, whereas the construction of $\mathbf{T}(Z)$ and $\mathbf{T}(Z')$ is closer in spirit to gadgets used for unweighted string edit distance such as [9] than to those for weighted TED.

Lower Bound for Dynamic RNA Folding and Dyck Edit Distance. Our starting point is once again the static lower bound [2]. Conceptually, the reduction is similar to the previous ones: the static lower bound encodes a $3k$ -Clique Detection instance into a single RNA Folding instance. For the dynamic lower bound, however, we instead encode a $4k$ -Clique Detection instance into a single RNA Folding instance, where k of the vertices are fixed for each round. Similar to the previous reductions, we must ensure that the part of our dynamic RNA Folding instance depending on the currently fixed vertices is small in size, so the number of updates stays small. Regarding Dyck Edit Distance, we do not need to build any new instance. We use the embedding of RNA Folding into Dyck Edit Distance from [18].

Organization of the Paper. First, we set up notation in Section 3. Section 4, Section 5 and Section 6 are dedicated to the proof of Theorem 19, Theorem 25, and Theorem 29, respectively. For the sake of completeness, the proof of Lemma 3 is presented in the appendix of the full version [40].

3 Preliminaries

Sets. For integers $i, j \in \mathbb{Z}$, we write $[i..j]$ to represent the set $\{i, \dots, j\}$, and $[i..j)$ to denote the set $\{i, \dots, j-1\}$. We define $(i..j]$ and $(i..j)$ similarly.

Strings. An *alphabet* Σ is a finite set of symbols. A *string* $X \in \Sigma^n$ of length n is written as $X = X[1]X[2]\cdots X[n]$, where each $X[i]$ denotes the i -th character of X , for $i \in \{1, \dots, n\}$. The length of string X is denoted as $|X|$. Given indices $1 \leq i \leq j \leq |X| + 1$, we define the *fragment* $X[i..j) := X[i]\cdots X[j-1]$. A *prefix* of X is any fragment starting at position 1, and a *suffix* is any fragment ending at position $|X|$. A string Y of length $m \in [0..n]$ is a *substring* of X if there exists an index i such that $Y = X[i..i+m)$. For two strings A and B , their *concatenation* is denoted by $A \circ B$ or simply AB . The notation A^k represents the string formed by concatenating k copies of A . The *reverse* of a string X is written as $X^R := X[n]\cdots X[1]$.

Finally, for strings A and B , we denote by $\text{ed}(A, B)$ the (*string*) *edit distance* (in short, ED) between A and B , defined as the minimum number of character insertions, deletions, and substitutions to transform A into B .

Tree Edit Distance (TED). We consider ordered, node-labeled trees from an alphabet Σ . Given a node v in a tree, we use $\ell(v)$ to denote its label, and we use $\text{sub}(v)$ to denote the subtree rooted at v .

Weighted Tree Edit Distance

Input: Two trees \mathbf{T}, \mathbf{T}' labeled from an alphabet Σ and a cost function $\delta : \Sigma \cup \{\varepsilon\} \times \Sigma \cup \{\varepsilon\} \mapsto \mathbb{R}$.

Output: $\text{ted}(\mathbf{T}, \mathbf{T}')$ defined as the minimum cost of transforming \mathbf{T} into \mathbf{T}' by performing a sequence of edit operations, which consist of the following three types:

- Changing the label of a node from ℓ to ℓ' , $\ell \neq \ell'$ at cost $\delta(\ell, \ell')$.
- Deleting a node v and attaching its children (if there are any) to the parent of v in their original order (if v is the root, then we obtain an ordered forest) at cost $\delta(\ell(v), \varepsilon)$.
- For an existing node v' inserting a new node v as a new leaf child in some position, or selecting consecutive children w_1, \dots, w_k of v' , inserting a new node v as the new parent of w_1, \dots, w_k , and placing v as a child of v' at cost $\delta(\varepsilon, \ell(v))$.

When all edit operations cost 1, we say the tree edit distance is *unweighted*. For our paper, it is more convenient to work with a different definition of TED, which is more similar to the definition used in [54] and [57], and is defined through the notion of *tree-alignment*.

► **Definition 4.** For two trees \mathbf{T}, \mathbf{T}' , we say the sequence $\mathcal{A} = \{(v_i, v'_i)\}_{i=1}^k \in \mathbf{T} \times \mathbf{T}'$ is a tree-alignment of \mathbf{T} onto \mathbf{T}' if for all $1 \leq i < j \leq k$:

- v_i is an ancestor of v_j in \mathbf{T} iff v'_i is an ancestor of v'_j in \mathbf{T}' ,
- if neither v_i nor v_j is the ancestor of the other, then v_i comes before v_j in the pre-order traversal of \mathbf{T} iff v'_i comes before v'_j in the pre-order traversal of \mathbf{T}' .

For a tree-alignment $\mathcal{A} = \{(v_i, v'_i)\}_{i=1}^k \in \mathbf{T} \times \mathbf{T}'$ and $i \in [1..k]$, we say \mathcal{A} aligns v_i with v'_i . If further $\ell(v_i) = \ell(v'_i)$, then we say \mathcal{A} matches v_i with v'_i , otherwise we say \mathcal{A} substitutes v_i with v'_i . The *cost* of a tree-alignment \mathcal{A} , denoted by $\text{cost}(\mathcal{A})$, is defined as

$$\text{cost}(\mathcal{A}) := \sum_{\substack{(v, v') \in \mathcal{A} \\ \ell(v) \neq \ell(v')}} \delta(\ell(v), \ell(v')) + \sum_{v \in \mathbf{T}: \exists_{v'} (v, v') \in \mathcal{A}} \delta(\ell(v), \varepsilon) + \sum_{v' \in \mathbf{T}': \exists_v (v, v') \in \mathcal{A}} \delta(\varepsilon, \ell(v')).$$

Finding the minimum cost of a tree-alignment is an equivalent definition of TED [54]⁵.

In our paper, we also view String Edit Distance through the lens of *string-alignment*. In this case, we can think of each string as a forest consisting of n single-node trees. Then Definition 4 recovers exactly their string-alignment.

Path Gadgets. To describe more concisely our lower bound instances for TED, we define a basic gadget $\mathbf{P}(S)$, which, given a string $S = s_1 \dots s_d$, constructs a path of length d with nodes labeled s_1, \dots, s_d from top to bottom. Moreover, given a path \mathbf{P} and a node $v \in \mathbf{P}$, we say that we *right-attach* a string $S = s_1, \dots, s_d$ to v when we attach d nodes as right children of v , labeled s_1, \dots, s_d from left to right. We define *left-attaching* analogously, for attaching nodes to the left of v .

RNA Folding. The RNA Folding Problem is defined over strings on the alphabet $\Sigma \cup \Sigma'$, where each character in Σ has a correspondingly paired character in Σ' . For $\sigma \in \Sigma$, we denote its paired character in Σ' by σ' . Moreover, for a string S over $\Sigma \cup \Sigma'$, we define $p(S)$ to be the string obtained by replacing each character with its paired character.

Two index pairs (i, j) and (i', j') , where $i < j$ and $i' < j'$, are said to *cross* if either $i < i' < j < j'$ or $i' < i < j' < j$ holds. Now, given a string $S \in (\Sigma \cup \Sigma')^n$, a set of index pairs $F = \{(i, j) : 1 \leq i < j \leq n\}$ is a *folding* of S if:

- (i) for all distinct $(i, j), (i', j') \in F$, the pairs do not cross; and
- (ii) for all $(i, j) \in F$, we have $S[i] = p(S[j])$.

This allows us to define the RNA Folding Problem, along with a weighted variant.

RNA Folding

Input: A string $S \in \{\Sigma \cup \Sigma'\}^n$.

Output: $\text{score}(S) := \max_{F \text{ folding of } S} |F|$.

⁵ To be precise, in [54] and [57], the complementary problem is considered, i.e., finding the maximum value of $\sum_{v \in \mathbf{T}} \delta(\ell(v), \varepsilon) + \sum_{v' \in \mathbf{T}'} \delta(\varepsilon, \ell(v')) - \text{cost}(\mathcal{A})$, taken over all tree-alignments \mathcal{A} . This is evidently equivalent to minimizing $\text{cost}(\mathcal{A})$, in which case we recover the tree edit distance.

Weighted RNA Folding

Input: A string $S \in \{\Sigma \cup \Sigma'\}^n$ and a weight function $w : \Sigma \cup \Sigma' \rightarrow [1..M]$ where $w(\sigma) = w(p(\sigma))$ for all $\sigma \in \Sigma \cup \Sigma'$.

Output: $\text{score}_w(S) := \max_{F \text{ folding of } S} \sum_{(i,j) \in F} w(S[i]).$

The weighted version is convenient because it can be reduced to the unweighted problem with only an overhead of M on the length of the new string, where M is the largest weight. This allows us to work with the weighted formulation, as long as we keep M reasonably small.

► **Lemma 5** ([2]). *Let $S \in \{\Sigma \cup \Sigma'\}^n$ be a string and $w : \Sigma \rightarrow [1..M]$ a weight function. Then, $\text{score}_w(S) = \text{score}(S')$ for the string $S' := s_1^{w(s_1)} \dots s_n^{w(s_n)}$.* ◻

Dyck Edit Distance. The Dyck Edit Distance Problem asks for the minimum number of edits required to transform a given string into a well-balanced string of parentheses. It can be thought of as a variant of RNA Folding, with three key differences:

- (1) it is formulated as a minimization problem rather than a maximization problem;
- (2) it also allows for substitution; and
- (3) a symbol σ may be matched with its corresponding closing symbol σ' only if σ appears before σ' in the string, but not the other way around.

Dyck Edit Distance

Input: A string $S \in \{\Sigma \cup \Sigma'\}^n$.

Output: $\text{dyck}(S) := \min_{S' \in \mathcal{L}_{\text{dyck}}(\Sigma)} \text{ed}(S, S')$, where $\mathcal{L}_{\text{dyck}}(\Sigma)$ is the language defined by the grammar with the rules $\mathbf{S} \rightarrow \mathbf{SS}$, $\mathbf{S} \rightarrow \varepsilon$, and $\mathbf{S} \rightarrow \sigma \mathbf{S} \sigma'$ for all $\sigma \in \Sigma$.

Regarding Dyck Edit Distance, no additional setup is required: for the purpose of proving lower bounds, one can directly reduce RNA Folding into Dyck Edit Distance.

► **Lemma 6** ([18]). *Let $S = s_1 \dots s_n$ be a string over $\Sigma \cup \Sigma'$. Let $\bar{\Sigma} := \{\#\} \cup \{0_\sigma, 0'_\sigma, 1_\sigma, 1'_\sigma\}_{\sigma \in \Sigma}$, and define the function $\phi : \Sigma \rightarrow \bar{\Sigma}^6$ such that $\phi(\sigma) = 0_\sigma \# 1'_\sigma 0_\sigma \# 1'_\sigma$ and $\phi(\sigma') = 1_\sigma 1_\sigma 0'_\sigma 0'_\sigma$, for $\sigma \in \Sigma$.*

Then, for $\Phi(S) := \phi(s_1) \dots \phi(s_n)$, we have $\text{dyck}(\Phi(S)) = 3|S| - 2\text{score}(S)^6$. ◻

Clique Gadgets. Central to the lower bound construction of [2], and also to the dynamic lower bounds presented here for Dyck Edit Distance, RNA Folding, and unweighted TED, are the gadget constructions from [2] that encode the neighborhoods of nodes and set of nodes. To describe such gadgets formally, given a graph $\mathbf{G} = (V, E)$ and a node $v \in V$, we use $N(v)$ to denote the set of neighbors of v . For a subset $V' \subseteq V$, we define the neighbor set of V' to be $N(V') = \bigcap_{v' \in V'} N(v')$.

► **Lemma 7** (Claim 4 in [2]). *Let $\mathbf{G} = (V, E)$ be a graph with $|V| = n$, and let $k \in \mathbb{Z}_+$. Then, there exist two string embeddings $\text{CLG} : V^k \rightarrow \Sigma^\lambda$ and $\text{CNG} : V^k \rightarrow \Sigma^\lambda$ of length⁷ $\lambda = \mathcal{O}(nk \log n)$ over an alphabet of size $|\Sigma| = \mathcal{O}(1)$, and a constant $C = C(n, k)$ such that for any $X, Y \in V^k$:*

$$\begin{aligned} \text{score}(\text{CLG}(X) \circ p(\text{CNG}(Y))^R) &= C && \text{if } X \subseteq N(Y), \\ \text{score}(\text{CLG}(X) \circ p(\text{CNG}(Y))^R) &< C && \text{otherwise.} \end{aligned}$$
 ◻

⁶ In [18], the author specifically considers the case $|\Sigma| = 2$, but its proof extends straightforwardly to alphabets of arbitrary size $|\Sigma| > 2$ by symmetry, as long as distinct letters are used for each symbol in the embedding.

⁷ In [2], the embeddings have length *at most* λ . We pad each one with two fresh symbols so that every embedding now has length *exactly* λ , without affecting the score.

Thus, these gadgets allow us to determine whether two k -cliques are fully connected to each other.

4 Lower Bounds for Dynamic Unweighted TED

In this section, we prove Theorem 19. As a first step, we construct gadgets that serve the same purpose as those in [2], but are adapted to work under edit distance.

► **Lemma 8.** *Let $\mathbf{G} = (V, E)$ be a graph, and let $k \in \mathbb{Z}_+$. Then, there exist two string embeddings $\text{CLG} : V^k \rightarrow \Sigma^{\lambda_1}$ and $\text{CNG} : V^k \rightarrow \Sigma^{\lambda_2}$ of lengths $\lambda_1, \lambda_2 = \mathcal{O}(nk^3 \log n)$ over an alphabet of size $|\Sigma| = \mathcal{O}(1)$, and a constant $C = C(n, k)$ such that for any $X, Y \in V^k$:*

$$\begin{aligned} \text{ed}(\text{CLG}(X), \text{CNG}(Y)) &= C && \text{if } X \subseteq N(Y), \\ \text{ed}(\text{CLG}(X), \text{CNG}(Y)) &> C && \text{otherwise.} \end{aligned}$$

Proof. In this proof, we use the alphabet $\Sigma := \{0, 1, 2, 3, \$, \#\}$. We set $\alpha = \mathcal{O}(\log n)$ to be the number of bits needed to encode values of $V = [1..n]$ in binary. Moreover, for a number $x \in V$, we denote with $\bar{x} \in \{0, 1\}^\alpha$ its binary encoding.

For nodes $v, u \in V = [1..n]$, we define the *node gadget* $\text{NG}(v)$ and *list gadget* $\text{LG}(v)$ as

$$\text{NG}(v) := (\bigcirc_{i=1}^{n-1} 2^\alpha \$) \bar{v} \$ (\bigcirc_{i=1}^{n-1} 2^\alpha \$) \quad \text{and} \quad \text{LG}(u) := (\bigcirc_{w=1}^n g_u(w) \$),$$

where $g_u(w) = \bar{w}$ if $w \in N(u)$ and $g_u(w) = 3^\alpha$, otherwise.

▷ **Claim 9.** Let $C' := (n-1)(2\alpha+1)$. Then, for nodes $u, v \in V$, we have:

$$\begin{aligned} \text{ed}(\text{NG}(v), \text{LG}(u)) &= C' && \text{if } v \in N(u), \\ \text{ed}(\text{NG}(v), \text{LG}(u)) &\in [C' + 1..C' + \alpha] && \text{otherwise.} \end{aligned}$$

Proof. We first prove that $\text{ed}(\text{NG}(v), \text{LG}(u)) \leq C' + \alpha$ (and if $v \in N(u)$, then $\text{ed}(\text{NG}(v), \text{LG}(u)) \leq C'$). To this end, construct the following (string) alignment \mathcal{A}_u :

- (a) Substitute $g_u(v)$ with \bar{v} (if $v \in N(u)$ then match exactly as $g_u(v) = \bar{v}$);
- (b) Align the $(\alpha+1)(v-1)$ characters coming before $g_u(v)$ and \bar{v} to each other;
- (c) Align the $(\alpha+1)(n-v)$ characters coming after $g_u(v)$ and \bar{v} to each other; and
- (d) delete the remaining characters.

We briefly argue that $\text{cost}(\mathcal{A}_u) \leq C' + \alpha = (n-1)(2\alpha+1) + \alpha$: in (a), we incur cost at most α (this becomes 0 if $v \in N(u)$), in (b) and (c) we incur cost $\alpha(v-1)$ and $\alpha(n-v)$, respectively, as out of the $(\alpha+1)(v-1)$ and $(\alpha+1)(n-v)$ characters, there are $v-1$ and $n-v$ perfectly matched $\$$. Finally, in (d) we incur cost $(\alpha+1)(n-1)$, as the remaining characters include a length- $(\alpha+1)(n-v)$ prefix and a length- $(\alpha+1)(v-1)$ suffix of $\text{NG}(u)$.

To conclude the proof, we consider an optimal alignment \mathcal{A} , and show that it has cost at least C' . Furthermore, we show that whenever \mathcal{A} has cost exactly C' , then $v \in N(u)$.

To this end, let m be the number of positions i in $\text{LG}(u)$ such that $\text{LG}(u)[i]$ is not matched perfectly by \mathcal{A} . Note that $m \geq \alpha(n-1)$, as out of the $n\alpha$ many non- $\$$ characters in $\text{LG}(u)$, no more than α many 0/1 characters can be matched exactly. Moreover, from $|\text{NG}(v)| \geq |\text{LG}(u)|$ follows that the cost of \mathcal{A} is at least

$$(|\text{NG}(v)| - |\text{LG}(u)|) + m \geq (\alpha+1)(2n-1) - (\alpha+1)n + \alpha(n-1) = (n-1)(2\alpha+1) = C'.$$

Now, if \mathcal{A} has cost exactly C' , this last inequality must be tight and $m = \alpha(n-1)$. In particular, this means that \mathcal{A} matches perfectly all $\$$ characters and α many 0/1 characters of $\text{LG}(u)$. Observe that these α many 0/1 characters must come from the same $g_u(w)$, for

78:10 Hardness of Dynamic Tree Edit Distance and Friends

some $w \in V$. (If not, then at least one $\$$ would be deleted, contradicting the fact that all $\$$ of $\text{LG}(v)$ are perfectly matched.) The only place where $g_u(w)$ can be matched perfectly to is \bar{v} , as it is the only part of $\text{NG}(v)$ having 0/1 characters. Since $g_u(w)$ and \bar{v} are matched perfectly, we have $g_u(w) = \bar{v}$, which can only happen when $w = v$ and $w \in N(u)$, i.e., $v \in N(u)$. \triangleleft

Next, for subsets $X, Y \subseteq V^k$, we define the *clique node gadget* $\text{CNG}(Y)$ and *clique list gadget* $\text{CLG}(X)$ as

$$\text{CNG}(Y) := \bigcirc_{v \in Y} (\text{NG}(v) \#^\ell)^k \quad \text{and} \quad \text{CLG}(X) := \left(\bigcirc_{v \in X} \text{LG}(v) \#^\ell \right)^k,$$

where $\ell := \lceil k^2 \cdot (C' + \alpha) / 2 \rceil$. Thus, λ_1, λ_2 satisfy $\lambda_1, \lambda_2 = \mathcal{O}(k \cdot \ell) = \mathcal{O}(nk^3 \log n)$.

For $i \in [1..k^2]$, let B_i and B'_i be the i -th substrings of the form $\text{NG}(v)$ and $\text{LG}(v)$ appearing in $\text{CNG}(Y)$ and $\text{CLG}(X)$, respectively. Moreover, let H_i and H'_i be the substrings $\#^\ell$ right after B_i and B'_i . Lastly, let \mathcal{A} be an optimal (string) alignment that aligns $\text{CLG}(X)$ onto $\text{CNG}(Y)$ and maximizes the number of i, j such that $H'_i[j]$ is matched to $H_i[j]$.

\triangleright **Claim 10.** There are no $i, j \in [1..k^2]$ such that $i \neq j$ and \mathcal{A} aligns any character of B_i and B'_j to each other.

Proof. For the sake of contradiction, assume there exist indices $i \neq j$ such that a character from B_i matches a character from B'_j . Observe that B_i is preceded by $(i-1)\ell$ and followed by $(k^2 - i - 1)\ell$ occurrences of $\#$, while B'_j is preceded by $(j-1)\ell$ and followed by $(k^2 - j - 1)\ell$ occurrences of $\#$. Since $i \neq j$, alignment \mathcal{A} necessarily incurs a cost of at least ℓ both before and after the matched character between B_i and B'_j . Hence, $\text{cost}(\mathcal{A}) \geq 2\ell$.

Now consider an alternative alignment \mathcal{A}' which matches every H_x perfectly with H'_x , and aligns each B_x optimally with B'_x for all $x \in [1..k^2]$. By Claim 10, the cost of such an alignment satisfies $\text{cost}(\mathcal{A}') \leq k^2 \cdot (C' + \alpha) \leq \text{cost}(\mathcal{A})$. This contradicts the assumed optimality of \mathcal{A} (and if $\text{cost}(\mathcal{A}') = \text{cost}(\mathcal{A})$, then \mathcal{A}' is strictly better in terms of the number of pairs (x, y) where $H'_x[y]$ is matched to $H_x[y]$). \triangleleft

\triangleright **Claim 11.** For all $i \in [1..k^2]$, H_i and H'_i are perfectly matched by \mathcal{A} .

Proof. For each $i \in [1..k^2]$, let c_i denote the number of characters in H_i that are deleted or substituted in the alignment \mathcal{A} , and define c'_i analogously for H'_i . Note that $\sum_i c_i = \sum_i c'_i$, since they are the only substrings that contain the character $\#$.

Now, consider modifying \mathcal{A} to \mathcal{A}' by perfectly matching each H_i to H'_i , and deleting any characters that were previously involved in substitutions with characters in H_i . By Claim 10, we have that \mathcal{A}' is still a valid alignment.

Let s denote the number of substitutions that are removed in this process. Then the change in cost is $\text{cost}(\mathcal{A}') - \text{cost}(\mathcal{A}) = s - \sum_i c_i - \sum_i c'_i$. Since $s \leq \sum_i c_i + \sum_i c'_i$, this change is not positive, so also \mathcal{A}' must be optimal. By our choice of \mathcal{A} , the number of pairs (i, j) such that $B_i[j]$ is matched to $B'_i[j]$ cannot increase. Therefore, if such a number remains unchanged, then $\sum_i c_i = 0$, which implies that \mathcal{A} already matches each H_i perfectly with H'_i . \triangleleft

Now observe that Claim 11 implies

$$\text{ed}(\text{CLG}(X), \text{CNG}(Y)) = \sum_i \text{ed}(B_i, B'_i) = \sum_{v \in Y, u \in X} \text{ed}(\text{CG}(v), \text{LG}(u)).$$

Since $X \subseteq N(Y)$ (or equivalently $Y \subseteq N(X)$) holds if and only if $v \in N(u)$ for all $u \in X$, $v \in Y$, it follows from Claim 9 that we can set $C := C'k^2$ to conclude the proof. \blacktriangleleft

Next, we present the trees in our lower-bound construction. As part of our construction, we use the label set Σ from Lemma 8 and introduce four additional labels: \S , $\$,$ and $\#_\ell$ for $\ell \in \{\mathbf{L}, \mathbf{R}\}$. We assume these new labels are disjoint from Σ . Further, in our lower-bound instance we use the gadgets CLG, CNG from Lemma 8 of lengths λ_1 and λ_2 , respectively. We set $\lambda := \max(\lambda_1, \lambda_2)$.

► **Definition 12.** Let $\mathbf{G} = (V, E)$ be a graph, let $k \in \mathbb{Z}_+$, and let \mathcal{X}, \mathcal{Y} , and \mathcal{Z} all denote the set of k -cliques in \mathbf{G} . Suppose $|\mathcal{X}| = |\mathcal{Y}| = |\mathcal{Z}| = N$. For $Z \in \mathcal{Z}$ define the two trees $\mathbf{T}_{\mathcal{X}}(Z)$ and $\mathbf{T}'_{\mathcal{Y}}(Z)$ as follows:

1. $\mathbf{T}_{\mathcal{X}}(Z)$ is obtained by attaching the following nodes to $\mathbf{P}(\$ \circ (\bigcirc_{X \in \mathcal{X}} \text{CLG}(X)) \circ \S^{16\lambda N + \lambda + 1})$:
 - (a) We left attach $\#_{\mathbf{L}}^{5\lambda N}$ to $\$,$
 - (b) For each $X \in \mathcal{X}$, we right attach $\#_{\mathbf{R}}^{4\lambda} \circ \text{CNG}(X)$ to the last node of $\text{CLG}(X)$ on the path, and
 - (c) We left attach $\text{CNG}(Z) \circ \#_{\mathbf{L}}^{5\lambda N}$ to the last node belonging to any $\text{CLG}(X)$ for $X \in \mathcal{X}$ (so to the spine node right before the first node with label \S).
2. $\mathbf{T}'_{\mathcal{Y}}(Z)$ is obtained by attaching the following nodes to $\mathbf{P}(\$ \circ (\bigcirc_{Y \in \mathcal{Y}} \text{CNG}(Y)) \circ \S^{16\lambda N + \lambda + 1})$:
 - (a) We right attach $\#_{\mathbf{R}}^{5\lambda N}$ to $\$,$
 - (b) For each $Y \in \mathcal{Y}$, we left attach $\text{CLG}(Y) \circ \#_{\mathbf{L}}^{5\lambda}$ to the last node of $\text{CNG}(Y)$ on the path, and
 - (c) We right attach $\#_{\mathbf{R}}^{5\lambda N} \circ \text{CLG}(Z)$ to the last node belonging to any $\text{CNG}(Y)$ for $Y \in \mathcal{Y}$ (so to the spine node right before the first node with label \S).

Refer to Figure 1 for a visualization of $\mathbf{T}_{\mathcal{X}}(Z)$ and $\mathbf{T}'_{\mathcal{Y}}(Z)$. We remark that in $\mathbf{T}_{\mathcal{X}}(Z)$ and $\mathbf{T}'_{\mathcal{Y}}(Z)$, each gadget $G(W)$ for $G \in \{\text{CLG}, \text{CNG}\}$ and $W \in \mathcal{X} \cup \mathcal{Y} \cup \{Z\}$ appears exactly once in each tree. Therefore, by a slight abuse of notation, we will also use $G(W)$ to refer to the set of nodes comprising that gadget.

► **Theorem 13.** We have $\text{ted}(\mathbf{T}_{\mathcal{X}}(Z), \mathbf{T}'_{\mathcal{Y}}(Z))$ equals to

$$\min_{Y \in \mathcal{Y}, X \in \mathcal{X}} \left(\text{ed}(\text{CLG}(X), \text{CNG}(Y)) + \text{ed}(\text{CNG}(X), \text{CLG}(Z)) + \text{ed}(\text{CLG}(Y), \text{CNG}(Z)) \right) + D,$$

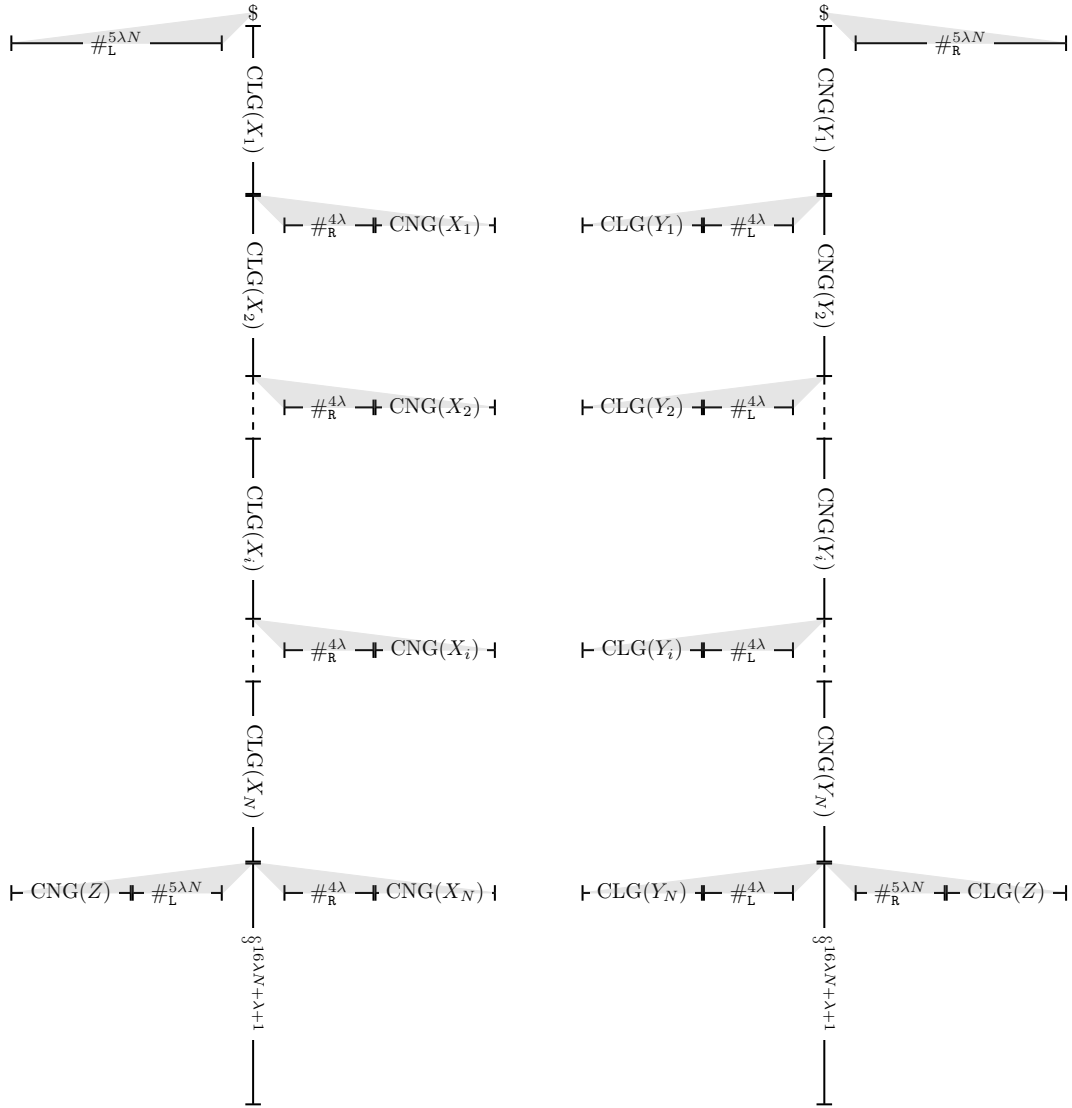
for $D = 12\lambda N - \lambda_2 N - \lambda_1 N - \lambda_2 - \lambda_1$.

Proof. Throughout this proof, we abbreviate $\mathbf{T}_{\mathcal{X}}(Z)$ and $\mathbf{T}'_{\mathcal{Y}}(Z)$ as \mathbf{T} and \mathbf{T}' , respectively. Additionally, let X_1, \dots, X_N denote the sets in \mathcal{X} such that $\text{CLG}(X_i)$ appears on the spine above $\text{CLG}(X_{i+1})$ in \mathbf{T} , and let \mathbf{L} (resp. \mathbf{R}) be the nodes on the left (resp. right) of the spine of \mathbf{T} . We define Y_1, \dots, Y_N and \mathbf{L}', \mathbf{R}' analogously for \mathcal{Y} and \mathbf{T}' .

Central in this proof is the alignment $\mathcal{A}_{a,b}$ of \mathbf{T} onto \mathbf{T}' , structured for $a, b \in [1..N]$ as follows:

- (i) The roots and the last $16\lambda N + \lambda + 1$ nodes of the spines of \mathbf{T}, \mathbf{T}' are perfectly matched.
- (ii) The three pairs of gadgets $\text{CLG}(X_a), \text{CNG}(Y_b)$, and $\text{CNG}(X_a), \text{CLG}(Z)$, and $\text{CLG}(Y_b), \text{CNG}(Z)$ are aligned as the corresponding strings would be under string ED.
- (iii) Match or substitute the nodes of $\mathbf{R} \setminus \text{CNG}(X_a)$ with nodes labeled $\#_{\mathbf{R}}$ as follows. The ones above the gadget $\text{CLG}(X_a)$ with the nodes of 2(a), while those below with the nodes of 2(c). Note that $|\mathbf{R}| = 5\lambda N$, so all nodes of $\mathbf{R} \setminus \text{CNG}(X_a)$ can be matched or substituted as described, as in 2(a) and 2(c) we add enough to do so.

78:12 Hardness of Dynamic Tree Edit Distance and Friends



■ **Figure 1** The trees $\mathbf{T}_X(Z)$ and $\mathbf{T}'_Y(Z)$ are depicted on the left and right, respectively.

- (iv) Similarly, match or substitute the nodes of $\mathbf{L}' \setminus \text{CLG}(Y_b)$ with nodes labeled $\#_L$ like so: The ones above the gadget $\text{CLG}(Y_b)$ with the nodes of 1(a), while those below with the nodes of 1(c).
- (v) All remaining nodes are deleted.

▷ **Claim 14.** The alignment $\mathcal{A}_{a,b}$ has cost

$$\text{ed}(\text{CLG}(X_a), \text{CNG}(Y_b)) + \text{ed}(\text{CNG}(X_a), \text{CLG}(Z)) + \text{ed}(\text{CLG}(Y_b), \text{CNG}(Z)) + D.$$

Proof. The cost can be directly verified. Refer to [40] for a detailed proof. ◁

In order to prove Theorem 13, it suffices to show that any optimal alignment \mathcal{A} transforming \mathbf{T} into \mathbf{T}' has cost at least as high as that of some alignment $\mathcal{A}_{a,b}$. To do so, we will demonstrate that \mathcal{A} satisfies all construction steps, (i) through (v), for some choice of $a, b \in [1 \dots N]$.

▷ **Claim 15.** For an optimal alignment \mathcal{A} , the nodes on the spine of \mathbf{T} must be matched or substituted by \mathcal{A} with nodes on the spine of \mathbf{T}' , and similarly, the nodes of \mathbf{R} (resp. \mathbf{L}) must be matched or substituted by \mathcal{A} with the nodes of \mathbf{R}' (resp. \mathbf{L}').

Proof. It suffices to argue that \mathcal{A} matches at least a pair of nodes labeled \S with each other, then by the rules of tree alignments, the structure specified in the claim follows. To this end, suppose for the sake of contradiction that every node labeled \S is deleted or substituted, each incurring a cost of at least 1 and resulting in a total cost of at least $2(16\lambda N + \lambda + 1)$. By matching the roots and all nodes labeled \S with each other and deleting the remaining nodes, we incur cost $4 \cdot 5\lambda N + \lambda_2 + \lambda_1 + 2 \cdot 4\lambda N + 2\lambda_2 N + 2\lambda_1 N < 2(16\lambda N + \lambda + 1)$, which contradicts the optimality of \mathcal{A} . ◁

By Claim 15, we may assume that \mathcal{A} satisfies the three following:

- (a) Let $v, v' \in \mathbf{T}$ be two spine nodes such that v is an ancestor of v' , and such that v and v' but none of $(\text{sub}(v) \setminus \{v\}) \setminus \text{sub}(v')$, are matched or substituted by \mathcal{A} . Then, none of the spine nodes between v and v' share the same label as v .
- (b) Let $v, v' \in \mathbf{L}$ be two nodes such that v occurs in the pre-order of \mathbf{T} before v' , and such that v and v' but none of the nodes in the pre-order between v and v' , are matched or substituted by \mathcal{A} . Then, none of the nodes in the pre-order between v and v' share the same label as v .
- (c) Similarly, let $v, v' \in \mathbf{R}$ be nodes such that v occurs in the post-order of \mathbf{T} after v' , and such that v and v' but none of the nodes in the post-order between v and v' , are matched or substituted by \mathcal{A} . Then, none of the nodes in the post-order between v and v' share the same label as v .

Whenever \mathcal{A} has this form, we say \mathcal{A} *favors matching downwards*. Note, if \mathcal{A} does not satisfy this structural form, then it can be modified (iteratively) so that it does, without changing its cost.

▷ **Claim 16.** The roots labeled $\$,$ as well as the last $16\lambda N + \lambda + 1$ spine nodes labeled $\S,$ in \mathbf{T} and \mathbf{T}' are matched with each other by \mathcal{A} . That is, construction step (i) is satisfied.

Proof. Clearly, the optimal alignment must match the two roots with each other.

Suppose, for contradiction, that for some $i \in [1..16\lambda N + \lambda + 1]$, the i -th last node on the spine of \mathbf{T} and the i -th last node on the spine of \mathbf{T}' are not matched with each other. Let i be the smallest such index. In the situation where either (1) both are deleted; or (2) exactly one is substituted with some other node whose label is not $\S,$ we can further reduce the cost of \mathcal{A} by matching these two nodes together, contradicting the optimality of \mathcal{A} . Thus, at most one of the two nodes is substituted with another node labeled $\S,$ but this contradicts our earlier assumption that \mathcal{A} favors matching downwards. ◁

Next, consider the largest $i \in [1..N]$ such that a node right-attached to $\text{CLG}(X_{i-1})$ is matched with a node labeled $\#_{\mathbf{R}}$ from 2(a) (if $i = 1$ then such a node does not exist). Moreover, consider the smallest $i' \in [1..N]$ such that a node right-attached to $\text{CLG}(X_{i'})$ is matched with a node labeled $\#_{\mathbf{R}}$ from 2(c). Note that i' is well-defined. Indeed, in \mathcal{A} the right-attached nodes to $\text{CLG}(X_N)$ with label $\#_{\mathbf{R}}$ must match with nodes from 2(c) because \mathcal{A} favors matching downwards. Symmetrically, we define j, j' indexing Y_1, \dots, Y_N .

In the remainder of the proof, we show that \mathcal{A} satisfies (ii), (iii), (iv), and (v), for $a = i$ and $b = j$.

▷ **Claim 17.** The indices i, i' satisfy $i' \geq i$.

Proof. First, note that since \mathcal{A} is a valid tree alignment, we must have $i' \geq i - 1$. To eliminate the remaining case $i' = i - 1$, for the sake of contradiction, suppose $i' = i - 1$, which we can only have if $i > 1$.

Observe that the condition $i > 1$ implies that among the $4\lambda N + \lambda_1 N$ nodes in \mathbf{R} , there exists a node $v \in \mathbf{R}$ that is not matched with a node labeled $\#_{\mathbf{R}}$ from 2(c), but instead is matched with a node $w \in \mathbf{R}'$ from 2(a). Let w be the first such matched node encountered in the pre-order traversal of \mathbf{T}' .

As a result, there must also exist a node $u \in \mathbf{R}'$ labeled $\#_{\mathbf{R}}$ from 2(c) that is neither matched nor substituted. By the assumption that \mathcal{A} favors downward matches, we may assume u is the last node visited in the pre-order traversal of \mathbf{T}' .

Finally, observe that the assumption $i' = i - 1$ implies that $\text{CLG}(Z)$ must be deleted. Since no nodes between w and u are matched or substituted, this contradicts the assumption that \mathcal{A} favors matching downward. \triangleleft

▷ **Claim 18.** The indices i, i' satisfy $i = i'$.

Proof. Suppose, for contradiction, that $i' > i$. Then all of the 4λ nodes labeled $\#_{\mathbf{R}}$, which are right-attached to $\text{CLG}(X_{i'-1})$, are deleted. This implies that also at least 4λ nodes labeled $\#_{\mathbf{R}}$ from 2(c) must be deleted. By the assumption that \mathcal{A} favors downward matching, we may assume these deletions occur among the last 4λ nodes visited in the pre-order traversal of \mathbf{T}' .

We now modify \mathcal{A} as follows: we match these 4λ last-visited nodes with the 4λ right-attached $\#_{\mathbf{R}}$ -labeled nodes of $\text{CLG}(X_{i'-1})$, and instead delete all nodes from $\text{CLG}(X_{i'-1})$, $\text{CNG}(X_{i'-1})$, as well as any nodes previously matched to the deleted gadgets (at most $2\lambda_2 + 2\lambda_1$ nodes in total). This modification reduces the overall cost of \mathcal{A} by at least $8\lambda - 2\lambda_2 + 2\lambda_1 \geq 4\lambda > 0$, contradicting its optimality. \triangleleft

It is now not difficult to see, by Claim 15 and Claim 18, that all nodes in $\mathbf{R} \setminus \text{CNG}(X_i)$ are matched to nodes labeled $\#_{\mathbf{R}}$ from 2(a) and 2(c), respectively, as specified in step (iii). By symmetry, we have $j = j'$, and construction step (iv) is satisfied as well. Moreover, by the optimality of \mathcal{A} , the three pairs of gadgets $\text{CLG}(X_i), \text{CNG}(Y_j), \text{CNG}(X_i), \text{CLG}(Z)$, and $\text{CLG}(Y_j), \text{CNG}(Z)$ must be aligned in the same way their corresponding strings would be under standard string ED, thus satisfying step (ii). All remaining nodes must be deleted, satisfying step (v). \blacktriangleleft

This allows us to conclude this section with the dynamic lower bound.

► **Theorem 19.** *Unless Conjecture 2 fails, for any $\varepsilon > 0$, there is no combinatorial algorithm for (unweighted) dynamic TED of size $\mathcal{O}(N)$ that satisfies $p(N) + N(u(N) + q(N)) = \mathcal{O}(N^{3-\varepsilon})$, where $p(N), u(N)$, and $q(N)$ denote the preprocessing, update, and query times of the algorithm, respectively. This lower bound holds even when the alphabet size is constant.*

Proof. Let k be a sufficiently large constant to be determined later. We show that an algorithm for unweighted dynamic TED, as described in the statement, can be used to construct a combinatorial algorithm for detecting a $3k$ -clique with running time $\mathcal{O}(n^{3k-\varepsilon'})$ for some $\varepsilon' > 0$.

Given an instance \mathbf{G} of $3k$ -Clique Detection, the algorithm proceeds as follows. We first discover the set of all k -cliques in \mathbf{G} , which we denote with \mathcal{S} . The algorithm performs $\mathcal{O}(n^k)$ rounds, one for each k -clique $Z \in \mathcal{S}$. In round Z , we dynamically maintain the two trees $\mathbf{T}_{\mathcal{X}}(Z)$ and $\mathbf{T}'_{\mathcal{Y}}(Z)$ where $\mathcal{X} = \mathcal{Y} = \mathcal{S}$ as defined in Definition 12. Note that between any two consecutive rounds Z and Z' , the corresponding trees $\mathbf{T}_{\mathcal{X}}(Z), \mathbf{T}_{\mathcal{X}}(Z')$ and $\mathbf{T}'_{\mathcal{Y}}(Z), \mathbf{T}'_{\mathcal{Y}}(Z')$ differ only in the gadgets $\text{CNG}(Z)$, and $\text{CLG}(Z)$, respectively. Both gadgets have length

$\lambda = \mathcal{O}(nk \log n)$, and can be transformed into one another using $\mathcal{O}(\lambda)$ updates. After constructing the instance for each round Z , we query the distance $d_Z := \text{ted}(\mathbf{T}_X(Z), \mathbf{T}_Y(Z))$ and keep track of the minimum value $d := \min_Z d_Z$ across all rounds. If $d - D \leq 3C$ for C, D as defined in Lemma 8 and Theorem 13, respectively, then we return YES. Otherwise, we return NO.

By Theorem 13, we have $d_Z - D \leq 3C$ if and only if there exist k -cliques $X, Y \in \mathcal{S}$ such that $X \cup Y \cup Z$ forms a $3k$ -clique. Since we iterate over all $Z \in \mathcal{S}$, the algorithm checks for the existence of a $3k$ -clique in \mathbf{G} .

The total size of the unweighted instance maintained throughout is $N = \mathcal{O}(n^{k+1} \log n)$, treating k as a constant. This is also the number of updates and queries performed. Therefore, if the total running time can be bounded by $\mathcal{O}(p(N) + N \cdot (u(N) + q(N))) = \mathcal{O}(N^{3(1-\varepsilon)})$, then for a large enough k and small enough $\varepsilon' > 0$, this evaluates to $\mathcal{O}(n^{3(k+1)(1-\varepsilon)} \log^{3(1-\varepsilon)} n) = \mathcal{O}(n^{3k-\varepsilon'})$, contradicting Conjecture 2. \blacktriangleleft

5 Lower Bounds for Dynamic Weighted TED

Our lower bound construction builds upon Bringmann et al. [13], where they reduce Min-Weight 3-Clique to TED. Their construction of the TED instance is in the style of the tree-alignment formulation as in Definition 4. Moreover, they do the following tweaks on the cost function: Denoting the old cost function with δ_{old} , they define

$$\begin{aligned} \delta_{new}(\ell(v), \varepsilon) &:= 0 \\ \delta_{new}(\varepsilon, \ell(v)) &:= 0 \\ \delta_{new}(\ell(v), \ell(v')) &:= \delta_{old}(\ell(v), \ell(v')) - \delta_{old}(\ell(v), \varepsilon) - \delta_{old}(\varepsilon, \ell(v')) \end{aligned}$$

for all v, v' . Note that this definition does not change the optimal alignment. The advantage that the new δ offers is that now we can ignore nodes that are not aligned by \mathcal{A} (deleting them has 0 cost), and $\text{cost}(\mathcal{A})$ consists solely of the aligned pairs.

Given a Min-Weight 3-Clique instance $\mathbf{G} = (V, E, w)$, to make the weight function complete, we additionally define $w(i, j) = \infty$ whenever $(i, j) \notin E$, and $w(i, i) = \infty$ for all nodes i .

► **Definition 20** (TED Instance [13]). *Given an instance $\mathbf{G} = (V, E, w)$ for Min-Weight-3-Clique with w augmented as above, we define an instance $(\mathbf{T}_1, \mathbf{T}_2, \Sigma, \delta)$ for TED as follows.*

- $\Sigma = \{a_1, \dots, a_n, b_1, \dots, b_{n+1}, a'_1, \dots, a'_n, b'_1, \dots, b'_{n+1}, c_1, \dots, c_n, d_1, \dots, d_{n+1}, c'_1, \dots, c'_n, d'_1, \dots, d'_{n+1}\}$.
- \mathbf{T}_1 is obtained from right-attaching the following nodes to $\mathbf{P}(a_1 \cdots a_n b_1 \cdots b_{n+1})$: For every $i \in [n]$, right-attach a leaf node a'_i to a_i , and right-attach b'_i to b_i .
- \mathbf{T}_2 is obtained from left-attaching the following nodes to $\mathbf{P}(c_1 \cdots c_n d_1 \cdots d_{n+1})$: For every $i \in [n]$, left-attach a leaf node c'_i to c_i , and left-attach d'_i to d_i .
- The cost $\delta(u, v)$ of matching nodes u and v is defined as follows, with M being a sufficiently large constant:
 1. $\delta(b'_k, d'_k) = -M^2 - 2M \cdot k$ for every $k \in [1..n]$
 2. $\delta(b_{k+1}, c'_j) = -M^2 + M \cdot k + M \cdot j + w(k, j)$ for every $k \in [1..n], j \in [1..n]$
 3. $\delta(a'_i, d_{k+1}) = -M^2 + M \cdot k + M \cdot i + w(i, k)$ for every $i \in [1..n], k \in [1..n]$
 4. $\delta(a_i, c_j) = -2M + w(i, j) - w(i-1, j-1)$ for every $i \in [2..n], j \in [2..n]$
 5. $\delta(a_i, c_1) = -M(i+1) + w(i, 1)$ for every $i \in [1..n]$
 6. $\delta(a_1, c_j) = -M(j+1) + w(1, j)$ for every $j \in [1..n]$
 7. All other costs are set to ∞ .

Hence, we can view each of \mathbf{T}_1 and \mathbf{T}_2 as having a “top” part (the a -nodes and c -nodes), and a “bottom” part (the b -nodes and d -nodes). Note that this instance has $\mathcal{O}(n)$ nodes and $\mathcal{O}(n)$ alphabet size. Bringmann et al. show that the minimum weight of a triangle can be extracted from the cost of an optimal tree alignment.

► **Definition 21.** *Given a TED instance $(\mathbf{T}_1, \mathbf{T}_2, \Sigma, \delta)$ as in Definition 20, consider a tree-alignment consisting of the following matchings, where $i, j, k \in [1..n]$.*

- (a) A matching (b'_k, d'_k)
- (b) A matching (b_{k+1}, c'_j)
- (c) A matching (a'_i, d_{k+1})
- (d) Consecutive matchings $(a_i, c_j), (a_{i-1}, c_{j-1}), \dots, (a_{i-j+2}, c_2)$ if $i \geq j$;
otherwise $(a_i, c_j), (a_{i-1}, c_{j-1}), \dots, (a_2, c_{j-i+2})$
- (e) A matching (a_{i-j+1}, c_1) if $i \geq j$; otherwise (a_1, c_{j-i+1})

We call a tree-alignment with the above structure $\mathcal{A}_{i,j,k}$. Some simple calculations show that $\mathcal{A}_{i,j,k}$ has cost $-3M^2 + w(i, k) + w(k, j) + w(i, j)$.

► **Theorem 22** ([13]). *Given a TED instance $(\mathbf{T}_1, \mathbf{T}_2, \Sigma, \delta)$ as in Definition 20, any optimal tree-alignment is of the form $\mathcal{A}_{i,j,k}$ for some $i, j, k \in [1..n]$. Therefore, an optimal tree-alignment selects the i, j, k to achieve the minimum cost*

$$-3M^2 + \min_{i,j,k} w(i, k) + w(k, j) + w(i, j)$$

And the minimum weight of a triangle can be extracted from the alignment cost. Moreover, each tree alignment not having the form $\mathcal{A}_{i,j,k}$ has cost at least $-3M^2 + M$, for large enough M . \square

Let us give some intuition on why these matchings encode a triangle. Note that weights (1), (2), and (3) have a “ $-M^2$ ” summand in them, so in order to minimize the total matching cost, the optimal alignment will for sure choose a matching of the form (b'_{k_1}, d'_{k_1}) , a matching of the form (b_{k_2+1}, c'_j) , and a matching of the form (a'_i, d_{k_3+1}) for some k_1, k_2, k_3, i, j . Moreover, there can only be *one* matching of each type in order to not violate the rules of tree alignment. In fact, one can show that $k_1 = k_2 = k_3$ achieves the lowest cost. We can think of matchings (k_2, j) and (k_3, i) as representing edges (k, j) and (k, i) for some k in the original graph. Finally, to bridge (i, j) , the spine-to-spine matchings form a telescoping sum, and the costs eventually cancel out to be $w(i, j)$.

Now, let us introduce our dynamic TED instance. See Figure 2 for an illustration.

► **Definition 23** (dynamic TED Instance). *Given an instance $\mathbf{G} = (V, E, w)$ for Min-Weight 4-Clique, and a node $x \in [1..n]$, build a TED instance $(\mathbf{T}_1^x, \mathbf{T}_2^x, \Sigma^x, \delta^x)$ from $(\mathbf{T}_1, \mathbf{T}_2, \Sigma, \delta)$ from Definition 20 as follows:*

- $\Sigma^x = \Sigma \cup \{a''_1, \dots, a''_n, c''_1, \dots, c''_n, b''_1, \dots, b''_n, \hat{b}_x, \hat{d}_x, \hat{c}_x, \perp\}$.
- For every $i \in [1..n]$, attach a''_i as the child of a'_i , c''_i as the children of c'_i , and b''_i as the first child of b_i (hence the left sibling of b'_i).
- Attach \hat{b}_x as the child of b_{n+1} , \hat{d}_x as the child of d_{n+1} , and add node \perp between c_n and d_1 . Attach \hat{c}_x as a left child of \perp .
- We define $\delta^x(\cdot, \cdot)$ out of δ . That is, we keep costs (1), (2), (3), (4), (5), and (6) from Definition 20, and additionally define the following costs:
 7. $\delta(\hat{b}_x, c''_j) = -M + w(x, j)$ for every $j \in [1..n]$.
 8. $\delta(\hat{d}_x, a''_i) = -M + w(x, i)$ for every $i \in [1..n]$.
 9. $\delta(\hat{c}_x, b''_k) = -M + w(x, k)$ for every $k \in [1..n]$.
 10. All other costs are set to ∞ .

Note that the alphabet size and number of nodes of our dynamic instance is still $\mathcal{O}(n)$. We call the part of the instance from Definition 20 the *underlying tree*.

► **Lemma 24.** *For a large enough M , an optimal tree-alignment in the above dynamic TED instance has the structure of $\mathcal{A}_{i,j,k}$ on the underlying trees from Definition 20, and three additional matchings $(\hat{b}_x, c'_j), (\hat{d}_x, a''_i), (\hat{c}_x, b''_k)$. Thus, an optimal alignment will choose i, j, k that gives the minimum cost*

$$-3M^2 - 3M + \min_{i,j,k} w(i, j) + w(i, k) + w(j, k) + w(x, i) + w(x, j) + w(x, k)$$

Proof. Let OPT denote the optimal alignment cost. First, note that the above alignment structure is a valid tree alignment. Thus OPT is upper-bounded by our desired cost.

It only remains to show that every alignment has at least this cost. For M large enough, in an optimal alignment, either we choose all three matchings $(b'_{k_1}, d'_{k_1}), (b_{k_2+1}, c'_j)$, and (a'_i, d_{k_3+1}) , or the total matching cost up to lower order terms is at least $-2M^2$. We can argue that $k_1 = k_2 = k_3 = k$ for some k . This is because if $k_1 < k_2$, then we can either increase k_1 or decrease k_2 to further lower the cost, and likewise for k_1 and k_3 . Now that we have fixed these three matchings, an optimal alignment will also match \hat{b}_x with c'_j , \hat{d}_x with a''_i , and \hat{c}_x with b''_k , since adding these matchings decreases the total cost by another $3M$, up to lower order terms. Note that by the definition of a tree alignment, these are the *only* possible matchings for \hat{b}_x, \hat{d}_x , and \hat{c}_x . Now, by our definition of δ , the only possible additional alignments are the spine-to-spine matchings. Note that our matchings can be partitioned into two groups: the matchings containing $\hat{d}_x, \hat{c}_x, \hat{d}_x$ and the remaining matchings of the underlying trees from Definition 20. If the spine-to-spine matchings do not have the form $(a_i, c_j), \dots, (a_{i-j+1}, c_1)$, then by Theorem 22, we have that the tree-alignment on the underlying tree has cost at least $-3M^2 + M$. Hence the tree-alignment on the *dynamic*TED instance has cost at least $-3M^2 - 2M$, which is not optimal for M large enough. ◀

This allows us to conclude this (sub)section with the dynamic lower bound.

► **Theorem 25.** *Unless Conjecture 1 fails, for any $\varepsilon > 0$, there is no algorithm for dynamic TED of size $\mathcal{O}(N)$ and alphabet size $\mathcal{O}(N)$ that satisfies $p(N) + N(u(N) + q(N)) = \mathcal{O}(N^{4-\varepsilon})$, where $p(N), u(N)$, and $q(N)$ denote the preprocessing, update, and query times of the algorithm, respectively.*

Proof. Given a Min-Weight 4-Clique instance $\mathbf{G} = (V, E, w)$ with n nodes, we construct our algorithm to perform n rounds, one for each node $x \in V$, while maintaining the dynamic TED instance $(\mathbf{T}_1^x, \mathbf{T}_2^x, \Sigma^x, \delta^x)$ throughout the rounds. Note that between consecutive rounds x and y , we only change the label of $\hat{b}_x, \hat{c}_x, \hat{d}_x$ into $\hat{b}_y, \hat{c}_y, \hat{d}_y$. In each round x , we query the tree edit distance between \mathbf{T}_1^x and \mathbf{T}_2^x , and keep track of the minimum value across all rounds (adding an offset of $3M^3 + 3M$ to each).

By Lemma 24, the result of round x allows us to identify nodes i, j, k such that i, j, k, x form a min-weight 4-clique containing x . After we iterate over all $x \in V$, we indeed return the min-weight 4-clique. The size and alphabet size of our dynamic TED instance is $\mathcal{O}(n)$, and there are $\mathcal{O}(n)$ updates and queries. Thus if there exists an algorithm for dynamic TED running in overall running time $\mathcal{O}(p(N) + N(u(N) + q(N))) = \mathcal{O}(N^{4-\varepsilon})$, we would solve Min-Weight 4-Clique in time $\mathcal{O}(n^{4-\varepsilon})$. So for a Min-Weight 4-Clique instance $\mathbf{G} = (V, E, w)$ satisfying the weight condition as in Conjecture 1, it would contradict Conjecture 1. ◀

Lower Bounds for Incremental/Decremental Dynamic TED. An easy modification of our construction allows us to also establish conditional lower bounds for incremental (resp., decremental) dynamic TED, where the updates are restricted to node insertions (resp., deletions). Let us briefly sketch the reduction to incremental dynamic TED here. The lower bound for decremental dynamic TED can be shown similarly.

First we assign some order to the n nodes x_1, \dots, x_n in the Min-Weight 4-Clique instance. Then we define δ_{incr} to contain all costs in δ from Definition 20, and additionally define

7. $\delta_{incr}(\hat{b}_{x_\ell}, c'_j) = -M \cdot \ell + w(x_\ell, j)$ for every $\ell, j \in [1..n]$
8. $\delta_{incr}(\hat{d}_{x_\ell}, a'_i) = -M \cdot \ell + w(x_\ell, i)$ for every $\ell, i \in [1..n]$
9. $\delta_{incr}(\hat{c}_{x_\ell}, b'_k) = -M \cdot \ell + w(x_\ell, k)$ for every $\ell, k \in [1..n]$
10. All other costs are set to ∞

Starting from an empty graph, we make $\mathcal{O}(n)$ insertions to arrive at the dynamic TED instance $\mathbf{T}_1^{x_1}, \mathbf{T}_2^{x_2}$ as in Definition 23. We make a query to our black-box incremental dynamic TED algorithm to get a min-weight 4-clique containing x_1 . This concludes the first round. For round ℓ , $2 \leq \ell \leq n$, we attach \hat{b}_{x_ℓ} as the child of $\hat{b}_{x_{\ell-1}}$, \hat{c}_{x_ℓ} as the child of $\hat{c}_{x_{\ell-1}}$, and \hat{d}_{x_ℓ} as the child of $\hat{d}_{x_{\ell-1}}$, and query for a min-weight 4-clique containing x_ℓ . Taking the minimum weight over all n rounds gives us a desired min-weight 4-clique.

To see the correctness, note that once $\hat{b}_{x_\ell}, \hat{c}_{x_\ell}, \hat{d}_{x_\ell}$ are inserted, they gain priority over all previously inserted $\hat{b}, \hat{c}, \hat{d}$ -nodes to be matched. So, essentially, we are considering each node $x \in V$ independently as in the previous reduction to dynamic TED. As there are $\mathcal{O}(n)$ updates and queries in this reduction, we can conclude that unless Conjecture 1 fails, for any $\varepsilon > 0$, there is no algorithm for incremental dynamic TED of size $\mathcal{O}(N)$ and alphabet size $\mathcal{O}(N)$ that satisfies $p(N) + N(u(N) + q(N)) = \mathcal{O}(N^{4-\varepsilon})$.

6 Lower Bounds for Dynamic RNA Folding and Dyck Edit Distance

We build upon the static lower bound instance from [2], summarized as follows.


► **Lemma 26** ([2]). *Let $\mathbf{G} = (V, E)$ be a graph on $n = |V|$ nodes, and let $k \in \mathbb{Z}_+$. Then, there exist embeddings $\text{CG}_A, \text{CG}_B, \text{CG}_C : V^k \rightarrow \{\Sigma \cup \Sigma'\}^{\leq \ell}$, with $\ell = \mathcal{O}(n)$ over a constant-size alphabet Σ and a weight function $w : \Sigma \rightarrow [M]$ with $M = \mathcal{O}(k^4 n \log n)$, such that for sets $\mathcal{X}, \mathcal{Y}, \mathcal{Z} \subseteq V^k$ of size N , the string*

$$S := \#_A^{2N} \left(\bigcirc_{X \in \mathcal{X}} \#'_A \text{CG}_A(X) \#'_A \right) \#_A^{2N} \circ \\ \#_B^{2N} \left(\bigcirc_{Y \in \mathcal{Y}} \#'_B \text{CG}_B(Y) \#'_B \right) \#_B^{2N} \circ \\ \#_C^{2N} \left(\bigcirc_{Z \in \mathcal{Z}} \#'_C \text{CG}_C(Z) \#'_C \right) \#_C^{2N}$$


satisfies the two following.

- (i) *There exists a positive constant C such that $\text{score}_w(S) = C$ if there are $X \in \mathcal{X}$, $Y \in \mathcal{Y}$, and $Z \in \mathcal{Z}$ such that $X \cup Y \cup Z$ forms a $3k$ -clique in \mathbf{G} . Otherwise, $\text{score}_w(S) < C$.*
- (ii) *For every optimal folding F of S , there exists a set $X \in \mathcal{X}$ such that if $\text{CG}_A(X)$ appears within the substring $S[x..x+\lambda)$, then for all pairs $(i, j) \in F$ with $\{S[i], S[j]\} = \{\#_A, \#'_A\}$, it holds that either $i < j < x$ or $x + \lambda \leq i < j$. Similarly, there exist sets $Y \in \mathcal{Y}$ and $Z \in \mathcal{Z}$ such that the same condition holds with respect to the delimiters $\#_B, \#'_B$ and $\#_C, \#'_C$, respectively. If $\text{score}_w(S) = C$, then this condition holds for all X, Y, Z forming a $3k$ -clique. \lrcorner*

We proceed to modify this instance to obtain a dynamic lower bound.



(a) The TED Instance from [13]. The numeric labels indicate the ranges of a certain type of matchings.



(b) Extension of the TED instance.

■ **Figure 2** The figure illustrates how to extend the TED instance from [13] to obtain a dynamic bound. Nodes newly introduced in Figure 2b are shown in full color, while those from Figure 2a are rendered translucent.

► **Lemma 27.** Consider the string S and the weight function w from Lemma 26:

■ For $W \in V^k$, modify the string S to get the string

$$S_W := \text{CLG}_A(W) \#_A^{2N} (\bigcirc_{X \in \mathcal{X}} \#'_A p(\text{CNG}_A(X))^R \text{CG}_A(X) \#'_A) \#_A^{2N} \circ \\ \text{CLG}_B(W) \#_B^{2N} (\bigcirc_{Y \in \mathcal{Y}} \#'_B p(\text{CNG}_B(Y))^R \text{CG}_B(Y) \#'_B) \#_B^{2N} \circ \\ \text{CLG}_C(W) \#_C^{2N} (\bigcirc_{Z \in \mathcal{Z}} \#'_C p(\text{CNG}_C(Z))^R \text{CG}_C(Z) \#'_C) \#_C^{2N},$$

where CLG_A and CNG_A are the strings CLG and CNG from Lemma 7, defined over a new alphabet Σ_A disjoint from Σ . Similarly, $\text{CLG}_B, \text{CNG}_B$ and $\text{CLG}_C, \text{CNG}_C$ are defined over new alphabets Σ_B and Σ_C , respectively, each disjoint from one another and from $\Sigma \cup \Sigma_A$.

■ Modify the weight function w from Lemma 26 to get $w' : \Sigma \cup \Sigma_A \cup \Sigma_B \cup \Sigma_C \rightarrow [1..4\lambda M]$ defined as

$$w'(\sigma) = \begin{cases} 4\lambda w(\sigma) & \sigma \in \Sigma \\ 1 & \sigma \in \Sigma_A \cup \Sigma_B \cup \Sigma_C \end{cases}$$

where λ is the length of CLG, CNG .

Then, there is a constant D such that $\text{score}_{w'}(S_W) = D$ if there are $X \in \mathcal{X}, Y \in \mathcal{Y}, Z \in \mathcal{Z}$ such that $X \cup Y \cup Z \cup W$ is a $4k$ -clique, and $\text{score}_{w'}(S_W) < D$, otherwise.

Proof. We first present Claim 28, and then we argue how Lemma 27 follows from it.

▷ **Claim 28.** There are $X \in \mathcal{X}, Y \in \mathcal{Y}, Z \in \mathcal{Z}$ for which the following holds:

$$\text{score}_{w'}(S_W) = 4\lambda \cdot \text{score}_w(S) + \sum_{(t,T) \in \{(A,X), (B,Y), (C,Z)\}} \text{score}(\text{CLG}_t(W) \circ p(\text{CNG}_t(T))^R). \quad (1)$$

Moreover, if $\text{score}_w(S) = C$ (where C is the constant from Lemma 26), the sets X, Y, Z are such that $X \cup Y \cup Z$ is a $3k$ -clique and such to maximize the summation of scores in (1).

Proof. Let F_W be an optimal folding of S_W with score $s_W = \sum_{(i,j) \in F_W} w(S_W[i])$, and let $s = \text{score}_w(S)$. Note that $s > 0$ if $N > 0$. Since S is a subsequence of S_W , we have $s_W \geq 4\lambda s$. Moreover, as the alphabet Σ is disjoint from Σ_A, Σ_B , and Σ_C , it follows that for any pair $(i, j) \in F_W$, either both positions i and j lie in S , or neither does.

We can therefore define a folding F of S by restricting F_W to those pairs whose endpoints are both in S . We claim that this folding F has score s , and must therefore be optimal. Assume for contradiction that this is not the case. Then the score s_W could be at most $4\lambda(s-1) + 3\lambda$, as the additional parts of S_W can contribute at most 3λ to the total score. This implies $s_W \leq 4\lambda(s-1) + 3\lambda = 4\lambda s - \lambda$, contradicting the earlier bound $s_W \geq 4\lambda s$.

Since F is optimal, it satisfies the property stated in Lemma 26(ii). Importantly, this property also holds for F_W : there is $X \in \mathcal{X}$ such that if $\text{CG}_A(X)$ appears within the substring $S_W[x..x+\lambda)$, then for any pair $(i, j) \in F_W$ with $\{S_W[i], S_W[j]\} = \{\#_A, \#'_A\}$, we have either $i < j < x$ or $x + \lambda \leq i < j$. The same holds for some $Y \in \mathcal{Y}$ and $Z \in \mathcal{Z}$.

At this point, it is not difficult to see that any optimal F would match the remaining characters of S_W (those not included in S) by including the matchings between the three strings $\text{CLG}_A(W) \circ p(\text{CNG}_A(X))^R$, and $\text{CLG}_B(W) \circ p(\text{CNG}_B(Y))^R$, and $\text{CLG}_C(W) \circ p(\text{CNG}_C(Z))^R$.

This completes the first part of the statement. For the second part, observe that if $\text{score}_w(S) = C$, then Lemma 26(ii) holds for all triples X, Y, Z that form a $3k$ -clique. Therefore, such triples must maximize the second part of the sum (1), otherwise, we would contradict the optimality of F_W . \triangleleft

Next, set $D = 4\lambda C + 3C'$, where C, C' are the constants from Lemma 26 and Lemma 7, respectively.

Now, suppose there exist $X \in \mathcal{X}$, $Y \in \mathcal{Y}$, and $Z \in \mathcal{Z}$ such that $X \cup Y \cup Z \cup W$ forms a $4k$ -clique. Since $X \cup Y \cup Z$ already forms a $3k$ -clique, Claim 28 implies that the first part of the sum (1) contributes $4\lambda C$. Additionally, by Lemma 7, the second part must contribute $3C'$. Therefore, we have $\text{score}_{w'}(S_W) = D$.

On the other hand, suppose $\text{score}_{w'}(S_W) = D$. Then we must have $\text{score}_w(S) = C$; otherwise, we would get $\text{score}_{w'}(S_W) \leq 4\lambda(C - 1) + 3\lambda = 4\lambda C - \lambda < D$, which is a contradiction. This implies that the sets $X \cup Y \cup Z$ form a $3k$ -clique. Consequently, the second part of the sum in (1) contributes $D - 4\lambda C = 3C'$, and by Lemma 7, this shows that $X \cup Y \cup Z \cup W$ forms a $4k$ -clique. ◀

This allows us to conclude this section with the dynamic lower bound.

► **Theorem 29.** *Unless Conjecture 2 fails, for any $\varepsilon > 0$, there is no combinatorial algorithms for dynamic RNA Folding or dynamic Dyck Edit Distance of size $\mathcal{O}(N)$ that satisfy $p(N) + N(u(N) + q(N)) = \mathcal{O}(N^{4-\varepsilon})$, where $p(N)$, $u(N)$, and $q(N)$ are the preprocessing, update, and query times of the algorithm, respectively. This lower bound holds even when the alphabet size is constant.*

Proof. The proof follows the same pattern as Theorem 19 and Theorem 25. Refer to the full version [40] for a detailed proof. ◀

References

- 1 Amir Abboud, Arturs Backurs, and Virginia Vassilevska Williams. Tight hardness results for LCS and other sequence similarity measures. In *Proceedings of the IEEE 56th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 59–78, 2015. doi:10.1109/FOCS.2015.14.
- 2 Amir Abboud, Arturs Backurs, and Virginia Vassilevska Williams. If the current clique algorithms are optimal, so is Valiant’s parser. *SIAM Journal on Computing*, 47(6):2527–2555, 2018. doi:10.1137/16M1061771.
- 3 Amir Abboud, Nick Fischer, Zander Kelley, Shachar Lovett, and Raghu Meka. New graph decompositions and combinatorial boolean matrix multiplication algorithms. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, STOC 2024, pages 935–943, New York, NY, USA, 2024. Association for Computing Machinery. doi:10.1145/3618260.3649696.
- 4 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. In *Proceedings of the 48th Annual ACM Symposium on Theory of Computing (STOC)*, pages 375–388, 2016. doi:10.1145/2897518.2897653.
- 5 Amir Abboud and Virginia Vassilevska Williams. Popular conjectures imply strong lower bounds for dynamic problems. In *Proceedings of the 55th Annual Symposium on Foundations of Computer Science (FOCS)*, 2014.
- 6 Ittai Abraham, Shiri Chechik, and Sebastian Krinninger. Fully dynamic all-pairs shortest paths with worst-case update-time revisited. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA ’17, pages 440–452, USA, 2017. Society for Industrial and Applied Mathematics. doi:10.1137/1.9781611974782.28.
- 7 Alfred V. Aho and Thomas G Peterson. A minimum distance error-correcting parser for context-free languages. *SIAM J. Comput.*, 1(4):305–312, 1972. doi:10.1137/0201022.
- 8 Josh Alman, Ran Duan, Virginia Vassilevska Williams, Yinzhan Xu, Zixuan Xu, and Renfei Zhou. More asymmetry yields faster matrix multiplication. In *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, page to appear, 2025.

- 9 Arturs Backurs and Piotr Indyk. Edit distance cannot be computed in strongly subquadratic time (unless SETH is false). In *Proceedings of the 47th Annual ACM Symposium on Theory of Computing (STOC)*, pages 51–58, 2015. doi:10.1145/2746539.2746612.
- 10 Grey Ballard, James Demmel, Olga Holtz, and Oded Schwartz. Graph expansion and communication costs of fast matrix multiplication. *J. ACM*, 59(6), January 2013. doi:10.1145/2395116.2395121.
- 11 John Bellando and Ravi Kothari. Region-based modeling and tree edit distance as a basis for gesture recognition. In *Proceedings of the 10th International Conference on Image Analysis and Processing (ICIAP)*, pages 698–703, 1999. doi:10.1109/ICIAP.1999.797676.
- 12 Mahdi Boroujeni, Mohammad Ghodsi, MohammadTaghi Hajiaghayi, and Saeed Seddighin. $1 + \varepsilon$ approximation of tree edit distance in quadratic time. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 709–720, 2019. doi:10.1145/3313276.3316388.
- 13 Karl Bringmann, Paweł Gawrychowski, Shay Mozes, and Oren Weimann. Tree edit distance cannot be computed in strongly subcubic time (unless APSP can). *ACM Trans. Algorithms*, 16(4), July 2020. doi:10.1145/3381878.
- 14 Karl Bringmann, Fabrizio Grandoni, Barna Saha, and Virginia Vassilevska Williams. Truly subcubic algorithms for language edit distance and RNA folding via fast bounded-difference min-plus product. *SIAM J. Comput.*, 48(2):481–512, 2019. doi:10.1137/17M112720X.
- 15 Karl Bringmann and Marvin Kunnemann. Quadratic Conditional Lower Bounds for String Problems and Dynamic Time Warping. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 79–97, Los Alamitos, CA, USA, October 2015. IEEE Computer Society. doi:10.1109/FOCS.2015.15.
- 16 Peter Buneman, Martin Grohe, and Christoph Koch. Path queries on compressed XML. In *Proceedings of the 29th International Conference on Very Large Data Bases (VLDB)*, pages 141–152, 2003. doi:10.1016/B978-012722442-8/50021-5.
- 17 Alejandro Cassis, Tomasz Kociumaka, and Philip Wellnitz. Optimal algorithms for bounded weighted edit distance. In *Proceedings of the 2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 2177–2187, 2023. doi:10.1109/FOCS57990.2023.00135.
- 18 Yi-Jun Chang. Hardness of RNA Folding Problem With Four Symbols. In Roberto Grossi and Moshe Lewenstein, editors, *27th Annual Symposium on Combinatorial Pattern Matching (CPM 2016)*, volume 54 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 13:1–13:12, Dagstuhl, Germany, 2016. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.CPM.2016.13.
- 19 Panagiotis Charalampopoulos, Tomasz Kociumaka, and Shay Mozes. Dynamic String Alignment. In *Proceedings of the 31st Annual Symposium on Combinatorial Pattern Matching (CPM)*, pages 9:1–9:13, 2020. doi:10.4230/LIPIcs.CPM.2020.9.
- 20 Sudarshan S. Chawathe. Comparing hierarchical data in external memory. In *Proceedings of the 25th International Conference on Very Large Data Bases (VLDB)*, pages 90–101, 1999. URL: <http://www.vldb.org/conf/1999/P8.pdf>.
- 21 Shiri Chechik and Tianyi Zhang. Faster deterministic worst-case fully dynamic all-pairs shortest paths via decremental hop-restricted shortest paths. In Nikhil Bansal and Viswanath Nagarajan, editors, *Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms, SODA 2023, Florence, Italy, January 22-25, 2023*, pages 87–99. SIAM, 2023. doi:10.1137/1.9781611977554.CH4.
- 22 Shucheng Chi, Ran Duan, and Tianle Xie. Faster algorithms for bounded-difference min-plus product. In Joseph (Seffi) Naor and Niv Buchbinder, editors, *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, Virtual Conference / Alexandria, VA, USA, January 9 - 12, 2022*, pages 1435–1447. SIAM, 2022. doi:10.1137/1.9781611977073.60.

- 23 Shucheng Chi, Ran Duan, Tianle Xie, and Tianyi Zhang. Faster min-plus product for monotone instances. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 1529–1542, 2022. doi:10.1145/3519935.3520057.
- 24 Debarati Das, Jacob Gilbert, MohammadTaghi Hajiaghayi, Tomasz Kociumaka, and Barna Saha. Weighted edit distance computation: Strings, trees, and Dyck. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, STOC 2023, pages 377–390, New York, NY, USA, 2023. Association for Computing Machinery. doi:10.1145/3564246.3585178.
- 25 Debarati Das, Jacob Gilbert, MohammadTaghi Hajiaghayi, Tomasz Kociumaka, Barna Saha, and Hamed Saleh. $\tilde{O}(n+\text{poly}(k))$ -time algorithm for bounded tree edit distance. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 686–697, 2022. doi:10.1109/FOCS54457.2022.00071.
- 26 Debarati Das, Jacob Gilbert, Tomasz Kociumaka, MohammadTaghi Hajiaghayi, and Barna Saha. Dynamic Dyck and tree edit distance: Decompositions and reductions to string edit distance. In *Proceedings of the 66th IEEE Symposium on Foundations of Computer Science (FOCS) 2025*, 2025. to appear.
- 27 Debarati Das, Tomasz Kociumaka, and Barna Saha. Improved Approximation Algorithms for Dyck Edit Distance and RNA Folding. In Mikołaj Bojańczyk, Emanuela Merelli, and David P. Woodruff, editors, *49th International Colloquium on Automata, Languages, and Programming (ICALP 2022)*, volume 229 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 49:1–49:20, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ICALP.2022.49.
- 28 Erik D. Demaine, Shay Mozes, Benjamin Rossman, and Oren Weimann. An optimal decomposition algorithm for tree edit distance. *ACM Trans. Algorithms*, 6(1), December 2010. doi:10.1145/1644015.1644017.
- 29 Camil Demetrescu and Giuseppe F. Italiano. Improved bounds and new trade-offs for dynamic all pairs shortest paths. In Peter Widmayer, Stephan Eidenbenz, Francisco Triguero, Rafael Morales, Ricardo Conejo, and Matthew Hennessy, editors, *Automata, Languages and Programming*, pages 633–643, Berlin, Heidelberg, 2002. Springer Berlin Heidelberg. doi:10.1007/3-540-45465-9_54.
- 30 Camil Demetrescu and Giuseppe F. Italiano. A new approach to dynamic all pairs shortest paths. In *Proceedings of the Thirty-Fifth Annual ACM Symposium on Theory of Computing*, STOC '03, pages 159–166, New York, NY, USA, 2003. Association for Computing Machinery. doi:10.1145/780542.780567.
- 31 Bartłomiej Dudek and Paweł Gawrychowski. Online Context-Free Recognition in OM_v Time. In Shunsuke Inenaga and Simon J. Puglisi, editors, *35th Annual Symposium on Combinatorial Pattern Matching (CPM 2024)*, volume 296 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 13:1–13:9, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.CPM.2024.13.
- 32 Anita Dürr. Improved bounds for rectangular monotone min-plus product and applications. *Inf. Process. Lett.*, 181(C), March 2023. doi:10.1016/j.ipl.2023.106358.
- 33 Paolo Ferragina, Fabrizio Luccio, Giovanni Manzini, and S. Muthukrishnan. Compressing and indexing labeled trees, with applications. *J. ACM*, 57(1):4:1–4:33, 2009. doi:10.1145/1613676.1613680.
- 34 Egor Gorbachev and Tomasz Kociumaka. Bounded edit distance: Optimal static and dynamic algorithms for small integer weights, 2025. doi:10.48550/arXiv.2404.06401.
- 35 Dan Gusfield. *Algorithms on Strings, Trees, and Sequences: Computer Science and Computational Biology*. Cambridge University Press, 1997. doi:10.1017/CB09780511574931.
- 36 Maximilian Probst Gutenberg and Christian Wulff-Nilsen. Fully-dynamic all-pairs shortest paths: Improved worst-case time and space bounds. In Shuchi Chawla, editor, *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020*, pages 2562–2574. SIAM, 2020. doi:10.1137/1.9781611975994.156.

- 37 Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. Dynamic approximate all-pairs shortest paths: Breaking the $o(mn)$ barrier and derandomization. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, pages 538–547, 2013. doi:10.1109/FOCS.2013.64.
- 38 Monika Henzinger, Sebastian Krinninger, Danupon Nanongkai, and Thatchaphol Saranurak. Unifying and strengthening hardness for dynamic problems via the online matrix-vector multiplication conjecture. In *Proceedings of the 47th Annual ACM on Symposium on Theory of Computing (STOC)*, 2015.
- 39 Matthias Höchsmann, Thomas Töller, Robert Giegerich, and Stefan Kurtz. Local similarity in RNA secondary structures. In *Proceedings of 2nd IEEE Computer Society Bioinformatics Conference (CSB)*, pages 159–168, 2003. doi:10.1109/CSB.2003.1227315.
- 40 Bingbing Hu, Jakob Nogler, and Barna Saha. Hardness of dynamic tree edit distance and friends, 2025. arXiv:2511.09842.
- 41 Bingbing Hu and Adam Polak. Non-boolean OMv: One more reason to believe lower bounds for dynamic problems. In Anne Benoit, Haim Kaplan, Sebastian Wild, and Grzegorz Herman, editors, *33rd Annual European Symposium on Algorithms, ESA 2025, September 15-17, 2025, Warsaw, Poland*, volume 351 of *LIPIcs*, pages 54:1–54:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025. doi:10.4230/LIPIcs.ESA.2025.54.
- 42 Yusuke Ishida, Shunsuke Inenaga, Ayumi Shinohara, and Masayuki Takeda. Fully incremental lcs computation. In *Proceedings of the 15th International Conference on Fundamentals of Computation Theory, FCT’05*, pages 563–574, Berlin, Heidelberg, 2005. Springer-Verlag. doi:10.1007/11537311_49.
- 43 Sung-Ryul Kim and Kunsoo Park. A dynamic edit distance table. *Journal of Discrete Algorithms*, 2(2):303–312, 2004. Combinatorial Pattern Matching. doi:10.1016/S1570-8667(03)00082-0.
- 44 V. King. Fully dynamic algorithms for maintaining all-pairs shortest paths and transitive closure in digraphs. In *40th Annual Symposium on Foundations of Computer Science (Cat. No.99CB37039)*, pages 81–89, 1999. doi:10.1109/SFFCS.1999.814580.
- 45 Valerie King and Mikkel Thorup. A space saving trick for directed dynamic transitive closure and shortest path algorithms. In Jie Wang, editor, *Computing and Combinatorics*, pages 268–277, Berlin, Heidelberg, 2001. Springer Berlin Heidelberg. doi:10.1007/3-540-44679-6_30.
- 46 Philip N. Klein. Computing the edit-distance between unrooted ordered trees. In *Proceedings of the 6th Annual European Symposium on Algorithms (ESA)*, pages 91–102, 1998. doi:10.1007/3-540-68530-8_8.
- 47 Philip N. Klein, Thomas B. Sebastian, and Benjamin B. Kimia. Shape matching using edit-distance: an implementation. In *Proceedings of the 12th Annual Symposium on Discrete Algorithms (SODA)*, pages 781–790, 2001. URL: <http://dl.acm.org/citation.cfm?id=365411.365779>.
- 48 Philip N. Klein, Srikanta Tirthapura, Daniel Sharvit, and Benjamin B. Kimia. A tree-edit-distance algorithm for comparing simple, closed shapes. In *Proceedings of the 11th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 696–704, 2000. URL: <http://dl.acm.org/citation.cfm?id=338219.338628>.
- 49 Tomasz Kociumaka, Anish Mukherjee, and Barna Saha. Approximating Edit Distance in the Fully Dynamic Model. In *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1628–1638, Los Alamitos, CA, USA, November 2023. IEEE Computer Society. doi:10.1109/FOCS57990.2023.00098.
- 50 Tomasz Kociumaka and Ali Shahali. Faster algorithm for bounded tree edit distance in the low-distance regime, 2025. To appear in the *European Symposium on Algorithms (ESA 2025)*. doi:10.48550/arXiv.2507.02701.
- 51 Michal Koucký and Michael E. Saks. Simple, deterministic, fast (but weak) approximations to edit distance and dyck edit distance. In Nikhil Bansal and Viswanath Nagarajan, editors, *Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms, SODA 2023, Florence, Italy, January 22-25, 2023*, pages 5203–5219. SIAM, 2023. doi:10.1137/1.9781611977554.CH188.

- 52 Gad M. Landau, Eugene W. Myers, and Jeanette P. Schmidt. Incremental string comparison. *SIAM J. Comput.*, 27(2):557–582, 1998. doi:10.1137/S0097539794264810.
- 53 Kasper Green Larsen and Ryan Williams. Faster online matrix-vector multiplication. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '17, pages 2182–2189, USA, 2017. Society for Industrial and Applied Mathematics. doi:10.1137/1.9781611974782.142.
- 54 Xiao Mao. Breaking the cubic barrier for (unweighted) tree edit distance. In *Proceedings of the 2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 792–803, 2022. doi:10.1109/FOCS52979.2021.00082.
- 55 Xiao Mao. Fully dynamic all-pairs shortest paths: Likely optimal worst-case update time. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, STOC 2024, pages 1141–1152, New York, NY, USA, 2024. Association for Computing Machinery. doi:10.1145/3618260.3649695.
- 56 Saul B. Needleman and Christian D. Wunsch. A general method applicable to the search for similarities in the amino acid sequence of two proteins. *Journal of Molecular Biology*, 48(3):443–453, March 1970. doi:10.1016/0022-2836(70)90057-4.
- 57 Jakob Nogler, Adam Polak, Barna Saha, Virginia Vassilevska Williams, Yinzhao Xu, and Christopher Ye. Faster weighted and unweighted tree edit distance and APSP equivalence. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, STOC '25, pages 2167–2178, New York, NY, USA, 2025. Association for Computing Machinery. doi:10.1145/3717823.3718116.
- 58 Ruth Nussinov and Ann B. Jacobson. Fast algorithm for predicting the secondary structure of single-stranded RNA. *Proc. Natl. Acad. Sci. U.S.A.*, 77(11):6309–6313, 1980.
- 59 Liam Roditty and Roei Tov. Approximating the girth. *ACM Trans. Algorithms*, 9(2), March 2013. doi:10.1145/2438645.2438647.
- 60 Barna Saha. The Dyck language edit distance problem in near-linear time. In *Proceedings of the 2014 IEEE 55th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 611–620, 2014. doi:10.1109/FOCS.2014.71.
- 61 Barna Saha. Language edit distance and maximum likelihood parsing of stochastic grammars: Faster algorithms and connection to fundamental graph problems. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pages 118–135, 2015. doi:10.1109/FOCS.2015.17.
- 62 Thomas B. Sebastian, Philip N. Klein, and Benjamin B. Kimia. Recognition of shapes by editing their shock graphs. *IEEE Trans. Pattern Anal. Mach. Intell.*, 26(5):550–571, 2004. doi:10.1109/TPAMI.2004.1273924.
- 63 Masoud Seddighin and Saeed Seddighin. $3 + \epsilon$ approximation of tree edit distance in truly subquadratic time. In *Proceedings of the 13th Innovations in Theoretical Computer Science Conference (ITCS)*, pages 115:1–115:22, 2022. doi:10.4230/LIPIcs.ITCS.2022.115.
- 64 Peter H. Sellers. On the theory and computation of evolutionary distances. *SIAM Journal on Applied Mathematics*, 26(4):787–793, 1974. doi:10.1137/0126070.
- 65 Bruce A. Shapiro and Kaizhong Zhang. Comparing multiple RNA secondary structures using tree comparisons. *Bioinformatics*, 6(4):309–318, October 1990. doi:10.1093/bioinformatics/6.4.309.
- 66 Volker Strassen. Gaussian elimination is not optimal. *Numerische Mathematik*, 13(4):354–356, 1969. doi:10.1007/BF02165411.
- 67 Kuo-Chung Tai. The tree-to-tree correction problem. *J. ACM*, 26(3):42–433, July 1979. doi:10.1145/322139.322143.
- 68 Mikkel Thorup. Fully-dynamic all-pairs shortest paths: Faster and allowing negative cycles. In Torben Hagerup and Jyrki Katajainen, editors, *Algorithm Theory - SWAT 2004*, pages 384–396, Berlin, Heidelberg, 2004. Springer Berlin Heidelberg. doi:10.1007/978-3-540-27810-8_33.

- 69 Mikkel Thorup. Worst-case update times for fully-dynamic all-pairs shortest paths. In *Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing, STOC '05*, pages 112–119, New York, NY, USA, 2005. Association for Computing Machinery. doi:10.1145/1060590.1060607.
- 70 Alexander Tiskin. Semi-local string comparison: Algorithmic techniques and applications. *Mathematics in Computer Science*, 1(4):571–603, 2008. doi:10.1007/s11786-007-0033-3.
- 71 Virginia Vassilevska Williams. *On some fine-grained questions in algorithms and complexity*, pages 3447–3487. World Scientific, 2018. doi:10.1142/9789813272880_0188.
- 72 Virginia Vassilevska Williams and R. Ryan Williams. Subcubic equivalences between path, matrix, and triangle problems. *J. ACM*, 65(5):1–38, 2018. doi:10.1145/3186893.
- 73 Virginia Vassilevska Williams and Yinzhan Xu. Truly subcubic min-plus product for less structured matrices, with applications. In *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 12–29, 2020. doi:10.1137/1.9781611975994.2.
- 74 Balaji Venkatachalam, Dan Gusfield, and Yelena Frid. Faster algorithms for RNA-folding using the four-russians method. *Algorithms for molecular biology : AMB*, 9:5, March 2014. doi:10.1186/1748-7188-9-5.
- 75 T. K. Vintsyuk. Speech discrimination by dynamic programming. *Cybernetics*, 4(1):52–57, 1968. doi:10.1007/BF01074755.
- 76 Robert A. Wagner and Michael J. Fischer. The string-to-string correction problem. *J. ACM*, 21(1):168–173, January 1974. doi:10.1145/321796.321811.
- 77 Michael S. Waterman. *Introduction to computational biology: maps, sequences and genomes*. CRC Press, 1995.
- 78 Kaizhong Zhang and Dennis Shasha. Simple fast algorithms for the editing distance between trees and related problems. *SIAM J. Comput.*, 18(6):1245–1262, 1989. doi:10.1137/0218082.