



An Improved Algorithm for The k -Dyck Edit Distance Problem

DVIR FRIED, Bar-Ilan University, Israel

SHAY GOLAN, Reichman University, Israel and University of Haifa, Israel

TOMASZ KOCIUMAKA, Max Planck Institute for Informatics, Germany

TSVI KOPELOWITZ and ELY PORAT, Bar-Ilan University, Israel

TATIANA STARIKOVSKAYA, DIENS, École normale supérieure, PSL Research University, France

A Dyck sequence is a sequence of opening and closing parentheses (of various types) that is balanced. The Dyck edit distance of a given sequence of parentheses S is the smallest number of edit operations (insertions, deletions, and substitutions) needed to transform S into a Dyck sequence. We consider the threshold Dyck edit distance problem, where the input is a sequence of parentheses S and a positive integer k , and the goal is to compute the Dyck edit distance of S only if the distance is at most k , and otherwise report that the distance is larger than k . Backurs and Onak [PODS'16] showed that the threshold Dyck edit distance problem can be solved in $O(n + k^{16})$ time.

In this work, we design new algorithms for the threshold Dyck edit distance problem which costs $O(n + k^{4.544184})$ time with high probability or $O(n + k^{4.853059})$ deterministically. Our algorithms combine several new structural properties of the Dyck edit distance problem, a refined algorithm for fast (min, +) matrix product, and a careful modification of ideas used in Valiant's parsing algorithm.

CCS Concepts: • **Theory of computation** → **Design and analysis of algorithms**;

Additional Key Words and Phrases: Dyck language, edit distance, fine-grained complexity

ACM Reference format:

Dvir Fried, Shay Golan, Tomasz Kociumaka, Tsvi Kopelowitz, Ely Porat, and Tatiana Starikovskaya. 2024. An Improved Algorithm for The k -Dyck Edit Distance Problem. *ACM Trans. Algor.* 20, 3, Article 26 (June 2024), 25 pages.

<https://doi.org/10.1145/3627539>

1 INTRODUCTION

The notion of balanced sequences of parentheses is a fundamental concept in the theory of computer science. Formally, a *Dyck sequence* is a sequence of opening and closing parentheses (of

S. Golan the work was done while the author was at Bar-Ilan University and the University of California, Berkeley.

T. Kociumaka the work was done while the author was at the University of California, Berkeley.

Authors' addresses: D. Fried, T. Kopelowitz, and E. Porat, Bar-Ilan University, Ramat Gan 5290002, Israel; e-mails: {friedvir1, kopelot}@gmail.com, porately@cs.biu.ac.il; S. Golan, Reichman University, Herzliya 4610101, Israel and University of Haifa, Haifa 3498838, Israel; e-mail: golansh1@biu.ac.il; T. Kociumaka, Max Planck Institute for Informatics, Saarbrücken 66123, Germany; e-mail: kociumaka@mimuw.edu.pl; T. Starikovskaya, DIENS, École normale supérieure, PSL Research University, 45 rue d'Ulm, Paris F-75230, France; e-mail: tat.starikovskaya@gmail.com.



This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2024 Copyright held by the owner/author(s).

1549-6325/2024/06-ART26

<https://doi.org/10.1145/3627539>

various types) that is balanced, and $\text{Dyck}(t)$ is the language containing all balanced sequences of parentheses over $\Sigma = \{(i, \cdot)_i \mid i \in [0..t-1]\}$. The $\text{Dyck}(t)$ language is a context-free language that is generated by the following grammar: $S \rightarrow \varepsilon$, where ε is the empty string, $S \rightarrow ({}_0S)_0S$, $S \rightarrow ({}_1S)_1S, \dots, S \rightarrow ({}_{t-1}S)_{t-1}S$.

The Dyck language has proven to be instrumental in the theory of context-free languages; in particular, Chomsky and Schotzenberger [9] proved that every context-free language can be mapped to a subset of the Dyck language (see also [16, 18]). The Dyck language also plays a role in practice since balanced parentheses are used to succinctly describe arbitrary rooted trees [23], and many programming languages have a balanced parentheses-like structure. Additionally, data files, such as XML files, which store data in a structured way, often utilize a notion of balanced parentheses. Moreover, the Dyck language is strongly related to DNA/RNA sequences since these sequences slightly deviate from the balancedness requirement, so understanding the behavior of the Dyck language is often an important building block for designing algorithms on such sequences [15, 26].

Language Edit Distance. The **language edit distance (LED)** problem is one of the most fundamental problems in formal language theory, generalizing both parsing and string edit distance problems. Formally, given a language L and a string S , the task is to compute the minimal edit distance between S and the strings in L .

The LED problem has been extensively studied for context-free languages. Consider a context-free language generated by a grammar G and a string S of length n . Aho and Peterson [2] showed a dynamic-programming algorithm with runtime $O(|G|^2n^3)$. Myers [24] improved the running time of the algorithm to $O(|G|n^3)$. In a recent breakthrough article, Bringmann et al. [6] bypassed the n^3 barrier and demonstrated $O(|G|^{O(1)}n^{2.8244})$ -time randomized and $O(|G|^{O(1)}n^{2.8603})$ -time deterministic algorithms. The algorithms of [6] are non-combinatorial, in the sense that they use **fast matrix multiplication (FMM)**. Unfortunately, this seems to be unavoidable: the lower bound result of Lee [22] implies that there is no algorithm that solves LED in time less than that of Boolean matrix multiplication. Saha studied the problem of approximating LED, developing a $(1 + \varepsilon)$ -factor approximation algorithm [28] followed by a solution providing an additive approximation of LED [29]. In another related work, Jayaram and Saha [17] studied LED for the class of linear grammars.

Dyck Edit Distance. In this work, we focus on the LED problem for the Dyck language. Formally, the Dyck edit distance of a given sequence of parentheses $S = S[0]S[1]S[2] \cdots S[n-1] \in \Sigma^n$, denoted by $\text{ed}_D(S)$, is the smallest number of edit operations (insertions, deletions, and substitutions) needed to transform S into a Dyck sequence. Notice that it is enough to consider only deletions and substitutions since insertions can be transformed into deletions.

One might hope that the LED problem for the Dyck language might be easier than for general context-free languages, but as Abboud et al. [1] showed, this is probably not the case: they proved that an efficient algorithm for the Dyck edit distance problem implies an algorithm for k -Clique whose runtime is faster than what is believed to be obtainable. In other words, an efficient algorithm for the Dyck edit distance problem must either provide approximate answers or focus on **the small distance regime**. Saha [27] presented a randomized algorithm with running time $O(n^{1+o(1)})$ and a $\tilde{O}(1)$ multiplicative approximation factor. Very recently, Debarati, Kociumaka, and Saha [10] showed a constant factor approximation algorithm with runtime $\tilde{O}(n^{1.971})$ and a $(1 + \varepsilon)$ -approximation algorithm with runtime $\tilde{O}(n^2/\varepsilon)$. *In this work, we turn to the small distance regime.*

1.1 The k -Dyck Edit Distance Problem.

For a string $S \in \Sigma^n$, define $\text{ed}_D(S)$ to be the *edit cost* of S which is the smallest number of deletions and substitutions of a parenthesis such that the resulting string belongs to the Dyck language.

In this article, we consider the k -Dyck edit distance problem, where the input is a sequence of parentheses $S \in \Sigma^n$ and a positive integer k , and the goal is to compute $\min\{\text{ed}_D(S), k + 1\}$. In other words, we want to compute the Dyck edit distance of S , as long as the distance is at most k (and return $k + 1$ otherwise). This problem was first considered by Backurs and Onak [5], who gave an $O(n + k^{16})$ -time algorithm. Krebs et al. [19] presented a randomized streaming algorithm for the Dyck language with two types of parentheses with running time $\tilde{O}(nk^{O(1)})$ and space complexity $O(k^{1+\varepsilon} + \sqrt{n \log n})$, for any fixed constant $\varepsilon > 0$. Debarati, Kociumaka, and Saha [10] showed a $(3 + \varepsilon)$ -approximation algorithm for the k -Dyck edit distance problem with runtime $\tilde{O}(kn/\varepsilon)$.

$O(n^3)$ -time dynamic programming algorithm. We begin by describing a simple folklore dynamic-programming algorithm that solves the Dyck edit distance problem in $O(n^3)$ time. This algorithm forms the basis of previous results, as well as the basis of our new algorithms for the k -Dyck edit distance problem.

Consider a sequence $S \in \Sigma^n$ and let \mathcal{D} be a matrix of size $(n + 1) \times (n + 1)$ such that $\mathcal{D}[i, j] = \text{ed}_D(S[i..j])$ (where $S[i..j]$ is the substring of S from position i to position $j - 1$) to the Dyck language. By definition, $\text{ed}_D(S) = \mathcal{D}[0, n]$. Notice that $\mathcal{D}[i, i] = 0$ for $i \in [0..n]$, $\mathcal{D}[i, i + 1] = 1$ for $i \in [0..n)$, and $\mathcal{D}[i, j]$ satisfies the following recursion for $i, j \in [0..n]$ with $j - i \geq 2$ ¹

$$\mathcal{D}[i, j] = \min \begin{cases} \mathcal{D}[i, m] + \mathcal{D}[m, j] & \text{for } m \in (i..j), \\ \text{ed}_D(S[i]S[j - 1]) + \mathcal{D}[i + 1, j - 1]. \end{cases} \quad (1)$$

One can use Equation (1) to compute $\mathcal{D}[0, n]$ in $O(n^3)$ time.

$O(n + k^{16})$ -time algorithm by Backurs and Onak [5]. Backurs and Onak [5] solved the k -Dyck edit distance problem by first showing that the problem can be reduced to an instance on a string S' of length $n' \leq n$ such that: (1) $\text{ed}_D(S) = \text{ed}_D(S')$, and (2) when solving k -Dyck edit distance on S' , it suffices to compute only $O(k^{14})$ values of \mathcal{D} when applied to S' . Moreover, by modifying the algorithm of Landau, Myers and Schmidt [20], after $O(n)$ -time preprocessing, each of the $O(k^{14})$ calls to \mathcal{D} is computed recursively in $O(k^2)$ time.

Our results. We begin by designing an algorithm for k -Dyck edit distance which costs $O(n^2k)$ time, and is thus faster than the algorithm of Backurs and Onak [5] whenever $k = \Omega(n^{2/15})$. We further improve on the algorithm of Backurs and Onak [5] by designing a combinatorial algorithm for k -Dyck edit distance which runs in $O(n + k^5)$ time. Our final contribution, summarized in Theorem 1.1, improves upon this runtime using FMM.

THEOREM 1.1. *Given a sequence S of parentheses of length n , $\min\{\text{ed}_D(S), k + 1\}$ can be computed in $O(n + k^{4.853059})$ deterministically or in $O(n + k^{4.544184})$ time with high probability.*

If $k > \sqrt{n}$, our algorithm is actually faster than claimed in Theorem 1.1. The precise bound, which never exceeds $O(n^2k)$, is provided in Section 1.2; it depends on the complexity of rectangular matrix multiplication and involves optimization over auxiliary parameters of the algorithm.

1.2 Algorithmic Overview and Organization.

Heights and valleys. In Section 2, we describe an $O(n^2k)$ -time algorithm for the k -Dyck edit distance problem which is based on Equation (1). The main idea is to reduce the number of options for m in Equation (1) from $O(n)$ to $O(k)$. To do so, we define the *height* of position $i \in [0..n]$ in S , denoted by $H(i)$, to be the difference between the number of opening parentheses and the number of closing parentheses in $S[0..i]$. (See Figure 1(a); we remark that our notion of height

¹For $i, j \in \mathbb{Z}$, we denote $[i..j] = \{k \in \mathbb{Z} : i \leq k \leq j\}$, $[i..j) = \{k \in \mathbb{Z} : i \leq k < j\}$, $(i..j] = \{k \in \mathbb{Z} : i < k \leq j\}$, and $(i..j) = \{k \in \mathbb{Z} : i < k < j\}$.




Fig. 1. (a) A plot example of the function $H(\cdot)$. (b) An example of a maximal tall trapezoid (a, b, c, d) , which covers the green indices. Notice that the red and yellow indices also form tall maximal trapezoids (a triangle is a degenerated instance of a trapezoid).

differs slightly from the notion of height in [5].) The notion of heights leads to a set of positions V , called *valleys* and defined as positions i such that $H(i-1) > H(i) < H(i+1)$. The result of Backurs and Onak [5] implies that one can convert S into a string S' of length at most n such that $\text{ed}_D(S) = \text{ed}_D(S')$ and S' has at most $2k$ valleys. The intuition is that if $\text{ed}_D(S[i-1]S[i]) = 0$ for some peak i (i.e., a position such that $H(i-1) < H(i) > H(i+1)$), then the two parentheses can be greedily matched; otherwise, they contribute at least 1 to $\text{ed}_D(S)$. Thus, we assume without loss of generality that S has at most $2k$ valleys. The main new insight behind our $O(n^2k)$ -time algorithm is that one can reduce the set of options for m in Equation (1) to the set $(i..j) \cap (M \cup \{i+1, i+2, j-2, j-1\})$, where $M = \bigcup_{v \in V} \{v-1, v, v+1\}$; see Lemma 2.1. Since $|V| \leq 2k$, this reduces the cost of each recursive call to $O(k)$ time.

Trapezoids. In order to obtain our $O(n+k^5)$ time algorithm, we introduce the notion of *trapezoids* and *clusters*, which are based on the heights of positions. Intuitively, a trapezoid is a quadruple (a, b, c, d) with $0 \leq a < b \leq c < d \leq n$, where $S[a..b)$ are all open parentheses, $S[c..d)$ are all closing parentheses, $H(a) = H(d)$, $H(b) = H(c)$, and for all $i \in (b..c)$ we have $H(i) \geq H(b)$; see Figure 1(b).

The motivation for defining trapezoids is that, in Section 3, we design an $O(k^2)$ -time algorithm for processing a trapezoid that, given values $\min\{\mathcal{D}[i, j], k+1\}$ for enough i, j in the vicinity of b and c , outputs the values $\min\{\mathcal{D}[i, j], k+1\}$ for enough i, j in the vicinity of a and d . Our algorithm utilizes the approach developed by Landau–Vishkin for computing the threshold edit distance [20, 21]. Thus, the trapezoids together with the trapezoid processing algorithm provide a method for shortcutting the evaluation of Equation (1). In order to leverage these shortcuts, and since the runtime per trapezoid does not depend on the distance between d and a , we choose to apply the trapezoid algorithm to the largest possible trapezoids. Thus, we define a *maximal trapezoid* to be a trapezoid (a, b, c, d) that cannot be *extended* to another trapezoid of the form either $(a-1, b, c, d+1)$ or $(a, b+1, c-1, d)$, and consider only *tall* maximal trapezoids which are maximal trapezoids where $b-a \geq 2k$ (and so the shortcutting is significant); see Figure 1(b).

Clusters. The computation of Equation (1) is not fully completed by just processing all tall maximal trapezoids. In order to address the remaining computation, we introduce the notion of *clusters*, which are defined as follows. Consider a cycle with vertices $[0..n]$ and edges $\{(i, i+1) : i \in [0..n]\} \cup \{(n, 0)\}$. For each tall maximal trapezoid (a, b, c, d) , delete edges $(i, i+1)$ for $i \in [a..b) \cup [c..d)$, delete vertices $(a..b) \cup (c..d)$, and add edges (a, d) and (b, c) . The connected components of the resulting graph are called *clusters*; see Figure 2(a).

From the way that clusters are constructed, the clusters and tall maximal trapezoids define a parent-child relationship as follows. Let $T = (a, b, c, d)$ be a tall maximal trapezoid and let C be a




Fig. 2. (a) An example of the partitioning to trapezoids and clusters. (b) The tree \mathcal{T} for the partitioning of Figure 2(a). (c) Focus on a cluster C_2 , the input contains the values $\min\{\mathcal{D}[i, j], k + 1\}$ for all pairs of red solid indices and for all pairs of yellow solid indices. The output of the processing of C_2 is the values $\min\{\mathcal{D}[i, j], k + 1\}$ for all pairs of green solid indices. (d) The values $\min\{\mathcal{D}[i, j], k + 1\}$ for pairs of red solid indices are the input for processing a trapezoid T_1 , and the values $\min\{\mathcal{D}[i, j], k + 1\}$ for pairs of blue solid indices are the output of the processing.

cluster. We say that C is the parent of T if $a, d \in C$, and that T is the parent of C if $b, c \in C$. This relationship defines a tree \mathcal{T} whose root is the cluster containing both 0 and n . Notice that every trapezoid in \mathcal{T} has a cluster parent and exactly 1 cluster child. Moreover, every cluster in \mathcal{T} has an arbitrary number of trapezoid children, and at most 1 trapezoid parent (the root of \mathcal{T} is a cluster and has no parent); see Figure 2(b).

By combining ideas based on the notions of heights and valleys, we design an algorithm for processing a cluster C with r children $T_q = (a_q, b_q, c_q, d_q)$ for $q \in [1 \dots r]$ and a parent $T = (a, b, c, d)$, which given the values $\min\{\mathcal{D}[i, j], k + 1\}$ for i, j in the vicinity of a_q and d_q for all $q \in [1 \dots r]$ (which are the outputs of the executions of the trapezoid processing algorithm on T_q for $q \in [1 \dots r]$), computes the values $\min\{\mathcal{D}[i, j], k + 1\}$ for i, j in the vicinity of b and c (which can later be used as input for processing T). Moreover, if C is the root of \mathcal{T} , then the cluster processing algorithm computes $\mathcal{D}[0, n]$. Similar to the $O(n^2k)$ -time algorithm, the intuition behind the cluster processing algorithm is to reduce the options for m when applying Equation (1) to the vicinity of the valleys contained in the cluster; see Figure 2(c) and (d).

Combining Trapezoids and Clusters. The tree \mathcal{T} provides a natural processing order, starting from the leaves (which could be either clusters or trapezoids) towards the root, in order to compute $\mathcal{D}[0, n]$ at the root. Whenever the processing order reaches a trapezoid, the input to the algorithm for processing the trapezoid is either trivial when the trapezoid is a leaf, or given from the cluster child of the trapezoid. Whenever the processing order reaches a cluster, the input to the algorithm for processing the cluster is either the output of the trapezoid processing algorithm when applied to the child trapezoid of the cluster, or trivial when the cluster is a leaf. Finally, the output $\mathcal{D}[0, n]$ is computed when processing the cluster that is the root of \mathcal{T} .

In Section 3, we show that the total runtime of the cluster processing algorithm on all clusters is $O(k^5)$. We also show that the number of tall maximal trapezoids is $O(k)$ and so the total time spent processing trapezoids is $O(k^3)$. In addition, constructing the tall maximal trapezoids, the

clusters, and the tree \mathcal{T} , together with some additional preprocessing, costs $O(n)$ time, for a total of $O(n + k^5)$ time.

Min-Plus Products, Fast Matrix Multiplication, and Valiant's Parser. Bringmann et al. [6] improved the running time of the dynamic programming algorithm based on Equation (1) from $O(n^3)$ to $O(n^{2.8603})$ (deterministically) and $O(n^{2.8244})$ (using randomization). Their approach is based on observations that \mathcal{D} satisfies a **column bounded difference (column-BD) property**² $|\mathcal{D}[i, j] - \mathcal{D}[i+1, j]| \leq 1$ and a **row bounded difference (row-BD) property** $|\mathcal{D}[i, j-1] - \mathcal{D}[i, j]| \leq 1$ all $i, j \in [0..n]$ with $i > j$. Moreover, the bottleneck of computing $\mathcal{D}[i, j]$ is the $(\min, +)$ product of vectors $(\mathcal{D}[i, m])_{m=i+1}^{j-1}$ and $(\mathcal{D}[m, j])_{m=i+1}^{j-1}$. If a matrix is both column-BD and row-BD, we say that the matrix is *fully-BD*. Their main technical contribution for the k -Dyck edit distance problem is an efficient algorithm for computing the $(\min, +)$ product³ of fully-BD square matrices. Building upon Valiant's parser, filling the entire dynamic-programming table \mathcal{D} is then reduced to computing the $(\min, +)$ products of the already constructed square submatrices of \mathcal{D} . It is worth mentioning that Bringmann et al. [6] introduce a subcubic $(\min, +)$ algorithm also for the case where only one of the matrices is either column-BD or row-BD. A subcubic $(\min, +)$ algorithm for even less structured matrices was later introduced by Vassilevska-Williams and Xu [31]. Recently, Chi et al. [7] improved upon the result of Bringmann et al. [6] for bounded difference $(\min, +)$ product, and, in follow up work, Chi et al. [8] introduced even more efficient algorithms that apply to a broader range of structured matrices. All of these recent improvements are randomized.

In the dynamic programming based on limiting the options for m , which results in the $O(n^2k)$ time algorithm, the bottleneck is determining the $(\min, +)$ product of vectors $(\mathcal{D}[i, m])_{m \in (i..j) \cap M}$ and $(\mathcal{D}[m, j])_{m \in (i..j) \cap M}$. Thus, instead of multiplying square matrices, it suffices to multiply smaller rectangular matrices (with columns of the first matrix and rows of the second matrix restricted to M). While this shrinks the two matrices, only the column-BD property is preserved for the left matrix in the $(\min, +)$ product, while only the row-BD property is preserved for the right matrix. Using the $(\min, +)$ product algorithm of Bringmann et al. [6] for the more general case, where only one of the matrices is either column-BD or row-BD would lead to inferior running times.

Nevertheless, in Section 6, we observe that the output matrix of the $(\min, +)$ product in our case is fully-BD and, accordingly, design a $(\min, +)$ product algorithm specialized for our case, which may be of independent interest. Formally, we prove the following theorem in Section 6.

THEOREM 1.2. *Let A and B be integer matrices of sizes $n \times n^\alpha$ and $n^\alpha \times n$, respectively, such that A is column-BD and B is row-BD. For any $\beta, s, \delta \in (0, 1]$, there exists a randomized algorithm that computes the $(\min, +)$ product $A \star B$ (correctly with high probability) in time*

$$\tilde{O}\left(n^\beta M(n, n^\alpha, n) + n^{2+\alpha-\beta}\right),$$

and a deterministic algorithm that runs in time

$$\tilde{O}\left(n^{\delta+s} M(n, n^{\alpha-s/3}, n) + n^\alpha \cdot M(n^{1-\delta}, n^s, n^{1-\delta}) + n^{2+\alpha-s/3} + n^\alpha \cdot M(n^{1-\delta}, n^{1-\delta}, n^{1-\delta})\right).$$

Here, $M(a, b, c)$ denotes the time needed for computing the Boolean product of an $a \times b$ matrix with a $b \times c$ matrix.

²The notion of bounded difference defined in [6] is parametrized by a natural number W . In our setting $W = 1$, so we only focus on this case.

³For two matrices A and B of size $a \times b$ and $b \times c$, respectively, the $(\min, +)$ -product of A and B is a matrix $C = A \star B$ of size $a \times c$ such that $C_{i,j} = \min_{\ell \in [1..b]} \{A_{i,\ell} + B_{\ell,j}\}$.

Henceforth, we assume that the running time of Theorem 1.2, after optimizing over β, s and δ , is $\tilde{O}(n^{\tilde{\omega}(\alpha)})$ (depending on context, this can either refer to the deterministic or the randomized runtime). In Section 6, we estimate $\tilde{\omega}(\frac{1}{2})$ (using the bounds of [3, 14]), deriving $\tilde{\omega}(\frac{1}{2}) < 2.426524$ in the deterministic case and $\tilde{\omega}(\frac{1}{2}) < 2.272092$ if randomization is allowed (for this case, a similar result was obtained independently and in parallel by [12]).

We emphasize that the techniques used to prove Theorem 1.2 mostly follow the paradigm and techniques of Bringmann et al. [6] and Chi et al. [7], with some adjustments to the details so that their paradigm fits our case. For sake of completeness, a full exposition is given in Section 6.

Proving Theorem 1.1. In Section 4, we define a generic Valiant-like recursion (whose structure resembles that of the dynamic-programming algorithm based on limiting the options for m) and show that it can be simulated in time proportional (up to polylogarithmic factors) to a single $(\min, +)$ product. As a straightforward corollary, our $O(n^2k)$ -time dynamic-programming algorithm can be sped up to $\tilde{O}(n^{\tilde{\omega}(\log_n k)})$ time. In particular, this proves Theorem 1.1 for $k \geq \sqrt{n}$ (if $k \gg \sqrt{n}$, the runtime is actually better than claimed in Theorem 1.1). In Section 5, we apply the result of Section 4 to convert our $O(n + k^5)$ -time algorithm into an $\tilde{O}(n + k^{2\tilde{\omega}(1/2)})$ -time solution, thereby completing the proof of Theorem 1.1 for the case of $k < \sqrt{n}$.

2 COMBINATORIAL OBSERVATIONS AND $O(N^2K)$ -TIME ALGORITHM

In this section, we introduce a new combinatorial property (Lemma 2.1) that allows the dynamic programming algorithm to spend only $O(k)$ time to compute each value of the matrix \mathcal{D} , and hence achieve $O(n^2k)$ runtime, which is an improvement over the runtime of Backurs and Onak's algorithm [5] for all $k = \Omega(n^{2/13})$.

Definition 2.1 (Heights). Define the function $h : \Sigma \rightarrow \{-1, 1\}$ so that $h(a) = 1$ if $a \in \Sigma$ is an opening parenthesis and $h(a) = -1$ otherwise. Given a sequence $S \in \Sigma^n$, define the height of a position i of S , $0 \leq i \leq n$, as $H(i) = \sum_{j=0}^{i-1} h(S[j])$.

Notice that $H(i)$ is the difference between the number of opening parentheses and the number of closing parentheses in $S[0 \dots i]$.

Definition 2.2 (Peaks and valleys). Let $S \in \Sigma^n$. We say that a position $i \in [1 \dots n]$ is a *peak* if $H(i-1) < H(i) > H(i+1)$ and a *valley* if $H(i-1) > H(i) < H(i+1)$.

The following claim allows us to assume, without loss of generality, that there are at most $2k$ valleys.

CLAIM 2.3 (COROLLARY OF [5, CLAIM 35]). *Let $S \in \Sigma^n$. There exists an algorithm that preprocesses S in $O(n)$ time, and either rejects S (meaning that $\text{ed}_D(S) > k$), or outputs a string S' of length at most n such that $\text{ed}_D(S) = \text{ed}_D(S')$ and S' has at most $2k$ valleys.*

From now on, we assume that the input string S has at most $2k$ valleys, and we denote the set of valleys by V . This assumption leads us to refine Equation (1) by observing that instead of considering all possible values of m for $i < m < j$, it is enough to consider only values of m that are at distance at most 1 from some valley or at most 2 from i or j .

LEMMA 2.1. *Let $M = \bigcup_{v \in V} \{v-1, v, v+1\}$ for $S \in \Sigma^n$. Let $i, j \in [0 \dots n]$ with $j - i \geq 2$. Then*

$$\mathcal{D}[i, j] = \min \begin{cases} \mathcal{D}[i, m] + \mathcal{D}[m, j] & \text{for } m \in (i \dots j) \cap (M \cup \{i+1, i+2, j-2, j-1\}), \\ \text{ed}_D(S[i]S[j-1]) + \mathcal{D}[i+1, j-1]. \end{cases} \quad (2)$$

PROOF. Notice that if $\mathcal{D}[i, j] = \text{ed}_D(S[i]S[j-1]) + \mathcal{D}[i+1, j-1]$ then, by the correctness of Equation (1), the equality holds since the options for m in Equation (2) are a subset of the options

for m in Equation (1). Thus, we focus on the case where $\mathcal{D}[i, j] = \mathcal{D}[i, m] + \mathcal{D}[m, j]$ for some $m \in (i \dots j)$. Let $p = \arg \min\{H(m) : m \in (i \dots j) \text{ and } \mathcal{D}[i, j] = \mathcal{D}[i, m] + \mathcal{D}[m, j]\}$. That is, p is an index from $(i \dots j)$ for which $\mathcal{D}[i, j] = \mathcal{D}[i, p] + \mathcal{D}[p, j]$ and $H(p)$ is minimized. In the following, we complete the proof by showing that $p \in M \cup \{i+1, i+2, j-2, j-1\}$. For a proof by contradiction, suppose that $p \notin M \cup \{i+1, i+2, j-2, j-1\}$. Then, $[p-2 \dots p+2] \subseteq (i \dots j)$ since $p \notin \{i+1, i+2, j-2, j-1\}$, and $\min\{H(p-2), H(p+2)\} = H(p) - 2$ since $p \notin M$. By symmetry, we assume without loss of generality that $H(p+2) = H(p) - 2$. In this case, $H(p+1) = H(p) - 1$, and $S[p], S[p+1]$ are both closing parentheses.

Let $q \in (p \dots j)$ be the smallest value such that $\mathcal{D}[p, j] = \mathcal{D}[p, q] + \mathcal{D}[q, j]$ (notice that $\mathcal{D}[p, j] = \mathcal{D}[p, j] + \mathcal{D}[j, j]$, so q is always well-defined). We consider two cases:

Case 1: $q \leq p+2$. In this case, $\mathcal{D}[i, j] = \mathcal{D}[i, p] + \mathcal{D}[p, q] + \mathcal{D}[q, j] \geq \mathcal{D}[i, q] + \mathcal{D}[q, j] \geq \mathcal{D}[i, j]$ and hence $\mathcal{D}[i, j] = \mathcal{D}[i, q] + \mathcal{D}[q, j]$. However, $H(q) < H(p)$, which contradicts the choice of p since $H(p)$ is not minimized.

Case 2: $q > p+2$. Recall that q is minimized. Therefore, there is no $q' \in (p \dots q)$ such that $\mathcal{D}[p, q] = \mathcal{D}[p, q'] + \mathcal{D}[q', q]$ as otherwise $\mathcal{D}[p, j] = \mathcal{D}[p, q] + \mathcal{D}[q, j] = \mathcal{D}[p, q'] + \mathcal{D}[q', q] + \mathcal{D}[q, j] \geq \mathcal{D}[p, q'] + \mathcal{D}[q', j] \geq \mathcal{D}[p, j]$ and hence $\mathcal{D}[p, j] = \mathcal{D}[p, q'] + \mathcal{D}[q', j]$, a contradiction. Consequently, we have $\mathcal{D}[p, q] = \text{ed}_D(S[p]S[q-1]) + \mathcal{D}[p+1, q-1]$ and $\mathcal{D}[i, j] = \mathcal{D}[i, p] + \text{ed}_D(S[p]S[q-1]) + \mathcal{D}[p+1, q-1] + \mathcal{D}[q, j]$. Let us define $r \in (p+1 \dots q-1)$ as the smallest value such that $\mathcal{D}[p+1, q-1] = \mathcal{D}[p+1, r] + \mathcal{D}[r, q-1]$. We consider two subcases:

Case 2a: $r = p+2$. Notice that $\text{ed}_D(S[p]S[q-1]) = \text{ed}_D(S[p+1]S[q-1])$ since $S[p]$ and $S[p+1]$ are closing parentheses. Recall that for $\ell \in [0 \dots n]$ we have $\mathcal{D}[\ell, \ell+1] = 1$. Consequently,

$$\begin{aligned} \mathcal{D}[i, j] &= \mathcal{D}[i, p] + \text{ed}_D(S[p]S[q-1]) + \mathcal{D}[p+1, p+2] + \mathcal{D}[p+2, q-1] + \mathcal{D}[q, j] \\ &= \mathcal{D}[i, p] + \mathcal{D}[p, p+1] + \text{ed}_D(S[p+1]S[q-1]) + \mathcal{D}[p+2, q-1] + \mathcal{D}[q, j] \\ &\geq \mathcal{D}[i, p+1] + \mathcal{D}[p+1, q] + \mathcal{D}[q, j] \\ &\geq \mathcal{D}[i, p+1] + \mathcal{D}[p+1, j] \\ &\geq \mathcal{D}[i, j]. \end{aligned}$$

Thus, $\mathcal{D}[i, j] = \mathcal{D}[i, p+1] + \mathcal{D}[p+1, j]$. However, $H(p+1) < H(p)$, which contradicts the choice of p since $H(p)$ is not minimized.

Case 2b: $r > p+2$. Notice that $\text{ed}_D(S[p]S[p+1]) \leq \text{ed}_D(S[p+1]S[r-1])$ since $S[p]$ and $S[p+1]$ are closing parentheses. Moreover, if $S[q-1]$ is a closing parenthesis, then $\text{ed}_D(S[r-1]S[q-1]) \leq 1 = \text{ed}_D(S[p]S[q-1])$. Otherwise, $S[q-1]$ is an opening parenthesis and $\text{ed}_D(S[r-1]S[q-1]) \leq 2 = \text{ed}_D(S[p]S[q-1])$. Notice that for $\ell \in [0 \dots n-1]$ we have $\mathcal{D}[\ell, \ell+2] = \text{ed}_D(S[\ell]S[\ell+1])$. Consequently,

$$\begin{aligned} \mathcal{D}[i, j] &= \mathcal{D}[i, p] + \text{ed}_D(S[p]S[q-1]) + \text{ed}_D(S[p+1]S[r-1]) + \mathcal{D}[p+2, r-1] + \mathcal{D}[r, q-1] \\ &\quad + \mathcal{D}[q, j] \\ &\geq \mathcal{D}[i, p] + \text{ed}_D(S[p]S[p+1]) + \mathcal{D}[p+2, r-1] + \text{ed}_D(S[r-1]S[q-1]) + \mathcal{D}[r, q-1] \\ &\quad + \mathcal{D}[q, j] \\ &\geq \mathcal{D}[i, p] + \mathcal{D}[p, p+2] + \mathcal{D}[p+2, r-1] + \mathcal{D}[r-1, q] + \mathcal{D}[q, j] \\ &\geq \mathcal{D}[i, p+2] + \mathcal{D}[p+2, j] \\ &\geq \mathcal{D}[i, j]. \end{aligned}$$

Thus, $\mathcal{D}[i, j] = \mathcal{D}[i, p+2] + \mathcal{D}[p+2, j]$. However, $H(p+2) < H(p)$, which contradicts the choice of p since $H(p)$ is not minimized. \square

By Lemma 2.1 it is straightforward to prove the following corollary.

COROLLARY 2.2. *There exists an algorithm for k -Dyck edit distance with $O(n^2k)$ running time.*

3 $O(N + K^5)$ -TIME ALGORITHM

In this section, we further refine the dynamic programming approach to achieve an $O(n + k^5)$ -time algorithm. To do so, the algorithm carefully selects a subset of $O(k^4)$ entries of the matrix \mathcal{D} based on Lemma 2.1, and the remaining entries are computed using a much faster method based on the approach developed by Landau and Vishkin for computing the threshold edit distance ([21], see also [20]).

3.1 Definitions and Combinatorial Observations.

We use the following simple fact that is an extension of [5, Fact 36].

FACT 3.1. *For $0 \leq i \leq j \leq n$, if $\mathcal{D}[i, j] \leq k$, then $H(m) \geq \max\{H(i), H(j)\} - 2k$ holds for every $m \in [i \dots j]$. In particular, $|H(i) - H(j)| \leq 2k$.*

Definition 3.2 (Trapezoid). For $0 \leq a < b \leq c < d \leq n$, we say that (a, b, c, d) is a *trapezoid* if the following properties hold: (1) $H(a) = H(d)$, (2) $b - a = H(b) - H(a) = H(c) - H(d) = d - c$ (this value is called the *height* of the trapezoid), and (3) for all $m \in [b \dots c]$, we have $H(m) \geq H(b) = H(c)$. See Figure 1(b).

Notice that all the characters in $S[a \dots b)$ are opening parentheses and all the characters in $S[c \dots d)$ are closing parentheses, respectively.

Definition 3.3 (Tall and maximal trapezoid). A trapezoid (a, b, c, d) is *maximal* if it cannot be extended to a different trapezoid (a', b', c', d') with $a' \leq a$, $b' \geq b$, $c' \leq c$, and $d' \geq d$. A trapezoid is *tall* if its height is at least $2k$.

FACT 3.4. *If (a, b, c, d) and (a', b', c', d') are distinct maximal trapezoids, then $[a \dots b) \cup [c \dots d)$ and $[a' \dots b') \cup [c' \dots d')$ are disjoint.*

PROOF. Suppose that these sets contain a common position i . By symmetry, assume without loss of generality that $i \in [a \dots b)$, which means that $S[i]$ is an opening parenthesis and $i \in [a' \dots b')$. Let $j \in (c \dots d)$ with $i + j = a + d$, and so $H(i) = H(j)$. Moreover, since (a, b, c, d) is a trapezoid, for $m \in (i \dots j)$ we have $H(m) > H(i)$. Symmetrically, define $j' \in (c' \dots d')$ so that $i + j' = a' + d'$. Thus, $H(i) = H(j')$ and for $m \in (i \dots j')$ we have $H(m) > H(i)$. This proves that $j = j'$ and that (a, b, c, d) and (a', b', c', d') both extend $(i, i + 1, j - 1, j)$. Finally, we note that $a' \geq a$ (otherwise $(a - 1, b, c, d + 1)$ is a trapezoid), $a' \leq a$ (otherwise $(a' - 1, b', c', d' + 1)$ is a trapezoid), $b' \leq b$ (otherwise $(a, b + 1, c + 1, d)$ is a trapezoid), and $b' \geq b$ (otherwise $(a', b' + 1, c' - 1, d')$ is a trapezoid). Hence, $(a, b, c, d) = (a', b', c', d')$ holds as claimed. \square

We process tall maximal trapezoids using a fast method based on [20, 21], but the computation of Equation (1) is not fully completed by just processing all tall maximal trapezoids. In order to address the remaining computation, The remaining positions are grouped into clusters, and processed using Lemma 2.1. Formally,

Definition 3.5 (Cluster). Consider a cycle with vertices $[0 \dots n]$ and edges $\{(i, i + 1) : i \in [0 \dots n]\} \cup \{(n, 0)\}$. For each tall maximal trapezoid (a, b, c, d) , delete edges $(i, i + 1)$ for $i \in [a \dots b) \cup [c \dots d)$, delete vertices $(a \dots b) \cup (c \dots d)$, and add edges (a, d) and (b, c) . The connected components of the resulting graph are called *clusters*. We say that a cluster C contains all positions $S[p]$ such that $(p, p + 1)$ is an edge of C . See Figure 2(a).

LEMMA 3.1. *If the input string S has at most $2k$ valleys and $\mathcal{D}[0, n] \leq k$, then the number of maximal trapezoids is $O(k)$ and the total number of positions in the clusters is $O(k^2)$.*

PROOF. Consider a maximal trapezoid (i, i', j', j) . First, note that $i \in V \cup \{0\}$ or $j \in V \cup \{n\}$ because $(i - 1, i', j', j + 1)$ is not a trapezoid. Moreover, by Fact 3.4, there is no other maximal trapezoid (a, b, c, d) with $a = i$ or $d = j$. Hence, there are at most two maximal trapezoids associated with each $v \in V$, and another maximal trapezoid associated with each extreme position (0 and $n - 1$). Thus, number of maximal trapezoids is bounded by $4k + 2$.

We now show that all positions in clusters, except for $O(k)$ of those positions, belong to some maximal trapezoid. Consider an opening parenthesis $S[p]$. If there is a position $r \in [p..n]$ such that $H(r) = H(p)$, then choose the leftmost such position. In this case, $S[r - 1]$ is a closing parenthesis and $(p, p + 1, r - 1, r)$ forms a trapezoid, which can be extended to a maximal trapezoid. Thus, an opening parenthesis $S[p]$ does not belong to any maximal trapezoid only if $H(p) < H(r)$ for all $r \in [p..n]$. The number of such positions does not exceed $H(n) - \min_{p \in [0..n]} H(p)$, since each height between $\min_{p \in [0..n]} H(p)$ and $H(n)$ can contribute at most once. By Fact 3.1, $H(n) - \min_{p \in [0..n]} H(p) \leq 2k$. Thus, the number of such positions is at most $2k$. Similarly, there are at most $2k$ closing parentheses that do not belong to any maximal trapezoid.

Finally, we note that if a maximal trapezoid is not tall, then it contains at most $4k$ positions. Since the number of maximal trapezoids is $O(k)$, this means that maximal trapezoid that are not tall contain at most $O(k^2)$ positions in total. Overall, we conclude that the total number of positions in the clusters is $O(k^2)$ as desired. \square

Definition 3.6 (Parent relation on clusters and trapezoids). Let $T = (a, b, c, d)$ be a tall maximal trapezoid and let C be a cluster. We say that C is the parent of T if $a, d \in C$, and that T is the parent of C if $b, c \in C$. See Figure 2(b).

The following lemma shows that clusters and trapezoids naturally define a tree structure.

LEMMA 3.2. *If $\mathcal{D}[0, n] \leq k$, then the parent relation on clusters and tall maximal trapezoids defines a rooted tree whose root is the cluster containing both 0 and n . We denote this tree by \mathcal{T} .*

Moreover, there exists an algorithm for constructing tree \mathcal{T} in $O(n)$ time.

PROOF. For a tall maximal trapezoid $T = (a, b, c, d)$, the graph G of Definition 3.5 does not contain any edge connecting $[b..c]$ with $[0..a] \cup [d..n]$. Consequently, the cluster C whose parent is T satisfies $\min C = b$ and $\max C = c$. At the same time, the parent C' of T satisfies $\min C' \leq a < b$ and $\max C' \geq d > c$. We conclude that the parent relation is acyclic.

It remains to prove that every cluster C (except for the one containing 0, n) has exactly one parent. Consider an edge $(\max C, \max C + 1)$ not present in G . Since $\max C \in G$, this edge must have been deleted while processing a tall maximal trapezoid $T = (a, b, c, d)$ such that $\max C \in \{a, c\}$. If $\max C = a$, then (a, d) is an edge of G , so $d \in C$ and $\max C \leq d < a = \max C$, a contradiction. Consequently, we must have $\max C = c$, which means that T is the parent of C . This trapezoid is unique due to Fact 3.4.

We now show an algorithm that builds \mathcal{T} in $O(n)$ time. First, for each opening parenthesis $S[p]$, we compute the value $B(p) = \min\{r \in (p..n) : H(p) = H(r)\}$, assuming $\min \emptyset = \infty$. For this, it suffices to scan S from left to right in linear time. Now, observe that (a, b, c, d) is a trapezoid if and only if $i + B(i) = a + d = b + c$ for all $i \in [a..b]$. In particular, maximal trapezoids correspond to maximal segments $S[a..b]$ of opening parentheses such that $i + B(i)$ is fixed for all $i \in [a..b]$. This allows constructing all maximal trapezoids in $O(n)$ time. Out of these trapezoids, we filter tall trapezoids (for which the height is at least $2k$).

We now show how to compute the clusters. Let G be a graph on nodes $[0..n]$ that contains edges $(i, i + 1)$ for each $i \in [0..n]$ and an edge $(n, 0)$ (i.e., G is a cycle on $[0..n]$). Next, for each tall maximal trapezoid (i, i', j', j) , modify G as specified in Definition 3.5. By Fact 3.4, no vertex

or edge is deleted twice. Finally, traverse the graph and mark each vertex with an identifier of its connected component (cluster). Thus, constructing all the clusters costs $O(n)$ time in total.

It remains to build the tree \mathcal{T} . For this, each tall maximal trapezoid $T = (a, b, c, d)$ is connected to its child (the cluster containing b and c) and its parent (the cluster containing a and d). Overall, the algorithm costs $O(n)$ time. \square

Notice that every trapezoid in \mathcal{T} has a cluster parent, and at most 1 child, which must be a cluster. Moreover, every cluster in \mathcal{T} has an arbitrary number of trapezoid children, and at most 1 trapezoid parent (the root of \mathcal{T} is a cluster and has no parent).

CLAIM 3.7. *Consider a trapezoid $(a, a + 2k, d - 2k, d)$ and indices $i \leq j$ such that $\mathcal{D}[i, j] \leq k$. If one of the indices i, j belongs to $(a + 2k \dots d - 2k)$, then the other belongs to $(a \dots d)$.*

PROOF. By symmetry, we assume without loss of generality that $i \in (a + 2k \dots d - 2k)$. Then, $H(i) > H(d - 2k) = H(d) + 2k$ and, by Fact 3.1, $d \notin [i \dots j]$, i.e., $j \in (a \dots d)$. \square

3.2 The Algorithm.

Our main algorithm constructs \mathcal{T} and then performs a post-order traversal of \mathcal{T} , with distinct procedures for processing nodes representing clusters and trapezoids.

Processing Clusters. When the algorithm processes a cluster C with r children $T_q = (a_q, b_q, c_q, d_q)$ for $q \in [1 \dots r]$, the algorithm is given the values $\min\{\mathcal{D}[i, j], k + 1\}$ for all $i, j \in [a_q \dots a_q + 2k] \cup [d_q - 2k \dots d_q]$ and $q \in [1 \dots r]$. If C has a parent $T = (a, b, c, d)$, then the outcome of processing C consists of the values $\min\{\mathcal{D}[i, j], k + 1\}$ for $i, j \in [b - 2k \dots b] \cup [c \dots c + 2k]$. See Figure 2(c). If C is the root of \mathcal{T} , then $0, n \in C$ by Lemma 3.2 and the output is $\mathcal{D}[0, n]$. Denote

$$\begin{aligned} E(C) &:= C \cup \left(\bigcup_{q=1}^r [a_q \dots a_q + 2k] \cup [d_q - 2k \dots d_q] \right) \cup ([b - 2k \dots b] \cup [c \dots c + 2k]) \\ &= [b - 2k \dots c + 2k] \setminus \bigcup_{q=1}^r (a_q + 2k \dots d_q - 2k), \end{aligned}$$

where the last term is included only if C is not the root of \mathcal{T} . Notice that $E(C)$ is the set of positions that are relevant for processing C . Let $M(C) = M \cap E(C)$, where M is defined as in Lemma 2.1. The following lemma is a simplification of Lemma 2.1, confined to values in $E(C)$ which are not given as input to cluster processing.

LEMMA 3.3. *Let $(i, j) \in E(C)^2 \setminus \bigcup_{q=1}^r [a_q \dots d_q]^2$. If $j \geq i + 2$ and $\mathcal{D}[i, j] \leq k$, then*

$$\mathcal{D}[i, j] = \min \begin{cases} \mathcal{D}[i, m] + \mathcal{D}[m, j] & \text{for } m \in (i \dots j) \cap (M(C) \cup \{i + 1, i + 2, j - 2, j - 1\}), \\ \text{ed}_D(S[i]S[j - 1]) + \mathcal{D}[i + 1, j - 1] & \text{if } i + 1, j - 1 \in E(C). \end{cases} \quad (3)$$

PROOF. If $\mathcal{D}[i, j] = \mathcal{D}[i + 1, j - 1] + \text{ed}_D(S[i]S[j - 1]) \leq k$, then observe that $i + 1, j - 1 \in E(C)$ unless $i + 1 \in (a_q + 2k \dots d_q - 2k)$ or $j - 1 \in (a_q + 2k \dots d_q - 2k)$ for some $q \in [1 \dots r]$. By Claim 3.7 applied to $(a_q, a_q + 2k, d_q - 2k, d_q)$, the assumption $\mathcal{D}[i, j] \leq k$ implies $\mathcal{D}[i + 1, j - 1] \leq k$ and so $(i + 1, j - 1) \in (a_q \dots d_q)^2$, i.e., $(i, j) \in [a_q \dots d_q]^2$, a contradiction.

Thus, by Lemma 2.1, we focus on the case where $\mathcal{D}[i, j] = \mathcal{D}[i, m] + \mathcal{D}[m, j]$ for some $m \in (i \dots j) \cap (M \cup \{i + 1, i + 2, j - 2, j - 1\})$. Assume by contradiction that $m \notin M(C)$, and so $m \in (a_q + 2k \dots d_q - 2k)$ for some $q \in [1 \dots r]$. By Claim 3.7 applied to $(a_q, a_q + 2k, d_q - 2k, d_q)$, the assumption $\mathcal{D}[i, m] + \mathcal{D}[m, j] \leq k$ implies $i, j \in (a_q \dots d_q)$. However, $(i, j) \notin [a_q \dots d_q]^2$, and so we obtain a contradiction. \square




Fig. 3. The submatrix of \mathcal{D} corresponding to a tall maximal trapezoid $T = (a, b, c, d)$. Green squares correspond to the values $\mathcal{D}[i, j]$ with $(i, j) \in [a..b]^2 \cup [c..d]^2$, for which the values are known. Except for the green squares, only the values in the red square, yellow diagonals, and the blue square can be smaller or equal to k . The processing of T must compute the values in the blue square given the values in the red square.

By Lemma 3.3, the values $\min\{\mathcal{D}[i, j], k + 1\}$ for all $i, j \in E(C)$ with $i \leq j$ can be computed in $O(|E(C)|^2(1 + |M(C)|))$ time. Notice that $|E(C)| \leq |C| + (4k + 2)(q + 1)$. Moreover, by Lemma 3.1, the sum of the number of children for all clusters is $O(k)$ and the total number of positions in the clusters is $O(k^2)$, and therefore $\sum_C |E(C)| = O(k^2)$. Furthermore, $|M(C) \setminus C| = O(q + 1)$ (because, for each trapezoid (a, b, c, d) , only $a, a + 1, d - 1, d$ may belong to M) and $\sum_C |M(C) \cap C| \leq |M| = O(k)$ as the clusters do not intersect, so $\sum_C |M(C)| = \sum_C |M(C) \setminus C| + \sum_C |M(C) \cap C| = O(k)$. Consequently, the total time required to process all clusters is $O((k^2)^2 \cdot (1 + k)) = O(k^5)$.

Processing Trapezoids. When we process a tall maximal trapezoid $T = (a, b, c, d)$, we are given values $\min\{\mathcal{D}[i, j], k + 1\}$ for all $i, j \in [b - 2k..b] \cup [c..c + 2k]$, and the output consists of the values $\min\{\mathcal{D}[i, j], k + 1\}$ for $i, j \in [a..a + 2k] \cup [d - 2k..d]$. See Figure 2(d).

For $(i, j) \in [a..b]^2 \cup [c..d]^2$, we have $\mathcal{D}[i, j] = \lceil (j - i) / 2 \rceil$ because all parentheses in $S[i..j]$ are of the same orientation. We therefore only need to compute entries $\mathcal{D}[i, j]$ such that $i \in [a..a + 2k]$ and $j \in [d - 2k..d]$. Consider a square submatrix $\mathcal{D}[a..b, c..d]$ of \mathcal{D} (see Figure 3). We start by showing that the values in the submatrix follow a simpler recursion:

CLAIM 3.8. For all $(i, j) \in [a..d]^2 \setminus [b - 2k..c + 2k]^2$ such that $i \leq j$ and $\mathcal{D}[i, j] \leq k$,

$$\mathcal{D}[i, j] = \min \left\{ \begin{array}{l} \mathcal{D}[i + 2, j] + 1, \mathcal{D}[i + 1, j] + 1, \mathcal{D}[i + 1, j - 1] + \text{ed}_D(S[i]S[j - 1]), \\ \mathcal{D}[i, j + 1] + 1, \mathcal{D}[i, j + 2] + 1 \end{array} \right\}.$$

PROOF. Note that $\mathcal{D}[i, i + 2] = \mathcal{D}[i, i + 1] = \mathcal{D}[j - 1, j] = \mathcal{D}[j - 2, j] = 1$, since for each computation all the parentheses that matter are of the same orientation. Hence, it suffices to prove that if $\mathcal{D}[i, j] = \mathcal{D}[i, m] + \mathcal{D}[m, j]$ then, without loss of generality, $m \in \{i + 1, i + 2, j - 2, j - 1\}$, since, for example, if $m = i + 1$ then $\mathcal{D}[i, j] = \mathcal{D}[i, i + 1] + \mathcal{D}[i + 1, j] = 1 + \mathcal{D}[i + 1, j]$. By Lemma 2.1, the only other possibility is that $m \in M$ and, since $(a..d) \cap V \subseteq (b..c)$, this implies that $m \in [b..c]$. However, $\mathcal{D}[i, m] \leq k$ would then contradict Claim 3.7 for a trapezoid $(b - 2k - 1, b - 1, c + 1, c + 2k + 1)$. \square

Consider an entry $\mathcal{D}[i, j] \leq k$, where $(i, j) \in [a..d]^2 \setminus [b - 2k..c + 2k]^2$. We have either $(i, j) \in [a, b - 2k]^2 \cup (c + 2k, d]^2$ or, by Fact 3.1, $i + j = b + c + \delta$, where $\delta \in [-2k..2k]$. Let us

denote $\delta_+ := \max(0, \delta)$ and $\delta_- := \min(0, \delta)$. For a fixed $\delta \in [-2k \dots 2k]$, we call the set of entries $(i, j) \in [a \dots d]^2 \setminus [b - 2k \dots c + 2k]^2$ such that $i + j = b + c + \delta$ the diagonal δ . Let us show that the values $\mathcal{D}[i, j]$ are monotone along each diagonal:

CLAIM 3.9. For all $(i, j) \in [a \dots d]^2 \setminus [b - 2k \dots c + 2k]^2$ such that $\mathcal{D}[i, j] \leq k$, we have $\mathcal{D}[i, j] \geq \mathcal{D}[i + 1, j - 1]$.

PROOF. Let $q \in (i \dots j)$ be the smallest value such that $\mathcal{D}[i, j] = \mathcal{D}[i, q] + \mathcal{D}[q, j]$. We consider four cases:

- (1) $q = j$. In this case, $\mathcal{D}[i, j] = \text{ed}_D(S[i]S[j - 1]) + \mathcal{D}[i + 1, j - 1] \geq \mathcal{D}[i + 1, j - 1]$.
- (2) $q = j - 1$. In this case, $\mathcal{D}[i, j] = \text{ed}_D(S[i]S[j - 2]) + \mathcal{D}[i + 1, j - 2] + \mathcal{D}[j - 1, j] \geq \mathcal{D}[i + 1, j - 2] + \mathcal{D}[j - 2, j - 1] \geq \mathcal{D}[i + 1, j - 1]$.
- (3) $q = i + 1$. Let $r \in [q \dots j]$ be the largest value such that $\mathcal{D}[q, j] = \mathcal{D}[q, r] + \mathcal{D}[r, j]$. We consider two subcases:
 - (a) $r = j - 1$. In this case, $\mathcal{D}[i, j] = \mathcal{D}[i, i + 1] + \mathcal{D}[i + 1, j - 1] + \mathcal{D}[j - 1, j] \geq \mathcal{D}[i + 1, j - 1]$.
 - (b) $r < j - 1$. In this case, $\mathcal{D}[i, j] = \mathcal{D}[i, i + 1] + \mathcal{D}[i + 1, r] + \text{ed}_D(S[r]S[j - 1]) + \mathcal{D}[r + 1, j - 1] \geq \mathcal{D}[i + 1, r] + \mathcal{D}[r, r + 1] + \mathcal{D}[r + 1, j - 1] \geq \mathcal{D}[i + 1, j - 1]$.
- (4) $i + 1 < q < j - 1$. Again, let $r \in [q \dots j]$ be the largest value such that $\mathcal{D}[q, j] = \mathcal{D}[q, r] + \mathcal{D}[r, j]$. We consider two subcases:
 - (a) $r = j - 1$. In this case, $\mathcal{D}[i, j] = \text{ed}_D(S[i]S[q - 1]) + \mathcal{D}[i + 1, q - 1] + \mathcal{D}[q, j - 1] + \mathcal{D}[j - 1, j] \geq \mathcal{D}[i + 1, q - 1] + \mathcal{D}[q - 1, q] + \mathcal{D}[q, j - 1] \geq \mathcal{D}[i + 1, j - 1]$.
 - (b) $r < j - 1$. In this case, $\mathcal{D}[i, j] = \text{ed}_D(S[i]S[q - 1]) + \mathcal{D}[i + 1, q - 1] + \mathcal{D}[q, r] + \text{ed}_D(S[r]S[j - 1]) + \mathcal{D}[r + 1, j - 1] \geq \mathcal{D}[i + 1, q - 1] + \text{ed}_D(S[q - 1]S[r]) + \mathcal{D}[q, r] + \mathcal{D}[r + 1, j - 1] \geq \mathcal{D}[i + 1, q - 1] + \mathcal{D}[q - 1, r + 1] + \mathcal{D}[r + 1, j - 1] \geq \mathcal{D}[i + 1, j - 1]$, where $\text{ed}_D(S[i]S[q - 1]) + \text{ed}_D(S[r]S[j - 1]) \geq 1 \geq \text{ed}_D(S[q - 1]S[r])$ follows from the fact that $S[q - 1]$ is an opening parenthesis or $S[r]$ is a closing parenthesis (otherwise, we would have $q, r \in (b \dots c)$, which contradicts $\mathcal{D}[i, q] \leq k$ and $\mathcal{D}[r, j] \leq k$ by Claim 3.7 applied to $(b - 2k, b, c, c + 2k)$).

□

For the diagonal δ and values $v \in [0 \dots k]$, we shall compute values $L_v[\delta] = \max\{j \in [c + 2k + \delta_- \dots d - 2k + \delta_+] : \mathcal{D}[b + c + \delta - j, j] \leq v\}$, assuming $\max \emptyset = -\infty$. From the resulting tables $L_v[\delta]$, we can determine $\min(\mathcal{D}[i, j], k + 1)$ for all $i, j \in [a \dots a + 2k] \cup [d - 2k \dots d]$ in $O(k^2)$ time by Claim 3.8. To compute the values $L_v[\delta]$, we rely on the fact that, after $O(n)$ -time preprocessing, the longest common prefix of any two substrings of $S\bar{S}$ can be computed in $O(1)$ -time [13].⁴ Specifically, we use an adaptation of the Landau–Vishkin method [20, 21] generalizing a similar subroutine present in [5]. Our procedure is implemented as Section 3.2, which assumes that $L_v[s] = -\infty$ for all out-of-bounds and uninitialized entries.

LEMMA 3.4. For all $(i, j) \in ([a \dots b] \times [c \dots d]) \setminus ([b - 2k \dots b] \times [c \dots c + 2k])$ and $v \in [0 \dots k]$, we have $\mathcal{D}[i, j] \leq v$ if and only if the value $L_v[(i + j) - (b + c)]$ computed by Section 3.2 is at least j .

PROOF. First, let us prove that $\mathcal{D}[i, j] \leq v$ implies $L_v[\delta] \geq j$, where $\delta = (i + j) - (b + c)$. We proceed by induction on v . Let $S[i \dots i'] = \overline{S[j' \dots j]}$ be the longest common prefix of $S[i \dots b - 2k + \delta_+]$ and $S[c + 2k + \delta_- \dots j]$. We shall prove $L'_v[\delta] \geq j'$, where $\delta = (i + j) - (b + c)$, by considering six cases:

- If $j' = c + 2k + \delta_-$ and $i' = b - 2k + \delta_+$, then $\mathcal{D}[i', j'] \leq v$ by Claim 3.9. Hence, $L'_v[\delta] \geq c + 2k + \delta_- = j'$.

⁴Here, \bar{S} denotes the reverse complement of S , obtained by reversing S and flipping the orientation of every parenthesis.

ALGORITHM 1: Processing trapezoids.

```

1: for  $v := 0$  to  $k$  do
2:   for  $\delta := -2v$  to  $2v$  do
3:      $L'_v[\delta] := \min(d - 2k + \delta_+, \max(L_{v-1}[\delta - 2] + 2, L_{v-1}[\delta - 1] + 2, L_{v-1}[\delta] + 1, L_{v-1}[\delta + 1], L_{v-1}[\delta + 2]))$ 
4:     if  $\mathcal{D}[b - 2k + \delta_+, c + 2k + \delta_-] \leq v$  then
5:        $L'_v[\delta] := \max(L'_v[\delta], c + 2k + \delta_-)$ 
6:     end if
7:      $L_v[\delta] := L'_v[\delta] + \text{lcp}(\overline{S[a..b + c + \delta - L'_v[\delta]]}, S[L'_v[\delta]..d])$ 
8:   end for
9: end for

```

Otherwise, by Claim 3.8, it suffices to consider one of the remaining five cases:

- If $\mathcal{D}[i', j'] = \mathcal{D}[i', j' - 2] + 1$, then the inductive assumption yields $L'_v[\delta] \geq L_{v-1}[\delta - 2] + 2 \geq (j' - 2) + 2 = j'$.
- If $\mathcal{D}[i', j'] = \mathcal{D}[i', j' - 1] + 1$, then the inductive assumption yields $L'_v[\delta] \geq L_{v-1}[\delta - 1] + 1 \geq (j' - 1) + 1 = j'$.
- If $\mathcal{D}[i', j'] = \mathcal{D}[i' + 1, j' - 1] + 1$, then the inductive assumption yields $L'_v[\delta] \geq L_{v-1}[\delta] + 1 \geq (j' - 1) + 1 = j'$.
- If $\mathcal{D}[i', j'] = \mathcal{D}[i' + 1, j'] + 1$, then the inductive assumption yields $L'_v[\delta] \geq L_{v-1}[\delta + 1] \geq j'$.
- If $\mathcal{D}[i', j'] = \mathcal{D}[i' + 2, j'] + 1$, then the inductive assumption yields $L'_v[\delta] \geq L_{v-1}[\delta + 2] \geq j'$.

In all cases $L_v[\delta] \geq j$ follows from $L'_v[\delta] \geq j'$ due to $\text{lcp}(\overline{S[a..i']}, S[j'..d]) \geq j - j'$.

The converse implication is also proved by induction on v . Let $\delta = (i + j) - (b + c)$, $j' = L'_v[\delta]$ and $i' = i + j - j'$. By Claim 3.9, we can assume $j = L_v[\delta]$ without loss of generality. Moreover, since $\overline{S[i..i']} = S[j'..j]$, we have $\mathcal{D}[i, j] \leq \mathcal{D}[i', j']$. Hence, it suffices to prove $\mathcal{D}[i', j'] \leq v$. For this, we consider six cases:

- If $j' = c + 2k + \delta_-$, then the algorithm explicitly checked $\mathcal{D}[i', j'] \leq v$ while setting j' .
- If $j' = L_{v-1}[\delta - 2] + 2$, then $\mathcal{D}[i', j'] \leq \mathcal{D}[i', j' - 2] + 1 \leq (v - 1) + 1 = v$ by the inductive assumption.
- If $j' = L_{v-1}[\delta - 1] + 1$, then $\mathcal{D}[i', j'] \leq \mathcal{D}[i', j' - 1] + 1 \leq (v - 1) + 1 = v$ by the inductive assumption.
- If $j' = L_{v-1}[\delta] + 1$, then $\mathcal{D}[i', j'] \leq \mathcal{D}[i' - 1, j' - 1] + 1 \leq (v - 1) + 1 = v$ by the inductive assumption.
- If $j' = L_{v-1}[\delta + 1]$, then $\mathcal{D}[i', j'] \leq \mathcal{D}[i' + 1, j'] + 1 \leq (v - 1) + 1 = v$ by the inductive assumption.
- If $j' = L_{v-1}[\delta + 2]$, then $\mathcal{D}[i', j'] \leq \mathcal{D}[i' + 2, j'] + 1 \leq (v - 1) + 1 = v$ by the inductive assumption. □

Lemma 3.4 shows correctness of Section 3.2. From the description of Section 3.2 it follows that the tables $L_v[\delta]$ and hence the entries $\min(\mathcal{D}[i, j], k + 1)$ for all $i, j \in [a..a + 2k] \cup [d - 2k..d]$ can be computed in $O(k^2)$ time.

THEOREM 3.5. *Given a sequence S of parentheses of length n , $\min\{\text{ed}_D(S), k + 1\}$ can be computed in $O(n + k^5)$ time.*

PROOF. We start with an $O(n)$ -time preprocessing of S according to Claim 2.3. After completing this preprocessing, we can assume that S contains at most $2k$ valleys. Next, we apply Lemma 3.2

to build a tree of clusters and trapezoids in $O(n)$ time, and then we start the main phase of the algorithm. The total time used to process the clusters is $O(k^5)$. The total time used to process the tall maximal trapezoids is $O(n + k \cdot k^2) = O(n + k^3)$. The theorem follows. \square

4 VALIANT-LIKE RECURSION

In this section, we describe how to speed up a dynamic-programming procedure that constructs a matrix $A[0 \dots n, 0 \dots n]$ so that, for some non-empty $M \subseteq [0 \dots n]$:

- (1) each entry $A[i, j]$ can be computed in $O(1)$ time given ($O(1)$ -time access to) the remaining entries $A[i', j']$ with $i \leq i' \leq j' \leq j$ and the value $\min\{A[i, m] + A[m, j] : m \in (i \dots j) \cap M\}$,
- (2) for every $m \in M$, subsequent entries $A[i, m]$ and subsequent entries $A[m, j]$ differ by at most 1.

Our procedure is an adaptation of Valiant's approach [30] to generalize computing a transitive closure by allowing a non-transitive operation. The description of our adaptation follows a more recent perspective of Valiant's recursion scheme as described by Okhotin [25].

The scheme recursively halves the domain $[0 \dots n]$, which yields a hierarchy of decompositions. We call intervals comprising the d th decomposition *level- d intervals* so that $[0 \dots n]$ is the unique level-0 interval and a level- d interval is partitioned into two disjoint level- $(d+1)$ intervals.

The computation of $A[0 \dots n, 0 \dots n]$ is based on two recursive procedures that utilize an auxiliary array $P[0 \dots n, 0 \dots n]$. Their definitions below are adapted to our setting.

compute(I): given a level- d interval I , compute the entries $A[i, j]$ for $i, j \in I$ with $i \leq j$.

complete(I, J): given two level- d intervals I, J (with I to the left of J), compute the entries $A[i, j]$ for $i \in I$ and $j \in J$, assuming that:

- the remaining entries $A[i, j]$ with $\min I \leq i \leq j \leq \max J$ have already been computed,
- P currently stores values $P[i, j] = \min\{A[i, m] + A[m, j] : m \in M \cap (\max I \dots \min J)\}$ for each $i \in I$ and $j \in J$.

The following procedure is the workhorse of the algorithm, allowing speed-ups through fast matrix multiplication.

update(I, K, J): given three distinct level- d intervals I, K, J , set $P[i, j] := \min\{P[i, j], \min\{A[i, m] + A[m, j] : m \in K \cap M\}\}$ for $i \in I$ and $j \in J$, assuming that all entries $A[i, m]$ and $A[m, j]$ with $i \in I, j \in J$, and $m \in K$ have already been computed.

In the original hierarchical decomposition, each level- d interval of length at least 2 is partitioned into two equal halves, whereas intervals of length 1 form the recursion base case. In this work, our goal is to make the algorithm more efficient when $|M|$ is much smaller than n , and we need to control both $|I|$ and $|I \cap M|$ for the intervals I in the decompositions. Moreover, we do not need to partition intervals I with $|I \cap M| = O(1)$. To facilitate these goals, we define the *weight* of an interval $I \subseteq [0 \dots n]$ as

$$w(I) := |I \cap M| + \frac{|I \setminus M| \cdot |M|}{n}.$$

Note that the weight of I is the sum of the singleton intervals constituting it. We also denote $W = w([0 \dots n])$; note that $|M| \leq W \leq 2|M|$. For $0 \leq d \leq \lfloor \log_2 W \rfloor$, we define a level- d decomposition of $[0 \dots n]$ into disjoint intervals. The only level-0 interval is $[0 \dots n]$. For $0 \leq d < \lfloor \log_2 W \rfloor$, each level- d interval I is decomposed into two level- $(d+1)$ intervals I_ℓ and I_r (the left and right subinterval) so that $w(I_\ell)$ and $w(I_r)$ are as balanced as possible. Since the weight of a singleton does not exceed 1, we can achieve $w(I_\ell), w(I_r) \leq \frac{1}{2}(w(I) + 1)$. A simple induction shows that, for every level- d interval

I , we have $w(I) \leq 2^{-d}(W-1)+1$. Consequently, $|I \cap M| \leq w(I) \leq 2^{1-d}W$, $|I \cap M| \leq w(I) \leq 2^{2-d}|M|$, and $|I| \leq \frac{n}{|M|}w(I) \leq 2^{2-d}n$.

LEMMA 4.1. *The update function for level- d intervals, $2 \leq d \leq \lfloor \log_2 W \rfloor$, costs $O((2^{-d}n)^{\bar{\omega}})$ time, where $\bar{\omega}$ is the exponent of Theorem 1.2 for $\alpha = \log_n |M|$.*

PROOF. This operation can be implemented via a min-plus product of a $|I| \times |K \cap M|$ matrix, storing $A[i, m]$ for $i \in I$ and $m \in K \cap M$, with a $|K \cap M| \times |J|$ matrix, storing $A[m, j]$ for $m \in K \cap M$ and $j \in J$. By Property (2), the first matrix is 1-column-BD, and the second matrix is 1-row-BD, so Theorem 1.2 can be applied to compute the product. Note that $|I|, |J| \leq 2^{2-d}n$ and $|K \cap M| \leq 2^{2-d}|M| \leq (2^{2-d}n)^{\log_n |M|}$. Consequently, the product can be computed in $O((2^{2-d}n)^{\bar{\omega}})$ time. \square

LEMMA 4.2. *The complete function for level- d intervals, $1 \leq d \leq \lfloor \log_2 W \rfloor$, costs $O((2^{-d}n)^{\bar{\omega}})$ time, where $\bar{\omega}$ is the exponent of Theorem 1.2 for $\alpha = \log_n |M|$.*

PROOF. If $d = \lfloor \log_2 W \rfloor$, then the naive implementation costs $O(|I| \cdot |J| \cdot |M \cap (I \cup J)|) = O((2^{-d}n)^2)$ time because $|M \cap (I \cup J)| = O(1)$. Otherwise, we proceed recursively based on the decompositions $I = I_\ell \cup I_r$, $J = J_\ell \cup J_r$ into level- $(d+1)$ intervals:

- (1) complete(I_r, J_ℓ),
- (2) update(I_ℓ, I_r, J_ℓ),
- (3) complete(I_ℓ, J_ℓ),
- (4) update(I_r, J_ℓ, J_r),
- (5) complete(I_r, J_r),
- (6) update(I_ℓ, I_r, J_r),
- (7) update(I_ℓ, J_ℓ, J_r),
- (8) complete(I_ℓ, J_r).

This sequence of steps is at the heart of Valiant's recursion [25, 30], and it is easy to check that the pre-conditions for each application of complete and update are satisfied. Since $I \times J = (I_r \times J_\ell) \cup (I_r \times J_r) \cup (I_\ell \times J_\ell) \cup (I_\ell \times J_r)$, the post-conditions of the four applications of complete guarantees that the resulting entries $A[i, j]$ for $i \in I$ and $j \in J$ are determined correctly. As for the running time, it suffices to observe that there are four applications of update and four recursive applications of complete, all on level- $(d+1)$ intervals. Since $\bar{\omega} > 2$, the runtime is $O((2^{-d}n)^{\bar{\omega}})$. \square

LEMMA 4.3. *The compute function for a level- d interval, $0 \leq d \leq \lfloor \log_2 W \rfloor$, costs $O((2^{-d}n)^{\bar{\omega}})$ time, where $\bar{\omega}$ is the exponent of Theorem 1.2 for $\alpha = \log_n |M|$.*

PROOF. If $d = \lfloor \log_2 W \rfloor$, then the naive implementation costs $O(|I| \cdot |I| \cdot |M \cap I|) = O((2^{-d}n)^2)$ time because $|M \cap I| = O(1)$. Otherwise, we proceed recursively based on the decomposition $I = I_\ell \cup I_r$ into level- $(d+1)$ intervals:

- (1) compute(I_ℓ)
- (2) compute(I_r)
- (3) complete(I_ℓ, I_r)

Again, it is easy to see that the pre-condition for complete is satisfied. Since each $(i, j) \in I^2$ with $i \leq j$ belongs to $I_\ell^2 \cup I_r^2 \cup (I_\ell \times I_r)$, the post-conditions of the two applications of compute and the one application of complete guarantee that all entries $A[i, j]$ with $(i, j) \in I^2$ and $i \leq j$ are computed correctly. As for the running time, it suffices to observe that there is one application of update and two recursive applications of complete, all on level- $(d+1)$ intervals. Since $\bar{\omega} > 1$, the runtime is $O((2^{-d}n)^{\bar{\omega}})$. \square

In particular, Lemma 4.3 applied to the initial recursive call $\text{compute}([0..n])$ yields the following.

COROLLARY 4.4. *The values $A[i, j]$ for $0 \leq i \leq j \leq n$ can be computed in $\tilde{O}(n^{\bar{\omega}})$ time, where $\bar{\omega}$ is the exponent of Theorem 1.2 for $\alpha = \log_n |M|$.*

5 FASTER PROCEDURE FOR CLUSTERS

In this section, we apply Corollary 4.4 to develop a faster implementation of the subroutine for processing clusters. For this, we shall prove that the dynamic-programming procedure following from Lemma 3.3 satisfies the two conditions of Valiant recursion. Formally, for a fixed cluster C , this procedure operates on an array $A[E(C), E(C)]$ resulting in $A[i, j] = \min\{\mathcal{D}[i, j], k + 1\}$ for all $i, j \in E(C)$ with $j \geq i$. Inspecting the recursion in Lemma 3.3, it is easy to see that this dynamic program satisfies condition (1) of Valiant recursion. As for condition (2), we need to prove that, for every $m \in M(C)$, the subsequent entries $A[i, m]$ (with $i \in E(C)$ and $i \leq m$) and $A[m, j]$ (with $j \geq m$ and $j \in E(C)$) differ by at most one. By symmetry, we focus without loss of generality on the former case. Let $A[i, m]$ and $A[i', m]$ be such subsequent entries with $i < i' \leq m$. If $i' = i + 1$, then $|\mathcal{D}[i, m] - \mathcal{D}[i + 1, m]| \leq 1$ because $S[i..m]$ and $S[i + 1..m]$ are at distance 1. Otherwise, we must have $i = a_q + 2k$ and $i' = d_q - 2k$ for some $q \in [1..r]$. If $\mathcal{D}[i, m] < k$ or $\mathcal{D}[i', m] < k$, then Claim 3.7 applied for a trapezoid $(a_q + 1, a_q + 2k - 1, d_q - 2k - 1, d_q - 1)$ implies $m \in (a_q + 1..d_q - 1)$, which contradicts $m \in M(C)$. Consequently, $A[i, m], A[i', m] \in \{k, k + 1\}$. In both case $|A[i, m] - A[i', m]| \leq 1$ holds as claimed.

Thus, Corollary 4.4 implies that C can be processed in time $\max\{|E(C)|, |M(C)|^2\}^{\bar{\omega}}$, where $\bar{\omega}$ is the exponent of Theorem 1.2 for $\alpha = \frac{1}{2}$. Due to $\sum_C |E(C)| = O(k^2)$ and $\sum_C |M(C)| = O(k)$, the total cost of processing all clusters is $\tilde{O}(k^{2\bar{\omega}})$, which is $O(k^{4.544184})$ with high probability or $O(k^{4.853059})$ deterministically. The remaining components of the algorithm from Section 3 cost $O(n + k^3)$ time, so this yields Theorem 1.1.

6 (MIN, +)-PRODUCT OF RECTANGULAR MATRICES WITH PARTIALLY BOUNDED DIFFERENCE

In this section we consider the problem of computing $(\min, +)$ product of rectangular matrices which are partially bounded difference. We first describe a deterministic algorithm that follows the ideas of Bringmann et al. [6]. We first describe a randomized algorithm follows the ideas of [6] and then we describe how to derandomize this algorithm to get a deterministic algorithm. Then, in Section 6.5 we describe a randomized algorithm which based on the recent result of Chi et al. [8], which improve upon the results of Bringmann et al. [6], but does not introduce any derandomization method.

Bringmann et al. [6] designed algorithms for $(\min, +)$ -product of two square matrices A and B of size $n \times n$ for two special cases. In the first case, both input matrices are fully-BD matrices, and the runtime for this case is $\tilde{O}(n^{2.8244})$ randomized or $\tilde{O}(n^{2.8603})$ deterministic. In the second case, only one of the input matrices is assumed to be either column-BD or row-BD, and the randomized runtime for the second case is $\tilde{O}(n^{2.9217})$ time.

In our setting, the two matrices are rectangular, the matrix A is column-BD, and the matrix B is row-BD. One could extend the solution of Bringmann et al. [6] for the second case to work on rectangular matrices. However, in our case, we take advantage of the additional structure of B in order to further reduce the runtime. Our algorithm is very similar to the algorithms of Bringmann et al. [6], and we follow their structure. We emphasize that our exposition is given here only for the sake of completeness; all the main ideas come from [6].

The (first) algorithm of Bringmann et al. [6] considers the matrices A, B , and $C = A \star B$ as composed of blocks of size $\Delta \times \Delta$, for some parameter Δ which is a small polynomial in n , and n is an integer multiple of Δ . The algorithm considers the right-bottom corner of each block as the *representative of the block*. The intuition is that, due to the bounded difference, the value of any two positions inside the same block differs by $O(\Delta)$, and so in order to compute a rough estimation of the output it suffices to consider only the representatives of the blocks.

Due to the following lemma, in our case where A is column-BD and B is row-BD, the matrix $C = A \star B$ is fully-BD.

LEMMA 6.1. *Let $\alpha > 0$. Let A and B be two integer matrices of sizes $n \times n^\alpha$ and $n^\alpha \times n$, respectively, such that A is column-BD and B is row-BD. Then, $C := A \star B$ is fully-BD.*

PROOF. Let $(i, j) \in [n-1] \times [n]$. By definition, $C[i, j] = \min_{\ell \in [n^\alpha]} \{A[i, \ell] + B[\ell, j]\}$. By the assumption on A , for any $(i, j) \in [n-1] \times [n^\alpha]$ we have $A[i, j] - 1 \leq A[i+1, j] \leq A[i, j] + 1$. Hence,

$$\begin{aligned} C[i+1, j] &= \min_{\ell \in [n^\alpha]} \{A[i+1, \ell] + B[\ell, j]\} \leq \min_{\ell \in [n^\alpha]} \{(A[i, \ell] + 1) + B[\ell, j]\} \\ &= \min_{\ell \in [n^\alpha]} \{A[i, \ell] + B[\ell, j]\} + 1 = C[i, j] + 1. \end{aligned}$$

Similarly,

$$\begin{aligned} C[i+1, j] &= \min_{\ell \in [n^\alpha]} \{A[i+1, \ell] + B[\ell, j]\} \geq \min_{\ell \in [n^\alpha]} \{(A[i, \ell] - 1) + B[\ell, j]\} \\ &= \min_{\ell \in [n^\alpha]} \{A[i, \ell] + B[\ell, j]\} - 1 = C[i, j] - 1. \end{aligned}$$

Thus, $|C[i, j] - C[i+1, j]| \leq 1$. The proof that $|C[i, j] - C[i, j+1]| \leq 1$ is symmetric (based on the assumption on B). \square

The algorithm of Bringmann et al. [6] is composed of three phases. We first describe the randomized algorithm, and in Section 6.4 we describe how to derandomize the algorithm.

6.1 Phase 1: Finding \tilde{C} – An Approximation of C with $O(\Delta)$ Additive Error.

Let Δ be a positive integer parameter (to be fixed later) and assume that n is divisible by Δ (we can always assume without loss of generality that both n and Δ are powers of 2). We partition $[n]$ into intervals of length Δ , that are defined as follows: for each $1 \leq i' \leq n$ which is divisible by Δ , let $I(i') = \{i \in [n] \mid i' - \Delta < i \leq i'\}$. For every i', j' which are divisible by Δ , the algorithm first computes $\tilde{C}[i', j'] = C[i', j']$ exactly (naively) and then, for any $(i, j) \in I(i') \times I(j')$ the algorithm sets $\tilde{C}[i, j] \leftarrow C[i', j']$. We emphasize that the definition of $\tilde{C}[i', j']$ in our algorithm differs from the definition in [6], since in our case A is not guaranteed to be row-BD and B is not guaranteed to be column-BD. Moreover, for pairs i', j' which are divisible by Δ we define $\tilde{C}[i', j']$ to be exactly $C[i', j']$, while in [6], $\tilde{C}[i', j']$ is an approximation of $C[i', j']$. Nevertheless, the following lemma is similar in flavor to [6, Lemma 2.1].

LEMMA 6.2. *For i', j' which are divisible by Δ and $(i, j) \in I(i') \times I(j')$ we have $|\tilde{C}[i, j] - C[i, j]| = |C[i', j'] - C[i, j]| \leq 2\Delta$.*

PROOF. Notice that $\tilde{C}[i, j] = C[i', j']$ by definition, so we only need to prove that $|C[i', j'] - C[i, j]| \leq 2\Delta$.

Notice that one can move from $C[i, j]$ to $C[i', j]$ in $i' - i < \Delta$ vertical steps. Moreover, one can move further from $C[i', j]$ to $C[i', j']$ in $j' - j < \Delta$ horizontal steps. Thus, one can move from $C[i, j]$ to $C[i', j']$ in less than 2Δ steps. By Lemma 6.1 each step increases the difference between $C[i, j]$ and the current entry by at most 1, hence the claim follows. \square

Runtime. Notice that computing $\tilde{C}[i', j']$ for all i', j' which are divisible by Δ costs $\tilde{O}((\frac{n}{\Delta})^2 n^\alpha) = \tilde{O}(\frac{n^{2+\alpha}}{\Delta^2})$ time. In addition, the process of filling the rest of the matrix \tilde{C} costs $O(n^2)$ time. Hence, the total running time of the first phase is $\tilde{O}(\frac{n^{2+\alpha}}{\Delta^2} + n^2)$ time.

6.2 Phase 2: Reduction to (min, +)-Product with Small Entries.

The following lemma is useful for the second phase.

LEMMA 6.3 (CF. [4, LEMMA 1]). *Let $R \in \mathbb{N}$. Let A and B be $n \times s$ and $s \times n$ matrices, respectively, with entries in $\{-R, -R + 1, \dots, R\} \cup \{\infty\}$. Then $A \star B$ can be computed in $\tilde{O}(R \cdot M(n, s, n))$ time.*

The second phase has ρ iterations (ρ will be fixed later as a small polynomial in n). In each iteration the algorithm chooses i' and j' independently and uniformly at random from $[n]$. Let A' be the matrix where $A'[i, \ell] = A[i, \ell] + B[\ell, j'] - \tilde{C}[i, j']$ and let B' be the matrix where $B'[\ell, j] = B[\ell, j] - B[\ell, j'] + \tilde{C}[i', j'] - \tilde{C}[i', j]$. Let $C^r = A' \star B'$. For each i, j , we have $C[i, j] = C^r[i, j] + \tilde{C}[i, j'] - \tilde{C}[i', j'] + \tilde{C}[i', j]$, and therefore one can compute C from C^r . However, computing C^r exactly seems to be inefficient, and so [6] proved that a partial computation of C^r can be done efficiently enough and still result in useful information regarding C .

Triples. By definition of (min, +)-product, for every (i, j) we have $C[i, j] = \min_{\ell \in [n^\alpha]} \{A[i, \ell] + B[\ell, j]\}$. Thus, we associate with each entry $C[i, j]$ the triples $(i, 1, j), (i, 2, j), \dots, (i, n^\alpha, j)$ which refer to the n^α different indices considered in the (min, +)-product definition. The value of $C[i, j]$ is $A[i, \ell] + B[\ell, j]$ for at least one value $\ell \in [n^\alpha]$; for such values ℓ we say that (i, ℓ, j) is *relevant*.

Let \hat{A}^r (\hat{B}^r) be the matrix A' (B') after replacing each entry whose absolute value is larger than 48Δ with ∞ , thereby effectively allowing a (min, +)-product to ignore those entries. Let $P^r = \hat{A}^r \star \hat{B}^r$ and define $\hat{C}^r[i, j] = P^r[i, j] + \tilde{C}[i, j'] - \tilde{C}[i', j'] + \tilde{C}[i', j]$. Notice that $\hat{C}^r[i, j] \geq C[i, j]$. For triples (i, ℓ, j) where both $|A'[i, \ell]| \leq 48\Delta$ and $|B'[\ell, j]| \leq 48\Delta$ we have $\hat{C}^r[i, j] \leq A[i, \ell] + B[\ell, j]$. We call such triples *covered* by the r th sample. Finally, \hat{C} is defined as the entry-wise minimum of \hat{C}^r for $r = 1, 2, \dots, \rho$. One can compute each P^r , and eventually \hat{C} , using Lemma 6.3. However, we employ a different approach which introduces additional speed-up.

Bounding the Number of Relevant and Not Covered Triples. Our goal is to show that with high probability, most of the relevant triples (i, ℓ, j) are covered by at least one of the ρ samples. Notice that for any pair (i, j) we have $\hat{C}[i, j] = C[i, j]$ if and only if there exist $\ell \in [n^\alpha]$ and $1 \leq r \leq \rho$ such that (i, ℓ, j) is relevant and is covered by the r th sample. Since the relevant triples are unknown, we use a weaker definition of relevance and also a weaker definition of being covered:

Definition 6.1 (cf. [6, Definition 2.2]). We call a triple $(i, \ell, j) \in [n] \times [n^\alpha] \times [n]$

- strongly relevant if $A[i, \ell] + B[\ell, j] = C[i, j]$;
- weakly relevant if $|A[i, \ell] + B[\ell, j] - C[i, j]| \leq 16\Delta$;⁵
- strongly r -uncovered if for all $1 \leq r' \leq r$ we have $|A^{r'}[i, \ell]| > 48\Delta$ or $|B^{r'}[\ell, j]| > 48\Delta$;
- weakly r -uncovered if for all $1 \leq r' \leq r$ we have $|A^{r'}[i, \ell]| > 40\Delta$ or $|B^{r'}[\ell, j]| > 40\Delta$.

A triple is called strongly (weakly) uncovered if it is strongly (weakly) ρ -uncovered. Finally, a triple is strongly (weakly) r -covered if it is not strongly (weakly) r -uncovered.

The following lemma is a straightforward generalization of [6, Lemma 3.1] to our case.

LEMMA 6.4 (BASED ON [6, LEMMA 3.1]). *With high probability for any $1 \leq r \leq \rho$, the number of weakly relevant triples that are also weakly r -uncovered is $\tilde{O}(n^{1.5+\alpha} + n^{2+\alpha}/r^{1/3})$.*

⁵We actually can use smaller constants, since by Lemma 6.2 the additive error is at most 2Δ , while in [6] the additive error in \tilde{C} is bounded by 4Δ . However, we prefer to use the original constants not to modify the claims and the proofs much.

SKETCH. The proof of [6, Lemma 3.1] shows that for every ℓ , with high probability, the number of weakly relevant triples (i, ℓ, j) that are also weakly r -uncovered is $\tilde{O}(n^{1.5} + n^2/r^{1/3})$. Since in our case there are only n^α possible values of ℓ , then, with high probability, the number of weakly relevant triples that are also weakly r -uncovered (among all possible values of ℓ) is $\tilde{O}(n^\alpha(n^{1.5} + n^2/r^{1/3})) = \tilde{O}(n^{1.5+\alpha} + n^{2+\alpha}/r^{1/3})$. \square

Efficient Method for Covering Relevant Triples. Instead of computing the $(\min, +)$ -product $P^r = \hat{A}^r \star \hat{B}^r$ in the r th round, we follow the method of [6] and use their additional insight to reduce the runtime of the algorithm. The proof of [6, Lemma 3.1] has the property that for any round r , the triples that are counted as covered by the r th round are exactly triples of the form (i, ℓ, j) such that the triple (i^r, ℓ, j^r) is weakly relevant and weakly $(r-1)$ -uncovered. Therefore, in order to cover the same triples at round r , after the algorithm chooses i^r and j^r , the algorithm computes L_r which is the set of ℓ values such that (i^r, ℓ, j^r) is both weakly relevant and weakly $(r-1)$ -uncovered. The algorithm removes all the columns $\ell \notin L_r$ from \hat{A}^r and all the rows $\ell \notin L_r$ from \hat{B}^r . Let $s_r = |L_r|$ be the number of surviving ℓ 's in round r . The algorithm computes $\hat{P}^r = \hat{A}^r \star \hat{B}^r$ using Lemma 6.3, in $O(\Delta \cdot M(n, s_r, n))$ time.

Runtime. For implementing the filtering in the r th iteration the algorithm checks n^α triples. For each triple the test takes $O(r)$ time (to check that the triple is $(r-1)$ -uncovered). Thus, the total time of the filtering takes in total $O(n^\alpha \cdot \rho^2) \leq O(\rho n^{1+\alpha})$ time. The runtime of computing the matrix products is $\sum_{r=1}^{\rho} \tilde{O}(\Delta M(n, s_r, n))$ time, which due to the analysis made in [6, Lemma 3.2] is bounded by $\tilde{O}(\rho \Delta \cdot M(n, \frac{n^\alpha}{\rho^{1/3}}, n))$ time with high probability. Thus, the total running time of the second phase is $\tilde{O}(\rho \Delta \cdot M(n, \frac{n^\alpha}{\rho^{1/3}}, n) + \rho \cdot n^{1+\alpha})$ time with high probability.

6.3 Phase 3: Complete the Relevant Uncovered Triples.

In the third phase, the goal is to find all the triples (i, ℓ, j) that are both strongly relevant and strongly uncovered, and use each such triple to update \hat{C} by setting $\hat{C}[i, j] = \min\{\hat{C}[i, j], A[i, \ell] + B[\ell, j]\}$. In order to identify all of these triples, we introduce the notions of approximately relevant and approximately covered.

Definition 6.2 (cf. [6, Definition 2.6]). We call a triple $(i, \ell, j) \in I(i') \times [n^\alpha] \times I(j')$

- approximately relevant if $|A[i', \ell] + B[\ell, j'] - \tilde{C}[i', j']| \leq 8\Delta$, and
- approximately r -uncovered if for all $1 \leq r' \leq r$ we have either $|A^{r'}[i', \ell]| > 44\Delta$ or $|B^{r'}[\ell, j']| > 44\Delta$ (or both).

A triple is called approximately uncovered if it is approximately ρ -uncovered. A triple that is both approximately relevant and approximately uncovered is called *interesting*.

A triple (i', ℓ, j') where i' and j' are divisible by Δ and $\ell \in [n^\alpha]$ is called a *representative triple*. A useful property of these definitions is that for a representative triple (i', ℓ, j') , all the triples $(i, \ell, j) \in I(i') \times \{\ell\} \times I(j')$ (and, in particular, the representative triple itself) are similar in the sense that either all of these triples are interesting, or all of these triples are not interesting. Thus, in order to find all the interesting triples, it suffices to test for every representative triple if the triple is interesting.

The following lemma describes the relationships between the different definitions of strongly, approximately and weakly relevant (uncovered) triples.

LEMMA 6.5 ([6, LEMMA 2.7]). *Any strongly relevant triple is also approximately relevant. Any approximately relevant triple is also weakly relevant. The same holds also with “relevant” replaced by “ r -uncovered”.*

To test whether a given (representative) triple is interesting, the algorithm tests if the triple is both approximately relevant and approximately uncovered. The test whether a given triple is approximately relevant takes $O(1)$ time (see Definition 6.2). In order to find all approximately uncovered representative triples, the algorithm employs rectangular Boolean matrix multiplication as follows. For each $\ell \in [n^\alpha]$ let U^ℓ be a matrix of size $\frac{n}{\Delta} \times \rho$ such that $U^\ell[x, r] = 1$ if and only if $|A^r[x\Delta, \ell]| \leq 44\Delta$ (see Definition 6.2). Similarly, let V^ℓ be a matrix of size $\rho \times \frac{n}{\Delta}$ such that $V^\ell[r, y] = 1$ if and only if $|B^r[\ell, y\Delta]| \leq 44\Delta$. Let Z^ℓ be the Boolean matrix product $U^\ell \cdot V^\ell$. Then, by Definition 6.2, for every $i' = x\Delta$, $j' = y\Delta$ and $\ell \in [n^\alpha]$, we have (i', ℓ, j') is approximately uncovered if and only if $Z^\ell[\frac{i'}{\Delta}, \frac{j'}{\Delta}] = 1$. Thus, the algorithm uses n^α rectangular Boolean matrix multiplications, in order to find all the interesting representative triples. At the last step of the third phase, for any (i', ℓ, j') which is known to be interesting from the previous step, the algorithm iterates over all $(i, \ell, j) \in I(i') \times \{\ell\} \times I(j')$ and updates $\hat{C}[i, j] \leftarrow \min\{\hat{C}[i, j], A[i, \ell] + B[\ell, j]\}$.

Runtime. At the first step of the phase, the algorithm finds all the interesting representative triples. The time cost of finding all the approximately relevant representative triples is $O(\frac{n^{2+\alpha}}{\Delta^2})$. The time cost of finding all the approximately uncovered representative triples is $O(n^\alpha \cdot M(\frac{n}{\Delta}, \rho, \frac{n}{\Delta}))$ time. Then, in the second step of the phase, the algorithm iterates over all the interesting triples, in time linear in the number of such triples which by Lemmas 6.4 and 6.5 is, with high probability, $\tilde{O}(n^{1.5+\alpha} + n^{2+\alpha}/\rho^{1/3})$. In total, the third phase takes $O(n^\alpha \cdot M(\frac{n}{\Delta}, \rho, \frac{n}{\Delta}) + n^{1.5+\alpha} + n^{2+\alpha}/\rho^{1/3})$ time with high probability.

6.4 Derandomization.

The only randomized part of the algorithm described above is the process of choosing i^r and j^r from $[n]$. The derandomization technique introduced by Bringmann et al. [6, Section 3.4] is based on the following observations.

Choosing Only Indices which are Divisible by Δ . Bringmann et al. [6, Section 3.4] show that the algorithm has guarantees similar to Lemma 6.4 even if i^r and j^r are chosen uniformly at random from the set of indices in $[n]$ which are divisible by Δ . The following lemma states that after this change in the algorithm, the number of interesting triples that will be covered in the third phase is still $\tilde{O}(n^{1.5+\alpha} + n^{2+\alpha}/\rho^{1/3})$. Hence, the running time of the third phase will not increase.

LEMMA 6.6 ([6, SECTION 3.4]). *With high probability, for any $1 \leq r \leq \rho$ the number of approximately relevant, approximately r -uncovered triples is $\tilde{O}(n^{1.5+\alpha} + n^{2+\alpha}/r^{1/3})$.*

Deterministic greedy choice. Let r be an iteration of the algorithm and let $\ell \in [n^\alpha]$. Bringmann et al. [6] showed that one could count for every pair (i', j') , where i' and j' are divisible by Δ , the number of triples (i, ℓ, j) which are approximately relevant and approximately $(r - 1)$ -uncovered but will be covered by choosing $(i^r, j^r) \leftarrow (i', j')$. The time cost of computing the counts for all pairs (i', j') is $O((\frac{n}{\Delta})^\omega)$ time for fixed ℓ and r . Thus, a greedy algorithm that computes in the r th iteration the pair (i', j') which maximizes the number of triples that will be covered, suffices to guarantee results which are at least as good as the expectation of the random process. The additional time in total for the derandomization of the second phase is $O(\rho n^\alpha (\frac{n}{\Delta})^\omega)$, and the guarantees of Lemma 6.6 hold.

6.4.1 Total Runtime. The additional runtime introduced by the derandomization process is exactly $O((\frac{n}{\Delta})^\omega \cdot n^\alpha) = O(\frac{n^{\omega+\alpha}}{\Delta^\omega})$. Summing up the running time of all the phases, the total runtime

of the algorithm is:

$$\begin{aligned} \tilde{O} \left(\frac{n^{2+\alpha}}{\Delta^2} + n^2 + \rho \Delta M(n, \frac{n^\alpha}{\rho^{1/3}}, n) + n^\alpha \cdot M(\frac{n}{\Delta}, \rho, \frac{n}{\Delta}) + n^{1.5+\alpha} + \frac{n^{2+\alpha}}{\rho^{1/3}} + \frac{n^{\omega+\alpha}}{\Delta^\omega} \right) \\ = \tilde{O} \left(\rho \Delta M(n, \frac{n^\alpha}{\rho^{1/3}}, n) + n^\alpha \cdot M(\frac{n}{\Delta}, \rho, \frac{n}{\Delta}) + \frac{n^{2+\alpha}}{\rho^{1/3}} + n^{1.5+\alpha} + \frac{n^{\omega+\alpha}}{\Delta^\omega} \right). \end{aligned}$$

We use the results of Le Gall and Urrutia [14]. We set $\Delta = n^\delta$ and $\rho = n^s$ for some s, δ positive integers. We will focus on $\alpha = \frac{1}{2}$. Hence, the running time is:

$$\tilde{O} \left(n^{\delta+s} M(n, n^{1/2-s/3}, n) + n^{1/2} \cdot M(n, n^{\frac{s}{1-\delta}}, n)^{1-\delta} + n^{5/2-s/3} + n^{1/2+(1-\delta)\omega} \right).$$

We optimize the running time by choosing $\delta = 0.1881$ and $s = 0.220425$. By interpolating the results of [14] the total running time of the algorithm is $O(n^{2.426524})$.

COROLLARY 6.7. *Let $0 < \alpha \leq \frac{1}{2}$. Let A and B be integer matrices of sizes $n \times n^\alpha$ and $n^\alpha \times n$, respectively, such that A is column-BD and B is row-BD. There exists a deterministic algorithm that computes $A \star B$ whose runtime is $\tilde{O}(n^{2.426524})$.*

6.5 Randomized Algorithm

In a recent article, Chi et al. [8] introduced an efficient randomized algorithm for computing the $(\min, +)$ -product of two square matrices A and B of size $n \times n$, for a broad range of structured matrices. In particular, their algorithm is applicable to our case, except for the fact that our matrices are *rectangular* and they described their algorithm for *squared* matrices. Notice that one can partition each rectangular matrix of size $n \times n^\alpha$ or $n^\alpha \times n$ into $n^{1-\alpha}$ sub-matrices of size $n^\alpha \times n^\alpha$, and so computing the $(\min, +)$ -product of the two original matrices is reduced to $n^{2(1-\alpha)}$ computations of $(\min, +)$ -products of squared matrices whose size is $n^\alpha \times n^\alpha$. Such an approach leads to an algorithm that costs $O(n^{2(1-\alpha)} n^{\alpha(3+\omega)/2}) = O(n^{2-\alpha/2+\alpha\omega/2}) = O(n^{2+(\omega-1)\alpha/2})$ time. Nevertheless, in this section we analyze the runtime obtained by adjusting the algorithm of Chi et al. [8] to our case in order to obtain a more efficient runtime.

We emphasize that the only change to the algorithm of [8] is replacing computations on squared matrices with computations on rectangular matrices. Thus, correctness follows immediately from [8], and our discussion focuses only on the runtime cost of the algorithm for our case. For complete definitions and details, the reader is referred to Chi et al. [8].

Overview. In our setting, the matrix A is column-BD and the matrix B is row-BD. As [8, Section 2] observes, one can reduce the computation of a min-plus product of two such matrices to the computation of a min-plus product of two matrices where the values in the second matrix are row monotone (the values never decrease when scanning a row from left to right), and all of the values in both matrices are integers from the range $[1..cn]$ for some constant $c \in \mathbb{N}$. Thus, for the rest of this section we assume that B is row monotone and all the entries in A and B are integers from the range $[1..cn]$.

Notice that [8] contains two algorithm, the first one is a basic algorithm that provides the main ideas of the article, and the second uses recursion to speedup the basic algorithm. Here, we refer directly to the recursive algorithm. Our algorithm follows the algorithm of [8] by replacing every part of the algorithm that computes a product of square matrices by an algorithm that computes a product of rectangular matrices.

The algorithm of Chi et al. [8] is composed of two parts that both depend on a random prime $p \in [40n^\beta .. 80n^\beta]$ (for some parameter β to be determined⁶). In the first part, the algorithm computes

⁶Chi et al. [8] uses ' α ' to denote this parameter. We use ' β ' instead since ' α ' is already taken.

a matrix C^* such that for every $i, j \in [1..n]$ we have $C_{i,j}^* = \lfloor C_{i,j}/p \rfloor$ for some prime number p . In the second part, the algorithm computes for every $i, j \in [1..n]$ the value $C_{i,j} \bmod p$. Using the fact that $C_{i,j} = \lfloor C_{i,j}/p \rfloor + C_{i,j} \bmod p$, after the execution of the two parts, the algorithm completes the computation of C . We describe the two parts below.

First part. The first part of the algorithm computes the matrix C^* defined as $C_{i,j}^* = \lfloor C_{i,j}/p \rfloor$.

Notice that it is straightforward to compute C^* in $O(n^3)$ time for squared matrices. The algorithm of [8] applies *segment trees* [11] to speedup the computation by a factor of $\tilde{O}(n^\beta)^7$ and so the runtime becomes $\tilde{O}(n^{3-\beta})$ (see [8, Lemma 3.2]). In our case, the straightforward algorithm for computing C^* runs in $O(n^{2+\alpha})$ time, and the application of segments trees in our case provides the same speedup, so the computation of C^* in our case becomes $\tilde{O}(n^{2+\alpha-\beta})$.

Second part. For $h = \lceil \log p \rceil + 1$, the algorithm iteratively computes a sequence of matrices $C^{(h)}, C^{(h-1)}, \dots, C^{(0)}$ where for every $\ell = h, h-1, \dots, 0$ the entry $C_{i,j}^{(\ell)}$ is an approximation of $\lfloor (C_{i,j} \bmod p)/2^\ell \rfloor$. One important property of the sequence of matrices is that $C_{i,j}^{(0)} = C_{i,j} \bmod p$. Each matrix $C^{(\ell)}$ is iteratively computed using $C^{(\ell+1)}$.

In order to describe the computation of the different matrices $C^{(\ell)}$, Chi et al. [8] define segments (unrelated to *segment trees*) as follows. Let $\ell \in [0..h]$, $B_{i,j}^* = \lfloor B_{i,j}/p \rfloor$ and $B_{i,j}^{(\ell)} = \lfloor (B_{i,j} \bmod p)/2^\ell \rfloor$. A *segment* $(i, k, [j_0, j_1])$ with respect to ℓ is a triple such that for every $j_0 \leq j \leq j_1$ we have $B_{k,j}^{(\ell)} = B_{k,j_0}^{(\ell)}$, $B_{k,j}^* = B_{k,j_0}^*$, $C_{k,j}^{(\ell)} = C_{k,j_0}^{(\ell)}$, and $C_{k,j}^* = C_{k,j_0}^*$. Chi et al. [8] showed that for every $\ell \in [0..h]$ there are $O(n^3/2^\ell)$ (maximal) segments. In our case, similar analysis shows that for every $\ell \in [0..h]$ there are $O(n^{2+\alpha}/2^\ell)$ segments. Throughout the iterations, for every ℓ the algorithm of Chi et al. [8] computes $O(1)$ sets of segments $T_b^{(\ell)}$ (for $b \in [-10..10]$), used for computing the matrix $C^{(\ell)}$. An important property proved by [8, Lemma 3.7] is that for every ℓ and b we have $\mathbb{E}[|T_b^{(\ell)}|] = O(n^{3-\beta})$. In particular, [8, Lemma 3.7] proved that the probability of every segment in level ℓ to be in $T_b^{(\ell)}$ is $\tilde{O}(2^\ell/n^\beta)$. In our settings this property still holds. Therefore, in our case $\mathbb{E}[|T_b^{(\ell)}|] = O(\frac{n^{2+\alpha}}{2^\ell}) \cdot \tilde{O}(\frac{2^\ell}{n^\beta}) = O(n^{2+\alpha-\beta})$.

Each iteration of the algorithm of Chi et al. [8] is composed of three phases. In the first phase, the algorithm computes a product of two polynomial matrices, by using fast matrix multiplication, where the entries of each matrix are polynomials of bounded degree $O(p) = O(n^\beta)$. The runtime of this phase is $O(n^{\omega+\beta})$. In our case, the two matrices are of sizes $n \times n^\alpha$ and $n^\alpha \times n$, and the polynomials are also of bounded degree $O(p) = O(n^\beta)$. Therefore, by using fast *rectangular* matrix multiplication, the running time of this phase in our case is generalized to $O(n^{n^\beta \cdot M(n, n^\alpha, n)})$ time. In the second phase, the algorithm of [8] computes $C^{(\ell)}$, using the result of the first phase and the sets $T_b^{(\ell+1)}$. The running time of this computation is dominated by two parts. The first part is reading and manipulating the result matrix of the first phase. The size of the matrix is $n \times n$ and each entry contains a polynomial of bounded degree $O(n^\beta)$. Thus, this part costs $O(n^{2+\beta})$ time, both in [8] and in our case. The second part is dominated by $\sum_b |T_b^{(\ell)}|$, which in [8] is $\tilde{O}(n^{3-\beta})$ and in our case is $\tilde{O}(n^{2+\alpha-\beta})$ in expectation. In the third phase of the ℓ th iteration, the algorithm computes the sets $T_b^{(\ell)}$ using the sets $T_b^{(\ell+1)}$ and the matrix $C^{(\ell)}$. The running time of this phase is $\tilde{O}(\sum_b |T_b^{(\ell+1)}|)$, which in our case is $\tilde{O}(n^{2+\alpha-\beta})$ in expectation.

Thus, the total time cost of the second part of the algorithm of [8] in our case is $\tilde{O}(n^\beta \cdot M(n, n^\alpha, n) + n^{2+\beta} + n^{2+\alpha-\beta}) = \tilde{O}(n^\beta \cdot M(n, n^\alpha, n) + n^{2+\alpha-\beta})$. For $\alpha = 1/2$ we optimize the run-

⁷The notation $\tilde{O}(\cdot)$ suppresses polylogarithmic factors.

ning time by choosing $\beta = 0.227908$. By interpolating the results of [14] the total running time of the algorithm is $\tilde{O}(n^{2.272092})$ in expectation. We mention that adapting the algorithm to compute $C = A \star B$ in $\tilde{O}(n^{2.272092})$ time with high probability is straightforward.

COROLLARY 6.8. *Let $0 < \alpha \leq \frac{1}{2}$. Let A and B be integer matrices of sizes $n \times n^\alpha$ and $n^\alpha \times n$, respectively, such that A is column-BD and B is row-BD. Then there exists a randomized algorithm that computes $A \star B$ in time $\tilde{O}(n^{2.272092})$ with high probability.*

Theorem 1.2 follows from Corollaries 6.7 and 6.8.

REFERENCES

- [1] Amir Abboud, Arturs Backurs, and Virginia Vassilevska Williams. 2015. If the current clique algorithms are optimal, so is Valiant’s parser. In *Proceedings of the 2015 IEEE 56th Annual Symposium on Foundations of Computer Science*. 98–117. DOI: <https://doi.org/10.1109/FOCS.2015.16>
- [2] Alfred V. Aho and Thomas G. Peterson. 1972. A minimum distance error-correcting parser for context-free languages. *SIAM Journal on Computing* 1, 4 (1972), 305–312. DOI: <https://doi.org/10.1137/0201022>
- [3] Josh Alman and Virginia Vassilevska Williams. 2021. A refined laser method and faster matrix multiplication. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms*. 522–539. DOI: <https://doi.org/10.1137/1.9781611976465.32>
- [4] Noga Alon, Zvi Galil, and Oded Margalit. 1997. On the exponent of the all pairs shortest path problem. *Journal of Computer and System Sciences* 54, 2 (1997), 255–262. DOI: <https://doi.org/10.1006/jcss.1997.1388>
- [5] Arturs Backurs and Krzysztof Onak. 2016. Fast algorithms for parsing sequences of parentheses with few errors. In *Proceedings of the 35th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*. 477–488. DOI: <https://doi.org/10.1145/2902251.2902304>
- [6] Karl Bringmann, Fabrizio Grandoni, Barna Saha, and Virginia Vassilevska Williams. 2019. Truly subcubic algorithms for language edit distance and RNA folding via fast bounded-difference min-plus product. *SIAM Journal on Computing* 48, 2 (2019), 481–512. DOI: <https://doi.org/10.1137/17M112720X>
- [7] Shucheng Chi, Ran Duan, Tianle Xie, and Tianyi Zhang. 2022. Faster min-plus product for monotone instances. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing (STOC’22)*. Association for Computing Machinery, New York, NY, 1529–1542. <https://doi.org/10.1145/3519935.3520057>
- [8] Shucheng Chi, Ran Duan, Tianle Xie, and Tianyi Zhang. 2022. Faster min-plus product for monotone instances. (2022), 1529–1542. DOI: <https://doi.org/10.1145/3519935.3520057>
- [9] Noam Chomsky and Marcel Paul Schützenberger. 1963. The algebraic theory of context-free languages. *Studies in Logic* 35 (1963), 118–161. DOI: [https://doi.org/10.1016/s0049-237x\(08\)72023-8](https://doi.org/10.1016/s0049-237x(08)72023-8)
- [10] Debarati Das, Tomasz Kociumaka, and Barna Saha. 2022. Improved approximation algorithms for Dyck edit distance and RNA folding. In *ICALP 2022*, Vol. 229. 49:1–49:20. DOI: <https://doi.org/10.4230/LIPIcs.ICALP.2022.49>
- [11] Mark de Berg, Otfried Cheong, Marc J. van Kreveld, and Mark H. Overmars. 2008. *Computational geometry: Algorithms and Applications, 3rd Edition*. Springer. Retrieved from <https://www.worldcat.org/oclc/227584184>
- [12] Anita Dürr. 2023. *Improved Bounds for Rectangular Monotone Min-Plus Product and Applications*.
- [13] Johannes Fischer and Volker Heun. 2006. Theoretical and practical improvements on the RMQ-problem, with applications to LCA and LCE. In *Proceedings of the Combinatorial Pattern Matching*. 36–48. DOI: https://doi.org/10.1007/11780441_5
- [14] Francois Le Gall and Florent Urrutia. 2018. Improved rectangular matrix multiplication using powers of the Coppersmith-Winograd tensor. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms*. 1029–1046. DOI: <https://doi.org/10.1137/1.9781611975031.67>
- [15] Robin R. Gutell, Jamie Cannone, Zhidi Shang, Yushi Du, and Martin J. Serra. 2000. A story: Unpaired adenosine bases in ribosomal RNAs. *Journal of Molecular Biology* 304, 3 (2000), 335–354. DOI: <https://doi.org/10.1006/jmbi.2000.4172>
- [16] Michael A. Harrison. 1978. *Introduction to Formal Language Theory*. Addison-Wesley Longman Publishing Co., Boston, MA, 1st edition.
- [17] Rajesh Jayaram and Barna Saha. 2017. Approximating language edit distance beyond fast matrix multiplication: Ultra-linear grammars are where parsing becomes hard!. In *Proceedings of the 44th International Colloquium on Automata, Languages, and Programming*. 19:1–19:15. DOI: <https://doi.org/10.4230/LIPIcs.ICALP.2017.19>
- [18] Dexter C. Kozen. 1997. *Automata and Computability* (1st ed.). Springer-Verlag, Berlin.
- [19] Andreas Krebs, Nutan Limaye, and Srikanth Srinivasan. 2011. Streaming algorithms for recognizing nearly well-parenthesized expressions. In *Proceedings of the International Symposium on Mathematical Foundations of Computer Science*. 412–423. DOI: https://doi.org/10.1007/978-3-642-22993-0_38

- [20] Gad M. Landau, Eugene W. Myers, and Jeanette P. Schmidt. 1998. Incremental string comparison. *SIAM Journal on Computing* 27, 2 (1998), 557–582. DOI : <https://doi.org/10.1137/S0097539794264810>
- [21] Gad M. Landau and Uzi Vishkin. 1988. Fast string matching with k differences. *Journal of Computer and System Sciences* 37, 1 (1988), 63–78. DOI: [https://doi.org/10.1016/0022-0000\(88\)90045-1](https://doi.org/10.1016/0022-0000(88)90045-1)
- [22] Lillian Lee. 2002. Fast context-free grammar parsing requires fast boolean matrix multiplication. *Journal of the ACM* 49, 1 (2002), 1–15. DOI : <https://doi.org/10.1145/505241.505242>
- [23] J. Ian Munro and Venkatesh Raman. 2001. Succinct representation of balanced parentheses and static trees. *SIAM Journal on Computing* 31, 3 (2001), 762–776. DOI : <https://doi.org/10.1137/S0097539799364092>
- [24] Gene Myers. 1995. Approximately matching context-free languages. *Information Processing Letters* 54, 2 (1995), 85–92. DOI : [https://doi.org/10.1016/0020-0190\(95\)00007-y](https://doi.org/10.1016/0020-0190(95)00007-y)
- [25] Alexander Okhotin. 2014. Parsing by matrix multiplication generalized to Boolean grammars. *Theoretical Computer Science* 516 (2014), 101–120. DOI : <https://doi.org/10.1016/j.tcs.2013.09.011>
- [26] Sanguthevar Rajasekaran, Sahar Al Seesi, and Reda Ammar. 2009. Improved algorithms for parsing ESLTAGs: A grammatical model suitable for RNA pseudoknots. In *Proceedings of the Bioinformatics Research and Applications*. 135–147. DOI : https://doi.org/10.1007/978-3-642-01551-9_14
- [27] Barna Saha. 2014. The dyck language edit distance problem in near-linear time. In *Proceedings of the 2014 IEEE 55th Annual Symposium on Foundations of Computer Science*. 611–620. DOI : <https://doi.org/10.1109/FOCS.2014.71>
- [28] Barna Saha. 2015. Language edit distance and maximum likelihood parsing of stochastic grammars: Faster algorithms and connection to fundamental graph problems. In *Proceedings of the 2015 IEEE 56th Annual Symposium on Foundations of Computer Science*. 118–135. DOI : <https://doi.org/10.1109/FOCS.2015.17>
- [29] Barna Saha. 2017. Fast space-efficient approximations of language edit distance and RNA folding: An amnesic dynamic programming approach. In *Proceedings of the 2017 IEEE 58th Annual Symposium on Foundations of Computer Science*. 295–306. DOI : <https://doi.org/10.1109/FOCS.2017.35>
- [30] Leslie G. Valiant. 1975. General context-free recognition in less than cubic time. *J. Comput. Syst. Sci.* 10, 2 (1975), 308–315. DOI : [https://doi.org/10.1016/S0022-0000\(75\)80046-8](https://doi.org/10.1016/S0022-0000(75)80046-8)
- [31] Virginia Vassilevska Williams and Yinzhan Xu. 2020. Truly subcubic min-plus product for less structured matrices, with applications. In *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 12–29. DOI: <https://doi.org/10.1137/1.9781611975994.2>

Received 8 June 2022; revised 14 June 2023; accepted 3 September 2023