

Tree Edit Distance Cannot be Computed in Strongly Subcubic Time (unless APSP can)

KARL BRINGMANN, Max Planck Institute for Informatics, Saarland Informatics Campus

PAWEŁ GAWRYCHOWSKI, University of Wrocław

SHAY MOZES, IDC Herzliya

OREN WEIMANN, University of Haifa

The edit distance between two rooted ordered trees with n nodes labeled from an alphabet Σ is the minimum cost of transforming one tree into the other by a sequence of elementary operations consisting of deleting and relabeling existing nodes, as well as inserting new nodes. Tree edit distance is a well known generalization of string edit distance. The fastest known algorithm for tree edit distance runs in cubic $O(n^3)$ time and is based on a similar dynamic programming solution as string edit distance. In this paper we show that a truly subcubic $O(n^{3-\epsilon})$ time algorithm for tree edit distance is unlikely: For $|\Sigma| = \Omega(n)$, a truly subcubic algorithm for tree edit distance implies a truly subcubic algorithm for the all pairs shortest paths problem. For $|\Sigma| = O(1)$, a truly subcubic algorithm for tree edit distance implies an $O(n^{k-\epsilon})$ algorithm for finding a maximum weight k -clique.

Thus, while in terms of upper bounds string edit distance and tree edit distance are highly related, in terms of lower bounds string edit distance exhibits the hardness of the strong exponential time hypothesis [Backurs, Indyk STOC'15] whereas tree edit distance exhibits the hardness of all pairs shortest paths. Our result provides a matching conditional lower bound for one of the last remaining classic dynamic programming problems.

CCS Concepts: • **Theory of computation** → **Pattern matching**;

Additional Key Words and Phrases: Conditional lower bound, String edit distance, Tree edit distance

ACM Reference Format:

Karl Bringmann, Paweł Gawrychowski, Shay Mozes, and Oren Weimann. 2018. Tree Edit Distance Cannot be Computed in Strongly Subcubic Time (unless APSP can). *ACM Trans. Algor.* 0, 0, Article 0 (2018), 22 pages. <https://doi.org/0000001.0000001>

1 INTRODUCTION

Tree edit distance is the most common similarity measure between labelled trees. Algorithms for computing the tree edit distance are being used in a multitude of applications in various domains including computational biology [24, 37, 52, 58], structured text and data processing (e.g., XML) [30, 31, 36], programming languages and compilation [38], computer vision [22, 43], character recognition [49], automatic grading [14], answer extraction [65], and the list goes on and on.

Let F and G be two rooted trees with a left-to-right order among siblings and where each vertex is assigned a label from an alphabet Σ . The edit distance between F and G is the minimum cost of transforming F into G by a sequence of elementary operations (at most one operation per node):

Authors' addresses: Karl Bringmann, Max Planck Institute for Informatics, Saarland Informatics Campus, kbringma@mpi-inf.mpg.de; Paweł Gawrychowski, University of Wrocław, gawry@cs.uni.wroc.pl; Shay Mozes, IDC Herzliya, smozes@idc.ac.il; Oren Weimann, University of Haifa, oren@cs.haifa.ac.il.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 0 Copyright held by the owner/author(s). Publication rights licensed to ACM.

1549-6325/2018/0-ART0 \$15.00
<https://doi.org/0000001.0000001>

changing the label of a node v , deleting a node v and setting the children of v as the children of v 's parent (in the place of v in the left-to-right order), and inserting a node v (the complement of delete¹). See Figure 1. The cost of these elementary operations is given by two functions: $c_{\text{del}}(a)$ is the cost of deleting or inserting a vertex with label a , and $c_{\text{match}}(a, b)$ is the cost of changing the label of a vertex from a to b .

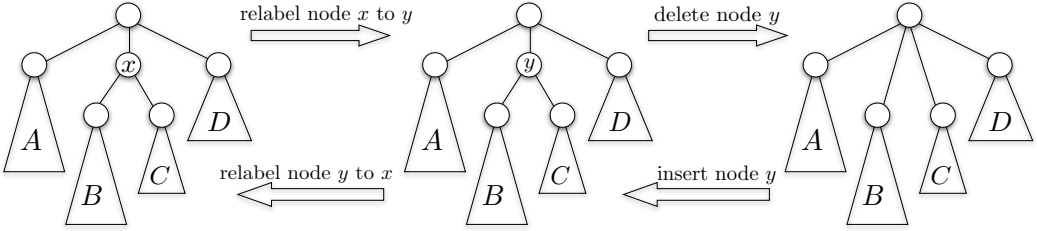


Fig. 1. The three editing operations on a tree with vertex labels.

The Tree Edit Distance (TED) problem was introduced by Tai in the late 70's [55] as a generalization of the well known string edit distance problem [57]. Since then it was extensively studied. Tai gave an $O(n^6)$ -time algorithm for TED which was subsequently improved to $O(n^4)$ in the late 80's [53], then to $O(n^3 \log n)$ in the late 90's [41], and finally to $O(n^3)$ in 2007 [34]. Many other algorithms have been developed for TED, see the popular survey of Bille [24] (this survey alone has more than 600 citations) and the books of Apostolico and Galil [16] and Valiente [56]. For example, Pawlik and Augsten [48] recently defined a class of dynamic programming algorithms that includes all the above algorithms for TED, and developed an algorithm whose performance on any input is not worse (and possibly better) than that of any of the existing algorithms. Other attempts achieved better running time by restricting the edit operations or the scoring schemes [31, 51, 54, 66], or by resorting to approximation [12, 17]. However, in the worst case no algorithm currently beats $\Omega(n^3)$ (not even by a logarithmic factor).

Due to their importance in practice, many of the algorithms described above, as well as additional heuristics and optimizations were studied experimentally [39, 48]. Solving tree edit distance in truly subcubic $O(n^{3-\epsilon})$ time is arguably one of the main open problems in pattern matching, and the most important one in tree pattern matching.

The fact that, despite the significant body of work on this problem, no truly subcubic time algorithm has been found, leads to the following natural conjecture that no such algorithm exists.

CONJECTURE 1. *For any $\epsilon > 0$ Tree Edit Distance on two n -node trees cannot be solved in $O(n^{3-\epsilon})$ time.*

In the same paper proving the $O(n^3)$ upper bound for TED [34], Demaine et al. prove that their algorithm is optimal within a certain class of dynamic programming algorithms for TED. However, proving Conjecture 1 seems to be beyond our current lower bound techniques.

A recent development in theoretical computer science suggests a more fine-grained classification of problems in P. This is done by showing lower bounds conditioned on the conjectured hardness of certain archetypal problems such as All Pairs Shortest Paths (APSP), 3-SUM, k -Clique, and Satisfiability, i.e., the Strong Exponential Time Hypothesis (SETH).

¹Since a deletion in F is equivalent to an insertion in G and vice versa, we can focus on finding the minimum cost of a sequence of just deletions and relabelings in both trees that transform F and G into isomorphic trees.

The APSP Conjecture. Given a directed or undirected graph with n vertices and integer edge weights, many classical algorithms for APSP (such as Dijkstra or Floyd-Warshall) run in $O(n^3)$ time. The fastest to date is the recent algorithm of Williams [60] that runs faster than $O(n^3/\log^C n)$ time for all constants C . Nevertheless, no truly subcubic $O(n^{3-\epsilon})$ time algorithm for APSP is known. This led to the following conjecture assumed in many papers, e.g. [5, 6, 8, 9, 11, 18, 21, 50, 61, 62].

CONJECTURE 2 (APSP). *For any $\epsilon > 0$ there exists $c > 0$, such that All Pairs Shortest Paths on n node graphs with edge weights in $\{1, \dots, n^c\}$ cannot be solved in $O(n^{3-\epsilon})$ time.*

The (Weighted) k -Clique Conjecture. The fundamental k -Clique problem asks whether a given undirected unweighted graph on n nodes and $O(n^2)$ edges contains a clique on k nodes. This is the parameterized version of the famously NP-hard Max-Clique [40]. k -Clique is amongst the most well-studied problems in theoretical computer science, and it is the canonical intractable (W[1]-complete) problem in parameterized complexity. A naive algorithm solves k -Clique in $O(n^k)$ time. A faster $O(n^{\omega k/3})$ -time algorithm (where $\omega < 2.373$ is the exponent of matrix multiplication) can be achieved via a reduction to Boolean matrix multiplication on matrices of size $n^{k/3} \times n^{k/3}$ if k is divisible by 3 [47]². Any improvement to this bound immediately implies a faster algorithm for MAX-CUT [59, 64]. It is a longstanding open question whether improvements to this bound are possible [46, 63]. The *k -Clique conjecture* asserts that for no $k \geq 3$ and $\epsilon > 0$ the problem has an $O(n^{\omega k/3-\epsilon})$ time algorithm, or an $O(n^{k-\epsilon})$ algorithm avoiding fast matrix multiplication, and has been used e.g. in [3, 28].

We work with a conjecture on a weighted version of k -Clique. In the *Max-Weight k -Clique* problem, the edges have integral weights and we seek the k -clique of maximum total weight. When the edge weights are small, one can obtain an $\tilde{O}(n^{k-\epsilon})$ time algorithm [13, 47]. However, when the weights are larger than n^k , the trivial $O(n^k)$ algorithm is the best known (ignoring $n^{o(k)}$ improvements). This gives rise to the following conjecture, which has been used e.g. in [10, 18, 21].

CONJECTURE 3 (MAX-WEIGHT k -CLIQUE). *For any $\epsilon > 0$ there exists a constant $c > 0$, such that for any $k \geq 3$ Max-Weight k -Clique on n -node graphs with edge weights in $\{1, \dots, n^{ck}\}$ cannot be solved in $O(n^{k(1-\epsilon)})$ time.*

In 2014, with the burst of the conditional lower bound paradigm, Abboud [1] highlighted seven main open problems in the field: The first two were to prove quadratic $n^{2-o(1)}$ lower bounds for *String Edit Distance* and *Longest Common Subsequence*, which were soon resolved in STOC'15 [19] and FOCS'15 [4, 29] conditional on SETH. The third problem was to show a cubic $n^{3-o(1)}$ lower bound for *RNA-Folding*. Surprisingly, in FOCS'16 [27] it was shown that RNA-Folding can actually be solved in truly subcubic time, thus ruling out the possibility of such a lower bound. The remaining four problems remain open. In fact, two of them, showing a cubic lower bound for *Graph Diameter* and an $n^{\lceil k/2 \rceil - o(1)}$ lower bound for *k -SUM*, have actually been used as hardness conjectures themselves, e.g., in SODA'15 [6] and ICALP'13 [8]. Until the present work, no progress has been made on the last problem posed by Abboud: A cubic lower bound for *Tree Edit Distance*. In the absence of progress on either upper bounds or conditional lower bounds for TED, one might think that Conjecture 1 is yet another fragment in the current landscape of fine grained complexity, and is unrelated to other common conjectures.

1.1 Our Results

In this paper we resolve the complexity of tree edit distance by showing a tight connection between edit distance of trees and all pairs shortest paths of graphs. We prove that Conjecture 2 implies Conjecture 1, and that Conjecture 3 implies Conjecture 1, even for constant alphabet.

²For the case that k is not divisible by 3 see [35].

THEOREM 1. *A truly subcubic algorithm for tree edit distance on alphabet size $|\Sigma| = \Omega(n)$ implies a truly subcubic algorithm for APSP. A truly subcubic algorithm for tree edit distance on sufficiently large alphabet size $|\Sigma| = O(1)$ implies an $O(n^{k(1-\varepsilon)})$ algorithm for Max-Weight k -Clique.*

Note that the known upper bounds for *string* edit distance and *tree* edit distance are highly related. The $O(n^2)$ algorithm for strings and the $O(n^3)$ algorithm for trees (and forests) are both based on a similar recursive solution: The recursive subproblems in strings (forests) are obtained by either deleting, inserting, or matching the rightmost or leftmost character (root). In strings, it is best to always consider the rightmost character. The recursive subproblems are then prefixes and the overall running time is $O(n^2)$. In trees however, sticking with the rightmost (or leftmost) root may result in an $O(n^4)$ running time. The specific way in which the recursion switches between leftmost and rightmost roots is exactly what enables the $O(n^3)$ solution. It is interesting that while the upper bounds for both problems are so similar, the lower bounds string edit distance exhibits the hardness of the SETH while tree edit distance exhibits the hardness of APSP.

While a considerable share of the recent conditional lower bounds is on string pattern matching problems [3, 4, 7, 10, 15, 19, 20, 28, 29, 32, 45], the only conditional lower bound for a tree pattern matching problem is the quadratic lower bound for *exact pattern matching* in unordered trees [2] (the problem of deciding whether one tree is a subtree of another). We solve the main remaining open problem in tree pattern matching, and one of the last remaining classic dynamic programming problems. Indeed, apart from the problems discussed above, for most of the classic dynamic programming problems a conditional lower bound or an improved algorithm have been found recently. This includes the Fréchet distance [25], bitonic TSP [33], context-free grammar parsing [3], maximum weight rectangle [18], and pseudopolynomial time algorithms for subset sum [26]. Tree edit distance was one of the few classic dynamic programming problems that so far resisted this approach. Two notable remaining dynamic programming problems without matching bounds are the optimal binary search tree problem ($O(n^2)$) [44] and knapsack (pseudopolynomial $O(nW)$) [23].

1.2 Our Reductions

APSP to TED. In order to prove APSP-hardness, by [61] it suffices to show a reduction from the negative triangle detection problem, where we are given an n -node graph G with edge weights $w(., .)$ and want to decide whether there are i, j, k with $w(i, j) + w(j, k) + w(i, k) < 0$. Our first result is a reduction from negative triangle detection to tree edit distance, which produces trees of size $O(n)$ over an alphabet of size $O(n)$. This yields the matching conditional lower bound of $O(n^{3-\varepsilon})$.

Our reduction constructs trees that are of a very special form: Both trees consist of a single path (called spine) of length $O(n)$ with a single leaf pending from every node (see Figure 2). Such instances already have been identified as difficult for a restricted class of algorithms based on a specific dynamic programming approach [34]. In our setting we cannot assume anything about the algorithm, and hence need a deeper insight on the structure of any valid sequence of edit operations (see Figure 2 and Lemma 1). Using this structural understanding, we then show that it is possible to carefully construct a cost function so that any optimal solution must obey a certain structure (Figure 3). Namely, for some i, j, k we match the two leaves in depth k , we match the right spine node in depth k to the left leaf in depth i (which encodes $w(i, k)$), we match the left spine node in depth k to the right leaf in depth j (which encodes $w(j, k)$), and we match as many spine nodes above depth i and j as possible (which together encode $w(i, j)$ by a telescoping sum).

Constant alphabet size. The drawback of the above reduction is the large alphabet size $|\Sigma|$, as essentially each node needs its own alphabet symbol. There are two major difficulties to improving this to constant alphabet size.

First, the instances identified as hard by the above reduction (and by Demaine et al. [34] for a restricted class of algorithms) are no longer hard for small alphabet! Indeed, in Section 4 we give an $O(n^2|\Sigma|^2 \log n)$ algorithm for these instances, which is truly subcubic for constant alphabet size. This algorithm is the first to break the barrier by Demaine et al. for such trees, and we believe it is of independent interest. Regarding lower bounds, this algorithm shows that for a reduction with constant alphabet size our trees necessarily need to be more complicated, making it much harder to reason about the structure of a valid edit sequence. We will construct new hard instances by taking the previous ones and attaching small subtrees to all nodes.

The second difficulty is that, since the input size of TED is $\tilde{O}(n + |\Sigma|^2)$, a reduction from negative triangle detection to TED with constant alphabet size would need to considerably compress the $\Omega(n^2)$ input size of negative triangle detection. It is a well-known open problem whether such compressing reductions exist. To circumvent this barrier, we assume the stronger Max-Weight k -Clique Conjecture, where the input size $\tilde{O}(n^2)$ is very small compared to the running time $O(n^k)$.

Max-Weight k -Clique to TED. Given an instance of Max-Weight k -Clique on an n -node graph G and weights bounded by $n^{O(k)}$ we construct a TED instance on trees of size $O(n^{k/3+2})$ over an alphabet of size $O(k)$. One can verify that an $O(n^{3-\epsilon})$ algorithm for TED now implies an $O(n^{k(1-\epsilon/6)})$ algorithm for Max-Weight k -Clique, for any sufficiently large $k = k(\epsilon)$.

We roughly follow the reduction from negative triangle detection; now each spine node corresponds to a $k/3$ -clique in G . To cope with the small alphabet, we simulate the previous matching costs with small gadgets. In particular, to each spine node, corresponding to some $k/3$ -clique U , we add a small subtree $T(U)$ of size $O(n^2)$ such that the edit distance between $T(U)$ and $T(U')$ encodes the total weight of edges between U and U' . This raises two issues. First, we need to represent a weight $w \in \{0, \dots, n^{O(k)}\}$ by trees over an alphabet of size $O(k)$ (that is, constant). This is solved by writing w in base n as $\sum_{i=0}^{O(k)} \alpha_i n^i$ and constructing α_i nodes of type i , such that the cost of matching two type i nodes is n^i . A second issue is that we need the small subtree $T(U)$ to interact with every other small subtree $T(U')$. So, in a sense, $T(U)$ needs to “prepare” for any possible U' , and yet its size needs to be small. We achieve this by creating in $T(U')$, for every node u in G , a separate component responsible for counting the total weight of all edges between u and all nodes in U' . Then, in $T(U)$ we have a separate component for every node $u \in U$, and make sure that it is necessarily matched to the appropriate component in $T(U')$.

The final and most intricate component of our reduction is to enforce that in any optimal solution we have some control on which small subtrees can be matched to which. A similar issue was present in the negative triangle reduction, when we require control over which spine nodes above depth i are matched to which spine nodes above depth j . This is handled in the negative triangle reduction by assigning a different matching cost depending on the node’s depth. Now however, we cannot afford so many different costs. We overcome this with yet another gadget, called an I -gadget, that achieves roughly the same goal, but in a more “distributed” manner.

Both of our reductions are highly non-trivial and introduce a number of new tricks that could be useful for other problems on trees. In particular, we show that TED probably requires cubic time for $|\Sigma| = O(1)$, but it remains an open problem to prove the same for the case of $|\Sigma| = 1$ (also known as the largest common subforest problem) or for the case where every elementary edit operation costs 1 (also known as the Levenshtein distance). Another intriguing open problem is whether or not TED is part of the APSP equivalence class. Namely, is there a (subcubic time) reduction from TED to APSP?

2 REDUCING ASPSP TO TED

We re-define the cost of matching two nodes to be the original cost minus the cost of deleting both nodes. Then, the goal of TED amounts to choosing a subset of *red nodes* in both trees, so that the subtrees defined by the red nodes are isomorphic (i.e., their left-right and ancestor-descendant relation is the same in both trees) and the total cost of matching the corresponding red nodes is minimized. See Figure 2. We work with this formulation from now on.

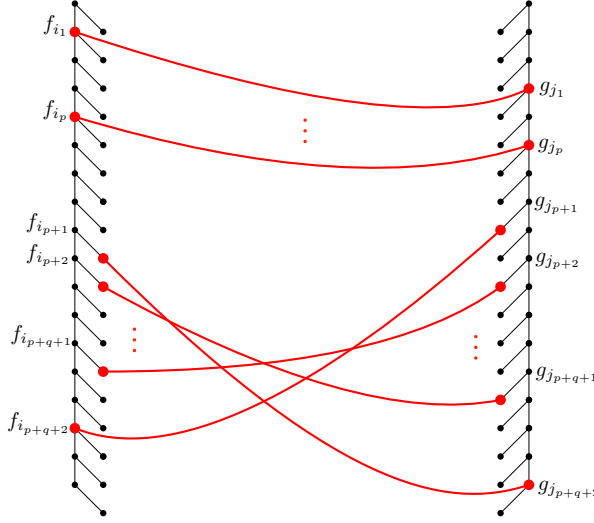


Fig. 2. Macro structure of the hard instance for TED: A tree F composed of a single spine with leaves hanging to the right and a tree G composed of a single spine with leaves hanging to the left.

It turns out that a hard instance for TED is given by two seemingly simple caterpillar trees. These two trees F and G , also called left and right, are shown in Figure 2. Each tree consists of *spine nodes* and *leaf nodes*. If u is a spine node then we denote by u' the (unique) leaf node attached to u . For any such hard instance of TED, the red nodes in any matching have the structure given by Lemma 1 below. Informally, it states that starting from the top of the left tree and ordering the nodes by depth, the matching consists of (1) a prefix of a matching subsequence of spine nodes in both trees, (2) a suffix of a matching subsequence of leaf nodes that are in reverse order in the other tree, and (3) at most one final spine node in each of the trees matching a leaf node in the other tree that is located between the prefix part (1) and the suffix part (2).

LEMMA 1. *Let f_1, f_2, \dots and g_1, g_2, \dots denote the spine nodes of F and G , respectively, ordered by the depth. Then, for some $p, q \geq 0$ and some $i_1 < i_2 < \dots < i_p < i_{p+1} < \dots < i_{p+q+1} < i_{p+q+2}$ and $j_1 < j_2 < \dots < j_p < j_{p+1} < \dots < j_{p+q+1} < j_{p+q+2}$ the set of red nodes consists of:*

- (1) *Spine nodes $f_{i_1}, f_{i_2}, \dots, f_{i_p}$ matched respectively to spine nodes $g_{j_1}, g_{j_2}, \dots, g_{j_p}$,*
- (2) *Leaf nodes $f'_{i_{p+2}}, f'_{i_{p+3}}, \dots, f'_{i_{p+q+1}}$ matched respectively to leaf nodes $g'_{j_{p+q+1}}, g'_{j_{p+q}}, \dots, g'_{j_{p+2}}$ (note the reversed order),*
- (3) *Optionally, a spine node $f_{i_{p+q+2}}$ matched to leaf node $g'_{j_{p+1}}$. Also optionally, a spine node $g_{j_{p+q+2}}$ matched to a leaf node $f'_{i_{p+1}}$.*

PROOF. Consider the subtree defined by the red nodes. It has two isomorphic copies, one in F and one in G . Its nodes are all the red nodes. The children of node u are all red nodes v_1, v_2, \dots, v_k

whose lowest red ancestor is u . The order is such that v_i precedes v_{i+1} in a left-to-right preorder traversal of F (or equivalently of G). Let u be a red node with two or more children v_1, v_2, \dots, v_k , $k \geq 2$. Observe that u must correspond to spine nodes in both F and G . Further observe that at most one v_i can correspond to a spine node (otherwise, for two spine nodes one must be an ancestor of the other). Consider any $\ell \in \{1, 2, \dots, k-1\}$. It is not hard to see that node $v_{\ell+1}$ must correspond to a leaf node in F and node v_ℓ must correspond to a leaf node in G . This implies that both v_ℓ and $v_{\ell+1}$ are leaves in the red subtree. Moreover, v_1 is the only node that may correspond to a spine node in F and v_k is the only node that may correspond to a spine node in G . Consequently, the red subtree has a particularly simple structure: it consists of nodes u_1, u_2, \dots, u_p such that for every $\ell = 1, 2, \dots, p-1$ the only child of u_ℓ is $u_{\ell+1}$, and nodes v_1, v_2, \dots, v_k (for some $k \geq 1$) that are all children of u_p .

For every $\ell = 1, 2, \dots, p$, the node u_ℓ must correspond to a spine node $f_{i_\ell} \in F$ and $g_{j_\ell} \in G$. We immediately obtain (1) that $i_1 < i_2 < \dots < i_p$ and that $j_1 < j_2 < \dots < j_p$. The nodes v_1, v_2, \dots, v_k are all children of u_p in the subtree. It is possible that all v_i are mapped to leaf nodes $f'_{i_{p+2}}, f'_{i_{p+3}}, \dots, f'_{i_{p+q+1}}$ and $g'_{j_{p+2}}, g'_{j_{p+3}}, \dots, g'_{j_{p+q+1}}$. In this case, they must be mapped in reverse order since a left-to-right preorder traversal visits the leaves of G in order of their depth and in reverse-depth order in F . This implies (2) that $i_p \leq i_{p+2} < \dots < i_{p+q+1}$ and $j_p \leq j_{p+2} < \dots < j_{p+q+1}$. Recall however that v_1 may be mapped to a spine node $f_{i_{p+q+2}}$ in F and a leaf node $g'_{j_{p+1}}$ in G . The requirement that $i_{p+q+2} > i_{p+q+1}$ and that $j_{p+2} > j_{p+1} \geq j_p$ follows from the fact that these nodes correspond to a leftmost leaf in the subtree. For symmetric reasons, v_k may be matched to a spine node $g_{j_{p+q+2}} \in G$ for some $j_{p+q+2} > j_{p+q+1}$ and $i_{p+2} > i_{p+1} \geq i_p$. This implies (3) and concludes the proof. \square

The above lemma characterizes the structure of a solution to what we call the *hard instance* of TED. We next show how to reduce the *negative triangle detection* problem to TED on the hard instance. Negative triangle detection is known to be subcubic equivalent to APSP [61]. Given a complete weighted n -node undirected graph, where $w(i, j)$ denotes the weight of the edge (i, j) , the problem asks whether there are i, j, k such that $w(i, j) + w(j, k) + w(i, k) < 0$. To solve negative triangle detection, we clearly only need to find i, j, k that minimize $w(i, j) + w(j, k) + w(i, k)$. We will show how to construct, given such a graph, a hard instance of TED of size $O(n)$, such that $\min_{i,j,k} w(i, j) + w(j, k) + w(i, k)$ can be extracted from the edit distance.

LEMMA 2. *Given a complete undirected n -node graph G with weights $w(\cdot, \cdot)$ in $\{1, \dots, n^\epsilon\}$, we construct, in linear time in the output size, an instance of TED of size $O(n)$ with alphabet size $|\Sigma| = O(n)$ such that the minimum weight of a triangle in G can be extracted from the edit distance.*

Consequently, an $O(n^{3-\epsilon})$ time algorithm for TED implies an $O(n^{3-\epsilon})$ algorithm for negative triangle detection, and thus an $O(n^{3-\epsilon/3})$ algorithm for APSP by a reduction in [61].

We create a hard instance of TED consisting of two trees F and G as in Figure 3. Every tree is divided into a *top* and a *bottom* part. The spine nodes of these parts are denoted by a_1, a_2, \dots, a_n for the top left part, b_1, b_2, \dots, b_{n+1} for the bottom left part, c_1, c_2, \dots, c_n for the top right part, and d_1, d_2, \dots, d_{n+1} for the bottom right part. The labels of all nodes are distinct (hence the alphabet size $|\Sigma|$ is $\Theta(n)$). We set the cost $c_{\text{match}}(u, v)$ of matching two nodes u and v as described below, where M denotes a sufficiently large number to be specified later. Intuitively, our assignment of costs ensures that any valid solution to TED must match b'_k to d'_k , b_{k+1} to c'_j , and d_{k+1} to a'_i for some i, j, k (as shown in Figure 3). Furthermore, the optimal solution (i.e., of minimum cost) must choose i, j, k that minimize $w(i, k) + w(k, j) + w(i, j)$. The costs are assigned as follows:

- (1) $c_{\text{match}}(b'_k, d'_k) = -M^2 - 2M \cdot k$ for every $k = 1, 2, \dots, n$.
- (2) $c_{\text{match}}(b_{k+1}, c'_j) = -M^2 + M \cdot k + M \cdot j + w(k, j)$ for every $k = 1, 2, \dots, n$ and $j = 1, 2, \dots, n$.

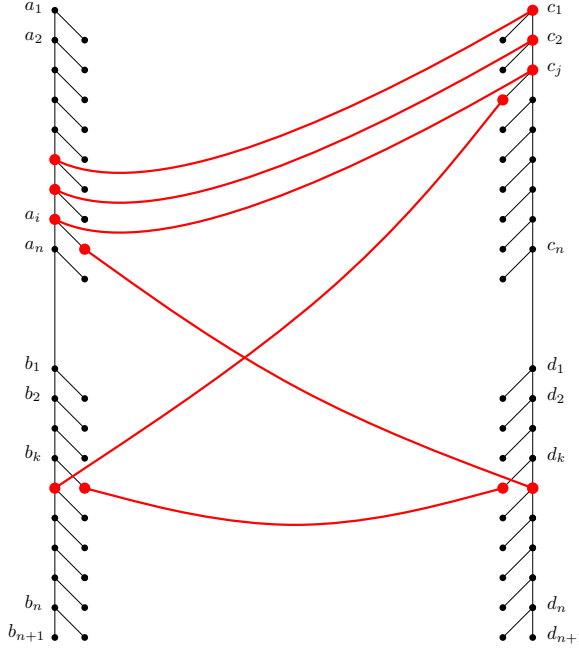


Fig. 3. A hard instance of TED constructed for a given instance of negative triangle detection. Appropriately chosen costs ensure that any optimal solution has a specific structure.

- (3) $c_{\text{match}}(a'_i, d_{k+1}) = -M^2 + M \cdot k + M \cdot i + w(i, k)$ for every $i = 1, 2, \dots, n$ and $k = 1, 2, \dots, n$.
- (4) $c_{\text{match}}(a_i, c_j) = -2M + w(i, j) - w(i-1, j-1)$ for every $i = 2, 3, \dots, n$ and $j = 2, 3, \dots, n$.
- (5) $c_{\text{match}}(a_i, c_1) = -M(i+1) + w(i, 1)$ for every $i = 1, 2, \dots, n$.
- (6) $c_{\text{match}}(a_1, c_j) = -M(j+1) + w(1, j)$ for every $j = 1, 2, \dots, n$.

All the remaining costs $c_{\text{match}}(u, v)$ are set to ∞ . The following theorem proves that these costs imply the required structure on the optimal solution as described above. Intuitively, by choosing sufficiently large M , because of the $-M^2$ addend in (1), (2) and (3) we can ensure that any optimal solution matches b'_k to d'_k , $b_{k'+1}$ to c'_j , and $d_{k''+1}$ to a'_i , for some i, j, k and $k', k'' \geq k$. Then, because of the $M \cdot k$ in (2) and (3), in any optimal solution actually $k = k' = k''$ and the total cost of all these matchings is $w(k, j) + w(i, k)$. Finally, because of the $-2M$ in (4), the $-M(i+1)$ in (5), and the $-M(j+1)$ in (6), in any optimal solution a_i is matched to c_j , a_{i-1} to c_{j-1} , a_{i-2} to c_{j-2} and so on. The total cost of these matching is $w(i, j)$ since the $w(i, j) - w(i-1, j-1)$ terms in (4) form a telescoping sum.

THEOREM 2. *For sufficiently large M , the total cost of an optimal matching in a hard instance with costs (1)-(6) is $-3M^2 + \min_{i,j,k} w(i, k) + w(k, j) + w(i, j)$.*

PROOF. Consider i, j, k minimizing $w(i, k) + w(k, j) + w(i, j)$. We assume without loss of generality that $i \geq j$. It is easy to see that it is possible to choose the following matching (see Figure 3):

1. b'_k to d'_k with cost $-M^2 - 2M \cdot k$.
2. b_{k+1} to c'_j with cost $-M^2 + M \cdot k + M \cdot j + w(k, j)$.
3. d_{k+1} to a'_i with cost $-M^2 + M \cdot k + M \cdot i + w(i, k)$.
4. a_i to c_j , a_{i-1} to c_{j-1} , a_{i-2} to c_{j-2} , \dots , a_{i-j+2} to c_2 with costs $-2M + w(i, j) - w(i-1, j-1)$, $-2M + w(i-1, j-1) - w(i-2, j-2)$, \dots , $-2M + w(i-j+2, 2) - w(i-j+1, 1)$.

5. a_{i-j+1} to c_1 with cost $-M \cdot (i - j + 2) + w(i - j + 1, 1)$.

Summing up and telescoping, the total cost is $-M^2 - 2M \cdot k - M^2 + M \cdot k + M \cdot j + w(k, j) - M^2 + M \cdot k + M \cdot i + w(i, k) - 2M \cdot (j - 1) + w(i, j) - M \cdot (i - j + 2)$ which is equal to $-3M^2 + w(i, k) + w(k, j) + w(i, j)$.

For the other direction, we need to prove that every solution has cost at least $-3M^2 + \min_{i,j,k} w(i, k) + w(k, j) + w(i, j)$. We first observe that, by Lemma 1, a solution can match b'_k to d'_k at most once for some k . Similarly, it can match $b_{k'+1}$ to c'_j at most once for some j and k' , and $d_{k''+1}$ to a'_i at most once for some i and k'' . Furthermore, for M large enough, either the cost up to the lower order terms is at least $-2M^2$ or all three such pairs of nodes are matched for some k, i, j , and $k', k'' \geq k$. Furthermore, if $k' > k$ and M is large enough then we can decrease k' by one thus decreasing the total cost, and similarly if $k'' > k$. It is enough to consider an optimal solution and hence we can assume that $k = k' = k''$.

Again by Lemma 1, the only possible additional matched pairs of nodes are a subsequence of spine nodes a_1, \dots, a_i and c_1, \dots, c_j . We show that an optimal solution matches a_i with c_j , a_{i-1} with c_{j-1} , ..., a_{i-j+1} with c_1 . To this end, suppose that a_{x_z} is matched to c_{y_z} , for every $z = 1, 2, \dots, L$, where $1 \leq x_1 < \dots < x_L \leq i$ and $1 \leq y_1 < \dots < y_L \leq j$. For every z this contributes, up to lower order terms less than M , $-2M$ if $x_z, y_z > 1$, or $-M(y_z + 1)$ if $x_z = 1$, or $-M(x_z + 1)$ if $y_z = 1$. In an optimal solution we have $x_1 = 1$ or $y_1 = 1$, as otherwise we can match a_1 to c_1 to decrease the total cost. First, assume that $y_1 = 1$. Then, if $x_{z'} + 1 < x_{z'+1}$ for some $z' \in \{1, \dots, L\}$ (where we define $x_{L+1} = i + 1$), we can increase all $x_1, x_2, \dots, x_{z'}$ by 1 to decrease the total cost by M , up to lower order terms. So $x_L = i, x_{L-1} = i - 1, \dots, x_1 = i - L + 1$. Now if $L < j$ then $x_1 > 1$ (recall that we assumed $i \geq j$) and also $y_{z'} + 1 < y_{z'+1}$ for some $z' \in \{1, \dots, L\}$ (again, we define $y_{L+1} = j + 1$). This means that we can increase all $y_1, y_2, \dots, y_{z'}$ by 1 and then additionally match a_{x_1-1} with c_1 to decrease the total cost by M , up to lower order terms. Second, if $x_1 = 1$ a symmetric argument applies. We obtain that indeed $L = \min(i, j) = j$, and a_{i-j+1} is matched to c_1 , a_{i-j+2} is matched to c_2 , ..., a_i is matched to c_j . Now, by the same calculations as in the previous paragraph, the total cost is $-3M^2 + w(i, k) + w(k, j) + w(i, j)$. \square

3 REDUCING MAX-WEIGHT k -CLIQUE TO TED

The drawback of the reduction described in Section 2 is the large size of the alphabet. That is, given a complete weighted n -node undirected graph it creates two trees of size $O(n)$ where labels of nodes are distinct, and therefore $|\Sigma| = \Theta(n)$. We would like to refine the reduction so that $|\Sigma| = O(1)$. However, as the input size of TED on n -node trees and alphabet Σ with $O(\log n)$ -bit integer weights is $\tilde{O}(n + |\Sigma|^2)$, such a reduction would need to compress the $\tilde{O}(n^2)$ input size of negative triangle detection considerably. To circumvent this barrier, we assume the stronger Max-Weight k -Clique Conjecture, where the input size $\tilde{O}(n^2)$ is very small compared to the running time bound $O(n^k)$.

LEMMA 3. *Given a complete undirected n -node graph G with weights in $\{1, \dots, n^{ck}\}$, we construct, in linear time in the output size, an instance of TED of size $O(n^{k/3+2})$ with alphabet size $|\Sigma| = O(ck)$ such that the maximum weight of a k -clique in G can be extracted from the edit distance.*

Thus, an $O(n^{3-\epsilon'})$ time algorithm for TED for sufficiently large $|\Sigma| = O(1)$ implies an $O(n^{(k/3+2)(3-\epsilon)})$ time algorithm for max-weight k -Clique. Setting $\epsilon = \epsilon'/6$, we obtain that, for every $c > 0$, there exists $k = \lceil 6/\epsilon \rceil$ such that max-weight k -Clique can be solved in time

$$O(n^{(k/3+2)(3-\epsilon)}) = O(n^{k-\epsilon'k/3+6-2\epsilon'}) = O(n^{k(1-\epsilon)-k\epsilon+6}) = O(n^{k(1-\epsilon)}),$$

so Conjecture 3 is violated.

The reduction starts with enumerating all $\frac{k}{3}$ -cliques in the graph and identifying them with numbers $1, 2, \dots, N$, where $N \leq n^{k/3}$. Let $Q(i)$ denote the set of nodes in the i -th clique. Then, for

i, j such that $Q(i) \cap Q(j) = \emptyset$, $W(i, j)$ is the total weight of all edges connecting two nodes in the i -th clique or a node in the i -th clique with a node in the j -th clique. Our goal is to calculate the maximum value of $W(i, z) + W(z, j) + W(j, i)$ over i, j, z such that $Q(i), Q(j)$ and $Q(z)$ are pairwise disjoint. If we define $w(u, u) = 0$ and increase every other weight $w(u, v)$ by $\Lambda := k^2 n^{ck}$, this is equivalent to maximising over all i, j, z . Indeed, if $Q(i), Q(j), Q(z)$ are pairwise disjoint, the total weight is at least $\binom{k}{2}\Lambda$, and otherwise it is at most $\left(\binom{k}{2} - 1\right)(\Lambda + n^{ck}) < \binom{k}{2}\Lambda$. Note that the new weights are still bounded by $n^{O(ck)}$.

Our construction of a hard instance of size $O(N \cdot \text{poly}(n))$ is similar to Section 2, however, the costs are set up differently and we attach small additional gadgets to some of the nodes (which is necessary, cf. Section 4). The original two trees (with some extra spine nodes without any leaves) are called the *macro structure* and all small gadgets are called the *micro structures*. With notation as in Section 2, the following micro structures are created for every $i = 1, 2, \dots, N$ (see Figure 4):

1. A'_i attached to the leaf a'_i ,
2. a copy of I attached as the left child of the leaf c'_i ,
3. a copy of I attached as the right child of the leaf a'_i ,
4. C'_i attached as the right child of the leaf c'_i ,
5. A_i attached to the spine node a_{i-1} between the previously existing children a_i and a'_{i-1} ,
6. B_i attached to the spine node b_i between the previously existing children b_{i+1} and b'_i ,
7. C_i attached to the spine node c_{i-1} as the rightmost child,
8. D_i attached to the spine node d_i between the previously existing children d'_i and d_{i+1} .

Notice that A_i is attached above a_i (and similarly C_j is attached above c_j). Therefore, we need to create dummy spine nodes a_0 and c_0 . We also insert an additional spine node b'_i between b_i and b_{i+1} and similarly d'_i between d_i and d_{i+1} , for every $i = 1, 2, \dots, N - 1$. See Figure 4.

The costs in the macro structure are chosen as follows, where again M is a sufficiently large value (picking $M = n^{O(ck)}$ will suffice):

1. $c_{\text{match}}(b_z, c'_j) = -M^8$ for every $z = 1, 2, \dots, N$ and $i = 1, 2, \dots, N$,
2. $c_{\text{match}}(a'_i, d_{z'}) = -M^8$ for every $i = 1, 2, \dots, N$ and $z' = 1, 2, \dots, N$,
3. $c_{\text{match}}(b'_z, d'_{z'}) = -M^7 \cdot 2$ for every $z = 1, 2, \dots, N$ and $z' = 1, 2, \dots, N$,
4. $c_{\text{match}}(a_i, c_j) = -M^3 \cdot 2 + M^2$ for every $i = 1, 2, \dots, N$ and $j = 1, 2, \dots, N$.

Additionally, the extra spine nodes b'_i and d'_i can be matched to some of the nodes of I . Each copy of I is a path consisting of $k/3$ segments $I_0, I_1, \dots, I_{k/3-1}$ of length $n - 1$, where the root of the whole I belongs to I_0 . The label of every $u \in I_i$ is the same and the costs are set so that $c_{\text{match}}(u, u) = -M^7 \cdot n^i$. The label of every b'_z (and also $d'_{z'}$) is chosen as the label of every $u \in I_m$, where n^m is the largest power of n dividing $N - z$. The cost of matching any other two labels used in the macro structure is set to infinity. For the nodes belonging to the other micro structures, the cost of matching is at least $-M^6$ and will be specified precisely later. This is enough to enforce the following structural property.

LEMMA 4. *For sufficiently large M , any optimal matching has the following structure: there exist i, j, z such that a'_i is matched to d_z , c'_j is matched to b_z , b'_1 is matched to d'_{z-1} , b'_2 is matched to d'_{z-2} , \dots , b'_{z-1} is matched to d'_1 . Furthermore, if $z < N$ then b'_z is matched to a descendant of c'_j and d'_z is matched to a descendant of a'_i . Ignoring the spine nodes $a_1, \dots, a_i, c_1, \dots, c_i$ and all micro structures that are not copies of I the cost of any such solution is $-M^8 \cdot 2 - M^7 \cdot 2(N - 1)$.*

PROOF. For sufficiently large M , any optimal solution must match a'_i to d_z and c'_j to $b_{z'}$, for some i, j, z, z' , as otherwise its cost is larger than $-M^8 \cdot 2$. By the reasoning described in Lemma 1, these i, j, z, z' are uniquely defined for any optimal solution.

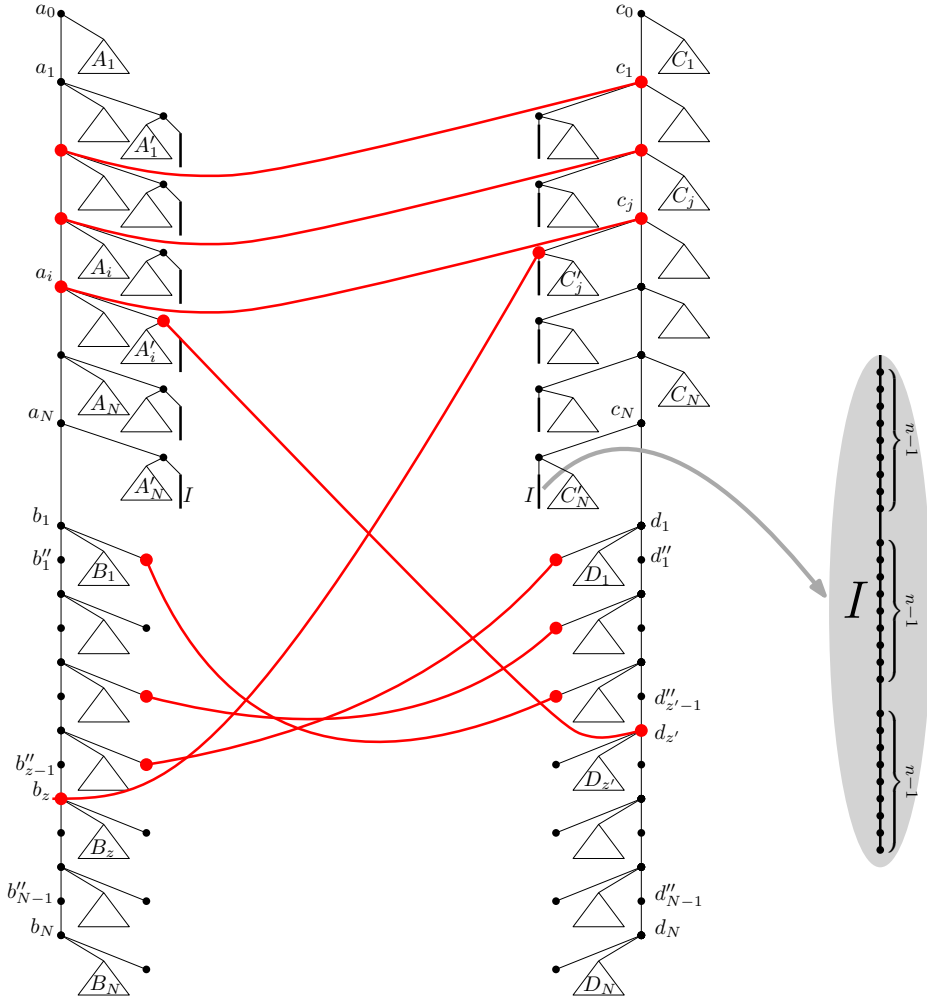


Fig. 4. A hard instance of TED constructed for a given instance of max-weight k -clique.

Nodes from the copy of I attached as the left child of the leaf c'_j can be matched to some spine nodes below b_z , nodes from the copy of I attached as the right child of the leaf a'_i can be matched to some spine nodes below $d_{z'}$, and no other nodes from the copies of I can be matched. We claim that the total contribution of these nodes is $-M^7(N - z)$ and $-M^7(N - z')$, respectively. By symmetry, it is enough to justify the former. Observe that the cost of matching a single $u \in I_i$ is smaller than the total cost of matching all nodes from $I_0 \cup \dots \cup I_{i-1}$, therefore an optimal solution must match as many nodes to nodes from $I_{k/3-1}$ as possible. Looking at the expansions of all numbers $N - z, N - (z + 1), \dots, N - (N - 1)$ in base n , where $N - z = \sum_{i=0}^{k/3-1} \alpha_i n^i$, we see that there are $\alpha_{k/3-1}$ such nodes, namely the nodes $b'_{N-z'}$ with $z \leq z' < N$ and $N - z'$ divisible by $n^{k/3-1}$. Then, an optimal solution must match as many nodes to nodes from $I_{k/3-2}$ as possible to nodes above the topmost node matched to a node from $I_{k/3-1}$. Looking again at the same expression, we see that there are $\alpha_{k/3-2}$ such nodes, namely the nodes $b'_{N-z'}$ with $z \leq z' < N - \alpha_{k/3-1} n^{k/3-1}$ and $N - z'$

divisible by $n^{k/3-2}$. Continuing in the same fashion, we obtain that there are α_i nodes matched to nodes from I_i , making the total cost $-M^7(N-z)$ as claimed.

We assume without loss of generality that $z \geq z'$. Then, an optimal solution must match $d'_{z'-1}$ to $b'_{x_{z'-1}}$, $d'_{z'-2}$ to $b'_{x_{z'-2}}$, \dots , and d'_1 to b'_{x_1} , for some $z > x_1 > \dots > x_{z'-1} \geq 1$, as otherwise its cost is larger than $-M^8 \cdot 2 - M^7(2N - z - z') - M^7 \cdot 2(z' - 1)$. Rewriting the cost we obtain $-M^8 \cdot 2 - M^7(2N - 2 - z + z')$, so recalling our assumption $z \geq z'$ we see that in fact $z = z'$ as otherwise its cost is larger than $-M^8 \cdot 2 - M^7 \cdot 2(N - 1)$. \square

We are now ready to state properties of the remaining micro structures. Let $c_{\text{match}}(T_1, T_2)$ denote the cost of matching two trees T_1 and T_2 . Then, we require that:

1. $c_{\text{match}}(A'_i, D_{z'}) = -M^6 - M^3(N - i) - W(i, z')$ for every $i = 1, 2, \dots, N$ and $z' = 1, 2, \dots, N$,
2. $c_{\text{match}}(B_z, C'_j) = -M^6 - M^3(N - j) - W(z, j)$ for every $z = 1, 2, \dots, N$ and $j = 1, 2, \dots, N$.
3. $c_{\text{match}}(A_i, C_j) = -M^2 - 2M(k/3)^2 - W(j, i) + W(j - 1, i - 1)$ for every $i = 2, 3, \dots, N$ and $j = 2, 3, \dots, N$.
4. $c_{\text{match}}(A_i, C_1) = -M^5 - M^3(i - 1) - W(1, i)$ for every $i = 1, 2, \dots, N$,
5. $c_{\text{match}}(A_1, C_j) = -M^5 - M^3(j - 1) - W(j, 1)$ for every $j = 1, 2, \dots, N$.

The labels of the nodes in the micro structures should be partitioned into disjoint subsets corresponding to the following micro structures:

1. $\{A'_1, A'_2, \dots, A'_N, D_1, D_2, \dots, D_N\}$,
2. $\{B_1, B_2, \dots, B_N, C'_1, C'_2, \dots, C'_N\}$,
3. $\{A_1, A_2, \dots, A_N, C_1, C_2, \dots, C_N\}$,

so that two nodes can be matched only if their labels belong to the same subset. The cost of matching any node of $A'_i, D_{z'}, B_z, C'_j$ should be at least $-M^6$. The cost of matching any node of A_i, C_j should be at least $-M^2$, except that the root of A_i (C_j) can be matched to the root of C_1 (A_1) with cost larger than $-M^5 - M$ but at most $-M^5$, and, for any non-root node of A_i (C_j) and for any non-root node of C_1 (A_1), the cost of matching is larger than $-M^4$. Finally, every A_i and C_j should consist of $O(n^2)$ nodes. Now we can show that, assuming these properties, any optimal solution has a specific structure.

LEMMA 5. *For sufficiently large M , the total cost of an optimal matching is*

$$-M^8 \cdot 2 - M^7 \cdot 2(N - 1) - M^6 \cdot 2 - M^5 - M^3 \cdot 2N + M^2 - 2M(k/3)^2 - \max_{i,j,z} \{W(i, z) + W(z, j) + W(j, i)\}.$$

PROOF. Consider i, j, z maximizing $W(i, z) + W(z, j) + W(j, i)$. We may assume that $i \geq j$. Then, it is possible to choose the following matching:

1. b_z to c'_j with cost $-M^8$,
2. some nodes from the copy of I being the left child of c'_j to some spine nodes below b_z with total cost $-M^7(N - z)$,
3. a'_i to d_z with cost $-M^8$,
4. some nodes from the copy of I being the right child of a'_i to some spine nodes below d_z with total cost $-M^7(N - z)$,
5. b'_1 to d'_{z-1} , b'_2 to d'_{z-2} , \dots , b'_{z-1} to d'_1 with cost $-M^7 \cdot 2$ each,
6. a_i to c_j , a_{i-1} to c_{j-1} , \dots , a_{i-j+1} to c_1 with cost $-M^3 \cdot 2 + M^2 + 2M(k/3)^2$ each,
7. A'_i to D_z with cost $-M^6 - M^3(N - i) - W(i, z)$,
8. B_z to C'_j with cost $-M^6 - M^3(N - j) - W(z, j)$,
9. A_i to C_j , A_{i-1} to C_{j-1} , \dots , A_{i-j+2} to C_2 with costs $-M^2 - 2M(k/3)^2 - W(j, i) + W(j - 1, i - 1)$, $-M^2 - 2M(k/3)^2 - W(j - 1, i - 1) + W(j - 2, i - 2)$, \dots , $-M^2 - 2M(k/3)^2 - W(2, i - j + 2) + W(1, i - j + 1)$,
10. A_{i-j+1} to C_1 with cost $-M^5 - M^3(i - j) - W(1, i - j + 1)$.

Summing up and telescoping, the total cost is

$$\begin{aligned}
& -M^8 \\
& -M^7(N-z) \\
& -M^8 \\
& -M^7(N-z) \\
& -M^7 \cdot 2(z-1) \\
& -M^3 \cdot 2j + M^2 \cdot j + M(k/3)^2 j \\
& -M^6 - M^3(N-i) - W(i, z) \\
& -M^6 - M^3(N-j) - W(z, j) \\
& -M^2(j-1) - M(k/3)^2(j-1) - W(j, i) - M^5 - M^3(i-j) \\
& = -M^8 \cdot 2 - M^7 \cdot 2(N-1) - M^6 \cdot 2 - M^5 - M^3 \cdot 2N + M^2 + M(k/3)^2 - W(i, z) - W(z, j) - W(j, i).
\end{aligned}$$

For the other direction, we need to argue that every solution has cost at least $-M^8 \cdot 2 - M^7 \cdot 2(N-1) - M^6 \cdot 2 - M^5 - M^3 \cdot 2N + M^2 + M(k/3)^2 - \max_{i,j,z} \{W(i, z) + W(z, j) + W(j, i)\}$. We start with invoking Lemma 4 and analyse the remaining small micro structures. Due to leaves $b'_1, \dots, b'_{z-1}, d'_1, \dots, d'_{z-1}$ being already matched, no node from $B_1, \dots, B_{z-1}, D_1, \dots, D_{z-1}$ can be matched (as they can in general only be matched to A'_* 's and C'_* 's). Then, due to b''_z and d''_z being already matched (or $z = N$) no node from $B_{z+1}, \dots, B_N, D_{z+1}, \dots, D_N$ can be matched, and nodes from B_z or D_z can be only matched to nodes from C'_j or A'_i , respectively. The cost incurred by all such nodes is $c_{\text{match}}(A'_i, D_z) + c_{\text{match}}(B_z, C'_j)$, making the total cost $-M^8 \cdot 2 - M^7 \cdot 2(N-1) - M^6 \cdot 2 - M^3(2N - i - j) - W(i, z) - W(z, j)$. It remains to analyse the contribution of all spine nodes $a_1, \dots, a_N, b_1, \dots, b_N$ and nodes from micro structures $A_1, \dots, A_N, C_1, \dots, C_N$.

Consider the micro structures C_1 and A_1 . Matching their roots to roots of some $A_{i'}$ and $C_{j'}$, respectively, decreases the total cost by at least $-M^5$, which is much smaller than the cost of matching the remaining nodes. Furthermore, it is not possible to match both the root of C_1 to the root of some $A_{i'}$ and the root of A_1 to the root of some $C_{j'}$ at the same time, unless the root of A_1 is matched to the root of C_1 . Therefore, an optimal solution matches exactly one of them or both to each other, say we match the root of C_1 to the root of some $A_{i'}$, thus adding $c_{\text{match}}(A_{i'}, C_1)$ to the total cost. Due to a'_i being matched to d_z , $i' \leq i$ holds. Now, unless $i' = 1$, no node from A_1 can be matched to a node from $C_{j'}$, so the cost of matching any $a_{i'}$ to $c_{j'}$ is now much smaller than the cost of matching nodes in the remaining micro structures (for each such node, the cost is at least $-M^2$, and there are at most $O(n^2)$ of them in a single micro structure, so the total cost contributed by a single micro structure is larger than $-M^3$ for M large enough) and, by Lemma 1, only nodes $a_1, \dots, a_i, c_1, \dots, c_j$ can be matched, so an optimal solution matches as many such pairs as possible. Due to the root of C_1 being matched to the root of $A_{i'}$, only nodes $a_{i'}, a_{i'+1}, \dots, a_i$ and c_1, \dots, c_j can be matched, so there are $\min(i - i' + 1, j)$ such matched pairs. If $i - i' + 1 < j$ and $i' > 1$ then C_1 can be matched with $A_{i'-1}$ instead of $A_{i'}$ which allows for an additional pair and decreases the total cost (because matching a pair (a_*, c_*) adds $-M^3 \cdot 2$ to the cost while decreasing i' by one adds M^3 to the cost $c_{\text{match}}(A_{i'}, C_1)$, up to lower order terms). If $i - i' + 1 < j$ and $i' = 1$ then A_1 can be matched with C_2 instead of C_1 while keeping the number of matched pairs intact to decrease the total cost. So $i - i' + 1 \geq j$ (implying $i \geq j$, which is due to our initial assumption that the root of C_1 is matched to the root of some $A_{i'}$). Then, if $i' < i - j + 1$, C_1 can be matched with $A_{i'+1}$ instead of $A_{i'}$ without changing the number of matched pairs to decrease the total cost. Thus, $i' = i - j + 1$ and a_i is matched to c_j , a_{i-1} to c_{j-1} , \dots , and a_{i-j+1} to c_1 ,

Then nodes from A_i can be only matched to nodes from C_j , nodes from A_{i-1} only to nodes from C_{j-1} , and so on. By the same calculations as in the previous paragraph, the total cost is therefore $-M^8 \cdot 2 - M^7 \cdot 2(N-1) - M^6 \cdot 2 - M^5 - M^3 \cdot 2N + M^2 + M(k/3)^2 - \max_{i,j,z} \{W(i,z) + W(z,j) + W(j,i)\}$. \square

To complete the proof we need to design the remaining micro structures. We start with describing some preliminary gadgets that will be later appropriately composed to obtain the micro structures $A_i, A'_i, B_z, C_j, C'_j, D_z$ with the required properties. Each such gadget consists of two trees, called left and right, and we are able to exactly calculate the cost of matching them. The main difficulty here is that we need to keep the size of the alphabet small, so for instance we are not able to create a distinct label for every node of the original graph. At this point it is also important to note that we can assume $M = n^{O(ck)}$, i.e., there is a constant $d = O(ck)$ such that all weights constructed above have absolute value less than n^d .

Decrease gadget $D(x)$. For any $x \in \{0, \dots, n^d - 1\}$, the cost of matching the left and right tree of $D(x)$ is $-x$, and furthermore the right tree does not depend on the value of x .

This is obtained by representing x in base n as $x = \sum_{i=0}^{d-1} \alpha_i n^i$. The left tree is a path composed of d segments, the i -th segment consisting of α_i nodes. The right tree is a path composed of d segments, each consisting of $n-1$ nodes. Nodes from the i -th segment of the left tree can be matched with nodes from the i -th segment of the right tree with cost $-n^i$, so the total cost is $-x$, see Figure 5 (left). We reuse the same set of distinct labels in every decrease gadget of the same type, hence we need only $O(d)$ distinct labels in total. Furthermore, note that the cost of matching the left tree of $D(x)$ with any tree is at least $-x$ and the cost of matching any node of $D(x)$ is $-n^i$ for some $i \in \{0, 1, \dots, d-1\}$.

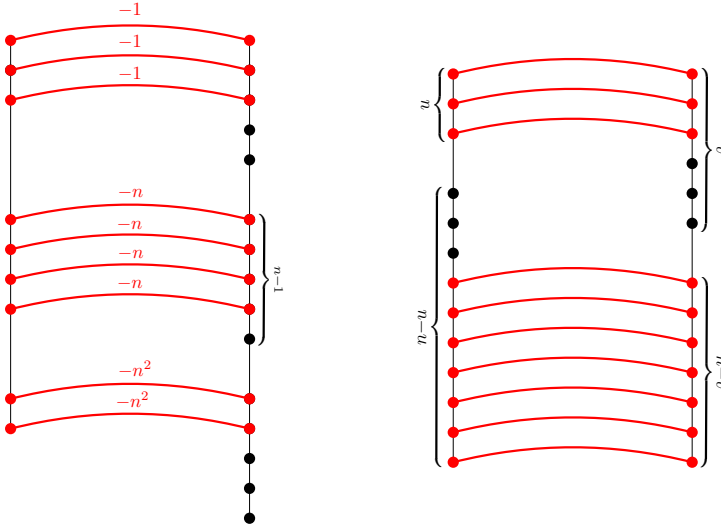


Fig. 5. Left: Decrease gadget built for $d = 3$, $n = 6$ and $x = n^2 \cdot 2 + n \cdot 4 + 3$. Right: Equality gadget for $u = 3$, $v = 6$.

Equality gadget $E(u, v, c_=-)$. For any $u, v \in \{1, \dots, n\}$ and $c_=- \in \{0, \dots, n^d - 1\}$, the cost of matching the left and right tree of $E(u, v, c_=-)$ is $-c_=- \cdot n$ if $u = v$ and at least $-c_=- \cdot n + c_=-$ otherwise. Also, the left tree does not depend on v and the right tree does not depend on u .

The left tree is a path composed of a segment of length u and a segment of length $n - u$. The right tree is a path composed of a segment of length v and a segment of length $n - v$. Nodes from the first segment of the left tree can be matched with nodes from the first segment of the right tree with cost $-c_{=}$, and similarly for the second segments. Then, if $u = v$ we can match all nodes in both trees, so the total cost is $-c_{=} \cdot n$. Otherwise, we can match at most $n - 1$ nodes, making the total cost at least $-c_{=} \cdot n + c_{=}$, see Figure 5 (right). Furthermore, note that the total cost of matching the left tree of $E(u, v, c_{=})$ with any tree is at least $-c_{=} \cdot n$ and the cost of matching any node of $E(u, v, c_{=})$ is $-c_{=}$.

Connection gadget $C(i, j, M)$. For any $i, j \in \{1, \dots, N\}$ and sufficiently large $M \in \{0, \dots, n^d - 1\}$, the cost of matching the left and right tree of $C(i, j, M)$ is $-M - W(i, j)$. The left tree does not depend on j and the right tree does not depend on i .

Let $\{u_1, \dots, u_{k/3}\}$ and $\{v_1, \dots, v_{k/3}\}$ be the $k/3$ -cliques corresponding to i and j , respectively, where $u_1 < u_2 < \dots, u_{k/3}$ and $v_1 < v_2 < \dots, v_{k/3}$. Recall that $W(i, j)$ denotes the total weight of all edges connecting two nodes in the i -th clique or a node in the i -th clique with a node in the j -th clique, where we assume that $w(u, u) = 0$. We construct the gadget $C(i, j, M)$ as follows. The root of the left tree has degree $1 + k/3$ and the root of the right tree has degree $1 + n$. Their rightmost children correspond to the root of the left and the right trees of $D(\sum_{x < y} w(u_x, u_y))$, respectively. Every other child of the left root can be matched with every other child of the right root with cost $-M_2$ (we fix M_1 and M_2 later). Intuitively, we would like the x -th child of the the left root to be matched with the u_x -th child of the right root, and then contribute $-\sum_y w(u_x, v_y)$ to the total cost, so that summing up over $x = 1, 2, \dots, k/3$ we obtain the desired sum. To this end, we attach the left tree of $E(u_x, \cdot, M_1)$ and the right tree of $D(\cdot)$ to the x -th child of the left root. Similarly, we attach the right tree of $E(\cdot, t, M_1)$ and the left tree of $D(\sum_y w(t, v_y))$ to the t -th child of the right root. Here we use \cdot to emphasise that a particular tree does not depend on the particular value of the parameter. All decrease gadgets are of the same type. See Figure 6.

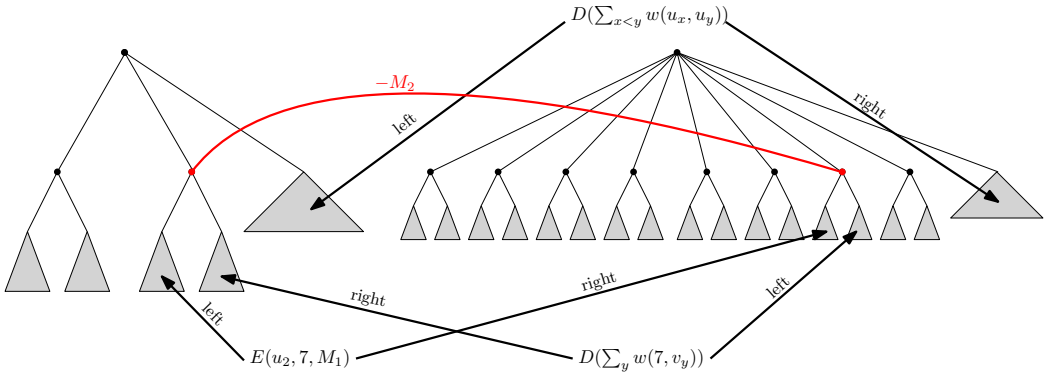


Fig. 6. Schematic illustration of a connection gadget for $k/3 = 2$ and $n = 8$.

We can clearly construct a solution with total cost $-M_2 \cdot k/3 - M_1 \cdot n \cdot k/3 - W(i, j)$ (because we have enumerated the clique corresponding to i so that $u_1 < u_2 < \dots < u_{k/3}$). We claim that, for appropriately chosen M_1 and M_2 , no better solution is possible. Let $W = \sum_{u, v} w(u, v)$. We fix $M_1 = W \cdot (k/3 + 1)$. This is enough to guarantee that the total cost contributed by nodes in all decrease gadgets is at least $-M_1$. The total cost contributed by nodes in all equality gadgets is at least $-M_1 \cdot n \cdot k/3$. Consequently, setting $M_2 = M_1 \cdot n \cdot k/3 + M_1$ guarantees that any optimal

solution must match all children of the left root, so in fact, for every $x = 1, 2, \dots, k/3$ we must match the x -th child of the left root to some child of the right root. Because matching the left tree of any decrease gadget contributes at least $-W$ to the total cost, by the choice of M_1 an optimal solution in fact must match the x -th child of the left root with the u_x -th child of the right root, as otherwise we lose at least M_1 due to the corresponding equality gadget that cannot be compensated by matching its accompanying decrease gadget. Finally, the corresponding decrease gadget adds $-\sum_y w(u_x, v_y)$ to the total cost. Therefore, as long as $M \geq M_2 \cdot k/3 + M_1 \cdot n \cdot k/3$ the total cost is indeed $-M - W(i, j)$ after choosing the cost of matching the roots to be $-M + M_2 \cdot k/3 + M_1 \cdot n \cdot k/3$. For any node in a decrease gadget, the cost of matching is at least $-W$, for any node in an equality gadget, the cost of matching is $-M_1$, and finally the cost of matching the children of the roots is $-M_2$, so the cost of matching any node of $C(i, j, W)$ is at least $-M$. For the correctness of the construction it is enough that M is at least

$$\begin{aligned} M_2 \cdot k/3 + M_1 \cdot n \cdot k/3 &= M_1((nk/3 + 1)k/3 + nk/3) \\ &= W(k/3 + 1)k/3(nk/3 + 1 + n) \\ &= W(k/3 + 1)k/3(n(k/3 + 1) + 1) \\ &\leq W \cdot n(k/3 + 1)^3 = n^{O(ck)}. \end{aligned}$$

Micro structures $A'_i, D_{z'}, B_z, C'_j$. We only explain how to construct A'_i and $D_{z'}$, for any $i = 1, 2, \dots, N$ and $z' = 1, 2, \dots, N$, as the construction of B_z and C'_j is symmetric. Recall that we require $c_{\text{match}}(A'_i, D_{z'}) = -M^6 - M^3(N - i) - W(i, z')$ and for every node in A'_i and $D_{z'}$ the cost of matching should be at least $-M^6$.

A'_i consists of a root to which we attach the left tree of $D(M^6 + M^3(N - i) - M)$ and the left tree of $C(i, \cdot, M)$, while $D_{z'}$ consists of a root to which we attach the right tree of $D(\cdot)$ and the right tree of $C(\cdot, z', M)$. All decrease gadgets attached as the left children of A'_i and $D_{z'}$ are of the same type, and all decrease gadgets used inside the connection gadgets attached as the right children are also of the same but other type. This guarantees that the total cost of matching A'_i to $D_{z'}$ is simply $-M^6 - M^3(N - i) + M - M - W(i, j) = -M^6 - M^3(N - i) - W(i, j)$. For sufficiently large $M \geq W \cdot n(k/3 + 1)^3$, the cost of matching any node in $D(M^6 + M^3(N - i) - M)$ is at least $-M^6$ and the cost of matching any node in $C(i, j, M)$ is at least $-M$.

Micro structures A_i, C_j . Here the situation is a bit more complex, as we simultaneously require that $c_{\text{match}}(A_i, C_j) = -M^2 - 2M(k/3)^2 - W(j, i) + W(j - 1, i - 1)$ for every $i = 2, 3, \dots, N$ and $j = 2, 3, \dots, N$ and $c_{\text{match}}(A_i, C_1) = -M^5 - M^3(i - 1) - W(1, i)$, and $c_{\text{match}}(A_1, C_j) = -M^5 - M^3(j - 1) - W(j, 1)$ for every $i = 1, 2, \dots, N$ and $j = 1, 2, \dots, N$. We must also make sure that the cost of matching a node of A_i to a node of C_j should be at least $-M^2$, except that the root of A_i (C_j) can be matched to the root of C_1 (A_1) with cost larger than $-M^5 - M$ but at most $-M^5$ and, for any non-root node of A_i (C_j) and for any non-root node of C_1 (A_1), the cost of matching is larger than $-M^4$.

For every $i > 1$ ($j > 1$), A_i (C_j) consists of two subtrees, called left and right, attached to the common root, while A_1 (C_1) consists of a single subtree connected to a root. For every $i > 1$ ($j > 1$), the left subtree of A_i (the right subtree of C_j) consists of a root with two subtrees, called left-right and left-right (right-left and right-right). For every $i > 1$, nodes of the right subtree of A_i can only be matched to nodes of C_1 and nodes of the left subtree of A_i can only be matched to nodes of the right subtree of C_j for any $j > 1$. For every $j > 1$, nodes of the left subtree of C_j can only be matched to nodes of A_1 and nodes of the right subtree of C_j can only be matched to nodes of the left subtree of A_i for any $i > 1$. Nodes of A_1 can be matched to nodes of the left subtree of C_j , for any $j > 1$. Nodes of C_1 can be matched to nodes of the right subtree of A_i , for any $i > 1$. Additionally, the root of A_1 can be matched to the root of C_1 with cost $-M^5 - W(1, 1) > -M^5 - M$, and for any $i > 1$

($j > 1$), the root of A_i (C_j) can be matched to the root of C_1 (A_1) with cost $-M^5$. The subtrees are constructed as follows:

1. the right subtree of A_i is the left tree of $D(M^3(i-1) + W(1, i))$,
2. the only subtree of A_1 is the right tree of $D(\cdot)$,
3. the left subtree of C_j is the left tree of $D(M^3(j-1) + W(j, 1))$,
4. the only subtree of C_1 is the right tree of $D(\cdot)$.

It remains to fully define the left subtree of every A_i and the right subtree of every C_j , for $i, j > 1$. Recall that the goal is to ensure $c_{\text{match}}(A_i, C_j) = -M^2 - 2M(k/3)^2 - W(j, i) + W(j-1, i-1)$. We define a new n -node graph with weight function $w'(u, v) = M - w(u, v)$ for any u, v (for sufficiently large M , the new weights are positive). Then, let $C'(i, j, M)$ denote the connection gadget $C(i, j, M)$ constructed for the new graph. Nodes of the left-left (left-right) subtree of A_i can be only matched to nodes of the right-left (right-right) subtree of C_j . The subtrees are constructed as follows:

1. the left-left subtree of A_i is the left tree of $C(i, \cdot, M \cdot (k/3)^2)$,
2. the right-left subtree of C_j is the right tree of $C(\cdot, j, M \cdot (k/3)^2)$,
3. the left-right subtree of A_i is the left tree of $C'(i-1, \cdot, M^2)$,
4. the right-right subtree of C_j is the right tree of $C'(\cdot, j-1, M^2)$.

See Figure 7. For the construction of $C(i, j, M \cdot (k/3)^2)$ and $C(i-1, j-1, M^2)$ to be correct we need that $M \cdot (k/3)^2 \geq W \cdot n(k/3 + 1)^3$ and $M^2 \geq M \cdot n^2 \cdot n(k/3 + 1)^3$, respectively, which holds for sufficiently large M .

Now we calculate $c_{\text{match}}(A_i, C_j)$. $C(i-1, j-1, M^2)$ contributes $-M^2$ minus the total cost of edges connecting two subsets of $k/3$ nodes in the new graph. As the weights in the new graph are defined as $w'(u, v) = M - w(u, v)$, this is exactly $-M^2 - (M \cdot (k/3)^2 - W(i-1, j-1))$. $C(i, j, M \cdot (k/3)^2)$ contributes $-M \cdot (k/3)^2 - W(i, j)$, so $c_{\text{match}}(A_i, C_j) = -M \cdot (k/3)^2 - W(i, j) - M^2 - (M \cdot (k/3)^2 - W(i-1, j-1)) = -M^2 - 2M(k/3)^2 - W(i, j) + W(i-1, j-1)$ as required.

It remains to bound the cost of matching nodes. Nodes in the left subtree of A_i (C_j) can be matched only to nodes of C_1 (A_1) with cost at least $-M^3 \cdot n > -M^4$, except that the roots can be matched with cost $-M^5$. The cost of matching a node of A_i to a node of C_j , for $i, j > 1$, is either at least $-M \cdot (k/3)^2$ (for the nodes of $C(i, j, M \cdot (k/3)^2)$) or at least $-M^2$ (for the nodes of $C(i, j, M^2)$), so for sufficiently large M at least $-M^2$.

Wrapping up. We have shown how to construct, given a complete undirected n -node graph G , two trees such that the weight of the max-weight k -clique in G can be extracted from the cost of an optimal matching (and, as mentioned in the beginning of Section 2, by a simple transformation this is equal to the edit distance). To complete the proof of Lemma 3, we need to bound the size of both trees and also the size of the alphabet used to label their nodes.

Initially, each tree consists of $2N$ original spine nodes, where $N \leq n^{k/3}$, $2N$ leaf nodes, and N additional spine nodes. Then, we attach appropriate microstructures to the original spine nodes and leaf nodes. The microstructures are $A'_i, I, C'_j, A_i, B_z, C_j, D_{z'}$. Every copy of I consists of $k/3 \cdot n$ nodes. To analyse the size of the remaining microstructures, first note that if $x \in \{0, \dots, n^d\}$ then the decrease gadget $D(x)$ consists of $O(d \cdot n)$ nodes. The equality gadget always consists of $O(n)$ nodes. Finally, the connection gadget $E(\cdot, \cdot, M)$ with $M \in \{0, \dots, n^d\}$ consists of $O(n(n + d \cdot n) + d \cdot n) = O(d \cdot n^2)$ nodes. Let $M = n^d$ with d to be specified later. Now, the size of the microstructures can be bounded as follows: A'_i and $D_{z'}$ (and also B_z and C'_j) consist of $O(6d \cdot n + d \cdot n^2) = O(d \cdot n^2)$ nodes. The right subtree of A_i (and the left subtree of C_j) consists of $O(3d \cdot n)$ nodes, while the left subtree of A_i (and the right subtree of C_j) consist of $O(k^2 d \cdot n^2 + 2d \cdot n^2) = O(k^2 d \cdot n^2)$ nodes. Thus, the total size of all microstructures is $O(N \cdot k^2 d \cdot n^2)$. It remains to bound M . Recall that we require $M \geq W \cdot n(k/3 + 1)^3$, $M \cdot (k/3)^2 \geq W \cdot n(k/3 + 1)^3$ and $M \geq n^3(k/3 + 1)^3$, where

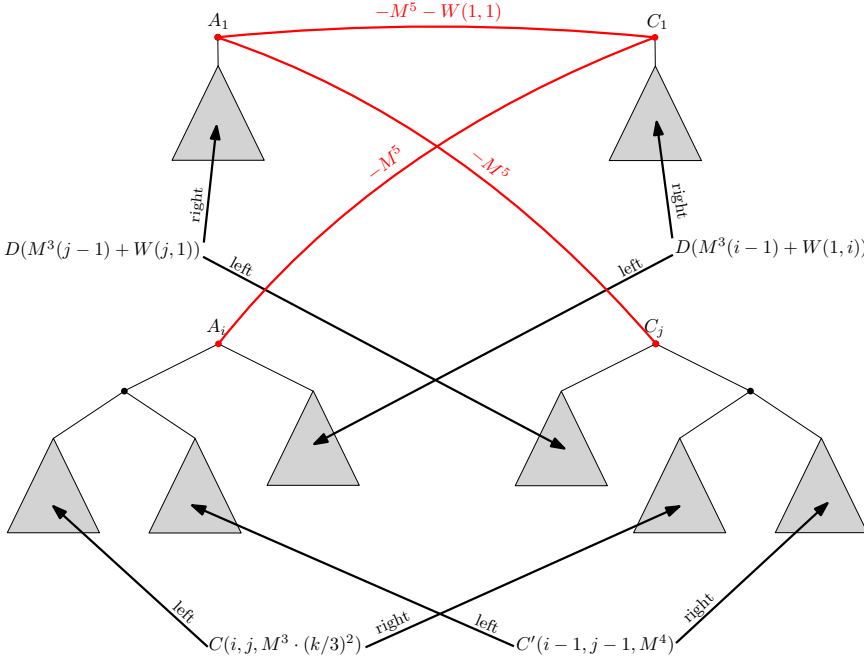


Fig. 7. Micro structures A_1, C_1 and A_i, C_j for $i, j > 1$.

$W \leq n^2 \cdot n^{O(ck)} = n^{O(ck)}$. Hence, it is sufficient that $M \geq 8W \cdot n^3 k^3$. Setting $d = \Theta(ck)$ is therefore enough. The size of the whole instance thus is $O(n^{k/3+2} \cdot ck) = O(n^{k/3+2})$.

We also have to bound the size of the alphabet. We need $k/3$ distinct labels for the nodes of I . We need $O(d)$ distinct labels for the nodes of all decrease gadgets of the same type. There is a constant number of types, and all other nodes require only a constant number of distinct labels (irrespectively on c and k), so the total size of the alphabet is $O(ck) = O(1)$.

4 ALGORITHM FOR CATERpillARS ON SMALL ALPHABET

In this section, we show that the hard instances of TED from Section 2 can be solved in time $O(n^2 |\Sigma|^2 \log n)$, where n is the size of the trees and Σ is the alphabet. Recall that in such an instance we are given two trees F and G both consisting of a single path (called spine) of length $O(n)$ with a single leaf pending from every node, and all these leaves are to the right of the path in F and to the left of the path in G (see Figure 2). In the following we use the same notation as in Lemma 1. At a high level, we want to guess the rootmost non-spine node in the left tree $f'_{i_{p+1}}$ and the rootmost non-spine node in the right tree $g'_{j_{p+1}}$. The optimal matching of spine nodes above these non-spine nodes can be precomputed in $O(n^2)$ total time with a simple dynamic programming algorithm. It might be tempting to say the same about the situation below, but this is much more complicated due to the fact that leaf nodes in this part are matched in reversed order. To overcome this difficulty, we need the following tool.

LEMMA 6. For strings $s[1..n]$ and $t[1..m]$ over alphabet Σ and matching cost $c_{\text{match}}(c, d)$ for any two letters $c, d \in \Sigma$, we define the optimal matching of s and the reverse of t as

$$\min \left\{ \sum_{\ell=1}^k c_{\text{match}}(s[i_\ell], t[j_\ell]) \mid k \geq 0, 1 \leq i_1 < \dots < i_k \leq n, 1 \leq j_k < \dots < j_1 \leq m \right\}.$$

Given two strings $s[1..n], t[1..n]$, in $O(n^2 \log n)$ total time we can calculate, for every i and j , the optimal matching of $s[1..i]$ and the reverse of $t[1..j]$.

PROOF. We construct an $(n+1) \times (n+1)$ grid graph on nodes $v_{i,j}$, where $i, j = 0, 1, \dots, n$ as follows. For every $i, j = 0, 1, \dots, n$, we create a directed edge from $v_{i,j}$ to $v_{i+1,j}$ and $v_{i,j+1}$ with length zero. Also, we create a directed edge from $v_{i,j}$ to $v_{i+1,j+1}$ with length $c_{\text{match}}(s[i], t[n+1-j])$. Then, paths from $v_{1,n+1-j}$ to $v_{i,n+1}$ are in one-to-one correspondence with matchings of $s[1..i]$ to the reverse of $t[1..j]$. Therefore, the cheapest such path corresponds to the optimal matching.

The grid is a planar directed graph, and all starting nodes $v_{1,n+1-j}$ lie on the external face, so we can use the multiple-source shortest paths algorithm of Klein [42] to compute, in $O(n^2 \log n)$ time, a representation of shortest paths from all starting nodes $v_{1,n+1-j}$ to all nodes of the grid.³ This representation can be then queried in $O(\log n)$ time to extract the length of any path from $v_{1,n+1-j}$ to $v_{i,n+1}$. Thus, the total time is $O(n^2 \log n)$. \square

To see how Lemma 6 can be helpful, consider the (simpler) case when there are no additional spine nodes $f_{i_{p+q+2}}$ and $g_{j_{p+q+2}}$. We construct two strings s and t by writing down the labels of leaf nodes $f'_n, f'_{n-1}, \dots, f'_1$ and $g'_n, g'_{n-1}, \dots, g'_1$, respectively, and preprocess them using Lemma 6. Then, to find the optimal matching we guess i_{p+2} and j_{p+2} . As explained above, optimal matching of spine nodes above $f'_{i_{p+2}}$ and $g'_{j_{p+2}}$ can be precomputed in $O(n^2)$ time in advance. Then, we need to match some of the leaf nodes $f'_{i_{p+2}}, f'_{i_{p+2}+1}, f'_{i_{p+2}+2}, \dots$ to some of the leaf nodes $g'_{j_{p+2}}, g'_{j_{p+2}+1}, g'_{j_{p+2}+2}, \dots$ in the reversed order. This exactly corresponds to matching $s[1, n+1-i_{p+2}]$ to the reverse of $t[1, n+1-j_{p+2}]$ and thus is also precomputed. Iterating over all possible i_{p+2} and j_{p+2} gives us the optimal matching in $O(n^2 \log n)$ total time.

Now consider the general case. We assume that both optional spine nodes $f_{i_{p+q+2}}$ and $g_{j_{p+q+2}}$ exist; if only one of them is present the algorithm is very similar. As in the simpler case, we iterate over all possible i_{p+1} and j_{p+1} . The natural next step would be to iterate over all possible i_{p+q+2} and j_{p+q+2} , but this is too expensive. However, because no spine nor leaf nodes below $f_{i_{p+q+2}}$ (or $g_{j_{p+q+2}}$) are matched, we can as well replace $f_{i_{p+q+2}}$ with the lowest spine node with the same label (and similarly for $g_{j_{p+q+2}}$). Thus, instead of guessing i_{p+q+2} we can guess the label of $f_{i_{p+q+2}}$ and choose the lowest spine node with such label (and similarly for j_{p+q+2}). Now we retrieve the precomputed optimal matching of spine nodes above $f'_{i_{p+1}}$ and $g'_{j_{p+1}}$. Then we need to find the optimal matching of leaf nodes $f'_{i_{p+1}+1}, f'_{i_{p+1}+2}, \dots, f'_{i_{p+q+2}-1}$ and $g'_{i_{p+1}+1}, g'_{i_{p+1}+2}, \dots, g'_{j_{p+q+2}-1}$. This can be precomputed in $O(n^2 |\Sigma|^2 \log n)$ time with Lemma 6. Indeed, there are only $|\Sigma|$ possibilities for $i_{p+q+2} - 1$ and also $|\Sigma|$ possibilities for $j_{p+q+2} - 1$, as both of them are defined by the lowest occurrence of a label among the spine nodes of the left and the right tree, respectively. For each such combination, we construct two strings s and t by writing down the labels of leaf nodes above $f'_{i_{p+q+2}}$ and $g'_{j_{p+q+2}}$ in the bottom-up order and preprocess them in $O(n^2 \log n)$ time. This allows us to retrieve the optimal matching of leaf nodes and then we only have to add $c_{\text{match}}(f'_{i_{p+1}}, g_{j_{p+q+2}})$ and $c_{\text{match}}(f_{i_{p+q+2}}, g'_{j_{p+1}})$ to

³In the presence of negative-length edges, Klein's algorithm requires an initial shortest paths tree from some node on the external face to all other nodes. In our case, computing this initial shortest path tree can easily be done in $O(n^2)$ time as our graph is a directed acyclic graph.

obtain the total cost. Thus, after $O(n^2|\Sigma|^2 \log n)$ preprocessing, we can find the optimal matching by iterating over $n^2|\Sigma|^2$ possibilities and checking each of them in constant time.

ACKNOWLEDGMENTS

The third and fourth authors are partially supported by the Israel Science Foundation grants 794/13 and 592/17.

REFERENCES

- [1] Amir Abboud. 2014. Hardness for Easy Problems. In *YR-ICALP Satellite Workshop of ICALP*.
- [2] Amir Abboud, Arturs Backurs, Thomas Dueholm Hansen, Virginia Vassilevska Williams, and Or Zamir. 2016. Subtree Isomorphism Revisited. In *SODA*. 1256–1271.
- [3] Amir Abboud, Arturs Backurs, and Virginia Vassilevska Williams. 2015. If the Current Clique Algorithms are Optimal, So is Valiant’s Parser. In *FOCS*. 98–117.
- [4] Amir Abboud, Arturs Backurs, and Virginia Vassilevska Williams. 2015. Tight Hardness Results for LCS and Other Sequence Similarity Measures. In *FOCS*. 59–78.
- [5] Amir Abboud and Søren Dahlgaard. 2016. Popular Conjectures as a Barrier for Dynamic Planar Graph Algorithms. In *FOCS*. 476–486.
- [6] Amir Abboud, Fabrizio Grandoni, and Virginia Vassilevska Williams. 2015. Subcubic Equivalences Between Graph Centrality Problems, APSP and Diameter. In *SODA*. 1681–1697.
- [7] Amir Abboud, Thomas Dueholm Hansen, Virginia Vassilevska Williams, and Ryan Williams. 2016. Simulating branching programs with edit distance and friends or: a polylog shaved is a lower bound made. In *STOC*. 375–388.
- [8] Amir Abboud and Kevin Lewi. 2013. Exact Weight Subgraphs and the k -SUM Conjecture. In *ICALP*. 1–12.
- [9] Amir Abboud and Virginia Vassilevska Williams. 2014. Popular Conjectures Imply Strong Lower Bounds for Dynamic Problems. In *FOCS*. 434–443.
- [10] Amir Abboud, Virginia Vassilevska Williams, and Oren Weimann. 2014. Consequences of Faster Alignment of Sequences. In *ICALP*. 39–51.
- [11] Amir Abboud, Virginia Vassilevska Williams, and Huacheng Yu. 2015. Matching Triangles and Basing Hardness on an Extremely Popular Conjecture. In *STOC*. 41–50.
- [12] Tatsuya Akutsu, Daiji Fukagawa, and Atsuhiko Takasu. 2006. Approximating Tree Edit Distance Through String Edit Distance. In *ISAAC*, Vol. 4288. 90–99.
- [13] Noga Alon, Zvi Galil, and Oded Margalit. 1997. On the Exponent of the All Pairs Shortest Path Problem. *JCSS* 54, 2 (1997), 255–262.
- [14] Rajeev Alur, Loris D’Antoni, Sumit Gulwani, Dileep Kini, and Mahesh Viswanathan. 2013. Automated Grading of DFA Constructions. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence (IJCAI ’13)*. 1976–1982. <http://dl.acm.org/citation.cfm?id=2540128.2540412>
- [15] Amihoud Amir, Timothy M. Chan, Moshe Lewenstein, and Noa Lewenstein. 2014. On Hardness of Jumbled Indexing. In *ICALP*. 114–125.
- [16] Alberto Apostolico and Zvi Galil (Eds.). 1997. *Pattern matching algorithms*. Oxford University Press, Oxford, UK.
- [17] Taku Aratsu, Kouichi Hirata, and Tetsuji Kuboyama. 2010. Approximating Tree Edit Distance through String Edit Distance for Binary Tree Codes. *Fundam. Inform.* 101, 3 (2010), 157–171.
- [18] Arturs Backurs, Nishanth Dikkala, and Christos Tzamos. 2016. Tight Hardness Results for Maximum Weight Rectangles. In *ICALP*. 81:1–81:13.
- [19] Arturs Backurs and Piotr Indyk. 2015. Edit Distance Cannot Be Computed in Strongly Subquadratic Time (unless SETH is false). In *STOC*. 51–58.
- [20] Arturs Backurs and Piotr Indyk. 2016. Which Regular Expression Patterns are Hard to Match?. In *FOCS*. 457–466.
- [21] Arturs Backurs and Christos Tzamos. 2017. Improving Viterbi is Hard: Better Runtimes Imply Faster Clique Algorithms. In *ICML*. 311–321.
- [22] John Bellando and Ravi Kothari. 1999. Region-based modeling and tree edit distance as a basis for gesture recognition. In *ICAP*. 698–703.
- [23] Richard E. Bellman. 1957. *Dynamic programming*. Princeton University Press.
- [24] Philip Bille. 2005. A survey on tree edit distance and related problems. *Theoretical Computer Science* 337, 1-3 (2005), 217–239.
- [25] Karl Bringmann. 2014. Why walking the dog takes time: Frechet distance has no strongly subquadratic algorithms unless SETH fails. In *FOCS*. 661–670.
- [26] Karl Bringmann. 2017. A near-linear pseudopolynomial time algorithm for subset sum. In *SODA*. 1073–1084.

- [27] Karl Bringmann, Fabrizio Grandoni, Barna Saha, and Virginia Vassilevska Williams. 2016. Truly Sub-cubic Algorithms for Language Edit Distance and RNA-Folding via Fast Bounded-Difference Min-Plus Product. In *57th FOCS*. 375–384.
- [28] Karl Bringmann, Allan Grønlund, and Kasper Green Larsen. 2016. A Dichotomy for Regular Expression Membership Testing. *CoRR* abs/1611.00918 (2016).
- [29] Karl Bringmann and Marvin Künnemann. 2015. Quadratic Conditional Lower Bounds for String Problems and Dynamic Time Warping. In *FOCS*. 79–97.
- [30] Peter Buneman, Martin Grohe, and Christoph Koch. 2003. Path queries on compressed XML. In *VLDB*. 141–152.
- [31] Sudarshan S. Chawathe. 1999. Comparing hierarchical data in external memory. In *VLDB*. 90–101.
- [32] Raphaël Clifford. 2009. Matrix multiplication and pattern matching under Hamming norm. <http://www.cs.bris.ac.uk/Research/Algorithms/events/BAD09/BAD09/Talks/BAD09-Hammingnotes.pdf>.
- [33] Mark de Berg, Kevin Buchin, Bart M. P. Jansen, and Gerhard Woeginger. 2016. Fine-Grained Complexity Analysis of Two Classic TSP Variants. In *ICALP*, Vol. 55. 5:1–5:14.
- [34] Erik Demaine, Shay Mozes, Benjamin Rossman, and Oren Weimann. 2009. An Optimal Decomposition Algorithm for Tree Edit Distance. *ACM Transactions on Algorithms* 6, 1 (2009), 1–19. Preliminary version in *ICALP* 2007.
- [35] Friedrich Eisenbrand and Fabrizio Grandoni. 2004. On the complexity of fixed parameter clique and dominating set. *Theoretical Computer Science* 326, 1-3 (2004), 57–67.
- [36] Paolo Ferragina, Fabrizio Luccio, Giovanni Manzini, and S. Muthukrishnan. 2009. Compressing and indexing labeled trees, with applications. *J. ACM* 57 (2009), 1–33. Issue 1.
- [37] Dan Gusfield. 1997. *Algorithms on strings, trees and sequences: computer science and computational biology*. Cambridge University Press.
- [38] Christoph M. Hoffmann and Michael J. O’Donnell. 1982. Pattern Matching in Trees. *J. ACM* 29, 1 (1982), 68–95.
- [39] Egor Ivkin. 2012. *Comparison of Tree Edit Distance Algorithms*. B.Sc. thesis. Charles University in Prague.
- [40] Richard M. Karp. 1972. Reducibility Among Combinatorial Problems. In *Complexity of Computer Computations*. 85–103.
- [41] Philip N. Klein. 1998. Computing the Edit-Distance between Unrooted Ordered Trees. In *ESA*. 91–102.
- [42] Philip N. Klein. 2005. Multiple-source shortest paths in planar graphs. In *SODA*. 146–155.
- [43] Philip N. Klein, Srikanta Tirhapura, Daniel Sharvit, and Benjamin B. Kimia. 2000. A tree-edit-distance algorithm for comparing simple, closed shapes. In *SODA*. 696–704.
- [44] Donald E. Knuth. 1971. Optimum binary search trees. *Acta Informatica* 1, 1 (1971), 14–25.
- [45] Kasper Green Larsen, J. Ian Munro, Jesper Sindahl Nielsen, and Sharma V. Thankachan. 2015. On hardness of several string indexing problems. *Theoretical Computer Science* 582 (2015), 74–82.
- [46] Miscellaneous Authors. 1984. Queries and problems. *SIGACT News* 16, 3 (1984), 38–47.
- [47] Jaroslav Nešetřil and Svatopluk Poljak. 1985. On the complexity of the subgraph problem. *Commentationes Math. Universitatis Carolinae* 026, 2 (1985), 415–419.
- [48] Mateusz Pawlik and Nikolaus Augsten. 2015. Efficient Computation of the Tree Edit Distance. *ACM Trans. Database Syst.* 40, 1, Article 3 (March 2015), 40 pages. <https://doi.org/10.1145/2699485>
- [49] Juan Ramón Rico-Juan and Luisa Micó. 2003. Comparison of AESA and LAESA search algorithms using string and tree-edit-distances. *Pattern Recognition Letters* 24, 9-10 (2003), 1417–1426. [https://doi.org/10.1016/S0167-8655\(02\)00382-3](https://doi.org/10.1016/S0167-8655(02)00382-3)
- [50] Liam Roditty and Uri Zwick. 2011. On Dynamic Shortest Paths Problems. *Algorithmica* 61, 2 (2011), 389–401.
- [51] Stanley M. Selkow. 1977. The tree-to-tree editing problem. *Inform. Process. Lett.* 6, 6 (1977), 184–186.
- [52] Bruce A. Shapiro and Kaizhong Zhang. 1990. Comparing multiple RNA secondary structures using tree comparisons. *Computer Applications in the Biosciences* 6, 4 (1990), 309–318. <https://doi.org/10.1093/bioinformatics/6.4.309>
- [53] Dennis E. Shasha and Kaizhong Zhang. 1989. Simple Fast Algorithms for the Editing Distance Between Trees and Related Problems. *SIAM J. Comput.* 18, 6 (1989), 1245–1262.
- [54] Dennis E. Shasha and Kaizhong Zhang. 1990. Fast algorithms for the unit cost editing distance between trees. *Journal of Algorithms* 11, 4 (1990), 581–621.
- [55] Kuo-Chung Tai. 1979. The Tree-to-Tree Correction Problem. *J. ACM* 26, 3 (1979), 422–433.
- [56] Gabriel Valiente. 2002. *Algorithms on Trees and Graphs*. Springer-Verlag.
- [57] Robert A. Wagner and Michael J. Fischer. 1974. The String-to-String Correction Problem. *J. ACM* 21, 1 (1974), 168–173.
- [58] Michael S. Waterman. 1995. *Introduction to computational biology: maps, sequences and genomes*, Chapters 13, 14. Chapman and Hall.
- [59] Ryan Williams. 2005. A new algorithm for optimal 2-constraint satisfaction and its implications. *Theor. Comput. Sci.* 348, 2-3 (2005), 357–365.
- [60] Ryan Williams. 2014. Faster all-pairs shortest paths via circuit complexity. In *STOC*. 664–673.
- [61] Virginia Vassilevska Williams and Ryan Williams. 2010. Subcubic equivalences between path, matrix and triangle problems. In *FOCS*. 645–654.
- [62] Virginia Vassilevska Williams and Ryan Williams. 2013. Finding, Minimizing, and Counting Weighted Subgraphs. *SIAM J. Comput.* 42, 3 (2013), 831–854.

- [63] Gerhard J. Woeginger. 2004. Space and Time Complexity of Exact Algorithms: Some Open Problems. In *IWPEC*. 281–290.
- [64] Gerhard J. Woeginger. 2008. Open problems around exact algorithms. *Discr. Appl. Math.* 156, 3 (2008), 397–405.
- [65] Xuchen Yao, Benjamin Van Durme, Chris Callison-Burch, and Peter Clark. 2013. Answer Extraction as Sequence Tagging with Tree Edit Distance. In *HLT-NAACL 2013*. 858–867. <http://aclweb.org/anthology/N/N13/N13-1106.pdf>
- [66] Kaizhong Zhang. 1995. Algorithms for the constrained editing distance between ordered labeled trees and related problems. *Pattern Recognition* 28, 3 (1995), 463–474.