# Small space and streaming pattern matching with $k$ edits 

Tomasz Kociumaka* ${ }^{* 1}$, Ely Porat ${ }^{2}$, and Tatiana Starikovskaya ${ }^{\dagger 3}$<br>${ }^{1}$ University of California, Berkeley, USA<br>kociumaka@berkeley.edu<br>${ }^{2}$ Bar-Ilan University, Israel<br>porately@cs.biu.ac.il<br>${ }^{3}$ DI/ENS, PSL Research University, France<br>tat.starikovskaya@gmail.com


#### Abstract

In this work, we revisit the fundamental and well-studied problem of approximate pattern matching under edit distance. Given an integer $k$, a pattern $P$ of length $m$, and a text $T$ of length $n \geq m$, the task is to find substrings of $T$ that are within edit distance $k$ from $P$. Our main result is a streaming algorithm that solves the problem in $\tilde{\mathcal{O}}\left(k^{5}\right)$ space and $\tilde{\mathcal{O}}\left(k^{8}\right)$ amortised time per character of the text, providing answers correct with high probability. (Hereafter, $\tilde{\mathcal{O}}(\cdot)$ hides a poly $(\log n)$ factor.) This answers a decade-old question: since the discovery of a poly $(k \log n)$-space streaming algorithm for pattern matching under Hamming distance by Porat and Porat [FOCS 2009], the existence of an analogous result for edit distance remained open. Up to this work, no poly $(k \log n)$-space algorithm was known even in the simpler semi-streaming model, where $T$ comes as a stream but $P$ is available for readonly access. In this model, we give a deterministic algorithm that achieves slightly better complexity.

Our central technical contribution is a new space-efficient deterministic encoding of two strings, called the greedy encoding, which encodes a set of all alignments of cost $\leq k$ with a certain property (we call such alignments greedy). On strings of length at most $n$, the encoding occupies $\tilde{\mathcal{O}}\left(k^{2}\right)$ space. We use the encoding to compress substrings of the text that are close to the pattern. In order to do so, we compute the encoding for substrings of the text and of the pattern, which requires read-only access to the latter.

In order to develop the fully streaming algorithm, we further introduce a new edit distance sketch parametrised by integers $n \geq k$. For any string of length at most $n$, the sketch is of size $\tilde{\mathcal{O}}\left(k^{2}\right)$ and it can be computed with an $\tilde{\mathcal{O}}\left(k^{2}\right)$-space streaming algorithm. Given the sketches of two strings, in $\tilde{\mathcal{O}}\left(k^{3}\right)$ time we can compute their edit distance or certify that it is larger than $k$. This result improves upon $\tilde{\mathcal{O}}\left(k^{8}\right)$-size sketches of Belazzougui and Zhu [FOCS 2016] and very recent $\tilde{\mathcal{O}}\left(k^{3}\right)$-size sketches of Jin, Nelson, and Wu [STACS 2021].


## 1 Introduction

In the pattern matching problem, given two strings, a pattern $P$ of length $m$ and a text $T$ of length $n$, one must find all substrings of the text equal to the pattern. This is a fundamental problem of string processing with a myriad of applications in such fields as computational biology, information retrieval, and signal processing, to mention just a few. However, in many

[^0]applications, retrieving substrings that are exactly equal to the pattern is not enough, and one must search for substrings merely similar to the pattern. This task, which is often referred to as approximate pattern matching problem, can be formalised in the following way: for each position $i$ in the text, compute the smallest distance $d_{i}$ between $P$ and any substring of $T$ that ends at position $i$. In string processing, the two most popular distances are the Hamming distance and the edit distance. Recall that the Hamming distance between two equal-length strings is the number of mismatching pairs of characters of the strings. The edit distance of two strings, not necessarily of equal lengths, is the smallest number of edits (character insertions, deletions, and substitutions) needed to transform one string into the other. Due to its practical importance, the approximate pattern matching problem has been extensively studied in the literature, originally in the classic setting with input strings explicitly stored in memory.

In general, computing Hamming distance is easier and can be considered as a preliminary step towards tackling edit distance. The first solution for approximate pattern matching under the Hamming distance was given by Abrahamson [2] and, independently, Kosaraju [40]; based on the fast Fourier transform, it spends $\mathcal{O}(n \sqrt{m \log m})$ time to compute the Hamming distance between the pattern and all the length $-m$ substrings of the text. Up to date, no algorithms improve upon this time complexity for the general version of approximate pattern matching under Hamming distance, but there are better solutions when one is interested only in the distances not exceeding a given $k$, a variant known as the $k$-mismatch problem. The first algorithm for the $k$-mismatch problem was given by Landau and Vishkin [42], who improved the running time to $\mathcal{O}(k n)$ via so-called "kangaroo jumps", a technique utilising the suffix tree to compute the longest common prefix of two suffixes of a string in constant time. This bound was further improved by Amir et al. [3] who showed two solutions, one with running time $\mathcal{O}(n \sqrt{k \log k})$ and one with running time $\tilde{\mathcal{O}}\left(n+k^{3} n / m\right)$-time algorithm ${ }^{1}$. Continuing this line of research, Clifford et al. [14] presented an $\tilde{\mathcal{O}}\left(n+k^{2} n / m\right)$-time algorithm, while Gawrychowski and Uznański [31] demonstrated a smooth trade-off between the latter and the solution of Amir et al. by designing an $\tilde{\mathcal{O}}(n+k n / \sqrt{m})$-time algorithm. Very recently, Chan et al. [11] shaved off most of the polylogarithmic factors and achieved the running time of $\mathcal{O}\left(n+\min \left(k^{2} n / m, k n \sqrt{\log m} / \sqrt{m}\right)\right)$ at the cost of Monte-Carlo randomization.

For the edit distance, a detailed survey of previous solutions can be found in [49], and here we only discuss the landmark results of the theoretical landscape. For the general variant of the problem, the first algorithm was given by Sellers [56]. The algorithm was based on dynamic programming and used $\mathcal{O}(n m)$ time. Masek and Paterson [46] improved the running time of the algorithm to $\mathcal{O}(n m / \log n)$ via the Four Russians technique. On the lower bound side, it is known that there is no solution with strongly subquadratic time complexity unless the Strong Exponential Time hypothesis [36] is false, even for the binary alphabet [5, 8]. Abboud et al. [1] gave a more precise bound under a weaker assumption: Namely, they showed that even shaving an arbitrarily large polylog factor would imply that NEXP does not have non-uniform $\mathrm{NC}^{1}$ circuits. Finally, Clifford et al. [15] showed that, in the cell-probe model with the word size $w=1$, any randomised algorithm that computes the edit distances between the pattern and the text online must spend $\Omega\left(\frac{\sqrt{\log n}}{(\log \log n)^{3 / 2}}\right)$ expected amortised time per character of the text.

Similarly to the Hamming distance, one can define the threshold variant of approximate pattern matching, which we refer to as approximate pattern matching with $k$ edits. The first algorithm for this variant of the problem was developed by Landau and Vishkin [43]; this by-now classical algorithm solves the problem in $\mathcal{O}(n k)$ time. The current best result was achieved by a series of work [55, 17] with the running time (for some range of parameters) $\mathcal{O}\left(n+k^{4} n / m\right)$. Very recently, Charalampopoulos et al. [12] studied the problem for both

[^1]distances in the grammar-compressed setting. Their result, in particular, implies an $\mathcal{O}\left(k^{4}\right)$ space and $\mathcal{O}\left(n k^{4}\right)$-time algorithm for the read-only model, where random access to characters in $P$ and $T$ is allowed and one accounts only for the "extra" space, the space required beyond the space needed to store the pattern and the text.

In this work, we focus on developing algorithms for approximate pattern matching with $k$ edits that use as little space as possible. In particular, we consider the streaming model of computation. In this model, we assume that the input arrives as a stream, one character at a time. We define the space complexity of the algorithm to be all the space used, in other words, we cannot store any information about the input without accounting for it.

The field of streaming algorithms for string processing is relatively recent but, because of its practical interest, it has received a lot of attention in the literature. It started with a seminal paper of Porat and Porat in FOCS 2009 [52], who showed streaming algorithms for exact pattern matching and for the $k$-mismatches problem. The result of Porat and Porat was followed by a series of works on streaming pattern matching [7, 13, 34, 14, 57, 33, 30, 16, 32, 53], search of repetitions in streams [20, 18, 19, 28, 48, 47, 29], and recognising formal languages in streams $[45,24,25,22,23,27,4,21,26]$.

All known streaming algorithms for approximate pattern matching under the Hamming distance $[52,14,16]$ are based on some rolling hash - a hash on strings of fixed length that can be efficiently updated when we delete the first character of a string and append a new one to the end of the string, such as, for example, the famous Karp-Rabin fingerprint, which allows computing the Hamming distance between two strings or certify that it exceeds $k$. The current best algorithm was demonstrated by Clifford et al. in SODA 2019 [16]. The algorithm uses only $\mathcal{O}\left(k \log \frac{m}{k}\right)$ space, which is optimal up to a logarithmic factor, and spends $\mathcal{O}\left(\log \frac{m}{k}(\sqrt{k \log k}+\right.$ $\left.\log ^{3} m\right)$ ) time per character. The algorithm is necessarily randomised and can err with high probability ${ }^{2}$. In other words, approximate pattern matching under the Hamming distance in the streaming model is essentially fully understood.

On the other hand, for the edit distance there are no small-space solutions, in particular, because there are no rolling hashes that allow to compute the edit distance between strings. When $k$ is small, the state-of-the-art solution [57] uses $\mathcal{O}\left(k^{8} \sqrt{m} \log ^{6} m\right)$ space and $\mathcal{O}\left(\left(k^{2} \sqrt{m}+\right.\right.$ $\left.k^{13}\right) \cdot \log ^{4} m$ ) worst-case time per symbol. Again, the algorithm is randomised and outputs all substrings at edit distance at most $k$ from the pattern with high probability. Another interesting result is that of Chakraborty et al. [9], who developed an algorithm for the general version of approximate pattern matching under edit distance in the model where the text is streaming, but the pattern is read-only. They showed a randomised algorithm that, for every position $i$ of the text $T$, computes the smallest edit distance $d_{i}$ between $P$ and a suffix of $T[1 \ldots i]$ with constant multiplicative and $m^{8 / 9}$-additive approximation (in other words, the algorithm returns a number between $d_{i}$ and $c \cdot d_{i}+m^{8 / 9}$, where $c \geq 1$ is a predetermined constant). The algorithm receives the text online and uses $\mathcal{O}\left(m^{1-1 / 54}\right)$ extra space, in addition to the space required to store the pattern.

Naturally, a question arises: is the true complexity of streaming approximate pattern matching under edit distance is on par with that of the Hamming distance? In this work, we answer this question affirmatively.

### 1.1 Our results

The main result of our work is a fully streaming algorithm for approximate pattern matching under the edit distance that uses $\tilde{\mathcal{O}}\left(k^{5}\right)$ space and $\tilde{\mathcal{O}}\left(k^{8}\right)$ amortized time per character of the text (Theorem 7.6). The algorithm is randomised and its answers are correct with high probability.

[^2]As a stepping stone, we also consider a simpler semi-streaming model introduced in [9]. In this model, we assume that the text arrives in a stream, but the pattern is read-only, which means that, at any moment, the algorithm can access any character of the pattern in constant time (but re-writing characters is prohibited). The space complexity of the algorithm is defined as the total space used on top of the read-only memory holding the pattern. In this setting, we show a deterministic algorithm for approximate pattern matching under the edit distance that uses $\tilde{\mathcal{O}}\left(k^{5}\right)$ space and $\tilde{\mathcal{O}}\left(k^{6}\right)$ amortized time per character of the text (Theorem 7.5).

Additionally, we design a new sketch for retrieving the exact edit distance (capped with a threshold $k$ ) between strings of length at most $n$ (Theorem 3.17). The sketch is of size $\tilde{\mathcal{O}}\left(k^{2}\right)$, and it can be built using a streaming algorithm that costs $\tilde{\mathcal{O}}(n k)$ total time and uses $\tilde{\mathcal{O}}\left(k^{2}\right)$ space. Given the sketches of two strings $X, Y$, in $\tilde{\mathcal{O}}\left(k^{3}\right)$ time and $\tilde{\mathcal{O}}\left(k^{2}\right)$ space, we can compute the edit distance between $X, Y$ or certify that it is larger than $k$. The answer is correct with a large constant probability (with standard amplification, we then achieve high probability of success). This improves upon the $\tilde{\mathcal{O}}\left(k^{8}\right)$-size sketches of Belazzougui and Zhang [6] and the $\tilde{O}\left(k^{3}\right)$-size sketches of Jin, Nelson, and Wu [37] developed independently.

The conceptual contribution of our work is described in the technical overview (Section 3).

## 2 Preliminaries

We assume an integer alphabet $\Sigma=\{1,2, \ldots, \sigma\}$ with $\sigma$ characters. A string $Y$ is a sequence of characters numbered from 1 to $n=|Y|$. By $Y[i]$ we denote the $i$-th symbol of $Y$. For a string $Y$ of length $n$, we denote its reverse $Y[n] Y[n-1] \ldots Y[1]$ by $\bar{Y}$. We define $Y[i \ldots j]$ to be equal to $Y[i] \ldots Y[j]$ which we call a fragment of $Y$ if $i \leq j$ and to the empty string $\varepsilon$ otherwise. We also use notations $Y[i . . j)$ and $Y(i . . j]$ which naturally stand for $Y[i] \ldots Y[j-1]$ and $Y[i+1] \ldots Y[j]$, respectively. We call a fragment $Y[1] \ldots Y[j]$ a prefix of $Y$ and use a simplified notation $Y[\ldots i]$, and a fragment $Y[i] \ldots Y[n]$ a suffix of $Y$ denoted by $Y[i .$.$] . We say that X$ is a substring of $Y$ if $X=Y[i \ldots j]$ for some $1 \leq i \leq j \leq|Y|$. The fragment $Y[i \ldots j]$ is called an occurrence of $X$.

For a string $Y$, we define $Y^{m}$ to be the concatenation of $m$ copies of $Y$. We also define $Y^{\infty}$ to be an infinite string obtained by concatenating infinitely many of copies of $Y$. We say that a string $X$ of length $x$ is a period of a string $T$ if $X=T[1 \ldots x]$ and $T[i]=T[i+x]$ for all $i=1, \ldots,|T|-x$. By $\operatorname{per}(T)$ we denote the length of the shortest period of $T$. The string $T$ is called periodic if $2 \operatorname{per}(T) \leq|T|$. For a string $Y \in \Sigma^{n}$, we define a forward rotation $\operatorname{rot}(Y)=Y[2] \cdots Y[n] Y[1]$. In general, a cyclic rotation $\operatorname{rot}^{s}(Y)$ with shift $s \in \mathbb{Z}$ is obtained by iterating rot or the inverse operation rot $^{-1}$. A non-empty string $X \in \Sigma^{n}$ is primitive if it is distinct from its non-trivial rotations, i.e., if $X=\operatorname{rot}^{s}(X)$ holds only when $s$ is a multiple of $n$.

We say that a fragment $X\left[i . . i+\ell\right.$ ) is a previous factor if $X[i \ldots i+\ell)=X\left[i^{\prime} \ldots i^{\prime}+\ell\right)$ holds for some $i^{\prime} \in[1 \ldots i)$. The LZ77 factorization of $X$ is a factorization $X=F_{1} \cdots F_{z}$ into nonempty phrases such tht the $j$ th phrase $F_{j}$ is the longest previous factor starting at position $1+\left|F_{1} \cdots F_{j-1}\right|$; if no previous factor starts there, then $F_{j}$ consists of a single character. In the underlying LZ77 representation, every phrase $F_{j}=T[j \ldots j+\ell)$ that is a previous fragment is encoded as $\left(i^{\prime}, \ell\right)$, where $i^{\prime} \in[1 \ldots i)$ satisfies $X[i \ldots i+\ell)=X\left[i^{\prime} \ldots i^{\prime}+\ell\right)$. The remaining length- 1 phrases are represented by the underlying character. We use $\operatorname{LZ}(X)$ to denote the underlying $L Z 77$ representation and $|\operatorname{LZ}(X)|$ to denote its size (the number of phrases).

### 2.1 Edit Distance Alignments

The edit distance ed $(X, Y)$ between two strings $X$ and $Y$ is defined as the smallest number of character insertions, deletions, and substitutions required to transform $X$ to $Y$. The Hamming distance $\operatorname{hd}(X, Y)$ allows substitutions only (and we assume hd $(X, Y)=\infty$ if $|X| \neq|Y|$ ).

Definition 2.1. A sequence $\left(x_{t}, y_{t}\right)_{t=1}^{m}$ is an alignment of $X, Y \in \Sigma^{*}$ if $\left(x_{1}, y_{1}\right)=(1,1)$, $\left(x_{m}, y_{m}\right)=(|X|+1,|Y|+1)$, and $\left(x_{t+1}, y_{t+1}\right) \in\left\{\left(x_{t}+1, y_{t}+1\right),\left(x_{t}+1, y_{t}\right),\left(x_{t}, y_{t}+1\right)\right\}$ for $t \in[1 \ldots m) .^{3}$

Given an alignment $\mathcal{A}=\left(x_{t}, y_{t}\right)_{t=1}^{m}$ of strings $X, Y \in \Sigma^{*}$, for every $t \in[1 \ldots m)$ :

- If $\left(x_{t+1}, y_{t+1}\right)=\left(x_{t}+1, y_{t}\right)$, we say that $\mathcal{A}$ deletes $X\left[x_{t}\right]$,
- If $\left(x_{t+1}, y_{t+1}\right)=\left(x_{t}, y_{t}+1\right)$, we say that $\mathcal{A}$ deletes $Y\left[y_{t}\right]$,
- If $\left(x_{t+1}, y_{t+1}\right)=\left(x_{t}+1, y_{t}+1\right)$, we say that $\mathcal{A}$ aligns $X\left[x_{t}\right]$ and $Y\left[y_{t}\right]$, denoted $X\left[x_{t}\right] \sim_{\mathcal{A}}$ $Y\left[y_{t}\right]$. If additionally $X\left[x_{t}\right]=Y\left[y_{t}\right]$, we say that $\mathcal{A}$ matches $X\left[x_{t}\right]$ and $Y\left[y_{t}\right]$, denoted $X\left[x_{t}\right] \simeq_{\mathcal{A}} Y\left[y_{t}\right]$. Otherwise, we say that $\mathcal{A}$ substitutes $X\left[x_{t}\right]$ for $Y\left[y_{t}\right]$.
The cost of an edit distance alignment $\mathcal{A}$ is the total number characters that $\mathcal{A}$ deletes or substitutes. We denote the cost by $\operatorname{cost}_{X, Y}(\mathcal{A})$, omitting the subscript if $X, Y$ are clear from context. The cost of an alignment $\mathcal{A}=\left(x_{t}, y_{t}\right)_{t=1}^{m}$ is at least its width width $(\mathcal{A})=$ $\max _{t=1}^{m}\left|x_{t}-y_{t}\right|$. Observe that ed $(X, Y)$ can be defined as the minimum cost of an alignment of $X$ and $Y$. An alignment of $X$ and $Y$ is optimal if its cost is equal to ed $(X, Y)$.

Given an alignment $\mathcal{A}=\left(x_{t}, y_{t}\right)_{t=1}^{m}$ of $X, Y \in \Sigma^{+}$, we partition the elements $\left(x_{t}, y_{t}\right)$ of $\mathcal{A}$ into matches (for which $X\left[x_{t}\right] \simeq_{\mathcal{A}} Y\left[y_{t}\right]$ ) and breakpoints (the remaining elements). We denote the set of matches and breakpoints by $\mathcal{M}_{X, Y}(\mathcal{A})$ and $\mathcal{B}_{X, Y}(\mathcal{A})$, respectively, omitting the subscripts if the strings $X, Y$ are clear from context. Observe that $\left|\mathcal{B}_{X, Y}(\mathcal{A})\right|=1+\operatorname{cost}(\mathcal{A})$.

We call $M \subseteq[1 \ldots|X|] \times[1 \ldots|Y|]$ a non-crossing matching of $X, Y \in \Sigma^{*}$ if $X[x]=Y[y]$ holds for all $(x, y) \in M$ and there are no distinct pairs $(x, y),\left(x^{\prime}, y^{\prime}\right) \in M$ with $x \leq x^{\prime}$ and $y \geq y^{\prime}$. Note that, for every alignment $\mathcal{A}$ of $X, Y$, the set $\mathcal{M}(\mathcal{A})$ is a non-crossing matching of $X, Y$.

Example 2.2. Consider strings $X=$ abbaabcb and $Y=$ acabaabab and a cost- 4 alignment

$$
\mathcal{A}:(1,1),(2,2),(3,3),(3,4),(4,5),(5,6),(6,7),(7,8),(8,8),(8,9),(9,10)
$$

The breakpoints are $\mathcal{B}(\mathcal{A})=\{(2,2),(3,3),(7,8),(8,8),(9,10)\}$; the first 4 breakpoints correspond to a substitution of $X[2]$ for $Y[2]$, a deletion of $Y[3]$, a deletion of $X[7]$, and a deletion of $Y[8]$, respectively. Graphically, the alignment is depicted on
 the right; the aligned pairs of characters are connected with an edge, and the substituted pair is highlighted.

Given an alignment $\mathcal{A}=\left(x_{t}, y_{t}\right)_{t=1}^{m}$ of $X$ and $Y$, for every $\ell, r \in[1 \ldots m]$ with $\ell \leq r$, we say that $\mathcal{A}$ aligns $X\left[x_{\ell} \ldots x_{r}\right)$ and $Y\left[y_{\ell} \ldots y_{r}\right)$, denoted $X\left[x_{\ell} \ldots x_{r}\right) \sim_{\mathcal{A}} Y\left[y_{\ell} \ldots y_{r}\right)$. If there is no breakpoint $\left(x_{t}, y_{t}\right)$ with $t \in[\ell \ldots r)$, we further say that $\mathcal{A}$ matches $X\left[x_{\ell} \ldots x_{r}\right)$ and $Y\left[y_{\ell} \ldots y_{r}\right)$, denoted $X\left[x_{\ell} \ldots x_{r}\right) \simeq_{\mathcal{A}} Y\left[y_{\ell} \ldots y_{r}\right)$.

An alignment $\mathcal{A}=\left(x_{t}, y_{t}\right)_{t=1}^{m}$ of $X, Y \in \Sigma^{*}$ naturally induces a unique alignment of any two fragments $X\left[x \ldots x^{\prime}\right)$ and $Y\left[y \ldots y^{\prime}\right)$. Formally, the induced alignment $\mathcal{A}_{\left[x \ldots x^{\prime}\right),\left[y . . y^{\prime}\right)}$ is obtained by removing repeated entries from $\left(\max \left(x, \min \left(x^{\prime}, x_{t}\right)\right)-x+1, \max \left(y, \min \left(y^{\prime}, y_{t}\right)\right)-y+1\right)_{t=1}^{m}$.

Fact 2.3. If an alignment $\mathcal{A}$ satisfies $X\left[x \ldots x^{\prime}\right) \sim_{\mathcal{A}} Y\left[y \ldots y^{\prime}\right)$, then $\left|x-x^{\prime}\right|,\left|y-y^{\prime}\right| \leq \operatorname{width}(\mathcal{A})$ and $\left|\left(x^{\prime}-x\right)-\left(y^{\prime}-y\right)\right| \leq \operatorname{ed}\left(X\left[x \ldots x^{\prime}\right), Y\left[y \ldots y^{\prime}\right)\right) \leq \operatorname{cost}\left(\mathcal{A}_{\left[x \ldots x^{\prime}\right),\left[y . y^{\prime}\right)}\right) \leq \operatorname{cost}(\mathcal{A})$.

## 3 Technical Overview

In this section, we provide an overview of our conceptual and technical contribution. Let us start with a formal statement of the pattern matching with $k$ edits problem. We say that

[^3]$T[\ell \ldots r]$ is a $k$-edit occurrence of $P$ if $\operatorname{ed}(P, T[\ell \ldots r]) \leq k$, and we denote the set of the right endpoints of the $k$-edit occurrences of $P$ in $T$ by $\mathrm{OCC}_{k}^{E}(P, T)$.
Problem 3.1 (Pattern matching with $k$ edits). Given a pattern $P$ of length $m$ over an alphabet $\Sigma$, a text $T$ of length $n$ over $\Sigma$, and an integer $k$, compute $\operatorname{OCC}_{k}^{E}(P, T)$.

We solve an online version of the problem, where the text arrives in a stream (character by character) and the algorithm must decide whether $r \in \operatorname{OCC}_{k}^{E}(P, T)$ while processing $T[r]$. The pattern is preprocessed in advance (consistently with [52], in the current version of this paper we do not account for this preprocessing in the complexity analysis). We consider two settings:

1. In the streaming setting, the algorithm can no longer access $P$ or $T[1 \ldots r$ ) while processing $T[r]$. In other words, all the information regarding these strings needs to be stored explicitly and accounted for in the space complexity of the algorithm.
2. In the semi-streaming setting, the algorithm can no longer access $T[1 \ldots r$ ) while processing $T[r]$, but it is given an oracle providing read-only constant-time access to individual characters of $P$. This oracle is not counted towards the space complexity of the algorithm.
For the semi-streaming setting, we provide a deterministic solution, whereas our solution for the streaming setting is Monte-Carlo randomized. Both algorithms are designed for the $w$-bit word RAM model, where $w=\Omega(\log n)$, and integer alphabets $\Sigma=\left[0 \ldots n^{\mathcal{O}(1)}\right)$.

## 3.1 (Semi-)Streaming Algorithm for Pattern Matching with $k$ Edits

Our algorithms solve a slightly stronger problem: every element $r \in \operatorname{OCC}_{k}^{E}(P, T)$ is augmented with the smallest integer $k^{\prime} \in[0 \ldots k]$ such that $r \in \mathrm{OCC}_{k^{\prime}}^{E}(P, T)$. At a very high-level, we reuse the structure of existing streaming algorithms for exact pattern matching and the $k$-mismatch problem [52, 7, 14, 16]. Namely, we consider $\mathcal{O}(\log m)$ prefixes $P_{i}=P\left[1 \ldots \ell_{i}\right]$ of exponentially increasing lengths $\ell_{i}$. The algorithms are logically decomposed into $\mathcal{O}(\log m)$ levels, with the $i$ th level receiving $\mathrm{OCC}_{k}^{E}\left(P_{i-1}, T\right)$ and producing $\mathrm{OCC}_{k}^{E}\left(P_{i}, T\right)$. In other words, the task of the $i$ th level is determine which $k$-edit occurrences of $P_{i-1}$ can be extended to $k$-edit occurrences of $P_{i}$. When the algorithm processes $T[r]$, the relevant positions $p \in \mathrm{OCC}_{k}^{E}\left(P_{i-1}, T\right)$ are those satisfying $\left|r-p-\left(\ell_{i}-\ell_{i-1}\right)\right| \leq k$. Since each $p \in \operatorname{OCC}_{k}^{E}\left(P_{i-1}, T\right)$ is reported when the algorithm processes $T[p]$, we need a buffer storing the active $k$-edit occurrences of $P_{i-1}$. We implement it using a recent combinatorial characterization of $k$-edit occurrences [12], which classifies strings based on the following notion of approximate periodicity:

Definition 3.2 ( $k$-periodic string). A string $X$ is $k$-periodic if there exists a primitive string $Q$ with $|Q| \leq|X| / 128 k$ such that the edit distance between $X$ and a prefix of $Q^{\infty}$ is at most $2 k$. We call $Q$ a $k$-period of $X$.

The main message of [12] is that only $k$-periodic strings may have many $k$-edit occurrences.
Corollary 3.3 (of [12, Theorem 5.1]). Let $X \in \Sigma^{m}, k \in[1 \ldots m]$, and $Y \in \Sigma^{n}$ with $n \leq 2 m$. If $X$ is not $k$-periodic, then $\left|\mathrm{OCC}_{k}^{E}(X, Y)\right|=\mathcal{O}\left(k^{2}\right)$.

In particular, if $P_{i-1}$ is not $k$-periodic, then it has $\mathcal{O}\left(k^{2}\right)$ active $k$-edit occurrences. For each active occurrence $p \in \mathrm{OCC}_{k}^{E}\left(P_{i-1}, T\right)$, we maintain an edit-distance sketch $\mathrm{sk}_{k}^{E}(T(p \ldots r])$ and combine it with a sketch $\mathrm{sk}_{k}^{E}\left(P\left(\ell_{i-1} \ldots \ell_{i}\right]\right)$ (constructed at preprocessing) in order to derive ed $\left(T(p \ldots r], P\left(\ell_{i-1} \ldots \ell_{i}\right]\right)$ or certify that this distance exceeds $k$. Since we have stored the smallest $k^{\prime} \in[0 \ldots k]$ such that $p \in \mathrm{OCC}_{k^{\prime}}^{E}(P, T)$, this lets us check whether any $k$-edit occurrence of $P_{i-1}$ ending at position $p$ extends to a $k$-edit occurrence of $P_{i}$ ending at position $r$. With existing $k$-edit sketches [6, 37], this already yields an $\operatorname{poly}(k \log n)$-space implementation in this case.

The difficulty lies in $k$-periodic strings whose occurrences form chains.

Definition 3.4 (Chain of occurrences). Consider strings $X, Y \in \Sigma^{*}$ and an integer $k \in \mathbb{Z} \geq 0$. An increasing sequence of positions $p_{1}, \ldots, p_{c}$ forms a chain of $k$-edit occurrences of $X$ in $Y$ if:

1. there is a difference string $D \in \Sigma^{*}$ such that $D=Y\left(p_{j} \ldots p_{j+1}\right]$ for $j \in[1 \ldots c)$, and
2. there is an integer $k^{\prime} \in[0 \ldots k]$ such that $p_{j} \in \operatorname{OCC}_{k^{\prime}}^{E}(X, Y) \backslash \operatorname{OCC}_{k^{\prime}-1}^{E}(X, Y)$ for $j \in[1 . . c]$.

Corollary 3.5 (of [12, Theorem 5.2, Claim 5.16, Claim 5.17]). Let $X \in \Sigma^{m}, k \in[1 . . m]$, and $Y \in \Sigma^{n}$ with $n \leq 2 m$. If $X$ is $k$-periodic with period $Q$, then $\operatorname{OCC}_{k}^{E}(X, Y)$ can be decomposed into $\mathcal{O}\left(k^{3}\right)$ chains whose differences are of the form $\operatorname{rot}^{s}(Q)$ with $|m-s| \leq 10 k$.

In the following discussion, assume that $P_{i-1}$ is $k$-periodic with period $Q_{i-1}$. Compared to the previous algorithm, we cannot afford maintaining a $\operatorname{sketch}^{\operatorname{sk}}{ }_{k}^{E}(T(p \ldots r])$ for all active $p \in \mathrm{OCC}_{k}^{E}(P, T)$. If $\mathrm{sk}_{k}^{E}$ were a rolling sketch (like the $k$-mismatch sketches of [16]), we would compute $\mathrm{sk}_{k}^{E}(D)$ at preprocessing time for all $\mathcal{O}(k)$ feasible chain differences $D$ and then, for any two subsequent positions $p_{j}, p_{j+1}$ in a chain with difference $D$, we could use $\mathbf{~ s k}_{k}^{E}(D)=$ $\mathbf{s k}_{k}^{E}\left(T\left(p_{j} . . p_{j+1}\right]\right)$ to transform $\operatorname{sk}_{k}^{E}\left(T\left(p_{j} \ldots r\right]\right)$ into $\mathbf{s k}_{k}^{E}\left(T\left(p_{j+1} . . r\right]\right)$. However, despite extensive research, no rolling edit distance sketch is known, which remains the main obstacle in designing streaming algorithms for approximate pattern matching with $k$ edits.

Our workaround relies on a novel encoding $\mathrm{q} G \mathrm{R}(X, Y)$ that, for a pair of strings $X, Y \in \Sigma^{*}$, represents a large class of low-distance edit distance alignments between $X, Y$. In the preprocessing phase of our algorithm, we build $\operatorname{qGR}\left(P\left(\ell_{i-1} \ldots \ell_{i}\right], D^{\infty}\left[1 \ldots \ell_{i}-\ell_{i-1}\right]\right)$ for every feasible chain difference $D$. In the main phase, for subsequent positions $p_{j} \in \operatorname{OCC}_{k}^{E}\left(P_{i-1}, T\right)$ in a chain with difference $D$, we aim to build $\mathrm{qGR}\left(T\left(p_{j} \ldots r\right], D^{\infty}\left[1 \ldots \ell_{i}-\ell_{i-1}\right]\right)$ when necessary, i.e., $\mid r-p_{j}-$ $\left(\ell_{i}-\ell_{i-1}\right) \mid \leq k$. We then combine the two encodings to derive $\mathrm{qGR}\left(P\left(\ell_{i-1} \ldots \ell_{i}\right], T\left(p_{j} \ldots r\right]\right)$ and $\operatorname{ed}\left(P\left(\ell_{i-1} \ldots \ell_{i}\right], T\left(p_{j} \ldots r\right]\right)$. Except for such products (transitive compositions), our encoding supports concatenations, i.e., $\mathrm{q} \mathrm{GR}\left(X_{1}, Y_{1}\right)$ and $\mathrm{q} \operatorname{GR}\left(X_{2}, Y_{2}\right)$ can be combined into $\mathrm{q} \operatorname{GR}\left(X_{1} X_{2}, Y_{1} Y_{2}\right)$. Consequently, it suffices to maintain $\mathrm{qGR}\left(T\left(p_{c} \ldots r\right], D^{\infty}\left[1 \ldots r-p_{c}-k\right]\right)$ (where $p_{c}$ is the rightmost element of the chain). When necessary, we prepend $(j-c)$ copies of $\mathrm{qGR}(D, D)$ (merged by doubling) and append $\mathrm{q} \operatorname{GR}\left(\varepsilon, D^{\infty}\left(r-p_{j}-k . . \ell_{i}-\ell_{i-1}\right]\right)$ to derive $\mathrm{q} \operatorname{GR}\left(P\left(\ell_{i-1} \ldots \ell_{i}\right], D^{\infty}\left[1 \ldots \ell_{i}-\ell_{i-1}\right]\right)$.

In the semi-streaming setting, we extend $\operatorname{qGR}\left(T\left(p_{c} . . r\right], D^{\infty}\left[1 . . r-p_{c}-k\right]\right)$ one character at a time using read-only random access to $D^{\infty}$. In the streaming setting, we cannot afford storing $D$, so we append the entire difference $D$ in a single step and utilize a new edit-distance sketch $\mathrm{sk}^{\mathrm{q}}$ that allows retrieving $\mathrm{qGR}(T(r-|D| \ldots r], D)$ from $\mathrm{sk}^{\mathrm{q}}(T(r-|D| \ldots r])$ and $\mathrm{sk}^{\mathrm{q}}(D)$. The sketch $\mathrm{sk}^{\mathrm{q}}(D)$ is constructed in the preprocessing phase, whereas $\mathrm{sk}^{\mathrm{q}}(T(r-|D| \ldots r])$ is built as the algorithm scans $T$. Similarly, we can (temporarily) append any of the $\mathcal{O}(k)$ prefixes that of $D$ may arise when $\mathrm{qGR}\left(T\left(p_{c} . . r\right], D^{\infty}\left[1 \ldots r-p_{c}-k\right]\right)$ is necessary. A complete presentation of our algorithms is provided in Section 7. Below, we outline the ideas behind our two main conceptual and technical contributions: the encoding $\mathrm{qGR}(\cdot, \cdot)$ and the sketch $\mathrm{sk}^{\mathrm{q}}(\cdot)$.

### 3.2 Greedy Alignments and Encodings

Recall that the encoding $\mathrm{qGR}(\cdot, \cdot)$ needs to support the following three operations:
Capped edit distance: given $\mathrm{q} G \mathrm{R}(X, Y)$, compute ed $(X, Y)$ or certify that ed $(X, Y)$ is large; Product: given $\mathrm{q} G \mathrm{R}(X, Y)$ and $\mathrm{q} \mathrm{GR}(Y, Z)$, retrieve $\mathrm{q} \mathrm{GR}(X, Z)$;
Concatenation: given $\mathrm{qGR}\left(X_{1}, Y_{1}\right)$ and $\mathrm{qGR}\left(X_{2}, Y_{2}\right)$, retrieve $\mathrm{qGR}\left(X_{1} X_{2}, Y_{1} Y_{2}\right)$.
Our encoding is parameterized with a threshold $k \in \mathbb{Z}_{+}$such that ed $(\cdot, \cdot)>k$ is considered large, and the goal is to achieve $\tilde{\mathcal{O}}\left(k^{\mathcal{O}(1)}\right)$ encoding size. In fact, whenever ed $(X, Y)>k$, we shall simply assume that $\mathrm{qGR}_{k}(X, Y)$ is undefined (formally, $\mathrm{qGR}_{k}(X, Y)=\perp$ ). Consequently, products and concatenations will require sufficiently large thresholds in the input encodings so that if either of them is undefined, the output encoding is also undefined.

In order to support concatenations alone, we could use so-called semi-local edit distances. For now, suppose that we only need to encode pairs of equal-length strings. ${ }^{4}$ Through a sequence of concatenations, we may only extend $\mathrm{qGR}_{k}(X, Y)$ to $\mathrm{qGR}_{k}\left(X^{\prime}, Y^{\prime}\right)$ so that $X=X^{\prime}[\ell \ldots r)$ and $Y=Y^{\prime}[\ell \ldots r)$. For any alignment $\mathcal{A}^{\prime}$ of $X^{\prime}, Y^{\prime}$ with $\operatorname{cost}\left(\mathcal{A}^{\prime}\right) \leq k$, consider the induced alignment $\mathcal{A}:=\mathcal{A}_{[\ell . . r),[\ell . r)}^{\prime}$. Note that $\mathcal{A}$ mimics the behavior of $\mathcal{A}^{\prime}$ except that it deletes some characters at the extremes of $X$ and $Y$ (which $\mathcal{A}^{\prime}$ aligns outside $Y$ and $X$, respectively). By Fact 2.3, we have $\operatorname{cost}(\mathcal{A}) \leq k$ and, in particular, $\mathcal{A}$ deletes at most $k$ characters at the extremes of $X$ and $Y$. If, after performing these deletions, we replace $\mathcal{A}$ with an optimal alignment between the remaining fragments of $X$ and $Y$, this modification may only decrease $\operatorname{cost}(\mathcal{A})$ and $\operatorname{cost}\left(\mathcal{A}^{\prime}\right)$. Consequently, it suffices to store the $\mathcal{O}(k)$ characters at the extremes of $X, Y$ and the $\mathcal{O}\left(k^{4}\right)$ edit distances ${ }^{5}$ between long fragments of $X$ and $Y$. In a sense, this encoding represents $\mathcal{O}\left(k^{4}\right)$ alignments between $X, Y$ that are sufficient to derive an optimal alignment of any extension.

The main challenge is to handle products, for which we develop a greedy encoding $\mathrm{GR}_{k}(X, Y)$ that compactly represents the following family $\mathrm{GA}_{k}(X, Y)$ of greedy alignments of $X, Y$.

Definition 3.6 (Greedy alignment). We say that an alignment $\mathcal{A}$ of two strings $X, Y \in \Sigma^{*}$ is greedy if $X[x] \neq Y[y]$ holds for every $(x, y) \in \mathcal{B}(\mathcal{A}) \cap([1 \ldots|X|] \times[1 \ldots|Y|])$. Given $k \geq 0$, we denote by $\mathrm{GA}_{k}(X, Y)$ the set of all greedy alignments $\mathcal{A}$ of $X, Y$ satisfying $\operatorname{cost}(\mathcal{A}) \leq k$.

Intuitively, whenever a greedy alignment encounters a pair of matching characters $X[x]$ and $Y[y]$, it must (greedily) match these characters (it cannot delete $X[x]$ or $Y[y]$ ). As stated below, this restriction does not affect the optimal cost. (All claims are proved in Section 5.)

Fact 3.7. For any two strings $X, Y \in \Sigma^{*}$, there is an optimal greedy alignment of $X, Y$.
For strings $X, Y \in \Sigma^{*}$ and an integer $k \geq e d(X, Y)$, we define a set $\mathcal{M}_{k}(X, Y)$ of common matches of all alignments $\mathcal{A} \in \mathrm{GA}_{k}(X, Y)$; formally $\mathcal{M}_{k}(X, Y)=\bigcap_{\mathcal{A} \in \mathrm{GA}_{k}(X, Y)} \mathcal{M}_{X, Y}(\mathcal{A})$. In our greedy encoding, we shall mask out all the characters involved in the common matches. Below, this transformation is defined for an arbitrary non-crossing matching of $X, Y$.

Definition 3.8. Let $M$ be a non-crossing matching of strings $X, Y \in \Sigma^{*}$. We define $X^{M}, Y^{M}$ to be the strings obtained from $X, Y$ by replacing $X[x]$ and $Y[y]$ with $\# \notin \Sigma$ for every $(x, y) \in M$. We refer to \# as a dummy symbol and to maximal blocks of \#'s as dummy segments.

The following lemma proves that masking out common matches does not affect ed $(X, Y)$ or $\mathcal{M}_{k}(X, Y)$ provided that we enumerate the dummy symbols, that is, any string $Z$ is transformed to num $(Z)$ by replacing the $i$ th leftmost occurrence of $\#$ with a unique symbol $\#_{i} \notin \Sigma$.

Lemma 3.9. Consider strings $X, Y \in \Sigma^{*}$, an integer $k \geq \operatorname{ed}(X, Y)$, and a set $M \subseteq \mathcal{M}_{k}(X, Y)$. Then, ed $(X, Y)=\operatorname{ed}\left(\operatorname{num}\left(X^{M}\right), \operatorname{num}\left(Y^{M}\right)\right)$ and $\mathcal{M}_{k}(X, Y)=\mathcal{M}_{k}\left(\operatorname{num}\left(X^{M}\right), \operatorname{num}\left(Y^{M}\right)\right)$.

At the same time, after masking out the common matches, the strings become compressible. Intuitively, this is because once two greedy alignments converge, they stay together until they encounter a mismatch. Moreover, when two alignments proceed in parallel without any mismatch, this incurs a small period (at most $2 k$ ) that is captured by the LZ factorization.

Lemma 3.10. Let $M=\mathcal{M}_{k}(X, Y)$ for strings $X, Y \in \Sigma^{*}$ and a positive integer $k \geq \operatorname{ed}(X, Y)$. Then, $\left|\operatorname{LZ}\left(X^{M}\right)\right|,\left|\operatorname{LZ}\left(Y^{M}\right)\right|=\mathcal{O}\left(k^{2}\right)$, and $X^{M}, Y^{M}$ contain $\mathcal{O}(k)$ dummy segments.

[^4]Consequently, for $k \geq \operatorname{ed}(X, Y)$, we could define the greedy encoding $\mathrm{GR}_{k}(X, Y)$ so that it consists of $\mathrm{LZ}\left(X^{M}\right)$ and $\mathrm{LZ}\left(Y^{M}\right)$. Instead, we use a more powerful compressed representation (developed in Section 4) that supports more efficient queries concerning $X^{M}$ and $Y^{M}$.

Even though $\mathrm{GR}_{k}(X, Y)$ is small, $\mathrm{GA}_{k}(X, Y)$ may consist of $2^{\Theta(k)}$ alignments, which is why constructing $\mathrm{GR}_{k}(X, Y)$ in poly $(k)$ time is far from trivial. The following combinatorial lemma lets us obtain an $\mathcal{O}\left(k^{5}\right)$-time algorithm in Section 5.2. Intuitively, the alignments in $\mathrm{GA}_{k}(X, Y)$ can be interpreted as paths in a directed acyclic graph with $\mathcal{O}\left(k^{5}\right)$ branching vertices.

Lemma 3.11. For all $X, Y \in \Sigma^{*}$ and $k \in \mathbb{Z}_{+}$, the set $\mathcal{B}_{k}(X, Y)=\underset{\mathcal{A} \in \operatorname{GA}_{k}(X, Y)}{ } \mathcal{B}(\mathcal{A})$ is of size $\mathcal{O}\left(k^{5}\right)$.
The reason why $\mathrm{GR}_{k}(X, Y)$ supports products is that every greedy alignment of $X, Z$ can be interpreted as a product of a greedy alignment of $X, Y$ and a greedy alignment of $Y, Z$.

Definition 3.12. Consider strings $X, Y, Z \in \Sigma^{*}$, an alignment $\mathcal{A}^{X, Y}$ of $X, Y$, an alignment $\mathcal{A}^{Y, Z}$ of $Y, Z$, and an alignment $\mathcal{A}^{X, Z}$ of $X, Z$. We say that $\mathcal{A}^{X, Z}$ is a product of $\mathcal{A}^{X, Y}$ and $\mathcal{A}^{Y, Z}$ if, for every $(x, z) \in \mathcal{A}^{X, Z}$, there is $y \in[1 \ldots|Y|+1]$ such that $(x, y) \in \mathcal{A}^{X, Y}$ and $(y, z) \in \mathcal{A}^{Y, Z}$.

Lemma 3.13. Consider strings $X, Y, Z \in \Sigma^{*}$ and $k \in \mathbb{Z}_{\geq 0}$. Every alignment $\mathcal{A}^{X, Z} \in \mathrm{GA}_{k}(X, Z)$ is a product of alignments $\mathcal{A}^{X, Y} \in \mathrm{GA}_{d}(X, Y)$ and $\mathcal{A}^{Y, Z} \in \mathrm{GA}_{d}(Y, Z)$, where $d=2 k+\operatorname{ed}(X, Y)$.

As a result, $\operatorname{GR}_{d}(X, Y)$ and $\mathrm{GR}_{d}(Y, Z)$ contain enough information to derive $\mathrm{GR}_{k}(X, Z)$. The underlying algorithm propagates the characters of $Y$ stored in $\mathrm{GR}_{d}(X, Y)$ and $\mathrm{GR}_{d}(Y, Z)$ along the matchings $\mathcal{M}_{d}(Y, Z)$ and $\mathcal{M}_{d}(X, Y)$, respectively. Then, $\mathrm{GR}_{k}(X, Z)$ is obtained by masking out all the characters corresponding to $\mathcal{M}_{k}(X, Y)$ (see Section 5.4 for details).

To support concatenations, we extend the family $\mathrm{GA}_{k}(X, Y)$ of greedy alignments to a family $\mathrm{qGA}_{k}(X, Y)$ of quasi-greedy alignments, which are allowed to delete a prefix of $X$ or $Y$ in violation of Definition 3.6. The quasi-greedy encoding $\mathrm{qGR}_{k}(X, Y)$ is defined analogously to $\mathrm{GR}_{k}(X, Y)$. Equivalently, $\mathrm{qGA}_{k}(X, Y)$ can be derived from $\mathrm{GA}_{k+1}\left(\$_{1} X, \$_{2} Y\right)$, where $\$_{1} \neq \$_{2}$ are sentinel symbols outside $\Sigma$, by taking the alignments induced by $X, Y$. The latter characterization makes all our claims regarding GA and GR easily portable to $q G A$ and $q G R$ (see Section 5.3). In particular, this is true for the sketches, described in the following subsection for GR only.

### 3.3 Edit Distance Sketches

Recall that we need an edit-distance sketch $\mathrm{sk}^{E}$ allowing to retrieve $\mathrm{GR}_{k}(X, Y)$ from $\mathrm{sk}_{k}^{E}(X)$ and $\mathbf{s k}_{k}^{E}(Y)$ for any strings $X, Y \in \Sigma^{\leq n}$ and any threshold $k \in[0 \ldots n]$. Furthermore, we need to make sure that $\mathbf{s k}_{k}^{E}(S)$ can be computed given streaming access to $S \in \Sigma^{\leq n}$, and that the encoding and decoding procedures use poly $(k, \log n)$ space. While the existing sketches $[6,37]$ are designed to compute the exact edit distance ed $(X, Y)$ capped with $k$, we believe that they could be adapted to output $\mathrm{qGR}(X, Y)$. Nevertheless, these sketches are relatively large (taking $\tilde{\mathcal{O}}\left(k^{8}\right)$ and $\tilde{\mathcal{O}}\left(k^{3}\right)$ bits, respectively), and we would need to strengthen the bulk of their analyses to prove that, in principle, they provide enough information to retrieve $\mathrm{qGR}_{k}(X, Y)$. Hence, to further demonstrate the power of our techniques, we devise a novel $\tilde{\mathcal{O}}\left(k^{2}\right)$-size sketch specifically designed to output $\mathrm{q} \mathrm{GR}_{k}(X, Y)$. We note that the $\tilde{\mathcal{O}}\left(k^{2}\right)$ size is optimal for $\mathrm{q} \mathrm{GR}_{k}(X, Y)$, but we are not aware of a matching lower bound for retrieving ed $(X, Y)$ capped with $k$.

Just like the sketches of $[6,37]$, ours relies on the embedding of Chakraborty, Goldenberg, and Koucký [10]. The CGK algorithm performs a random walk over the input string (with forward and stationary steps only). In abstract terms, such walk can be specified as follows:
Definition 3.14 (Complete walk). For a string $S \in \Sigma^{*}$, we say that $\left(s_{t}\right)_{t=1}^{m+1}$ is an $m$-step complete walk over $S$ if $s_{1}=1, s_{m+1}=|S|+1$, and $s_{t+1} \in\left\{s_{t}, s_{t}+1\right\}$ for $t \in[1 \ldots m]$.

For any two strings $X, Y \in \Sigma^{*}$, the two walks underlying $\operatorname{CGK}(X)$ and $\operatorname{CGK}(Y)$ can be interpreted as an edit distance alignment using the following abstract definition:

Definition 3.15 (Zip alignment). The zip alignment of $m$-step complete walks $\left(x_{t}\right)_{t=1}^{m+1}$ and $\left(y_{t}\right)_{t=1}^{m+1}$ over $X, Y \in \Sigma^{*}$ is obtained by removing repeated entries in $\left(x_{t}, y_{t}\right)_{t=1}^{m+1}$.

The key result of [10] is that the cost of the zip alignment of CGK walks over $X, Y \in \Sigma^{*}$ is $\mathcal{O}\left(\operatorname{ed}(X, Y)^{2}\right)$ with good probability, which is then exploited to derive a metric embedding (mapping edit distance to Hamming distance) with quadratic distortion. In our sketch, we also need to observe that the CGK alignment is greedy and that its width is $\mathcal{O}(\operatorname{ed}(X, Y))$ with good probability. The following proposition, proved in Section 6.1, provides a complete black-box interface of the properties of the CGK algorithm utilized in our sketches. It also encapsulates Nisan's pseudorandom generator [50] that reduces the number of (shared) random bits.

Proposition 3.16. For every constant $\delta \in(0,1)$, there exists a constant $c$ and an algorithm W that, given an integer $n$, a seed $r$ of $\mathcal{O}\left(\log ^{2} n\right)$ random bits, and a string $S \in \Sigma^{\leq n}$, outputs a $3 n$-step complete walk $\mathrm{W}(n, r, S)$ over $S$ satisfying the following property for all $X, Y \in \Sigma^{\leq n}$ and the zip alignment $\mathcal{A}_{\mathrm{W}}$ of $\mathrm{W}(n, r, X)$ and $\mathrm{W}(n, r, Y)$ :

$$
\operatorname{Pr}_{r}\left[\mathcal{A}_{\mathrm{W}} \in \mathrm{GA}_{c \cdot \operatorname{ed}(X, Y)^{2}}(X, Y) \text { and } \text { width }\left(\mathcal{A}_{\mathrm{W}}\right) \leq c \cdot \operatorname{ed}(X, Y)\right] \geq 1-\delta .
$$

Moreover, W is an $\mathcal{O}\left(\log ^{2} n\right)$-bit streaming algorithm that $\operatorname{costs} \mathcal{O}(n \log n)$ time and reports any element $s_{t} \in[1 . .|S|]$ of $\mathrm{W}(n, r, S)$ while processing the corresponding character $S\left[s_{t}\right]$.

Next, we analyze the structural similarity between $\mathcal{A}_{\mathrm{w}}$ and any alignment $\mathcal{A} \in \mathrm{GA}_{k}(X, Y)$. Based on Proposition 3.16, we may assume that $\mathcal{A}_{\mathrm{w}} \in \mathrm{GA}_{\mathcal{O}\left(k^{2}\right)}(X, Y)$ and width $\left(\mathcal{A}_{\mathrm{W}}\right)=\mathcal{O}(k)$. Consider the set $M=\mathcal{M}(\mathcal{A}) \cap \mathcal{M}\left(\mathcal{A}_{\mathrm{W}}\right)$ of the common matches of $\mathcal{A}$ and $\mathcal{A}_{\mathrm{W}}$ and the string $X^{M}$ obtained by masking out the underlying characters of $X$. Whereas Lemma 3.10 immediately implies that the LZ factorization of $X^{M}$ consists of $\mathcal{O}\left(k^{4}\right)$ phrases, a more careful application of the same technique provides a refined bound of $\mathcal{O}\left(k^{2}\right)$ phrases. Furthermore, there are $\mathcal{O}(k)$ dummy segments in $X^{M}$ and, if $X[x]$ is not masked out in $X^{M}$ (for some $x \in[1 . .|X|]$ ), then there is a breakpoint $\left(x^{\prime}, y^{\prime}\right) \in \mathcal{B}_{X, Y}\left(\mathcal{A}_{\mathrm{W}}\right)$ with $x^{\prime} \in[1 \ldots x]$ and $\left|\operatorname{LZ}\left(X\left[x^{\prime} \ldots x\right]\right)\right|=\mathcal{O}(k)$. Intuitively, this means that $\mathcal{A}$ and $\mathcal{A}_{\mathrm{W}}$ diverge only within highly compressible regions following the breakpoints $\mathcal{B}_{X, Y}\left(\mathcal{A}_{\mathrm{W}}\right)$. We call these regions forward contexts (formally, a forward context is the longest fragment starting at a given position and satisfying certain compressiblity condition). Since our choice of $\mathcal{A} \in \mathrm{GA}_{k}(X, Y)$ was arbitrary, any two alignments $\mathcal{A}, \mathcal{A}^{\prime} \in \mathrm{GA}_{k}(X, Y)$ diverge only within these forward contexts. Hence, in order to reconstruct $\mathrm{GR}_{k}(X, Y)$ and, in particular, $X^{\mathcal{M}_{k}(X, Y)}$, the sketch should be powerful enough to retrieve all characters in forward contexts of breakpoints $\mathcal{B}_{X, Y}\left(\mathcal{A}_{\mathrm{W}}\right)$. Even though $\mathcal{B}_{X, Y}\left(\mathcal{A}_{\mathrm{W}}\right)$ could be of size $\Theta\left(k^{2}\right)$, due to the aforementioned bounds on $\left|\operatorname{LZ}\left(X^{M}\right)\right|$ and the number of dummy segments in $X^{M}$, it suffices to take $\mathcal{O}(k)$ among these forward contexts to cover the unmasked regions of $X^{M}$ and $X^{\mathcal{M}_{k}(X, Y)}$. Each context can be encoded in $\tilde{\mathcal{O}}(k)$ bits, so this paves a way towards sketches of size $\tilde{\mathcal{O}}\left(k^{2}\right)$.

Nevertheless, while processing a string $X \in \Sigma^{\leq n}$, we only have access to the string $X$ and the $m$-complete walk $\left(x_{t}\right)_{t=1}^{m}=\mathrm{W}(n, r, X)$ over $X$. In particular, depending on $Y$, any position in $X$ could be involved in a breakpoint. A naive strategy would be to build a context encoding $\mathrm{CE}(X)[1 \ldots m]$ that stores at $t \in[1 \ldots m]$ (a compressed representation of) the forward context starting at $X\left[x_{t}\right]$, and then post-process it using a Hamming-distance sketch. This is sufficiently powerful because $X\left[x_{t}\right] \neq Y\left[y_{t}\right]$ holds for any $\left(x_{t}, y_{t}\right) \in \mathcal{B}\left(\mathcal{A}_{\mathrm{W}}\right)$ (recall that $\mathcal{A}_{\mathrm{W}}$ is greedy). Unfortunately, this construction does not guarantee any upper bound on hd $(\operatorname{CE}(X), \mathrm{CE}(Y))$ in terms of $k$. (For example, if $X$ is compressible, modifying its final character of $X$ affects
the entire $\operatorname{CE}(X)$.) Hence, we sparsify $\operatorname{CE}(X)$ by placing a blank symbol $\perp$ at some positions $\mathrm{CE}(X)[t]$ so that just a few forward contexts stored in $\mathrm{CE}(X)[t]$ cover any single position in $X$.

This brings two further challenges. First, if $X[x]$ is involved in a breakpoint, then we are only guaranteed that it is covered by some forward context $X[p \ldots q)$ of $\mathrm{CE}(X)[t]$ (i.e., $x \in[p \ldots q)$ ). In particular, the forward context starting at position $x$ could extend beyond $X[x \ldots q)$. Hence, the string $\operatorname{CE}(X)[t]$ actually stores double forward contexts $X[p \ldots r)=X[p . . q) X[q \ldots r)$ defined as the concatenation of the forward contexts of $X[p]$ and $X[q]$. We expect this double forward context $X[p \ldots r)$ to cover the entire forward context of $X[x]$. Unfortunately, this is not necessarily true if we use the Lempel-Ziv factorization to quantify compressibility: we could have $|\operatorname{LZ}(X[x \ldots r))|<|\operatorname{LZ}(X[q \ldots r))|$ because $\mathrm{LZ}(\cdot)$ is not monotone. Instead, we use an ad-hoc compressibility measure defined as $\overline{\operatorname{maxLZ}}(S)=\max _{[\ell . r) \subseteq[1 .|S|]} \mathrm{LZ}(\overline{S[\ell \ldots r)})$. Here, maximization over substrings guarantees monotonicity whereas reversal helps designing an efficient streaming algorithm constructing contexts (beyond the scope of this overview).

Another challenge is that the sparsification needs to be consistent between $\operatorname{CE}(X)$ and $\mathrm{CE}(Y)$ : assuming ed $(X, Y) \leq k$, we should have $\operatorname{hd}(\operatorname{CE}(X), \mathrm{CE}(Y))=\tilde{\mathcal{O}}(k)$, which also accounts for mismatches between $\perp$ and a stored double forward context. This rules out a naive strategy of covering $X$ from left to right using disjoint forward contexts: any substitution at the beginning of $X$ could then have a cascade of consequences throughout $\operatorname{CE}(X)$. Hence, we opt for a memoryless strategy that decides on $\operatorname{CE}(X)[t]$ purely based on the forward contexts $X\left[x_{t-1} \ldots x_{t-1}^{\prime}\right)$ and $X\left[x_{t} \ldots x_{t}^{\prime}\right)$. For example, we could set $\operatorname{CE}(X)[t]=\perp$ unless the smallest dyadic interval containing $\left[x_{t-1} \ldots x_{t-1}^{\prime}\right)$ differs from the smallest dyadic interval containing $\left[x_{t} \ldots x_{t}^{\prime}\right)$ (a dyadic interval is of the form $\left[i 2^{j} \ldots(i+1) 2^{j}\right)$ for some integers $\left.i, j \geq 0\right)$. With this approach, each position of $X$ is covered by at least one and at most $\mathcal{O}(\log |X|)$ forward contexts. Furthermore, substituting any character in $X$ does not have far-reaching knock-on effects. Unfortunately, insertions and deletions are still problematic as they shift the positions $x_{t}$. Thus, the decision concerning $\operatorname{CE}(X)[t]$ should be independent of the numerical value of $x_{t}$. Consequently, instead of looking at the smallest dyadic interval containing $\left[x_{t} \ldots x_{t}^{\prime}\right)$, we choose the largest $t^{\prime}$ such that $\left[x_{t} \ldots x_{t}^{\prime}\right)=\left[x_{t} \ldots x_{t^{\prime}}\right)$, and we look at the smallest dyadic interval containing $\left[t \ldots t^{\prime}\right)$.

With each forward context $X\left[x_{t} \ldots x_{t}^{\prime}\right.$ ) retrieved, we also need to determine the value $x_{t}$ (so that we know which fragment of $X$ we can learn from $\operatorname{CE}(X)[t]$. To avoid knock-on effects, we actually store differences $x_{t}-x_{u}$ with respect to the previous index satisfying $\operatorname{CE}(X)[u] \neq \perp$. This completes the intuitive description of the context encoding CE studied in Section 6.2.

Our edit-distance sketch contains the Hamming-distance sketch of $\operatorname{CE}(X)$. For this, we use an existing construction [16], augmented in a black-box manner to support large alphabets (recall that the each forward contexts takes $\tilde{\mathcal{O}}(k)$ bits). Furthermore, to retrieve the starting positions $x_{t}$ (rather than just the differences $x_{t}-x_{u}$ ), we use a hierarchical Hamming-distance sketch similar to those used in $[6,37]$. This way, given sketches of $X$ and $Y$, we can recover all characters that remain unmasked in $X^{\mathcal{M}_{k}(X, Y)}$ and $Y^{\mathcal{M}_{k}(X, Y)}$. The tools of Section 5 are then used to compute ed $(X, Y)$ (or certify ed $(X, Y)>k$ ) and retrieve the greedy encoding $\mathrm{GR}_{k}(X, Y)$ (see Section 6.4). We summarize the properties of the edit distance sketches below:

Theorem 3.17. For every constant $\delta \in(0,1)$, there is a sketch $\mathbf{s k}_{k}^{E}$ (parametrized by integers $n \geq k \geq 1$, an alphabet $\Sigma=\left[0 \ldots n^{\mathcal{O}(1)}\right)$, and a seed of $\mathcal{O}\left(\log ^{2} n\right)$ random bits) such that:
(a) The sketch $\mathbf{s k}_{k}^{E}(S)$ of a string $\Sigma^{\leq n}$ takes $\mathcal{O}\left(k^{2} \log ^{3} n\right)$ bits. Given streaming access to $S$, it can be constructed in $\tilde{\mathcal{O}}(n k)$ time using $\tilde{\mathcal{O}}\left(k^{2}\right)$ space.
(b) There exists an $\tilde{\mathcal{O}}\left(k^{2}\right)$-space decoding algorithm that, given $\mathbf{s k}_{k}^{E}(X), \mathbf{s k}_{k}^{E}(Y)$ for $X, Y \in$ $\Sigma^{\leq n}$, with probability at least $1-\delta$ outputs $\mathrm{GR}_{k}(X, Y)$ and $\min (\mathrm{ed}(X, Y), k+1)$. Retrieving $\mathrm{GR}_{k}(X, Y)$ costs $\tilde{\mathcal{O}}\left(k^{5}\right)$ time, whereas computing $\min (\mathrm{ed}(X, Y), k+1)$ costs $\tilde{\mathcal{O}}\left(k^{3}\right)$ time.

## 4 Compressed String Representation

In this section, we develop a data structure that stores a string $X$ in $\mathcal{O}\left(|\operatorname{LZ}(X)| \log ^{2}|X|\right)$ space, supporting various operations in a relatively efficient way (the $\log |X|$ factors are not optimized).

Our data structure can be constructed not only from $\operatorname{LZ}(X)$, but from any $L Z$-like representation describing a factorization $X=F_{1} \cdots F_{f}$ into non-empty phrases such that each phrase $F_{j}$ with $\left|F_{j}\right|>1$ is a previous factor (unlike LZ77, the phrases do not need to be a maximal).

Observation 4.1. Every LZ-like factorization of a string $X \in \Sigma^{*}$ has at least $|\operatorname{LZ}(X)|$ phrases. Moreover, $|\mathrm{LZ}(X)| \leq|\mathrm{LZ}(X Y)| \leq|\mathrm{LZ}(X)|+|\mathrm{LZ}(Y)|$ holds for all strings $X, Y \in \Sigma^{*}$.

Proposition 4.2. In the $w$-bit word RAM model, every string $X \in \Sigma^{n}$ satisfying $w=\Omega(\log n+$ $\log |\Sigma|)$, $z=|\operatorname{LZ}(X)|$, and $\bar{z}=|\operatorname{LZ}(\bar{X})|$ has a representation $\mathrm{D}(X)$ that uses $\mathcal{O}\left(z \log ^{2} n\right)$ space and supports the following queries:
(a) retrieve $n=|X|$, in $\mathcal{O}(1)$ time;
(b) given $i \in[1 \ldots n]$, retrieve $X[i]$, in $\mathcal{O}(\log n)$ time;
(c) given $i, j \in[1 . . n]$, compute $\operatorname{LCE}(X[i \ldots n], X[j \ldots n])$, in $\mathcal{O}(\log n)$ time;
(d) given $i, j \in[1 \ldots n]$, compute $\operatorname{LCE}(\bar{X}[i \ldots n], \bar{X}[j \ldots n])$, in $\mathcal{O}(\log n)$ time;
(e) compute $\mathrm{LZ}(X)$, in $\mathcal{O}\left(z \log ^{4} n\right)$ time;
(f) compute $\mathrm{LZ}(\bar{X})$, in $\mathcal{O}\left(\bar{z} \log ^{4} n\right)$ time;
(g) given $i, j \in[1 \ldots n]$ with $i \leq j$, compute $\mathrm{D}(X[i \ldots j])$, in $\mathcal{O}\left(z \log ^{4} n\right)$ time.

Moreover, $\mathrm{D}(X)$ can be constructed:
(h) in $\mathcal{O}\left(f \log ^{2} n\right)$ time given an $f$-phrase LZ-like representation of $X$ or $\bar{X}$;
(i) in $\mathcal{O}\left(z \log ^{4} n\right)$ time given $\mathrm{D}(X[1 \ldots i))$ and $\mathrm{D}(X[i \ldots n])$ for some $i \in[1 \ldots n]$.

Before proving Proposition 4.2, we need to introduce several more compression schemes. These concepts are not used any of the subsequent sections.

A straight-line grammar is a tuple $\mathcal{G}=(\mathcal{S}, \Sigma$, rhs,$S)$, where $\mathcal{S}$ is a finite set of symbols, $\Sigma \subseteq \mathcal{S}$ is a set of terminal symbols, rhs: $(\mathcal{S} \backslash \Sigma) \rightarrow \mathcal{S}^{*}$ is the production (or right-hand side) function, and $S \in \mathcal{S}$ is the start symbol. We further require existence of an order $\prec$ on $\mathcal{S}$ such that $B \prec A$ if $B$ occurs in $\operatorname{rhs}(A)$.

The expansion function $\exp : \mathcal{S} \rightarrow \Sigma^{+}$is defined recursively:

$$
\exp (A)= \begin{cases}A & \text { if } A \in \Sigma, \\ \exp \left(B_{1}\right) \cdots \exp \left(B_{k}\right) & \text { if } \operatorname{rhs}(A)=B_{1} \cdots B_{k},\end{cases}
$$

We say that $\mathcal{G}$ is a grammar-compressed representation of $\exp (S)$.
The parse tree $\mathcal{T}(A)$ of a symbol $A \in \mathcal{S}$ is a rooted ordered tree with each node $\nu$ associated to a symbol $s(\nu) \in \mathcal{S}$. The root of $\mathcal{T}(A)$ is a node $\rho$ with $s(\rho)=A$. If $A \in \Sigma$, then $\rho$ has no children. Otherwise, if $\operatorname{rhs}(A)=B_{1} \cdots B_{k}$, then $\rho$ has $k$ children, and the subtree rooted at the $i$ th child is (a copy of) $\mathcal{T}\left(B_{i}\right)$. The parse tree of a grammar $\mathcal{G}$ is defined as the parse tree $\mathcal{T}(S)$ of the starting symbol $S$, and the height of $\mathcal{G}$ is defined as the height of $\mathcal{T}(S)$.

Each node $\nu$ of $\mathcal{T}(A)$ is associated with an occurrence $\exp (\nu)$ of $\exp (s(\nu))$ in $\exp (A)$. For the root $\rho$, we define $\exp (\rho)=\exp (A)[1 \ldots|\exp (A)|]$ to be the whole $\exp (A)$. Moreover, if $\exp (\nu)=\exp (A)[\ell . . r), \operatorname{rhs}(s(\nu))=B_{1} \cdots B_{k}$, and $\nu_{1}, \ldots, \nu_{k}$ are the children of $\nu$, then $\exp \left(\nu_{i}\right)=\exp (A)\left[r_{i-1} \ldots r_{i}\right)$, where $r_{i}=\sum_{j=1}^{i}\left|\exp \left(B_{j}\right)\right|$ for $0 \leq i \leq k$. This way, the fragments $\exp \left(\nu_{i}\right)$ form a partition of $\exp (\nu)$, and $\exp \left(\nu_{i}\right)$ is indeed an occurrence $\exp \left(s\left(\nu_{i}\right)\right)$ in $\exp (A)$.

A straight-line program (SLP) is a straight-line grammar $\mathcal{G}=(\mathcal{S}, \Sigma$, rhs,$S)$ such that $|\operatorname{rhs}(A)|=2$ holds for each nonterminal $A \in \mathcal{S} \backslash \Sigma$. In a run-length straight-line program (RLSLP), productions of the form $\operatorname{rhs}(A)=B^{k}$, with $B \in \mathcal{S}$ and $k \in \mathbb{Z}_{\geq 2}$, are also allowed.

An RLSLP $\mathcal{G}$ of size $g$ (with $g$ symbols) representing a string of length $n$ can be stored in $\mathcal{O}(g)$ space $(\mathcal{O}(g \log (n+g))$ bits) with each non-terminal $A \in \mathcal{S} \backslash \Sigma$ storing rhs $(A)$ (represented by $(B, k)$ if $\left.\operatorname{rhs}(A)=B^{k}\right)$ and $|\exp (A)|$. This representation allows for efficiently traversing the parse tree $\mathcal{T}_{\mathcal{G}}$ : given a node $\nu$ represented as a pair $(s(\nu), \exp (\nu))$, it is possible to retrieve in constant time an analogous representation of any child of $\nu$.

Fact 4.3. Let $\mathcal{G}$ be an RLSLP of size $g$ representing a string $X$.
(a) An LZ-like factorization of $X$ can be constructed in $\mathcal{O}(g)$ time.
(b) For every fragment $X[i \ldots j]$, an RLSLP representing $X[i \ldots j]$ can be constructed in $\mathcal{O}(g)$ time.

Proof. We traverse the parse tree $\mathcal{T}(S)$ in pre-order, skipping some subtrees depending on the application.
(a) For each symbol $A$, we maintain an already visited node $\nu$ with symbol $s(\nu)=A$, if any. When visiting a node $\nu$, we first check whether any node $\nu^{\prime}$ with $s\left(\nu^{\prime}\right)=s(\nu)$ has already been visited. If so, we retrieve the expansion $\exp \left(\nu^{\prime}\right)=X\left[i^{\prime} . . i^{\prime}+\ell\right.$ ), and we output $\left(i^{\prime}, \ell\right)$ as a part of the LZ-like representation. Otherwise, we save $\nu$ as an already visited node with symbol $s(\nu)$ and proceed as follows:

- If $\nu$ is a leaf, then we output $s(\nu) \in \Sigma$ as a part of the LZ-like representation.
- If $\nu$ has two children, we process them recursively.
- If $\nu$ has $k \geq 3$ children, we process the first child $\nu_{1}$ recursively, retrieve $\exp \left(\nu_{1}\right)=$ $X[i \ldots i+\ell)$, and output $(i,(k-1) \ell)$ as a part of the LZ-like representation.
The overall running time $\mathcal{O}(g)$ is amortized by the number of distinct symbols visited.
(b) Here, an intermediate goal is to construct a sequence $\left(A_{1}, m_{1}\right), \ldots,\left(A_{t}, m_{t}\right)$ with $A_{p} \in \mathcal{S}$ and $m_{p} \in \mathbb{Z}_{\geq 1}$ such that $\exp \left(A_{1}\right)^{m_{1}} \cdots \exp \left(A_{t}\right)^{m_{t}}=X[i \ldots j]$ and $t=\mathcal{O}(g)$. We only visit nodes $\nu$ such that $\exp (\nu)$ intersects $X[i . . j]$.
- If $\exp (\nu)$ is contained in $X[i . . j]$, we simply output $(s(\nu), 1)$ and skip the subtree of $\nu$.
- Otherwise, we determine the children $\nu_{\ell}, \ldots, \nu_{r}$ of $\nu$ whose expansions intersect $X[i \ldots j]$. Then, we recurse into $\nu_{\ell}$, output $\left(s\left(\nu_{\ell}\right), r-\ell-1\right)$ if $r>\ell+1$, and recurse into $\nu_{r}$ if $r>\ell$.
The number of visited nodes is proportional to the height of $\mathcal{G}$ and thus this traversal takes $\mathcal{O}(g)$ time. In the post-processing, for each pair $\left(A_{p}, m_{p}\right)$, we create a new symbol $B_{p}$ with $\operatorname{rhs}\left(B_{p}\right)=A_{p}^{m_{p}}$ (if $m_{p} \geq 2$ ) or set $B_{p}$ as an alias of $A_{p}$ (if $m_{p}=1$ ). Next, we set $C_{1}$ to be an alias of $B_{1}$ and, for $p \in[2 \ldots t]$, create a new symbol $C_{p}$ with $\operatorname{rhs}\left(C_{p}\right)=C_{p-1} B_{p-1}$. Finally, we return the extended RLSLP with $C_{t}$ as the new starting symbol.

Theorem 4.4 (I [35]). Let $\mathcal{G}$ be a size-g SLP representing a string $X \in \Sigma^{n}$ and let $g^{*}$ be the minimum size of an $S L P$ representing $X$. Given $\mathcal{G}$, in $\mathcal{O}(g \log n)$ time, one can construct a size- $\mathcal{O}\left(g^{*} \log n\right)$ RLSLP representing $X$ and a size- $\mathcal{O}\left(g^{*} \log n\right)$ data structure that answers the following queries in $\mathcal{O}(\log n)$ time:
Access: Given $i \in[1 \ldots n]$, retrieve $S[i]$;
LCE queries: Given $i, j \in[1 \ldots n]$, compute $\operatorname{LCE}(S[i \ldots n], S[j \ldots n])$.
Proof of Proposition 4.2. We define $\mathrm{D}(X)$ so that it consists of two instances of the data structure of Theorem 4.4, one for $X$ and one for $\bar{X}$, including the RLSLPs representing $X$ and $\bar{X}$. Thus, the size of $\mathrm{D}(X)$ is $\mathcal{O}\left(g^{*} \log n\right)$, where $g^{*}$ is the minimum size of an SLP representing $X$ (it is the same for $\bar{X}$ ).

To implement (h), we first use [39, Theorem 6.1] and build a size- $\mathcal{O}(f \log n)$ SLP representing $X$ or $\bar{X}$. By reversing the right-hand sides of all the productions, we obtain an analogous SLP representing $\bar{X}$ or $X$, respectively. Finally, we pass these SLPs to the construction algorithm of Theorem 4.4. Overall, the running time is $\mathcal{O}\left(f \log ^{2} n\right)$.

In (i), given $\mathrm{D}(X[1 \ldots i))$ and $\mathrm{D}(X[i \ldots n])$, we convert the underlying RLSLPs to LZ-like representations (Fact 4.3), concatenate them into an $\mathcal{O}\left(g^{*} \log n\right)$-phrase LZ-like representation of $X$, and apply (h). Overall, this takes $\mathcal{O}\left(g^{*} \log ^{3} n\right)=\mathcal{O}\left(z \log ^{4} n\right)$ time.

As for (a), we note that $|X|$ can be trivially retrieved from the RLSLP representing $X$. Queries (b), (c), and (d) are implemented directly using Theorem 4.4.

In (e), we construct an index of [39, Theorem 6.11] that, given a fragment $X[i . . i+\ell$ ), in $\mathcal{O}\left(\log ^{3} n\right)$ time locates the leftmost position $i^{\prime}$ with $X\left[i^{\prime} . . i^{\prime}+\ell\right)=X[i . . i+\ell)$. This index can be constructed in $\mathcal{O}\left(g^{*} \log ^{3} n\right)$ time: it is already based on the RLSLP of Theorem 4.4, and the extra construction time is $\mathcal{O}\left(\log ^{2} n\right)$ per RLSLP symbol; see [39, Lemma 6.9]. This way, each phrase $F_{j}$ of the LZ77 parsing of $X$ can be constructed, by binary search, with $\mathcal{O}(\log n)$ queries to the index, for a total of $\mathcal{O}\left(z \log ^{4} n\right)$ time (see also [51, Section 5]). The algorithm for (f) is symmetric.

As for (g), we construct an $\mathcal{O}\left(g^{*} \log n\right)$-phrase LZ-like representation of $X[i \ldots j]$ using Fact 4.3, and then apply (h). Overall, this takes $\mathcal{O}\left(g^{*} \log ^{3} n\right)=\mathcal{O}\left(z \log ^{4} n\right)$ time.

## 5 Greedy Alignments and Encodings

Definition 3.6 (Greedy alignment). We say that an alignment $\mathcal{A}$ of two strings $X, Y \in \Sigma^{*}$ is greedy if $X[x] \neq Y[y]$ holds for every $(x, y) \in \mathcal{B}(\mathcal{A}) \cap([1 . .|X|] \times[1 . .|Y|])$. Given $k \geq 0$, we denote by $\mathrm{GA}_{k}(X, Y)$ the set of all greedy alignments $\mathcal{A}$ of $X, Y$ satisfying $\operatorname{cost}(\mathcal{A}) \leq k$.

Fact 3.7. For any two strings $X, Y \in \Sigma^{*}$, there is an optimal greedy alignment of $X, Y$.
Proof. Let $\mathcal{A}=\left(x_{t}, y_{t}\right)_{t=1}^{m}$ be an optimal alignment maximizing the sum $\sum_{t=1}^{m}\left(x_{t}+y_{t}\right)$. For a proof by contradiction, suppose that $\mathcal{A}$ is not greedy. Then, there exists $i \in[1 \ldots m]$ such that $\left(x_{i}, y_{i}\right) \in \mathcal{B}(A) \cap([1 \ldots|X|] \times[1 \ldots|Y|])$ and $X\left[x_{i}\right]=Y\left[y_{i}\right]$. Observe that $x_{i+1}=x_{i}$ or $y_{i+1}=y_{i}$ and, by symmetry, assume $x_{i+1}=x_{i}$ without loss of generality. Let us define $j>i$ so that $x_{i}=\cdots=x_{j}<x_{j+1}$ and consider two cases depending on whether $y_{j}=y_{j+1}$ :

- If $y_{j}=y_{j+1}$, we define an alignment $\mathcal{A}^{\prime}$ obtained from $\mathcal{A}$ by deleting $\left(x_{j}, y_{j}\right)$. It is easy to see that $\mathcal{A}^{\prime}$ is a valid alignment. Moreover, $\mathcal{B}\left(\mathcal{A}^{\prime}\right) \subseteq \mathcal{B}(\mathcal{A})$ and $\mathcal{B}(\mathcal{A}) \backslash \mathcal{B}\left(\mathcal{A}^{\prime}\right)=\left\{\left(x_{j}, y_{j}\right)\right\}$, so $\operatorname{cost}\left(\mathcal{A}^{\prime}\right)<\operatorname{cost}(\mathcal{A})$, contradicting the choice of $\mathcal{A}$ as an optimal alignment.
- If $y_{j}<y_{j+1}$, we define an alignment $\mathcal{A}^{\prime}=\left(x_{t}^{\prime}, y_{t}^{\prime}\right)_{t=1}^{m}$ obtained from $\mathcal{A}$ by incrementing $x_{t}$ for $t \in(i . . j]$. Moreover, $\mathcal{B}(\mathcal{A}) \backslash \mathcal{B}\left(\mathcal{A}^{\prime}\right) \subseteq\left\{\left(x_{t}, y_{t}\right): t \in[i . . j)\right\}$ and $\mathcal{B}\left(\mathcal{A}^{\prime}\right) \backslash \mathcal{B}(\mathcal{A})=$ $\left\{\left(x_{t}^{\prime}, y_{t}^{\prime}\right): t \in(i \ldots j]\right\}$, so $\operatorname{cost}\left(\mathcal{A}^{\prime}\right) \leq \operatorname{cost}(\mathcal{A})$. Furthermore, $\sum_{t=1}^{m}\left(x_{t}^{\prime}+y_{t}^{\prime}\right)=j-i+$ $\sum_{t=1}^{m}\left(x_{t}+y_{t}\right)$, contradicting the choice of $\mathcal{A}$ as an optimal alignment maximizing $\sum_{t=1}^{m}\left(x_{t}+\right.$ $y_{t}$ ).

For strings $X, Y \in \Sigma^{*}$ and an integer $k \geq e d(X, Y)$, we define a set

$$
\mathcal{M}_{k}(X, Y)=\bigcap_{\mathcal{A} \in \mathrm{GA}_{k}(X, Y)} \mathcal{M}_{X, Y}(\mathcal{A})
$$

of common matches of all alignments $\mathcal{A} \in \mathrm{GA}_{k}(X, Y)$; note that it forms a non-crossing matching of $X, Y$.

Definition 3.8. Let $M$ be a non-crossing matching of strings $X, Y \in \Sigma^{*}$. We define $X^{M}, Y^{M}$ to be the strings obtained from $X, Y$ by replacing $X[x]$ and $Y[y]$ with $\# \notin \Sigma$ for every $(x, y) \in M$. We refer to \# as a dummy symbol and to maximal blocks of \#'s as dummy segments.

Additionally, for a string $Z \in(\Sigma \cup \#)^{*}$, we define num $(Z)$ as a string obtained from $Z$ by replacing the $i$ th leftmost occurrence of $\#$ in it with a unique dummy symbol $\#_{i} \notin \Sigma$.

Fact 5.1. Let $X, Y \in \Sigma^{*}$. For each non-crossing matching $M$ of $X, Y$, we have ed $(X, Y) \leq$ ed (num $\left(X^{M}\right)$, num $\left.\left(Y^{M}\right)\right)$.

Proof. For every $(x, y) \in[1 \ldots|X|] \times[1 \ldots|Y|]$, if num $\left(X^{M}\right)[x]=\operatorname{num}\left(Y^{M}\right)[y]$, then $X[x]=Y[y]$. Thus, $\operatorname{cost}_{X, Y}(\mathcal{A}) \leq \operatorname{cost}_{\text {num }\left(X^{M}\right), \text { num }\left(Y^{M}\right)}(\mathcal{A})$ holds for all alignments $\mathcal{A}$ of $X, Y$.

Lemma 3.9. Consider strings $X, Y \in \Sigma^{*}$, an integer $k \geq \operatorname{ed}(X, Y)$, and a set $M \subseteq \mathcal{M}_{k}(X, Y)$. Then, $\operatorname{ed}(X, Y)=\operatorname{ed}\left(\operatorname{num}\left(X^{M}\right), \operatorname{num}\left(Y^{M}\right)\right)$ and $\mathcal{M}_{k}(X, Y)=\mathcal{M}_{k}\left(\operatorname{num}\left(X^{M}\right), \operatorname{num}\left(Y^{M}\right)\right)$.

Proof. Denote $X^{\prime}=\operatorname{num}\left(X^{M}\right)$ and $Y^{\prime}=\operatorname{num}\left(Y^{M}\right)$, and let $\mathcal{A} \in \mathrm{GA}_{k}(X, Y)$. We shall first prove that $\mathcal{M}_{X, Y}(\mathcal{A}) \subseteq \mathcal{M}_{X^{\prime}, Y^{\prime}}(\mathcal{A})$. Suppose that $(x, y) \in \mathcal{M}_{X, Y}(\mathcal{A})$, i.e., $X[x] \simeq_{\mathcal{A}} Y[y]$. Note that $\mathcal{M}_{X, Y}(\mathcal{A})$ is a non-crossing matching of $X, Y$ and a superset of $M$, so either $(x, y) \in M$, and then $X^{\prime}[x]=Y^{\prime}[y]$ is a dummy symbol, or $M$ contains no pair involving $x$ or $y$, and then $X^{\prime}[x]=X[x]=Y[y]=Y^{\prime}[y]$. In both cases, we have $X^{\prime}[x] \simeq_{\mathcal{A}} Y^{\prime}[y]$, i.e., $(x, y) \in \mathcal{M}_{X^{\prime}, Y^{\prime}}(\mathcal{A})$. This completes the proof that $\mathcal{M}_{X, Y}(\mathcal{A}) \subseteq \mathcal{M}_{X^{\prime}, Y^{\prime}}(\mathcal{A})$, from which we derive $\mathcal{B}_{X^{\prime}, Y^{\prime}}(\mathcal{A}) \subseteq$ $\mathcal{B}_{X, Y}(\mathcal{A})$ and $\operatorname{cost}_{X^{\prime}, Y^{\prime}}(\mathcal{A}) \leq \operatorname{cost}_{X, Y}(\mathcal{A}) \leq k$.

Next, we shall prove that $\mathcal{A} \in \operatorname{GA}_{k}\left(X^{\prime}, Y^{\prime}\right)$. Suppose that $X^{\prime}[x]=Y^{\prime}[y]$ holds for some $(x, y) \in \mathcal{A}$. Then, we have $X[x]=Y[y]$ because either $(x, y) \in M$ or $X[x]=X^{\prime}[x]=Y^{\prime}[y]=$ $Y[y]$. Since $\mathcal{A}$ is a greedy alignment of $X, Y$, this implies $(x, y) \in \mathcal{M}_{X, Y}(\mathcal{A})$, i.e., $X[x] \simeq_{\mathcal{A}} Y[y]$. Due to the assumption $X^{\prime}[x]=Y^{\prime}[y]$, we conclude that $X^{\prime}[x] \simeq_{\mathcal{A}} Y^{\prime}[y]$, i.e., $(x, y) \in \mathcal{M}_{X^{\prime}, Y^{\prime}}(\mathcal{A})$. This proves $\mathcal{A} \in \mathrm{GA}_{k}\left(X^{\prime}, Y^{\prime}\right)$.

Now, let $\mathcal{A} \in \mathrm{GA}_{k}\left(X^{\prime}, Y^{\prime}\right)$. We shall first prove that $\mathcal{M}_{X^{\prime}, Y^{\prime}}(\mathcal{A}) \subseteq \mathcal{M}_{X, Y}(\mathcal{A})$. Suppose that $(x, y) \in \mathcal{M}_{X^{\prime}, Y^{\prime}}(\mathcal{A})$, i.e., $X^{\prime}[x] \simeq_{\mathcal{A}} Y^{\prime}[y]$. In particular, $X^{\prime}[x]=Y^{\prime}[y]$, which implies $X[x]=Y[y]$ because either $(x, y) \in M$ or $X[x]=X^{\prime}[x]=Y^{\prime}[y]=Y[y]$. Hence, $X[x] \simeq_{\mathcal{A}} Y[y]$, i.e., $(x, y) \in \mathcal{M}_{X, Y}(\mathcal{A})$. This completes the proof that $\mathcal{M}_{X^{\prime}, Y^{\prime}}(\mathcal{A}) \subseteq \mathcal{M}_{X, Y}(\mathcal{A})$, from which we derive $\mathcal{B}_{X, Y}(\mathcal{A}) \subseteq \mathcal{B}_{X^{\prime}, Y^{\prime}}(\mathcal{A})$ and $\operatorname{cost}_{X, Y}(\mathcal{A}) \leq \operatorname{cost}_{X^{\prime}, Y^{\prime}}(\mathcal{A}) \leq k$.

Next, we shall prove that $\mathcal{A} \in \mathrm{GA}_{k}(X, Y)$. For a proof by contradiction, suppose that $X[x]=$ $Y[y]$ holds for some $(x, y) \in \mathcal{B}_{X, Y}(\mathcal{A})$; if there are several such breakpoints, let us choose the leftmost one. Note that $X[1 \ldots x) \sim_{\mathcal{A}} Y[1 \ldots y)$ and $X[x \ldots|X|] \sim_{\mathcal{A}} Y[y \ldots|Y|]$. Let us construct an alignment $\mathcal{A}^{\prime}$ obtained from $\mathcal{A}$ by replacing the induced alignment $\mathcal{A}_{[x .|X|,|y .|Y|]}$ with an optimum greedy alignment of $X[x . .|X|], Y[y . .|Y|]$ (see Fact 3.7). By the choice of $(x, y)$, the induced alignment $\mathcal{A}_{[1 . . x),[1 . . y)}^{\prime}=\mathcal{A}_{[1 \ldots x),[1 . y)}$ is a greedy alignment of $X[1 \ldots x), Y[1 \ldots y)$ and thus $\mathcal{A}^{\prime}$ is a greedy alignment of $X, Y$. Due to $\operatorname{cost}_{X, Y}\left(\mathcal{A}^{\prime}\right) \leq \operatorname{cost}_{X, Y}(\mathcal{A}) \leq k$, this implies $\mathcal{A}^{\prime} \in \mathrm{GA}_{k}(X, Y)$. Hence, $\mathcal{M}_{X, Y}\left(\mathcal{A}^{\prime}\right)$ is a non-crossing matching of $X, Y$ and a superset of $M$. The construction of $\mathcal{A}^{\prime}$ further guarantees $(x, y) \in \mathcal{M}_{X, Y}\left(\mathcal{A}^{\prime}\right)$, so either $(x, y) \in M$, and then $X^{\prime}[x]=Y^{\prime}[y]$ is a dummy symbol, or $M$ contains no pair involving $x$ or $y$, and then $X^{\prime}[x]=X[x]=Y[y]=Y^{\prime}[y]$. In both cases, we have $X^{\prime}[x]=Y^{\prime}[y]$ and, since $\mathcal{A}$ is a greedy alignment of $X^{\prime}, Y^{\prime}$, we derive $X^{\prime}[x] \simeq_{\mathcal{A}} Y^{\prime}[y]$ and $(x, y) \in \mathcal{M}_{X, Y}(\mathcal{A})$. This contradicts the choice of $(x, y)$, completing the proof that $\mathcal{A} \in \mathrm{GA}_{k}(X, Y)$.

Overall, we conclude that $\mathrm{GA}_{k}(X, Y)=\mathrm{GA}_{k}\left(X^{\prime}, Y^{\prime}\right)$ and that every alignment $\mathcal{A}$ in this family satisfies both $\mathcal{M}_{X, Y}(\mathcal{A})=\mathcal{M}_{X^{\prime}, Y^{\prime}}(\mathcal{A})$ and $\operatorname{cost}_{X, Y}(\mathcal{A})=\operatorname{cost}_{X^{\prime}, Y^{\prime}}(\mathcal{A})$. Consequently, $\mathcal{M}_{k}(X, Y)=\mathcal{M}_{k}\left(X^{\prime}, Y^{\prime}\right)$ and, due to Fact 3.7, ed $(X, Y)=\operatorname{ed}\left(X^{\prime}, Y^{\prime}\right)$.

Proposition 5.2. Consider a string $S \in(\Sigma \cup\{\#\})^{n}$ with s dummy segments. Given the sorted list of dummy segments in $S$, one can in $\mathcal{O}(s)$ time construct a data structure $\mathrm{RS}_{\#}(S)$ supporting the following queries in $\mathcal{O}(\log s)$ time:

1. $\operatorname{rank}_{\#}(S, i):$ Given $i \in[1 \ldots n+1]$, return $\left|\left\{i^{\prime} \in[1 \ldots i): S\left[i^{\prime}\right]=\#\right\}\right|$.
2. $\operatorname{select}_{\#}(S, j)$ : Given $j \in\left[1 \ldots \operatorname{rank}_{\#}(S, n)\right]$, return the $j$ th smallest element of $\{i \in[1 \ldots n]$ : $S[i]=\#\}$.

Proof. For each dummy segment $S[\ell \ldots r)$, the data structure stores a tuple $\left(\ell, r, \operatorname{rank}_{\#}(S, \ell)\right)$. These tuples are stored in a sorted array (note that the order is the same for all coordinates). To compute $\operatorname{rank}_{\#}(S, i)$, we binary search for the rightmost segment $S\left[\ell^{*} \ldots r^{*}\right)$ such that $\ell^{*} \leq i$. If there is no such segment, we return $\operatorname{rank}_{\#}(S, i)=0$. Otherwise, $\operatorname{rank}_{\#}(S, i)=\operatorname{rank}_{\#}\left(S, \ell^{*}\right)+$ $\min \left(i, r^{*}\right)-\ell^{*}$.

To compute select $\#(S, j)$, we binary search for the rightmost segment $S\left[\ell^{*} \ldots r^{*}\right)$ such that $\operatorname{rank}_{\#}\left(S, \ell^{*}\right)<j$. Then, $\operatorname{select}_{\#}(S, j)=\ell^{*}+j-1-\operatorname{rank}_{\#}\left(S, \ell^{*}\right)$.

Definition 5.3. For a non-crossing matching $M$ of strings $X, Y \in \Sigma^{*}$ we set (cf. Proposition 4.2):

$$
\mathrm{E}^{M}(X, Y):=\left(\mathrm{D}\left(X^{M} Y^{M}\right), \mathrm{RS}_{\#}\left(X^{M}\right), \mathrm{RS}_{\#}\left(Y^{M}\right)\right)
$$

Lemma 5.4. Let $X^{\prime}=\operatorname{num}\left(X^{M}\right)$ and $Y^{\prime}=\operatorname{num}\left(Y^{M}\right)$ for a non-crossing matching $M$ of strings $X, Y \in \Sigma^{\leq n}$. The encoding $\mathrm{E}^{M}(X, Y)$ allows answering LCE queries on the suffixes of $X^{\prime}, Y^{\prime}$ and on the suffixes of $\overline{X^{\prime}}, \overline{Y^{\prime}}$ in $\mathcal{O}(\log n)$ time.

Proof. Consider a query asking for $\operatorname{LCE}\left(X^{\prime}[x .],. Y^{\prime}[y .].\right)$. Observe that $\operatorname{LCE}\left(X^{\prime}[x .],. Y^{\prime}[y .].\right)=$ $\operatorname{LCE}\left(X^{M}[x .],. Y^{M}[y .].\right)$ if $\operatorname{rank}_{\#}\left(X^{M}, x\right)=\operatorname{rank}_{\#}\left(Y^{M}, y\right)$. Otherwise, $\operatorname{LCE}\left(X^{\prime}[x .],. Y^{\prime}[y .].\right)=$ $\min \left(\operatorname{LCE}\left(X^{M}[x .],. Y^{M}[y .].\right), \operatorname{select}_{\#}\left(X^{M}, \operatorname{rank}_{\#}\left(X^{M}, x\right)+1\right)-x, \operatorname{select}_{\#}\left(Y^{M}, \operatorname{rank}_{\#}\left(Y^{M}, y\right)+\right.\right.$ $1)-y)$. In either case, the query can be answered in $\mathcal{O}(\log n)$ time using Propositions 5.2 and $4.2(\mathrm{c})$. The values $\operatorname{LCE}\left(\overline{X^{\prime}}[x .],. \overline{Y^{\prime}}[y .].\right)$ are answered in a similar way.

Corollary 5.5. Given an integer $k>0$ and the encoding $\mathrm{E}^{M}(X, Y)$ for a non-crossing matching $M$ of strings $X, Y \in \Sigma^{\leq n}$, one can in $\mathcal{O}\left(k^{2} \log n\right)$ time compute a integer $d \in[0 \ldots k+1]$ such that

$$
d= \begin{cases}\operatorname{ed}(X, Y) & \text { if } \operatorname{ed}(X, Y) \leq k \text { and } M \subseteq \mathcal{M}_{k}(X, Y) \\ k+1 & \text { if } \operatorname{ed}(X, Y)>k\end{cases}
$$

Proof. We compute $d:=\min \left(k+1, \operatorname{ed}\left(\operatorname{num}\left(X^{M}\right), \operatorname{num}\left(Y^{M}\right)\right)\right)$ using the Landau-Vishkin algorithm [42], with $\mathcal{O}\left(k^{2}\right)$ LCE queries on suffixes of num $\left(X^{M}\right)$, num $\left(Y^{M}\right)$ implemented using Lemma 5.4. Due to Fact 5.1, we have $d \geq \min (k+1$, ed $(X, Y))$ and, in particular, $d=k+1$ if $\operatorname{ed}(X, Y)>k$. If $\operatorname{ed}(X, Y) \leq k$ and $M \subseteq \mathcal{M}_{k}(X, Y)$, then Lemma 3.9 yields $d=\min (k+1, \operatorname{ed}(X, Y))=\mathrm{ed}(X, Y)$.

### 5.1 Greedy Encoding and Its Size

Definition 5.6. For strings $X, Y \in \Sigma^{*}$ and an integer $k$, we define the greedy encoding

$$
\mathrm{GR}_{k}(X, Y)= \begin{cases}\mathrm{E}^{\mathcal{M}_{k}(X, Y)}(X, Y) & \text { if } k \geq \operatorname{ed}(X, Y) \\ \perp & \text { otherwise }\end{cases}
$$

Lemma 3.10. Let $M=\mathcal{M}_{k}(X, Y)$ for strings $X, Y \in \Sigma^{*}$ and a positive integer $k \geq \operatorname{ed}(X, Y)$. Then, $\left|\operatorname{LZ}\left(X^{M}\right)\right|,\left|\operatorname{LZ}\left(Y^{M}\right)\right|=\mathcal{O}\left(k^{2}\right)$, and $X^{M}, Y^{M}$ contain $\mathcal{O}(k)$ dummy segments.

Proof. We show the claim of the lemma for $X^{M}$, the claim for $Y^{M}$ follows by symmetry.
As ed $(X, Y) \leq k$, by Fact $3.7 \mathrm{GA}_{k}(X, Y) \neq \emptyset$, and therefore there is an alignment $\mathcal{A} \in$ $\mathrm{GA}_{k}(X, Y)$ of cost at most $k$ between $X$ and $Y$. We consider yet another graphical representation of an alignment. Namely, we represent $\mathcal{A}$ as a set of at most $k$ horizontal segments, where a horizontal segment $I=[i \ldots j ; \Delta]$ from $(i, \Delta)$ to $(j, \Delta)$ means that $X[i \ldots j] \simeq_{\mathcal{A}} Y[i+\Delta \ldots j+\Delta]$ (see Fig. 1). This representation induces a partitioning of $X=f_{1} \cdots f_{z}$, where $z=\mathcal{O}(k)$, and each factor $f_{\ell}$ is either a single character deleted under $\mathcal{A}$ or a fragment $X[i \ldots j]$ corresponding to a horizontal segment $I=[i \ldots j ; \Delta]$.


Alignment $\mathcal{A}$.

Alignment $\mathcal{A}^{\prime \prime}$.


Figure 1: Graphical representations of two greedy alignments $\mathcal{A}, \mathcal{A}^{\prime \prime} \in \mathrm{GA}_{5}(X, Y)$, where $X=$ abababcbcbc and $Y=$ abbabcbcbc. The alignment $\mathcal{A}$ (red) has cost 1 (we delete $X[3]$ ), and the alignment $\mathcal{A}^{\prime \prime}$ cost 5 (we delete $X[5], X[10], X[11], Y[3], Y[6]$ ). The alignments imply that $X[6 \ldots 9]=Y[7 \ldots 10]=X[8 \ldots 11]$.

If $\left|\mathrm{GA}_{k}(X, Y)\right|=1$, then every factor $f_{\ell}$ with $\left|f_{\ell}\right| \geq 2$ is replaced with an equal length string of dummy symbols. The upper bound on $\left|\mathrm{LZ}\left(X^{M}\right)\right|$ follows by sub-additivity of the Lempel-Ziv encoding, and the upper bound on the number of dummy segments is trivial.

Suppose now that $\left|\mathrm{GA}_{k}(X, Y)\right|>1$. We show that in this case the Lempel-Ziv encoding of any factor $f_{\ell}$ has size $\mathcal{O}(k)$. Fix $\ell$. If $f_{\ell}$ is a single character, the claim obviously holds. Suppose now that the factor $f_{\ell}$ of the partitioning defined above corresponds to a horizontal segment $I=[i \ldots j ; \Delta]$. Let us show that there is $i \leq q \leq j$ such that $X^{M}[i \ldots q]=X[i \ldots q]$ and $X^{M}[q+1 \ldots j]=\# \# \ldots \#$. This follows immediately from the following observation. Let $\mathcal{A}^{\prime} \neq \mathcal{A}$ be another alignment $\in \mathrm{GA}_{k}(X, Y)$. The alignment $\mathcal{A}^{\prime}$ is greedy, and therefore, if for some $i \leq p \leq j$ the alignment $\mathcal{A}^{\prime}$ matches $X[p]$ and $Y[p+\Delta]$, then it also matches $X[p+1]$ and $Y[p+1+\Delta], \ldots, X[j]$ and $Y[j+\Delta]$. We obtain an upper bound on the number of dummy segments as an immediate corollary, but an upper bound on $\operatorname{LZ}\left(X^{M}\right)$ requires more work.

Consider an alignment $\mathcal{A}^{\prime \prime} \in \mathrm{GA}_{k}(X, Y)$ realising $q$. We have $X^{M}[q+1 \ldots j]=\# \# \ldots \#$, and therefore $\left|\operatorname{LZ}\left(X^{M}[q+1 \ldots j]\right)\right|=\mathcal{O}(1)$. Consider now $X^{M}[i \ldots q]=X[i \ldots q]$. Consider the set of the horizontal segments corresponding to $\mathcal{A}^{\prime \prime}$ that contain positions in $[i \ldots q]$. We claim that either all segments in the set have height less than $\Delta$, or all segments have height larger than $\Delta$. Suppose that there are two consecutive horizontal segments $\left[i_{1} \ldots j_{1} ; \Delta_{1}\right]$ and $\left[i_{2} \ldots j_{2} ; \Delta_{2}\right]$ such that one of them is above $I$ and the other is below. $\mathcal{A}^{\prime \prime}$ must delete the characters $X\left[j_{1}+1\right], X\left[j_{1}+2\right], \ldots, X\left[i_{2}-1\right]$ and $Y\left[j_{1}+\Delta_{1}\right], Y\left[j_{1}+\Delta_{1}+1\right], \ldots, Y\left[i_{2}+\Delta_{2}-1\right]$ in some order. As it deletes the characters one by one, at some moment it arrives to a pair $X[x], Y[x+\Delta]$, where $i<x<q$. However, since $\mathcal{A}^{\prime \prime}$ is greedy, from this pair it must follow the segment $I$, a contradiction.

We now use this property to show that $|\mathrm{LZ}(X[i \ldots q])|=\mathcal{O}(k)$. First, consider the case when the segments corresponding to $\mathcal{A}^{\prime \prime}$ are below $I$. Let $I^{\prime}=\left[i^{\prime} \ldots j^{\prime} ; \Delta^{\prime}\right]$ be one such segment, where $i \leq i^{\prime} \leq j^{\prime} \leq q$, and $-k \leq \Delta^{\prime}<\Delta$. Let $i^{\prime \prime}=\max \left\{i+2 k, i^{\prime}\right\}$. If $i^{\prime \prime} \leq j^{\prime}$, we have $X\left[i^{\prime \prime} \ldots j^{\prime}\right]=Y\left[i^{\prime \prime}+\Delta^{\prime} \ldots j^{\prime}+\Delta^{\prime}\right]$, and the alignment $\mathcal{A}$ implies that $Y\left[i^{\prime \prime}+\Delta^{\prime} \ldots j^{\prime}+\Delta^{\prime}\right]=$ $X\left[i^{\prime \prime}+\left(\Delta^{\prime}-\Delta\right) \ldots j^{\prime}+\left(\Delta^{\prime}-\Delta\right)\right]$. Consequently, $X\left[i^{\prime \prime} \ldots j^{\prime}\right]=X\left[i^{\prime \prime}+\left(\Delta^{\prime}-\Delta\right) \ldots j^{\prime}+\left(\Delta^{\prime}-\Delta\right)\right]$ is a previous factor as $\Delta^{\prime}-\Delta<0$. It follows that $\mathcal{A}^{\prime \prime}$ defines a partitioning of $X[i \ldots q]=t_{1} \cdots t_{z}$, where $z=\mathcal{O}(k)$ and each factor $t_{i}$ is either a single character, or a previous factor. By Fact 4.1, we obtain that $|\operatorname{LZ}(X[i \ldots q])|=\mathcal{O}(k)$.

The proof for the complementary case is similar. There are at most $k$ horizontal segments
corresponding to $\mathcal{A}^{\prime}$ that contain positions in $[i . . q]$, denote them by $\left[i_{\ell}^{\prime} . . j_{\ell}^{\prime} ; \Delta_{\ell}^{\prime}\right]$. For all $\ell$, we have $k \geq \Delta_{\ell}^{\prime}>\Delta$. Let $j_{\ell}^{\prime \prime}=\min \left\{j_{\ell}^{\prime}, q-2 k\right\}$. If $i_{\ell}^{\prime} \leq j_{\ell}^{\prime \prime}$, we have $X\left[i_{\ell}^{\prime} \ldots j_{\ell}^{\prime \prime}\right]=Y\left[i_{\ell}^{\prime}+\Delta_{\ell}^{\prime} . . j_{\ell}^{\prime \prime}+\Delta_{\ell}^{\prime}\right]$ (from the alignment $\mathcal{A}^{\prime \prime}$ ), and $Y\left[i_{\ell}^{\prime \prime}+\Delta_{\ell}^{\prime} . . j_{\ell}^{\prime \prime}+\Delta_{\ell}^{\prime}\right]=X\left[i_{\ell}^{\prime}+\left(\Delta_{\ell}^{\prime}-\Delta\right) \ldots j_{\ell}^{\prime \prime}+\left(\Delta_{\ell}^{\prime}-\Delta\right)\right]$. Consequently, $X\left[i_{\ell}^{\prime}+\left(\Delta_{\ell}^{\prime}-\Delta\right) \ldots j_{\ell}^{\prime \prime}+\left(\Delta_{\ell}^{\prime}-\Delta\right)\right]=X\left[i_{\ell}^{\prime} . . j_{\ell}^{\prime \prime}\right]$ is a previous factor as $\Delta_{\ell}^{\prime}-\Delta>$ 0 . As the cost of $\mathcal{A}^{\prime \prime}$ is bounded from above by $k$, we have $\left|[i . . q]-\cup_{\ell}\left[i_{\ell}^{\prime} . . j_{\ell}^{\prime \prime}\right]\right| \leq k$ and $\sum_{\ell}\left|\Delta_{\ell}^{\prime}-\Delta_{\ell+1}^{\prime}\right| \leq k$. Therefore, $\left|[i \ldots q]-\cup_{\ell}\left[i_{\ell}^{\prime}+\left(\Delta_{\ell}^{\prime}-\Delta\right) \ldots j_{\ell}^{\prime \prime}+\left(\Delta_{\ell}^{\prime}-\Delta\right)\right)\right|=\mathcal{O}(k)$, and we can conclude as above.

The claim follows.
Corollary 5.7. For every $X, Y \in \Sigma^{\leq n}$ and $k \in \mathbb{Z}_{+}$, the encoding $\mathrm{GR}_{k}(X, Y)$ takes $\mathcal{O}\left(k^{2} \log ^{2} n\right)$ space.

Proof. Follows immediately from Propositions 4.2 and 5.2 and Lemma 3.10.

### 5.2 Construction of greedy encodings

In this section, we give a deterministic approach to construction of $\mathrm{GR}_{k}(X, Y)$. We start with a combinatorial lemma that is a key to this approach.

Lemma 3.11. For all $X, Y \in \Sigma^{*}$ and $k \in \mathbb{Z}_{+}$, the set $\mathcal{B}_{k}(X, Y)=\underset{\mathcal{A} \in \operatorname{GA}_{k}(X, Y)}{ } \mathcal{B}(\mathcal{A})$ is of size $\mathcal{O}\left(k^{5}\right)$.
Proof ${ }^{6}$. We prove $\mathcal{B}_{k}(X, Y) \leq 136 k^{5}$ by induction on $|X|$. In the base case of $|X| \leq 42 k^{4}$, we have $\left|\mathcal{B}_{k}(X, Y)\right| \leq(2 k+1)|X| \leq 136 k^{5}$ since each $(x, y) \in \mathcal{B}_{k}(X, Y)$ satisfies $|x-y| \leq k$. Thus, we assume that $|X|>42 k^{4}, \mathrm{GA}_{k}(X, Y) \neq \emptyset$ (otherwise, $\mathcal{B}_{k}(X, Y)=\emptyset$ ), and the claim holds for shorter strings.

Case 1. If $\mathcal{M}_{k}(X, Y)$ contains two adjacent pairs $(x-1, y-1),(x, y)$, then we construct strings $X^{*}$ and $Y^{*}$ by deleting $X[x]$ and $Y[y]$, respectively, aiming to prove that $\left|\mathcal{B}_{k}(X, Y)\right| \leq$ $\left|\mathcal{B}_{k}\left(X^{*}, Y^{*}\right)\right| \leq 136 k^{5}$. For this, consider an alignment $\mathcal{A} \in \mathrm{GA}_{k}(X, Y)$. Due to $(x-1, y-$ $1),(x, y) \in \mathcal{M}(\mathcal{A})$, the pairs $(x-1, y-1),(x, y)$, and $(x+1, y+1)$ are three consecutive elements of $\mathcal{A}$. Let $X^{*}$ be the string obtained by removing $X[x]$ from $X$, and $Y^{*}$ be the string obtained by removing $Y[y]$ from $Y$. We define an alignment $\mathcal{A}^{*}$ of $X^{*}, Y^{*}$ by removing $(x, y)$ and shifting all pairs to the right of $(x, y)$ by one position to the left. Each breakpoint $\left(x^{\prime}, y^{\prime}\right) \in \mathcal{B}(\mathcal{A})$ is preserved (along with $\left.X\left[x^{\prime}\right] \neq Y\left[y^{\prime}\right]\right)$, if it is located to the left of ( $x, y$ ), or shifted by 1 position to the left (along with $X\left[x^{\prime}\right] \neq Y\left[y^{\prime}\right]$ ), if it is located to the right of $(x, y)$. Consequently, $\mathcal{A}^{*} \in \mathrm{GA}_{k}\left(X^{*}, Y^{*}\right)$ and, since the shifts are independent of $\mathcal{A}$, we derive $\left|\mathcal{B}_{k}\left(X^{*}, Y^{*}\right)\right| \geq\left|\mathcal{B}_{k}(X, Y)\right|$.

Case 2: If $\mathcal{M}_{k}(X, Y)$ does not contain any two adjacent pairs, we first prove the following claim:
Claim 5.8. There are alignments $\mathcal{A}, \mathcal{A}^{\prime} \in \mathrm{GA}_{k}(X, Y)$ such that $X\left[a \ldots a+9 k^{2}\right) \simeq_{\mathcal{A}} Y[i . . i+$ $\left.9 k^{2}\right) \simeq_{\mathcal{A}^{\prime}} X\left[b . . b+9 k^{2}\right)$ holds for some positions $a \neq b$.

Proof. Consider $k+1$ disjoint fragments of $X$, each of length $21 k^{3} \leq|X| /(k+1)$. Each $\mathcal{A} \in \mathrm{GA}_{k}(X, Y)$ aligns at least one fragment without mismatches: $X[s \ldots t] \simeq_{\mathcal{A}} Y\left[s^{\prime} \ldots t^{\prime}\right]$ for some positions $s, t, s^{\prime}, t^{\prime}$. Due to $X[t]=Y\left[t^{\prime}\right]$, if $\left(t-1, t^{\prime}-1\right) \in \mathcal{M}_{k}(X, Y)$, then $\left(t, t^{\prime}\right) \in \mathcal{M}_{k}(X, Y)$ holds by the greedy nature of the considered alignments. Consequently, there is $\mathcal{A}^{\prime} \in \mathrm{GA}_{k}(X, Y)$ such that $X[t-1] \not{\not ㇒ \mathcal{A}^{\prime}} Y\left[t^{\prime}-1\right]$.

Next, consider $k+1$ disjoint parts of $Y[s+k . . t-k)$, each of length $9 k^{2} \leq\left(21 k^{3}-1-2 k\right) /(k+$ 1). By Fact $2.3, \mathcal{A}$ aligns the entire $Y[s+k \ldots t-k)$ without mismatches to a fragment of $X[s, . t)$.

[^5]Moreover, $\mathcal{A}^{\prime}$ aligns at least one part without mismatches. Denote this part $Y\left[i . . i+9 k^{2}\right)$, and fix positions $a, b$ so that $X\left[a \ldots a+9 k^{2}\right) \simeq_{\mathcal{A}} Y\left[i \ldots i+9 k^{2}\right) \simeq_{\mathcal{A}^{\prime}} X\left[b \ldots b+9 k^{2}\right)$. Note that $a \in[s . . t)$ and thus $X[a . . t) \simeq_{\mathcal{A}} Y\left[i . . t^{\prime}\right)$. Now, if $a=b$, then we would have $X[a] \simeq_{\mathcal{A}^{\prime}} Y[i]$ and the greedy nature of $\mathcal{A}^{\prime}$ would guarantee $X[a \ldots t) \simeq_{\mathcal{A}^{\prime}} Y\left[i \ldots t^{\prime}\right)$, contradicting $X[t-1] \not 千_{\mathcal{A}^{\prime}} Y\left[t^{\prime}-1\right]$. Consequently, $a \neq b$.

By Fact 2.3, we have $p:=\operatorname{per}\left(Y\left[i . . i+9 k^{2}\right)\right) \leq|a-b| \leq|i-a|+|i-b| \leq 2 k$. Moreover, due to $7 k^{2} \geq 2 p$, the fragment $X\left[i+k^{2} \ldots i+8 k^{2}\right.$ ) has the same period (up to cyclic rotation).
Claim 5.9. For every $\mathcal{A}^{\prime \prime} \in \mathrm{GA}_{k}(X, Y)$, we have $X\left[i+3 k^{2} \ldots i+8 k^{2}\right) \simeq \simeq_{\mathcal{A}^{\prime \prime}} Y\left[v . . v+5 k^{2}\right)$ for some $v \in\left[i+2 k^{2} \ldots i+4 k^{2}\right]$.

Proof. Consider $k+1$ disjoint parts of $X\left[i+k^{2} \ldots i+3 k^{2}+p\right.$ ), each of length $p$ (note that $\left.(k+1) p \leq 2 k^{2}+p\right)$. Each $\mathcal{A}^{\prime \prime} \in \mathrm{GA}_{k}(X, Y)$ matches at least one part without mismatches: $X[j \ldots j+p) \simeq_{\mathcal{A}^{\prime \prime}} Y[u \ldots u+p)$ for some $j \in\left[i+k^{2} . . i+3 k^{2}\right]$ and $u \in[j-k \ldots j+k]$. Note that $X\left[j \ldots i+8 k^{2}\right)$ and $Y\left[u \ldots i+9 k^{2}\right)$ have, up to cyclic rotation, the same string period of length $p$. This string period is primitive, so $X[j \ldots j+p)=Y[u \ldots u+p)$ implies that the periods synchronize, i.e., $X\left[j \ldots i+8 k^{2}\right)=Y\left[u \ldots u+\left(i+8 k^{2}-j\right)\right]$. By the greedy nature of $\mathcal{A}^{\prime \prime}$, we must have $X\left[j \ldots i+8 k^{2}\right) \simeq_{\mathcal{A}^{\prime \prime}} Y\left[u \ldots u+\left(i+8 k^{2}-j\right)\right]$, and therefore $X\left[i+3 k^{2} \ldots i+8 k^{2}\right) \simeq_{\mathcal{A}^{\prime \prime}}$ $Y\left[v . . v+5 k^{2}\right]$ for $v=u+i+3 k^{2}-j \in\left[i+3 k^{2}-k \ldots i+3 k^{2}+k\right]$.

Consider fragments $Q:=X\left[i+3 k^{2} . . i+8 k^{2}\right.$ ) and $F:=Y\left[i+4 k^{2} \ldots i+7 k^{2}\right)$, where $F$ is contained in all fragments $Y\left[v \ldots v+5 k^{2}\right)$ with $v \in\left[i+2 k^{2} . . i+4 k^{2}\right]$. We observe that $|Q| \geq|F|=3 k^{2}>p$ and construct $X^{*}$ and $Y^{*}$ by removing $X\left[i+8 k^{2}-p \ldots i+8 k^{2}\right.$ ) and $Y\left[i+7 k^{2}-p \ldots i+7 k^{2}\right.$ ). For every $\mathcal{A}^{\prime \prime} \in \mathrm{GA}_{k}(X, Y)$, we derive $\mathcal{A}^{*}$ by removing the pairs $(x, y)$ with $y \in\left[i+7 k^{2}-p \ldots i+7 k^{2}\right.$ ), and shifting all the subsequent pairs by $p$ positions to the left. We conclude from Claim 5.9 that the breakpoints $\left(x^{\prime}, y^{\prime}\right) \in \mathcal{B}(\mathcal{A})$ are preserved (along with $\left.X\left[x^{\prime}\right] \neq Y\left[y^{\prime}\right]\right)$, if located to the left of $\left(i+3 k^{2}, i+4 k^{2}\right)$, or shifted by $p$ positions to the left (along with $\left.X\left[x^{\prime}\right] \neq Y\left[y^{\prime}\right]\right)$, if located to the right of $\left(i+8 k^{2}, i+7 k^{2}\right)$. Consequently, $\mathcal{A}^{*} \in \mathrm{GA}_{k}\left(X^{*}, Y^{*}\right)$ and, since the shifts are independent of $\mathcal{A}^{\prime \prime}$, we derive $\left|\mathcal{B}_{k}(X, Y)\right| \leq\left|\mathcal{B}_{k}\left(X^{*}, Y^{*}\right)\right| \leq 136 k^{5}$ from the inductive assumption.

Algorithm 1 is the key procedure for computing greedy encodings.
Lemma 5.10. Given two strings $X, Y \in \Sigma^{*}$ and integer $k \geq \operatorname{ed}(X, Y)$, Algorithm 1 computes the dummy segments in $X^{\mathcal{M}_{k}(X, Y)}$. Moreover, it can be implemented in $\mathcal{O}\left(\left|\mathcal{B}_{k}(X, Y)\right|(t+\log k)\right)$ time and $\tilde{\mathcal{O}}(k)$ space, provided with an oracle that, given $x \in[1 \ldots|X|+1]$ and $y \in[1 . .|Y|+1]$, in $\mathcal{O}(t)$ time computes $\operatorname{LCE}(X[x .],. Y[y .]$.$) and \min (k+1, \operatorname{ed}(X[x .],. Y[y .])$.$) .$

Proof. To show correctness of Algorithm 1, consider an (imaginary set) of alignments $\mathcal{G}$ generated by the algorithm. First, associate with each point in $[1 . .|X|] \times[1 \ldots|Y|]$ an empty set of alignments. Let $(x, y)$ be the latest point extracted from $Q$ and $\mathcal{G}[(x, y)]$ the set of alignments associated with it. In line 7 (if the algorithm executes it), we create a new set of alignments containing all alignments in $\mathcal{G}[(x, y)]$ extended by $(x+1, y)$ (deletion of $X[x]$ ), $(x+1, y),(x+2, y+1), \ldots(x+\ell, y+\ell-1)$, where $\ell=\operatorname{LCE}(X[x+1 .],. Y[y .]$.$) (matching of$ $X[x+1 \ldots x+\ell]$ and $Y[y \ldots y+\ell-1])$. We union the resulting set of alignments with $\mathcal{G}[x+\ell, y+\ell-1]$ and push $(x+\ell, y+\ell-1)$ to $Q$. In lines 9 and 11 we process $\mathcal{G}[(x, y)]$ analogously. Finally, we define $\mathcal{G}=G[|X|,|Y|]$. We claim that $\mathcal{G}=\mathrm{GA}_{k}(X, Y)$. First note that all alignments in $\mathcal{G}$ are greedy by construction: an alignment can only delete $X[x], Y[y]$, or substitute $X[x]$ for $Y[y]$ if $X[x] \neq Y[y]$. In other words, for every breakpoint $(x, y)$ we have $X[x] \neq Y[y]$. Additionally, the cost of every alignment in $\mathcal{G}$ is at most $k$ as at each call of GreedyMatch credit decreases by one. Hence, $\mathcal{G}=\mathrm{GA}_{k}(X, Y)$.

```
Greedy ( \(k\) )
    \(Q=\emptyset ;\)
    GreedyMatch ( \(Q, 1,1, k\) );
    while \(Q\) is not empty do
        \((x, y)\), credit \(=Q\). extract-min ()\(;\)
        if \(x \leq|X|\) then
            GreedyMatch \((Q, x+1, y\), credit -1\()\); // delete \(X[x]\)
        if \(y \leq|Y|\) then
            GreedyMatch \((Q, x, y+1\), credit -1\() ; \quad\) // delete \(Y[y]\)
        if \(x \leq|X|\) and and \(y \leq|Y|\) then
                GreedyMatch \((Q, x+1, y+1\), credit -1\() ; \quad / /\) substitute \(X[x]\) for \(Y[y]\)
    GreedyMatch ( \(Q, x, y\), credit):
        if \(\operatorname{ed}(X[x .],. Y[y .]\).\() \leq credit then\)
            \(\ell=\operatorname{LCE}(X[x .],. Y[y .].) ; / /\) match \(X[x]\) and \(Y[y], \ldots, X[x+\ell-1]\) and \(Y[y+\ell-1]\)
            \(x=x+\ell ; y=y+\ell ;\)
            Push ( \(Q,(x, y)\), credit);
            if \(|Q|=1\) then output \([x-\ell, x-1]\);
    Push \((Q,(x, y)\), credit)
        if \((x, y) \notin Q\) then \(Q[(x, y)]=\) credit;
        else \(Q[(x, y)]=\max \{Q[(x, y)]\), credit \(\} ; \quad / /\) points are ordered lexicographically.
```

Algorithm 1: The algorithm receives as an input strings $X, Y$, and integer $k$. Lexicographic order on points is defined in the following way: $\left(x_{1}, y_{1}\right)<\left(x_{2}, y_{2}\right)$ if either $x_{1}<x_{2}$ or $x_{1}=x_{2}$ and $y_{1}<y_{2}$.

We define successors and predecessors in the lexicographic order in a standard way. For an alignment $\mathcal{A} \in \mathcal{G}$ and a point $(x, y) \in \mathcal{B}(\mathcal{A})$, define $\operatorname{credit}_{\mathcal{A}}(x, y)$ as the difference between $k$ and the number of edits that $\mathcal{A}$ makes to reach $(x, y)$. Algorithm 1 satisfies the following properties:

1. The set $\mathcal{P}$ of the points that are ever added to $Q$ is equal to $\mathcal{B}(\mathcal{G})$, and when we process a point $(x, y)$, the credit associated with it equals $\max _{\mathcal{A} \in \mathcal{G}:(x, y) \in \mathcal{B}(\mathcal{A})} \operatorname{credit}_{\mathcal{A}}(x, y)$;
2. When we process a point $(x, y), Q=\cup_{\mathcal{A} \in \mathcal{G}}\left\{\left(x^{\prime}, y^{\prime}\right)\right.$ is the successor of $(x, y)$ in $\left.\mathcal{B}(\mathcal{A})\right\}$.
3. If $Q$ contains points $p_{1}=\left(x_{1}, y_{1}\right)$ and $p_{2}=\left(x_{2}, y_{2}\right)$ such that $x_{1}-y_{1}=x_{2}-y_{2}$, then $p_{1}=p_{2}$.
To show the properties, we exploit the fact that the points are processed in the lexicographical order: When we process a point $q \in Q$, it is the lexicographically smallest point in $Q$, and all the points that we add to $Q$ while processing $q$ are larger than it.

Property 1. We first prove that $\mathcal{P} \supseteq \mathcal{B}(\mathcal{G})$ and that when we process a point $(x, y)$, the credit associated with it is at least $\max _{\mathcal{A} \in \mathcal{G}:(x, y) \in \mathcal{B}(\mathcal{G})} \operatorname{credit}_{\mathcal{A}}(x, y)$. Let $(x, y) \in \mathcal{B}(\mathcal{G})$ be the lexicographically smallest point which either does not belong to $\mathcal{P}$ or such that the credit associated with it is smaller than $\max _{\mathcal{A} \in \mathcal{G}:(x, y) \in \mathcal{B}(\mathcal{A})} \operatorname{credit}_{\mathcal{A}}(x, y)$. Consider an alignment $\mathcal{A} \in \mathcal{G}$ such that $(x, y) \in \mathcal{B}(\mathcal{A})$ and let $\left(x^{\prime}, y^{\prime}\right)$ be the predecessor of $(x, y)$ in $\mathcal{B}(\mathcal{A})$. When we process $\left(x^{\prime}, y^{\prime}\right)$, which happens before we process $(x, y)$ as we process the points in the lexicographic order, the credit associated with it is at least $\operatorname{credit}_{\mathcal{A}}\left(x^{\prime}, y^{\prime}\right) \geq \operatorname{ed}\left(X\left[x^{\prime} ..\right], Y\left[y^{\prime} ..\right]\right)$, and therefore $(x, y)$ is added to $Q$ with credit at least $\operatorname{credit}_{\mathcal{A}}\left(x^{\prime}, y^{\prime}\right)-1=\operatorname{credit}_{\mathcal{A}}(x, y)$. As it holds for any $\mathcal{A} \in \mathcal{G}$, we obtain a contradiction.

We now show the opposite direction. Consider a point $(x, y) \in Q$ associated with the edit distance credit $c$. Suppose that $c$ is achieved when we process a point $\left(x^{\prime}, y^{\prime}\right)$. By the induction assumption, there is a greedy alignment $\mathcal{A} \in \mathcal{G}$ such that the credit $c+1$ associated with
$\left(x^{\prime}, y^{\prime}\right)$ equals $\operatorname{credit}_{\mathcal{A}}\left(x^{\prime}, y^{\prime}\right)$. Consider a greedy alignment $\mathcal{A}^{\prime}$ that behaves as $\mathcal{A}$ until ( $x^{\prime}, y^{\prime}$ ) ( $k-\operatorname{credit}_{\mathcal{A}}\left(x^{\prime}, y^{\prime}\right)$ edits), then proceeds to $\left(x^{\prime}, y^{\prime}\right)$ (one edit), and ends as an arbitrary optimal greedy alignment between $X\left[x^{\prime} ..\right]$ and $Y\left[y^{\prime} ..\right]$. The cost of the latter is ed $\left(X\left[x^{\prime} ..\right], Y\left[y^{\prime} ..\right]\right) \leq c$. Therefore, the cost of $\mathcal{A}^{\prime}$ is at most $k$ and it is in $\mathcal{G}$. It follows that $\left(x^{\prime}, y^{\prime}\right) \in \mathcal{B}\left(\mathcal{A}^{\prime}\right)$ and that $c=\operatorname{credit}_{\mathcal{A}^{\prime}}\left(x^{\prime}, y^{\prime}\right)$, finishing the proof.

Property 2. Let $(x, y)$ be the lexicographically smallest point in $\mathcal{B}(\mathcal{G})$ such that when we process it $Q \ni(u, v) \notin \cup_{\mathcal{A} \in \mathcal{G}}\left\{\left(x^{\prime}, y^{\prime}\right)\right.$ is the successor of $(x, y)$ in $\left.\mathcal{B}(\mathcal{A})\right\}$. Suppose that $(u, v)$ was added to $Q$ when we processed a point ( $u^{\prime}, v^{\prime}$ ). Similarly to above, we can construct an alignment $\mathcal{A} \in \mathcal{G}$ such that $\left(u^{\prime}, v^{\prime}\right),(u, v) \in \mathcal{B}(\mathcal{A})$. As we processed $\left(u^{\prime}, v^{\prime}\right)$ before $(x, y)$, we must have that $\left(u^{\prime}, v^{\prime}\right)$ is lexicographically smaller than $(x, y)$, while $(u, v)$ is lexicographically larger than $(x, y)$, and therefore $(u, v)$ is the successor of $(x, y)$ in $\mathcal{B}(\mathcal{A})$, a contradiction. On the other hand, if $(u, v)$ is the successor of $(x, y)$ in $\mathcal{B}(\mathcal{A})$ for some $\mathcal{A} \in \mathcal{G}$, then it must be in $Q$ when we process $(x, y)$ : if $\left(u^{\prime}, v^{\prime}\right)$ is the predecessor of $(u, v)$ in $\mathcal{B}(\mathcal{A})$, then $\left(u^{\prime}, v^{\prime}\right)$ is at most $(x, y)$, and when we process it, we add $(u, v)$ to $Q$.

Property 3. We show the property by induction. It is true at initialisation. Suppose that the property is violated when we process a point $(x, y) \in Q$. Recall that we extract $(x, y)$ and call GreedyMatch for a subset of points $\{(x+1, y+1),(x+1, y),(x, y+1)\}$. The point $(x+1, y+1)$ is in the same diagonal as $(x, y)$ that we extract from $Q$, and therefore the call for it does not violate the property. Consider now the point $(x+1, y)$ (resp., $(x, y+1)$ ) and assume that the diagonal containing it also contains a point $\left(x^{\prime}, y^{\prime}\right) \in Q$. If we added ( $x^{\prime}, y^{\prime}$ ) to $Q$ when we called GreedyMatch for a point $\left(x^{\prime \prime}, y^{\prime \prime}\right)$, we have that $\left(x^{\prime \prime}, y^{\prime \prime}\right)$ is in the same diagonal, $X\left[x^{\prime \prime} . . x^{\prime}-1\right]=Y\left[y^{\prime \prime} . . y^{\prime}-1\right]$, and $X\left[x^{\prime}\right] \neq Y\left[y^{\prime}\right]$. As we process the points in the lexicographic order, at least one of the points $\left(x^{\prime \prime}-1, y^{\prime \prime}-1\right),\left(x^{\prime \prime}-1, y^{\prime \prime}\right),\left(x^{\prime \prime}, y^{\prime \prime}-1\right)$ is strictly smaller than $(x, y)$. By considering each of the cases, from the definition we obtain that ( $x^{\prime \prime}, y^{\prime \prime}$ ) is at most $(x+1, y)$ (resp., $(x, y+1)$ ). As all three points $\left(x^{\prime \prime}, y^{\prime \prime}\right),(x+1, y)$ (resp., $(x, y+1)$ ), $\left(x^{\prime}, y^{\prime}\right)$ are on the same diagonal, it follows that GreedyMatch for $(x+1, y)$ outputs $\left(x^{\prime}, y^{\prime}\right)$.

Correctness and complexity. $\left[x, x^{\prime}\right]$ is a dummy segment in $X^{\mathcal{M}_{k}(X, Y)}$ iff there exists $-k \leq$ $\delta \leq k$ such that the following three conditions are satisfied, where $y=x+\delta$ and $y^{\prime}=x^{\prime}+\delta$ :

1. $\mathcal{B}(\mathcal{G})$ does not contain any points with the $x$-coordinate in $\left[x \ldots x^{\prime}\right]$ or the $y$-coordinate in $\left[y . . y^{\prime}\right]$;
2. $\mathcal{B}(\mathcal{G})$ contains $q=\left(x^{\prime}+1, y^{\prime}+1\right)$;
3. The predecessor $p$ of $q$ in $\mathcal{B}(\mathcal{G})$ is one of the points $(x-1, y-1),(x-1, y),(x, y-1)$.

Therefore, if $\left[x, x^{\prime}\right]$ is a maximal dummy segment, then by Property 2 , after we have processed $p, Q$ contains a single point $q$ that is obtained by calling GreedyMatch for the point $(x, y)$. Hence, the implementation creates the dummy segment $\left[x, x^{\prime}\right]$. On the other hand, if the implementation creates a segment $\left[x, x^{\prime}\right]$, then $\mathcal{B}(\mathcal{G})$ satisfies all three conditions and hence $\left[x, x^{\prime}\right]$ is a maximal dummy segment.

Algorithm 1 implements the set $Q$ as a binary search tree. By Lemma 3.11 and Property 1 , we process $|\mathcal{B}(\mathcal{G})|=\left|\mathcal{B}_{k}(X, Y)\right|$ points, spending $\mathcal{O}(\log k+t)$ time per point, which gives the desired time complexity. To show the space complexity, it suffices to show that at any moment $|Q|=\mathcal{O}(k)$. By Fact 2.3, for any $(x, y) \in Q$ we have $y=x+\delta$ for some $-k \leq \delta \leq k$, in other words, the point belongs to one of $2 k+1$ diagonals of the grid. By Property 3, each of the diagonals contains at most one point from $Q$, which concludes the proof.

We proceed with an auxiliary claim that allows constructing an edit distance oracle.
Claim 5.11. Given two strings $U, V$, and a data structure of size $s$ that can answer the LCE queries on the suffixes of $\bar{U}, \bar{V}$ in $t$ time. We can build in $\mathcal{O}\left(k^{2} t\right)$ time a data structure that occupies $\mathcal{O}\left(k^{2}\right)$ space and can retrieve $\min (k+1, \operatorname{ed}(U[u .],. V[v .])$.$) in \mathcal{O}(\log n)$ time.

Proof. The data structure is based on the Landau-Myers-Schmidt algorithm [41]. Consider a table $D$ of size $|U| \times|V|$ such that $D[i, j]=\operatorname{ed}(U[|U|-i+1 .],. V[|V|-j+1 .])=$. ed $(\bar{U}[1 \ldots i], \bar{V}[1 \ldots j])$. By Fact 2.3, if $D[i, i+\delta] \leq h$, then $-h \leq \delta \leq h$. Let $L^{h}(\delta)=$ $\max \{i: D[i, i+\delta] \leq h\}$. The data structure consists of $2 k+1$ sorted arrays containing $L^{0}(\delta), L^{1}(\delta), \ldots, L^{k}(\delta),-k \leq \delta \leq k$. The arrays occupy $\mathcal{O}\left(k^{2}\right)$ space and allow to retrieve $\operatorname{ed}(U[u .],. V[v .]$.$) in \tilde{\mathcal{O}}(1)$ time by simple binary search. The data structure can be computed via the following recurrence, where $\operatorname{SLIDE}_{\delta}(u)=u+\operatorname{LCE}(\bar{U}[u .],. \bar{V}[u+\delta .]$.$) .$

$$
L^{h}(\delta)=\operatorname{SLIDE}_{\delta}\left(\max \left\{\begin{array}{ll}
L^{h-1}(\delta-1), & \delta>-h ;  \tag{1}\\
L^{h-1}(\delta)+1, & \text { always; } \\
L^{h-1}(\delta+1)+1, & \delta<h .
\end{array}\right)\right.
$$

Therefore, given $L^{h-1}, L^{h}(\delta)$ can be computed via three LCE queries on the suffixes of $\bar{U}, \bar{V}$. As a corollary, the arrays $\left\{L^{0}(\delta), L^{1}(\delta), \ldots, L^{k}(\delta)\right\},-k \leq \delta \leq k$, can be built in $\tilde{\mathcal{O}}\left(k^{2}\right)$ time.

Proposition 5.12. Consider a non-crossing matching $M$ of strings $X, Y \in \Sigma^{\leq n}$ and an integer $k \in \mathbb{Z}_{+}$such that $M \subseteq \mathcal{M}_{k}(X, Y)$ holds if $\operatorname{ed}(X, Y) \leq k$. Given $k$ and $\mathrm{E}^{M}(X, Y)$, the greedy encoding $\mathrm{GR}_{k}(X, Y)$ can be computed in $\tilde{\mathcal{O}}\left(z k+k^{2}+\left|\mathcal{B}_{k}(X, Y)\right|\right)$ time and $\tilde{\mathcal{O}}\left(z+k^{2}\right)$ space, where $z=\left|\operatorname{LZ}\left(X^{M} Y^{M}\right)\right|$.

Proof. First, we pass $k$ and $\mathrm{E}^{M}(X, Y)$ to the procedure of Corollary 5.5. Observe that the returned value $d$ satisfies $d \leq k$ if and only if ed $(X, Y) \leq k$. If ed $(X, Y)>k$, then we simply return $\mathrm{GR}_{k}(X, Y)=\perp$. In this case, the running time and the space complexity are $\tilde{\mathcal{O}}\left(k^{2}\right)$.

Otherwise, our strategy is to compute $\mathcal{M}_{k}(X, Y)$ and then mask out the corresponding characters of $X^{M}$ and $Y^{M}$ to obtain $X^{\mathcal{M}_{k}(X, Y)}$ and $Y^{\mathcal{M}_{k}(X, Y)}$. By Lemma 3.9, we have $\mathcal{M}_{k}(X, Y)=\mathcal{M}_{k}\left(X^{\prime}, Y^{\prime}\right)$, where $X^{\prime}=\operatorname{num}\left(X^{M}\right)$ and $Y^{\prime}=\operatorname{num}\left(Y^{M}\right)$. Hence, we shall use Lemma 3.11 to compute the dummy segments corresponding to $\mathcal{M}_{k}\left(X^{\prime}, Y^{\prime}\right)$. For this, we need to support LCE queries and edit distance queries on the suffixes of $X^{\prime}, Y^{\prime}$. As for LCE queries, we rely on Lemma 5.4, which provides $\mathcal{O}(\log n)$-time LCE queries on the suffixes of $X^{\prime}, Y^{\prime}$ and on the suffixes of $\overline{X^{\prime}}, \overline{Y^{\prime}}$. We use the latter queries to build a component of Claim 5.11 for the edit distance queries. After $\mathcal{O}\left(k^{2} \log n\right)$-time preprocessing, these queries can be answered in $\mathcal{O}(\log k)$ time. Overall, constructing the dummy segments corresponding to $\mathcal{M}_{k}\left(X^{\prime}, Y^{\prime}\right)$ costs $\tilde{\mathcal{O}}\left(k^{2}+\left|\mathcal{B}_{k}(X, Y)\right|\right)$ time and $\mathcal{O}\left(k^{2}\right)$ working space.

Our next goal is to convert $\mathrm{D}\left(X^{M} Y^{M}\right)$ to $\mathrm{D}\left(X^{\mathcal{M}_{k}(X, Y)} Y^{\mathcal{M}_{k}(X, Y)}\right)$. For this, we need to place \#s within each of the computed dummy segments. Consider updating a working string $Z$ by setting $Z[\ell . . r):=\#^{r-\ell}$. To implement this operation, we build $L Z\left(\#^{r-\ell}\right)$ (of size at most 2), derive $\mathrm{D}\left(\#^{r-\ell}\right)$ (Proposition $4.2(\mathrm{~h})$ ), extract $\mathrm{D}(Z[1 \ldots \ell)$ ) and $\mathrm{D}(Z[r \ldots|Z|])$ (Proposition $4.2(\mathrm{~g})$ ), and finally concatenate into $\mathrm{D}\left(Z[1 . . \ell) \#^{r-\ell} Z[r . .|Z|]\right)$ (Proposition 4.2(i)). Overall, each dummy segment is processed in $\mathcal{O}\left((z+k) \log ^{4} n\right)$ time and space, for a total of $\tilde{\mathcal{O}}\left(z k+k^{2}\right)$ time and $\tilde{\mathcal{O}}(z+k)$ space across all segments.

Finally, $\mathrm{RS}_{\#}\left(X^{\mathcal{M}_{k}(X, Y)}\right)$ and $\mathrm{RS}_{\#}\left(Y^{\mathcal{M}_{k}(X, Y)}\right)$ are constructed in $\mathcal{O}(k)$ time and space using Proposition 5.2.

Corollary 5.13. Given an integer $k \in \mathbb{Z}_{+}$and two strings $X, Y \in \Sigma \leq k^{2}$, one can compute $\mathrm{GR}_{k}(X, Y)$ in $\tilde{\mathcal{O}}\left(k^{3}\right)$ time and $\tilde{\mathcal{O}}\left(k^{2}\right)$ space.

Proof. We shall construct $\mathrm{E}^{\emptyset}(X, Y)$ and then derive $\mathrm{GR}_{k}(X, Y)$ from Proposition 5.12. Hence, we build a trivial LZ-like representation of $X Y$ (with length-1 phrases) and derive $\mathrm{D}(X Y)=$ $\mathrm{D}\left(X^{\emptyset} Y^{\emptyset}\right)$ using Proposition $4.2(\mathrm{~h})$. This costs $\mathcal{O}\left(k^{2} \log ^{2} k\right)$ time and space. The components $\mathrm{RS}_{\#}\left(X^{\emptyset}\right)$ and $\mathrm{RS}_{\#}\left(Y^{\emptyset}\right)$ are trivial (there are no dummy segments). Using Proposition 5.12 costs $\tilde{\mathcal{O}}\left(k^{3}+\left|\mathcal{B}_{k}(X, Y)\right|\right)$ time and $\tilde{\mathcal{O}}\left(k^{2}\right)$ space. To finish the proof, note that $\mathcal{B}_{k}(X, Y)$ can
only contain pairs $(x, y)$ such that $1 \leq x \leq|X|, 1 \leq y \leq|Y|$, and $|x-y| \leq k$, which implies $\left|\mathcal{B}_{k}(X, Y)\right|=\mathcal{O}(|X| k)=\mathcal{O}\left(k^{3}\right)$.

### 5.3 Concatenations and Quasi-greedy Alignments

Definition 5.14 (Quasi-greedy alignment). We say that an alignment $\mathcal{A}$ of two strings $X, Y \in$ $\Sigma^{*}$ is quasi-greedy if it satisfies at least one of the following symmetric conditions:

- there exists $\delta \in[1 \ldots|X|+1]$ such that $(\delta, 1) \in \mathcal{A}$ and $X[x] \neq Y[y]$ holds for every $(x, y) \in$ $\mathcal{B}(\mathcal{A}) \cap([\delta \ldots|X|] \times[1 \ldots|Y|]) ;$
- there exists $\delta \in[1 \ldots|Y|+1]$ such that $(1, \delta) \in \mathcal{A}$ and $X[x] \neq Y[y]$ holds for every $(x, y) \in$ $\mathcal{B}(\mathcal{A}) \cap([1 \ldots|X|] \times[\delta \ldots|Y|]) ;$
Given $k \geq 0$, we denote by $\mathrm{qGA}_{k}(X, Y)$ the set of all quasi-greedy alignments $\mathcal{A}$ of $X, Y$ with $\operatorname{cost}(\mathcal{A}) \leq k$.

Lemma 5.15. Consider an alignment $\mathcal{A} \in \mathrm{qGA}_{k}\left(X_{p} X_{s}, Y_{p} Y_{s}\right)$, where $X_{p}, X_{s}, Y_{p}, Y_{s} \in \Sigma^{*}$ and $k \in \mathbb{Z}_{\geq 0}$. Then, $\mathcal{A}_{\left[1 .\left|\left|X_{p}\right|+1\right),\left[1 .\left|Y_{p}\right|+1\right)\right.} \in \mathrm{qGA}_{k+\left|\left|X_{s}\right|-\left|Y_{s}\right|\right|}\left(X_{p}, Y_{p}\right)$ and $\mathcal{A}_{\left[\left|X_{p}\right|+1 .\right],\left[\left|Y_{p}\right|+1 .\right]} \in$ $\mathrm{qGA}_{k+\left|\left|X_{p}\right|-\left|Y_{p}\right|\right|}\left(X_{s}, Y_{s}\right)$.

Proof. For brevity, let $X=X_{p} X_{s}, Y=Y_{p} Y_{s} ; n_{p}=\left|X_{p}\right|, n_{s}=\left|X_{s}\right|, n=|X|, m_{p}=\left|Y_{p}\right|$, $m_{s}=\left|Y_{s}\right|, m=|Y|$.

Let $\mathcal{A}=\left(x_{t}, y_{t}\right)_{t=1}^{q}$ and let $\left(x_{i_{p}}, y_{i_{p}}\right)$ be the leftmost pair $(x, y) \in \mathcal{A}$ such that $x>n_{p}$ or $y>$ $m_{p}$. By symmetry, we assume without loss of generality that $x_{i_{p}}=n_{p}+1$; see Fig. 2. Observe that $\mathcal{A}_{p}:=\mathcal{A}_{\left[1 . . n_{p}+1\right),\left(1 . . m_{p}+1\right)}$ is the union of $\left(x_{t}, y_{t}\right)_{t=1}^{i_{p}-1}$ and $\left(n_{p}+1, y\right)_{y=y_{i_{p}}}^{m_{p}+1}$. Furthermore, $\mathcal{B}\left(\mathcal{A}_{p}\right)=\left\{\left(x_{t}, y_{t}\right) \in \mathcal{B}(\mathcal{A}): t \in[1 \ldots p)\right\} \cup\left\{\left(n_{p}+1, y\right): y \in\left[y_{p} \ldots m_{p}+1\right]\right\}$. Let $k_{p}$ be the size of the former component. Then, $\left.k \geq \operatorname{cost}(\mathcal{A}) \geq k_{p}+\operatorname{ed}\left(X_{s}, Y_{p}\left[y_{i_{p}}.\right]\right] Y_{s}\right) \geq k_{p}+\left|Y_{p}\left[y_{i_{p}} ..\right]\right|+m_{s}-n_{s}$, so that $\operatorname{cost}\left(\mathcal{A}_{p}\right)=k_{p}+\left|Y_{p}\left[y_{p} ..\right]\right| \leq k+n_{s}-m_{s} \leq k-\left|\left|X_{s}\right|-\left|Y_{s}\right|\right|$. It remains to prove that $\mathcal{A}_{p}$ is a quasi-greedy alignment. By Definition 3.6 applied to $\mathcal{A}$, we have two possibilities:

- There exists $\delta \in[1 \ldots n+1]$ such that $(\delta, 1) \in \mathcal{A}$ and $X[x] \neq Y[y]$ holds for every $(x, y) \in$ $\mathcal{B}(\mathcal{A}) \cap([\delta \ldots n] \times[1 \ldots m])$. If $\delta \leq n_{p}+1$, then we also have that $(\delta, 1) \in \mathcal{A}_{p}$ and $X_{p}[x] \neq Y_{p}[y]$ holds for every $(x, y) \in \mathcal{B}\left(\mathcal{A}_{p}\right) \cap\left(\left[\delta \ldots n_{p}\right] \times\left[1 \ldots m_{p}\right]\right)$. Otherwise, $\left(x_{i_{p}}, y_{i_{p}}\right)=\left(n_{p}+1,1\right) \in$ $\mathcal{A}_{p}$, so $\mathcal{A}_{p}$ is trivially quasi-greedy.
- There exists $\delta \in[1 \ldots m+1]$ such that $(1, \delta) \in \mathcal{A}$ and $X[x] \neq Y[y]$ holds for every $(x, y) \in \mathcal{B}(\mathcal{A}) \cap([1 \ldots n] \times[\delta \ldots m])$. If $\delta \leq y_{i_{p}}$, then we also have that $(1, \delta) \in \mathcal{A}_{p}$ and $X_{p}[x] \neq Y_{p}[y]$ holds for every $(x, y) \in \mathcal{B}\left(\mathcal{A}_{p}\right) \cap\left(\left[1 \ldots n_{p}\right] \times\left[\delta \ldots m_{p}\right]\right)$. Otherwise, $n_{p}=0$ because $(1, \delta)$ cannot cross $\left(n_{p}+1, y_{p}\right)$. Hence, $\mathcal{A}_{p}$ is trivially quasi-greedy.
Now, let $\left(x_{i_{s}}, y_{i_{s}}\right)$ be the leftmost pair $(x, y) \in \mathcal{A}$ such that $x>n_{p}$ and $y>m_{p}$. By symmetry, we assume without loss of generality that $y_{i_{s}}=m_{p}+1$; see Fig. 2. Observe that $\mathcal{A}_{s}:=\mathcal{A}_{\left[n_{p}+1 .\right],\left[m_{p}+1 .\right]}$ is the union of $(x, 1)_{x=1}^{x_{i_{s}}-n_{p}-1}$ and $\left(x_{t}-n_{p}, y_{t}-m_{p}\right)_{t=i_{s}}^{q}$. Furthermore, $\mathcal{B}\left(\mathcal{A}_{s}\right)=\left\{\left(x_{t}-n_{s}, y_{t}-m_{s}\right) \in \mathcal{B}(\mathcal{A}): t \in\left[i_{s} \ldots q\right]\right\} \cup\left\{(x, 1): x \in\left[1 \ldots x_{i_{s}}-n_{p}\right)\right\}$. Let $k_{s}$ be the size of the former component. Then, $k \geq \operatorname{cost}(\mathcal{A}) \geq k_{s}-1+\operatorname{ed}\left(X_{p} X_{s}\left[1 \ldots x_{i_{s}}-n_{p}\right), Y_{p}\right) \geq k_{s}-1+x_{i_{s}}-m_{p}$, so that $\operatorname{cost}\left(\mathcal{A}_{s}\right)=k_{s}-1+x_{i_{s}}-n_{p} \leq k+m_{p}-n_{p} \leq k-\left|\left|X_{p}\right|-\left|Y_{p}\right|\right|$. It remains to prove that $\mathcal{A}_{s}$ is a quasi-greedy alignment. By Definition $3.6 \operatorname{applied}$ to $\mathcal{A}$, we have two possibilities:
- There exists $\delta \in[1 \ldots n+1]$ such that $(\delta, 1) \in \mathcal{A}$ and $X[x] \neq Y[y]$ holds for every $(x, y) \in$ $\mathcal{B}(\mathcal{A}) \cap([\delta \ldots n] \times[1 \ldots m])$. If $\delta \leq x_{i_{s}}$, then $\left(x_{i_{s}}-n_{p}, 1\right) \in \mathcal{A}_{s}$ and $X_{s}[x] \neq Y_{s}[y]$ holds for every $(x, y) \in \mathcal{B}\left(\mathcal{A}_{s}\right) \cap\left(\left[x_{i_{s}}-n_{p} . . n_{s}\right] \times\left[1 \ldots m_{s}\right]\right)$. Otherwise, $\left(\delta-n_{p}, 1\right) \in \mathcal{A}_{s}$ and $X_{s}[x] \neq Y_{s}[y]$ holds for every $(x, y) \in \mathcal{B}\left(\mathcal{A}_{s}\right) \cap\left(\left[\delta-n_{p} . . n_{s}\right] \times\left[1 \ldots m_{s}\right]\right)$.
- There exists $\delta \in[1 \ldots m+1]$ such that $(1, \delta) \in \mathcal{A}$ and $X[x] \neq Y[y]$ holds for every $(x, y) \in \mathcal{B}(\mathcal{A}) \cap([1 \ldots n] \times[\delta \ldots m])$. If $\delta \leq m_{p}+1$, then $\left(x_{i_{s}}-n_{p}, 1\right) \in \mathcal{A}_{s}$ and $X_{s}[x] \neq Y_{s}[y]$ holds for every $(x, y) \in \mathcal{B}\left(\mathcal{A}_{s}\right) \cap\left(\left[x_{i_{s}}-n_{p} \ldots n_{s}\right] \times\left[1 \ldots m_{s}\right]\right)$. Otherwise, $x_{i_{s}}=1$ and $n_{p}=0$
because $(1, \delta)$ cannot $\operatorname{cross}\left(x_{i_{s}}, m_{p}+1\right)$. Hence, $\left(1, \delta-m_{p}\right) \in \mathcal{A}$ and $X_{s}[x] \neq Y_{s}[y]$ holds for every $(x, y) \in \mathcal{B}\left(\mathcal{A}_{s}\right) \cap\left(\left[1 \ldots n_{s}\right] \times\left[\delta-m_{p} \ldots m_{s}\right]\right)$.


Figure 2: Alignment $\mathcal{A} \in \mathrm{qGA}_{k}\left(X_{p} X_{s}, Y_{p} Y_{s}\right)$.
For strings $X, Y \in \Sigma^{*}$ and an integer $k \geq \operatorname{ed}(X, Y)$, we define a set

$$
\mathrm{q} \mathcal{M}_{k}(X, Y)=\bigcap_{\mathcal{A} \in \mathrm{qGA}_{k}(X, Y)} \mathcal{M}_{X, Y}(\mathcal{A})
$$

of common matches of all alignments $\mathcal{A} \in \mathrm{qGA}_{k}(X, Y)$.
Definition 5.16. For strings $X, Y \in \Sigma^{*}$ and an integer $k$, we define the quasi-greedy encoding

$$
\mathrm{qGR}_{k}(X, Y)= \begin{cases}\mathrm{E}^{\mathrm{q} \mathcal{M}_{k}(X, Y)}(X, Y) & \text { if } \operatorname{ed}(X, Y) \leq k \\ \perp & \text { otherwise }\end{cases}
$$

Corollary 5.17. Consider strings $X, Y \in \Sigma^{\leq n}$. Define $X^{\prime}=\$_{1} X$ and $Y^{\prime}=\$_{2} Y$, where $\$_{1} \neq \$_{2}$ are special symbols not in $\Sigma$. Let $M^{\prime}=\mathcal{M}_{k+1}\left(X^{\prime}, Y^{\prime}\right)$ and $M=\mathrm{q} \mathcal{M}_{k}(X, Y)$. The following properties hold:
(a) Given the dummy segments for one of the strings $X^{M},\left(X^{\prime}\right)^{M^{\prime}}$, the dummy segments for the other string can be constructed in $\mathcal{O}(k)$ time and space;
(b) Given one of the encodings $\mathrm{qGR}_{k}(X, Y), \mathrm{GR}_{k+1}\left(X^{\prime}, Y^{\prime}\right)$, the other encoding can be constructed in $\mathcal{O}\left(k^{2} \log ^{4} n\right)$ time and space;
(c) The quasi-greedy encoding $\mathrm{qGR}_{k}(X, Y)$ occupies $\mathcal{O}\left(k^{2} \log ^{4} n\right)$ space.

Proof. Let us first show that $\mathcal{A}^{\prime} \mapsto \mathcal{A}_{[2 .],[2 . .]}^{\prime}$ bijectively maps $\mathrm{GA}_{k+1}\left(X^{\prime}, Y^{\prime}\right)$ to $\mathrm{qGA}_{k}(X, Y)$.
Consider a greedy alignment $\mathcal{A}^{\prime} \in \mathrm{GA}_{k+1}\left(X^{\prime}, Y^{\prime}\right)$ and let $\mathcal{A}=\mathcal{A}_{[2.1],[2 .]}^{\prime}$. Lemma 5.15 yields $\mathcal{A} \in \mathrm{qGA}_{k+1}(X, Y)$. However, we actually have $\operatorname{cost}(\mathcal{A}) \leq k$ because $\mathcal{M}\left(\mathcal{A}^{\prime}\right)=\{(x, y):$ $(x+1, y+1) \in \mathcal{M}(\mathcal{A})\}$ holds due to $X^{\prime}[1] \neq Y^{\prime}[1]$.

We now show the opposite direction. Let $\mathcal{A} \in \mathrm{qGA}_{k}(X, Y)$. Suppose that there is $\delta \in$ $[1 \ldots|Y|]$ such that $(1, \delta) \in \mathcal{A}$ and $X[x] \neq Y[y]$ holds for every $(x, y) \in \mathcal{B}(\mathcal{A}) \cap([1 \ldots|X|] \times$ $[\delta \ldots|Y|])$ (the other case is symmetrical). It follows in particular that $\mathcal{A}$ contains elements $\{(1,1),(1,2), \ldots,(1, \delta)\}$. We start by replacing each element $(x, y)$ of $\mathcal{A}$ with $(x+1, y+1)$. We add an element $(1,1)$ to $\mathcal{A}$ and replace each element $(2, i), 2 \leq i \leq \delta+1$, with $(1, i)$. Finally, we add an element $(2, \delta+1)$. We claim that the resulting alignment $\mathcal{A}^{\prime}$, that we treat as an alignment of $X^{\prime}$ and $Y^{\prime}$, is greedy. For every element $(x, y) \in \mathcal{B}\left(\mathcal{A}^{\prime}\right) \cap\left(\left[2 \ldots\left|X^{\prime}\right|\right] \times\left[\delta+1 \ldots\left|Y^{\prime}\right|\right]\right)$ we have $(x-1, y-1) \in \mathcal{B}(\mathcal{A}) \cap([1 \ldots|X|] \times[\delta \ldots|Y|])$ and hence from quasi-greediness of $\mathcal{A}$ we obtain $X^{\prime}[x]=X[x-1] \neq Y[y-1]=Y^{\prime}[y]$. For every element $(1, i) \in \mathcal{A}^{\prime}$ we have $X^{\prime}[1]=\$_{1} \neq Y[i]$. It follows that $\mathcal{A}^{\prime}$ is greedy. In addition, the cost of $\mathcal{A}^{\prime}$ equals the cost of $\mathcal{A}$ plus one and hence is bounded by $k+1$. Finally, note that $\mathcal{M}\left(\mathcal{A}^{\prime}\right)=\{(x, y):(x-1, y-1) \in \mathcal{M}(\mathcal{A})\}$.

Let $M^{\prime}=\mathcal{M}_{k+1}\left(X^{\prime}, Y^{\prime}\right)$ and $M=\mathrm{q} \mathcal{M}_{k}(X, Y)$. From above, we obtain $M^{\prime}=\{(x+1, y+1)$ : $(x, y) \in M\}$. Therefore, $\left(X^{\prime}\right)^{M^{\prime}}[2 .]=.X^{M}$ and $\left(Y^{\prime}\right)^{M^{\prime}}[2 .]=.Y^{M}$.

It follows that we can obtain the dummy segments for $X^{M}$ by subtracting one from the endpoints of each dummy segment of $\left(X^{\prime}\right)^{M^{\prime}}$, and analogously for $Y^{M}$. The reverse claim is obtained analogously, by adding one to the endpoints. As the number of the dummy segments in $X^{M}$ and $\left(X^{\prime}\right)^{M^{\prime}}$ is $\mathcal{O}(k)$ by Lemma 3.10, (a) follows.

From (a) it follows that given the rank data structures for one pair of strings $X^{M}, Y^{M}$ and $\left(X^{\prime}\right)^{M^{\prime}},\left(Y^{\prime}\right)^{M^{\prime}}$, the rank data structures for the other pair can be built in $\mathcal{O}(k)$ time and space. We now must show that given the D data structure for one of the two strings $X^{M} Y^{M},\left(X^{\prime}\right)^{M^{\prime}}\left(Y^{\prime}\right)^{M^{\prime}}$, we can construct the D data structure for the other string efficiently. Assume that we are given $\mathrm{D}\left(\left(X^{\prime}\right)^{M^{\prime}}\left(Y^{\prime}\right)^{M^{\prime}}\right)$, the other case is analogous. We extract all nondummy segments using Proposition $4.2(\mathrm{~g})$ and concatenate them using Proposition 4.2(i). In total, it takes $\mathcal{O}\left(k^{2} \log ^{4} n\right)$ time and space, giving (b).
(c) follows immediately from (b).

From Corollaries 5.13 and 5.17 it immediately follows that:
Claim 5.18. Given an integer $k \in \mathbb{Z}_{+}$and strings $X, Y \in \Sigma \leq k^{2}$, one can compute $\mathrm{qGR}_{k}(X, Y)$ in $\tilde{\mathcal{O}}\left(k^{3}\right)$ time and $\tilde{\mathcal{O}}\left(k^{2}\right)$ space.

Corollary 5.19. Consider a non-crossing matching $M$ of strings $X, Y \in \Sigma^{\leq n}$ and an integer $k \in \mathbb{Z}_{+}$such that $M \subseteq \mathrm{q} \mathcal{M}_{k}(X, Y)$ if $\operatorname{ed}(X, Y) \leq k$. Given $k$ and $\mathrm{E}^{M}(X, Y)$, the quasigreedy encoding $\mathrm{qGR}_{k}(X, Y)$ can be computed in $\tilde{\mathcal{O}}\left(k^{5}+z k\right)$ time and $\tilde{\mathcal{O}}\left(k^{2}+z\right)$ space, where $z=\left|\operatorname{LZ}\left(X^{M} Y^{M}\right)\right|$.

Proof. Let $\$_{1} \neq \$_{2}$ be auxiliary symbols not in $\Sigma \cup\{\#\}$. We construct $M^{\prime}=\{(x+1, y+$ 1) : $(x, y) \in M\}$ and $\mathrm{E}^{M^{\prime}}\left(\$_{1} X, \$_{2} Y\right)$. Note that $M^{\prime}$ is a non-crossing matching of $\$_{1} X, \$_{2} Y$. Moreover, if $M \subseteq \mathrm{q} \mathcal{M}_{k}(X, Y)$, then $M^{\prime} \subseteq \mathcal{M}_{k+1}\left(\$_{1} X, \$_{2} Y\right)$ holds by Corollary 5.17. Hence, we can construct $\mathrm{GR}_{k+1}\left(\$_{1} X, \$_{2} Y\right)$ using Proposition 5.12. Then, we derive $\mathrm{qGR}_{k}(X, Y)$ using Corollary 5.17 again. By Lemma 3.11, we have $\left|\mathcal{B}_{k+1}\left(\$_{1} X, \$_{2} Y\right)\right|=\mathcal{O}\left(k^{5}\right)$, which gives the desired time bound.

Observation 5.20. Let $k^{\prime} \leq k$. Given $\mathrm{qGR}_{k}(X, Y)$, the encoding $\mathrm{qGR}_{k^{\prime}}(X, Y)$ can be computed in $\tilde{\mathcal{O}}\left(k^{5}\right)$ time and $\tilde{\mathcal{O}}\left(k^{2}\right)$ space.

Proof. For $k^{\prime} \leq k$, we have $\mathrm{q} \mathcal{M}_{k}(X, Y) \subseteq \mathrm{q} \mathcal{M}_{k^{\prime}}(X, Y)$ by definition. The claim follows from Corollary 5.19.

We now show that the quasi-greedy encodings are concatenatable.
Lemma 5.21. Consider strings $X_{p}, Y_{p}, X_{s}, Y_{s} \in \Sigma^{\leq n}$ and $k \in \mathbb{Z}_{+}$. Assume $\max \left\{\left|\left|X_{p}\right|-\right.\right.$ $\left|Y_{p}\right|\left|,\left|\left|X_{s}\right|-\left|Y_{s}\right|\right|\right\}=\mathcal{O}(k)$. Given $\mathrm{qGR}_{k+\left|\left|X_{s}\right|-\left|Y_{s}\right|\right|}\left(X_{p}, Y_{p}\right)$ and $\mathrm{qGR}_{k+\left|\left|X_{p}\right|-\left|Y_{p}\right|\right|}\left(X_{s}, Y_{s}\right)$, one can compute $\mathrm{qGR}_{k}\left(X_{p} X_{s}, Y_{p} Y_{s}\right)$ in $\tilde{\mathcal{O}}\left(k^{5}\right)$ time and $\tilde{\mathcal{O}}\left(k^{2}\right)$ space.

Proof. For brevity, let $M_{p}=\mathrm{q} \mathcal{M}_{k+\left|\left|X_{s}\right|-\left|Y_{s}\right|\right|}\left(X_{p}, Y_{p}\right)$ and $M_{s}=\mathrm{q} \mathcal{M}_{k+\left|\left|X_{p}\right|-\left|Y_{p}\right|\right|}\left(X_{s}, Y_{s}\right)$. Let also $X=X_{p} X_{s}, Y=Y_{p} Y_{s}$.

By Lemma 5.15 , if ed $(X, Y) \leq k$, then $M:=M_{p} \cup\left\{\left(x+\left|X_{p}\right|, y+\left|Y_{p}\right|\right):(x, y) \in M_{s}\right\} \subseteq$ q $\mathcal{M}_{k}(X, Y)$. Hence, we shall construct $\mathrm{E}^{M}(X, Y)$ and then apply Corollary 5.19.

For this, we extract $\mathrm{D}\left(X_{p}^{M_{p}}\right), \mathrm{D}\left(X_{s}^{M_{s}}\right), \mathrm{D}\left(Y_{p}^{M_{p}}\right), \mathrm{D}\left(Y_{s}^{M_{s}}\right)$ using Proposition $4.2(\mathrm{~g})$, and then concatenate them to $\mathrm{D}\left(X^{M} Y^{M}\right)=\mathrm{D}\left(X_{p}^{M_{p}} X_{s}^{M_{s}} Y_{p}^{M_{p}} Y_{s}^{M_{s}}\right)$ using Proposition 4.2(i). Overall, this takes $\mathcal{O}\left(k^{2} \log ^{4} n\right)$ time. Then, we use Proposition 5.2 to build $\mathrm{RS}_{\#}\left(X^{M}\right)$ and $\mathrm{RS}_{\#}\left(Y^{M}\right)$ in $\tilde{\mathcal{O}}(k)$ time.

Finally, we note that using Corollary $5.19 \operatorname{costs} \tilde{\mathcal{O}}\left(k^{5}\right)$ time and $\tilde{\mathcal{O}}\left(k^{2}\right)$ space (here we use the fact that $\left.\left|\left|X_{p}\right|-\left|Y_{p}\right|\right|=\mathcal{O}(k)\right)$.

Finally, as a corollary we derive an algorithm that can compute the quasi-greedy encoding of arbitrarily long strings.

Corollary 5.22. Assuming constant-time random access to a string $X \in \Sigma^{\ell}$ and streaming access to a string $Y \in \Sigma^{\ell}$, where $\ell \leq n$, there is an algorithm that constructs $\mathrm{q}_{\mathrm{GR}}(X, Y)$ with a delay of at most $k^{2}$ characters in $\tilde{\mathcal{O}}\left(k^{3}\right)$ amortised time per character and $\tilde{\mathcal{O}}\left(k^{2}\right)$ space. The delay means that at the moment when the $y$-th character of $Y$ arrives, the algorithm knows $\mathrm{qGR}{ }_{k}\left(X\left[. . y^{\prime}\right], Y\left[. . y^{\prime}\right]\right)$, where $\left|y-y^{\prime}\right| \leq k^{2}$.

Proof. If $\ell \leq k^{2}$, construct $\mathrm{qGR}_{k}(X, Y)$ via Claim 5.18 in $\tilde{\mathcal{O}}\left(k^{3}\right)$ (total) time and $\tilde{\mathcal{O}}\left(k^{2}\right)$ space. Otherwise, partition the strings into non-overlapping blocks $X=X_{1} \cdots X_{p}$ and $Y=Y_{1} \cdots Y_{p}$ so that $\left|X_{i}\right|=\left|Y_{i}\right|=k^{2}$ for all $i \in[1 \ldots p)$ and $\left|X_{p}\right|=\left|Y_{z}\right|=\ell \bmod k^{2}$. Suppose that we have computed $\mathrm{qGR}_{k}\left(X_{1} \cdots X_{i-1}, Y_{1} \cdots Y_{i-1}\right)$ for $i \in[1 \cdots p)$. Compute $\mathrm{qGR}_{k}\left(X_{1} \cdots X_{i}, Y_{1} \cdots Y_{i}\right)$ in the following manner: first, compute $\mathrm{qGR}_{k}\left(X_{i}, Y_{i}\right)$ in $\mathcal{O}\left(k^{5}\right)$ time and $\mathcal{O}\left(k^{2}\right)$ space via Claim 5.18, and then compute $\mathrm{qGR}_{k}\left(X_{1} \cdots X_{i}, Y_{1} \cdots Y_{i}\right)$ in $\tilde{\mathcal{O}}\left(k^{5}\right)$ time and $\tilde{\mathcal{O}}\left(k^{2}\right)$ space via Lemma 5.21.

### 5.4 Products of Greedy Alignments

Definition 3.12. Consider strings $X, Y, Z \in \Sigma^{*}$, an alignment $\mathcal{A}^{X, Y}$ of $X, Y$, an alignment $\mathcal{A}^{Y, Z}$ of $Y, Z$, and an alignment $\mathcal{A}^{X, Z}$ of $X, Z$. We say that $\mathcal{A}^{X, Z}$ is a product of $\mathcal{A}^{X, Y}$ and $\mathcal{A}^{Y, Z}$ if, for every $(x, z) \in \mathcal{A}^{X, Z}$, there is $y \in[1 \ldots|Y|+1]$ such that $(x, y) \in \mathcal{A}^{X, Y}$ and $(y, z) \in \mathcal{A}^{Y, Z}$.
Lemma 3.13. Consider strings $X, Y, Z \in \Sigma^{*}$ and $k \in \mathbb{Z}_{\geq 0}$. Every alignment $\mathcal{A}^{X, Z} \in \mathrm{GA}_{k}(X, Z)$ is a product of alignments $\mathcal{A}^{X, Y} \in \mathrm{GA}_{d}(X, Y)$ and $\mathcal{A}^{Y, Z} \in \mathrm{GA}_{d}(Y, Z)$, where $d=2 k+\operatorname{ed}(X, Y)$.

Proof. We proceed by induction on $|X|+|Y|+|Z|$. In the base case, when at least one of the strings $X, Y, Z$ is empty, we set $\mathcal{A}^{X, Y}$ and $\mathcal{A}^{Y, Z}$ to be any greedy optimal alignments of $X, Y$ and $Y, Z$, respectively, so that $\operatorname{cost}\left(\mathcal{A}^{X, Y}\right)=\operatorname{ed}(X, Y)$ and $\operatorname{cost}\left(\mathcal{A}^{Y, Z}\right)=\operatorname{ed}(Y, Z) \leq$ $\mathrm{ed}(X, Y)+\operatorname{ed}(X, Z) \leq \operatorname{ed}(X, Y)+\operatorname{cost}\left(\mathcal{A}^{X, Z}\right)$. Moreover, it easy to check that $\mathcal{A}^{X, Z}$ is a product of $\mathcal{A}^{X, Y}$ and $\mathcal{A}^{Y, Z}$ because two out of these three alignments simply delete all characters of the non-empty string.

In the inductive step, we assume that all strings $X, Y, Z$ are non-empty, and we consider several cases.

1. $\mathbf{X}[\mathbf{1}]=\mathbf{Z}[\mathbf{1}]=\mathbf{Y}[\mathbf{1}]$.

We recurse on $X^{\prime}=X[2 .],. Y^{\prime}=Y[2 .],. Z^{\prime}=Z[2 .$.$] , and \mathcal{A}^{X^{\prime}, Z^{\prime}}=\mathcal{A}_{[2, .1,][2 .]}^{X,[2,]}$, which is greedy due to $X[1] \simeq_{\mathcal{A}^{X, Z}} Z[1]$. This yields greedy alignments $\mathcal{A}^{X^{\prime}, Y^{\prime},} A^{Y^{\prime}, Z^{\prime}}$ of cost at most $d^{\prime}=2 \operatorname{cost}\left(\mathcal{A}^{X, Z}\right)+\operatorname{ed}(X, Y)=d$. We extend them so that $X[1] \simeq_{\mathcal{A}^{X}, Y} Y[1]$ and $Y[1] \simeq_{\mathcal{A}^{Y, Z}} Z[1]$, obtaining alignments of cost up to $d$.
2. $\mathbf{X}[\mathbf{1}]=\mathbf{Z}[\mathbf{1}] \neq \mathbf{Y}[\mathbf{1}]$.

In this case, we have ed $(X, Y)>\min (\operatorname{ed}(X[2 .],. Y[2 .]),. \operatorname{ed}(X[2 .], Y$.$) , ed (X, Y[2 .])$.$) .$
(a) If $\operatorname{ed}(X, Y)>\operatorname{ed}(X[2 .],. Y[2 .]$.$) , we recurse on X^{\prime}=X[2 .],. Y^{\prime}=Y[2 \ldots], Z^{\prime}=$ $Z[2 .$.$] , and \mathcal{A}^{X^{\prime}, Z^{\prime}}=\mathcal{A}_{[2,1,[2 .]]}^{X, Z}$, which is greedy due to $X[1] \simeq_{\mathcal{A}^{X, Z}} Z[1]$. This yields greedy alignments $\mathcal{A}^{X^{\prime}, Y^{\prime}}, A^{Y^{\prime}, Z^{\prime}}$ of cost at most $d^{\prime}=2 \operatorname{cost}\left(\mathcal{A}^{X, Z}\right)+\operatorname{ed}(X, Y)-1=$ $d-1$. We extend them so that $X[1] \sim_{\mathcal{A}^{X, Y}} Y[1]$ and $Y[1] \sim_{\mathcal{A}^{Y, Z}} Z[1]$, obtaining alignments of cost up to $d$.
(b) If ed $(X, Y)>\operatorname{ed}(X[2 .], Y$.$) , we recurse on X^{\prime}=X[2 .],. Y^{\prime}=Y, Z^{\prime}=Z[2 .$.$] ,$ and $\mathcal{A}^{X^{\prime}, Z^{\prime}}=\mathcal{A}_{[2 . .],[2 .]}^{X, Z}$, which is greedy due to $X[1] \simeq_{\mathcal{A}^{X, Z}} Z[1]$. This yields greedy alignments $\mathcal{A}^{X^{\prime}, Y^{\prime},} A^{Y^{\prime}, Z^{\prime}}$ of cost at most $d^{\prime}=2 \operatorname{cost}\left(\mathcal{A}^{X, Z}\right)+\operatorname{ed}(X, Y)-1=d-1$. We extend them so that $\mathcal{A}^{X, Y}$ deletes $X[1]$ and $\mathcal{A}^{Y, Z}$ deletes $Z[1]$, obtaining alignments of cost up to $d$.
(c) If $\operatorname{ed}(X, Y)>\operatorname{ed}(X, Y[2 \ldots])$, we recurse on $X^{\prime}=X, Y^{\prime}=Y[2 \ldots], Z^{\prime}=Z$, and $\mathcal{A}^{X^{\prime}, Z^{\prime}}=\mathcal{A}^{X, Z}$. This yields greedy alignments $\mathcal{A}^{X^{\prime}, Y^{\prime}}, A^{Y^{\prime}, Z^{\prime}}$ of cost at most $d^{\prime}=$ $2 \operatorname{cost}\left(\mathcal{A}^{X, Z}\right)+\operatorname{ed}(X, Y)-1=d-1$. We extend them so that $\mathcal{A}^{X, Y}$ and $\mathcal{A}^{Y, Z}$ both delete $Y[1]$, obtaining alignments of cost up to $d$.
3. $\mathbf{X}[\mathbf{1}] \neq \mathbf{Z}[\mathbf{1}]=\mathbf{Y}[\mathbf{1}]$.
(a) If $\mathcal{A}^{X, Z}$ deletes $X[1]$, we recurse on $X^{\prime}=X[2 \ldots], Y^{\prime}=Y, Z^{\prime}=Z$, and $\mathcal{A}^{X^{\prime}, Z^{\prime}}=$ $\mathcal{A}_{[2 . .],[1 .]}^{X, Z}$. This yields greedy alignments $\mathcal{A}^{X^{\prime}, Y^{\prime}}, A^{Y^{\prime}, Z^{\prime}}$ of cost $\leq d^{\prime}=2 \operatorname{cost}\left(\mathcal{A}^{X, Z}\right)-$ $2+\operatorname{ed}\left(X^{\prime}, Y^{\prime}\right) \leq d-1$. We derive $\mathcal{A}^{Y, Z}=\mathcal{A}^{Y^{\prime}, Z^{\prime}}$ and extend $\mathcal{A}^{X^{\prime}, Y^{\prime}}$ so that $\mathcal{A}^{X, Y}$ deletes $X[1]$, obtaining an alignment of cost up to $d$.
(b) If $\mathcal{A}^{X, Z}$ deletes $Z[1]$, we recurse on $X^{\prime}=X, Y^{\prime}=Y[2 \ldots], Z^{\prime}=Z[2 .$.$] , and \mathcal{A}^{X^{\prime}, Z^{\prime}}=$
 $2+\operatorname{ed}\left(X^{\prime}, Y^{\prime}\right) \leq d-1$. We extend them so that $\mathcal{A}^{X, Y}$ deletes $Y[1]$ and $Y[1] \simeq_{\mathcal{A}^{Y, Z}} Z[1]$, obtaining alignments of cost up to $d$.
(c) If $X[1] \sim_{\mathcal{A}^{x, Z}} Z[1]$, we recurse on $X^{\prime}=X[2 \ldots], Y^{\prime}=Y[2 \ldots], Z^{\prime}=Z[2 .$.$] , and$ $\mathcal{A}^{X^{\prime}, Z^{\prime}}=\mathcal{A}_{[2 . .],[2 . .]}^{X, Z}$. This yields greedy alignments $\mathcal{A}^{X^{\prime}, Y^{\prime}}, A^{Y^{\prime}, Z^{\prime}}$ of cost $\leq d^{\prime}=$ $2 \operatorname{cost}\left(\mathcal{A}^{X, Z}\right)-2+\operatorname{ed}\left(X^{\prime}, Y^{\prime}\right) \leq d-1$. We extend them so that $X[1] \sim_{\mathcal{A}^{X, Y}} Y[1]$ and $Y[1] \simeq_{\mathcal{A}^{Y, Z}} Z[1]$, obtaining alignments of cost up to $d$.
4. $\mathbf{X}[\mathbf{1}] \neq \mathbf{Z}[\mathbf{1}] \neq \mathbf{Y}[\mathbf{1}]$. (Note that this case allows both $X[1]=Y[1]$ and $X[1] \neq Y[1]$.)
(a) If $\mathcal{A}^{X, Z}$ deletes $X[1]$, we recurse on $X^{\prime}=X[2 .],. Y^{\prime}=Y[2 .],. Z^{\prime}=Z$, and $\mathcal{A}^{X^{\prime}, Z^{\prime}}=\mathcal{A}_{[2 . .],[1 . .]}^{X, Z}$. This yields greedy alignments $\mathcal{A}^{X^{\prime}, Y^{\prime}}, A^{Y^{\prime}, Z^{\prime}}$ of cost at most $d^{\prime}=2 \operatorname{cost}\left(\mathcal{A}^{X, Z}\right)-2+\operatorname{ed}\left(X^{\prime}, Y^{\prime}\right) \leq d-2$. We extend them so that $X[1] \sim_{\mathcal{A}^{X, Y}} Y[1]$ and $\mathcal{A}^{Y, Z}$ deletes $Y[1]$, obtaining alignments of cost up to $d-1$.
(b) If $\mathcal{A}^{X, Z}$ deletes $Z[1]$, we recurse on $X^{\prime}=X, Y^{\prime}=Y, Z^{\prime}=Z[2 \ldots]$, and $\mathcal{A}^{X^{\prime}, Z^{\prime}}=$ $\mathcal{A}_{[1 . .],[2 .]}^{X, Z}$. This yields greedy alignments $\mathcal{A}^{X^{\prime}, Y^{\prime},} A^{Y^{\prime}, Z^{\prime}}$ of $\operatorname{cost} \leq d^{\prime}=2 \operatorname{cost}\left(\mathcal{A}^{X, Z}\right)-$ $2+\operatorname{ed}(X, Y)=d-2$. We derive $\mathcal{A}^{X, Y}=\mathcal{A}^{X^{\prime}, Y^{\prime}}$ and extend $\mathcal{A}^{Y^{\prime}, Z^{\prime}}$ so that $\mathcal{A}^{Y, Z}$ deletes $Z[1]$, obtaining an alignment of cost up to $d-1$.
(c) If $X[1] \sim_{\mathcal{A}^{x, Z}} Z[1]$, we recurse on $X^{\prime}=X[2 \ldots], Y^{\prime}=Y[2 \ldots], Z^{\prime}=Z[2 .$.$] , and$ $\mathcal{A}^{X^{\prime}, Z^{\prime}}=\mathcal{A}_{[2 . .],[2 . .]}^{X, Z}$. This yields greedy alignments $\mathcal{A}^{X^{\prime}, Y^{\prime}}, A^{Y^{\prime}, Z^{\prime}}$ of $\operatorname{cost} \leq d^{\prime}=$ $2 \operatorname{cost}\left(\mathcal{A}^{X, Z}\right)-2+\operatorname{ed}\left(X^{\prime}, Y^{\prime}\right) \leq d-2$. We extend them so that $X[1] \sim_{\mathcal{A}^{X, Y}} Y[1]$ and $Y[1] \sim_{\mathcal{A}^{Y, Z}} Z[1]$, obtaining alignments of cost up to $d-1$.
In all the cases above, $\mathcal{A}^{X, Y}$ is greedy because $\mathcal{A}^{X^{\prime}, Y^{\prime}}$ is greedy and $\mathcal{A}^{X, Y}$ matches $X[1]$ with $Y[1]$ whenever $X[1]=Y[1]$. Similarly, $\mathcal{A}^{Y, Z}$ is greedy because $\mathcal{A}^{Y^{\prime}, Z^{\prime}}$ is greedy and $\mathcal{A}^{Y, Z}$ matches $Y[1]$ with $Z[1]$ whenever $Y[1]=Z[1]$. Moreover, $\mathcal{A}^{X, Z}$ is a product of $\mathcal{A}^{X, Y}$ and $\mathcal{A}^{Y, Z}$ because each alignment starts with $(1,1)$ and since $\mathcal{A}^{X^{\prime}, Z^{\prime}}$ is a product of $\mathcal{A}^{X^{\prime}, Y^{\prime}}$ and $\mathcal{A}^{Y^{\prime}, Z^{\prime}}$.

Assume that we are given three strings $X, Y, Z$. Let $d=\operatorname{ed}(X, Y)+2 k$, and define $\mathcal{G}_{X}=\mathrm{GR}_{d}(X, Y)$ and $\mathcal{G}_{Z}=\mathrm{GR}_{d}(Y, Z)$. We show that given $\mathcal{G}_{X}$ and $\mathcal{G}_{Z}$, we can compute an optimal alignment between $X$ and $Z$ efficiently if its cost is at most $k$. Let $M=\{(x, z)$ : $\exists y$ such that $(x, y) \in \mathcal{M}_{d}(X, Y)$ and $\left.(y, z) \in \mathcal{M}_{d}(Y, Z)\right\}$.

Lemma 5.23. If $\operatorname{ed}(X, Z) \leq k$, then $M \subseteq \mathcal{M}_{k}(X, Z)$.
Proof. Suppose that $(x, y) \in \mathcal{M}_{d}(X, Y)$ and $(y, z) \in \mathcal{M}_{d}(Y, Z)$. For a proof by contradiction, suppose that $X[x] \not \chi_{\mathcal{A}} Z[z]$ for some $\mathcal{A} \in \mathrm{GA}_{k}(X, Z)$. Then, there exists $\left(x^{\prime}, z^{\prime}\right) \in \mathcal{A}$ such that either $x^{\prime} \leq x$ and $z^{\prime}>z$, or $x^{\prime}>x$ and $z^{\prime} \leq z$. By symmetry, without loss of generality, we consider the first alternative. By Lemma $3.13, \mathcal{A}$ is a product $\mathcal{A}^{X, Y} \in \mathrm{GA}_{d}(X, Y)$ and $\mathcal{A}^{Y, Z} \in \mathrm{GA}_{d}(Y, Z)$. According to Definition 3.12, this means that there exists $y^{\prime}$ such that $\left(x^{\prime}, y^{\prime}\right) \in \mathcal{A}^{X, Y}$ and $\left(y^{\prime}, z^{\prime}\right) \in \mathcal{A}^{Y, Z}$. If $y^{\prime} \leq y$, then $(y, z),\left(y^{\prime}, z^{\prime}\right) \in \mathcal{A}^{Y, Z}$ implies that $\mathcal{A}^{Y, Z}$
deletes $Z[z]$, contradicting $(y, z) \in \mathcal{M}\left(\mathcal{A}^{Y, Z}\right)$. Similarly, if $y^{\prime}>y$, then $(x, y),\left(x^{\prime}, y^{\prime}\right) \in \mathcal{A}^{X, Y}$ implies that $\mathcal{A}^{X, Y}$ deletes $Y[y]$, contradicting $(x, y) \in \mathcal{M}\left(\mathcal{A}^{X, Y}\right)$. This completes the proof that $X[x] \sim_{\mathcal{A}} Z[z]$ for every $\mathcal{A} \in \operatorname{GA}_{k}(X, Z)$. Due to $X[x]=Y[y]=Z[z]$, we also have $X[x] \simeq_{\mathcal{A}} Z[z]$ for every $\mathcal{A} \in \mathrm{GA}_{k}(X, Z)$, i.e., $(x, y) \in \mathcal{M}_{k}(X, Z)$ holds as claimed.

Lemma 5.24. If $\operatorname{ed}(X, Y)=\mathcal{O}(k)$ and $\mathcal{G}_{X}, \mathcal{G}_{Z} \neq \perp, \mathrm{E}^{M}(X, Z)$ can be computed in $\tilde{\mathcal{O}}\left(k^{3}\right)$ time and $\tilde{\mathcal{O}}\left(k^{2}\right)$ space.

Proof. Let us first explain how $X^{M}$ can be constructed. Consider a position $x \in[1 \ldots|X|]$ such that $(x, z) \notin M$ for every $z \in[1 \ldots|Z|]$. If $X^{\mathcal{M}_{d}(X, Y)}[x] \neq \#$, then $X^{M}[x]=X[x]=$ $X^{\mathcal{M}_{d}(X, Y)}[x]$. Otherwise, $(x, y) \in \mathcal{M}_{d}(X, Y)$ for some $y \in[1 . .|Y|]$. By Lemma 5.23, we have $(y, z) \notin \mathcal{M}_{d}(Y, Z)$ for every $z \in[1 \ldots|Z|]$. Hence, $Y^{\mathcal{M}_{d}(Y, Z)}[y] \neq \#$, and we have $X[x]=$ $Y[y]=Y^{\mathcal{M}_{d}(Y, Z)}[y]$. In words, non-dummy characters of $X^{M}$ can be retrieved from non-dummy characters of $X^{\mathcal{M}_{d}(X, Y)}$ and non-dummy characters of $Y^{\mathcal{M}_{d}(Y, Z)}$.

Thus, for each dummy segment $[\ell \ldots r)$ in $X^{\mathcal{M}_{d}(X, Y)}$ and its counterpart $\left[\ell^{\prime} \ldots r^{\prime}\right)$ in $Y^{\mathcal{M}_{d}(X, Y)}$, we need to set $X^{M}[\ell \ldots r):=Y^{\mathcal{M}_{d}(Y, Z)}\left[\ell^{\prime} \ldots r^{\prime}\right)$. Such a copy-paste operation can be implemented using Proposition $4.2(\mathrm{~g})(\mathrm{i})$, in $\mathcal{O}\left(k^{2} \log ^{4} n\right)$ time per dummy segment. We can also keep track of the dummy segments in $X^{M}$ within the same procedure. The algorithm for $Z^{M}$ is symmetric, and thus we can construct $\mathrm{D}\left(X^{M} Z^{M}\right)$ along with the dummy segments in $\tilde{\mathcal{O}}\left(k^{3}\right)$ time and $\tilde{\mathcal{O}}\left(k^{2}\right)$ space. Finally, we build $\mathrm{RS}_{\#}\left(X^{M}\right)$ and $\mathrm{RS}_{\#}\left(Z^{M}\right)$ in $\mathcal{O}(k)$ time.

Corollary 5.25. Given $\mathcal{G}_{X}$ and $\mathcal{G}_{Z}$. If $\operatorname{ed}(X, Y)=\mathcal{O}(k)$, then we can compute $\min (k+$ $1, \operatorname{ed}(X, Z)$ ) in $\tilde{\mathcal{O}}\left(k^{3}\right)$ time and $\tilde{\mathcal{O}}\left(k^{2}\right)$ space.

Proof. If $\mathcal{G}_{X}$ or $\mathcal{G}_{Z}$ equals to $\perp$, then $\operatorname{ed}(X, Z)>k$ by Lemma 3.13. Otherwise, construct $\mathrm{E}^{M}(X, Z)$ using Lemma 5.24. Finally, we pass $k$ and $\mathrm{E}^{M}(X, Z)$ to the algorithm of Corollary 5.5, and we return the resulting value $d$.

Note that $M$ is a non-crossing matching of $X, Z$, so $d=k+1$ is correctly returned if ed $(X, Z)>k$. If ed $(X, Z) \leq k$, then Lemma 5.23 implies $M \subseteq \mathcal{M}_{k}(X, Z)$, and thus $d=\operatorname{ed}(X, Z)$ holds as claimed.

The overall runtime and space complexity are dominated by the procedure of Lemma 5.24.

Remark 5.26. Using Proposition 5.12 instead of Corollary 5.5, we could construct $\mathrm{GR}_{k}(X, Z)$.
Corollary 5.27. Consider three strings $X, Y, Z$. Let $d=\operatorname{ed}(X, Y)+2 k$, and assume that we are given $\mathrm{qGR}_{d}(X, Y)$ and $\mathrm{qGR}_{d}(Y, Z)$. If $\operatorname{ed}(X, Y)=\mathcal{O}(k)$, then we can compute $\min (k+$ $1, \operatorname{ed}(X, Z)$ ) in $\tilde{\mathcal{O}}\left(k^{3}\right)$ time and $\tilde{\mathcal{O}}\left(k^{2}\right)$ space.

Proof. We compute $\mathcal{G}_{X}=\operatorname{GR}_{d+1}\left(\$_{1} X, \$_{2} Y\right)$ and $\mathrm{qGR}_{d+1}\left(\$_{2} Y, \$_{3} Z\right)$ in $\tilde{\mathcal{O}}\left(k^{2}\right)$ time and space via Corollary 5.17, and apply Corollary 5.25 to compute $\min \left(k+2\right.$, ed $\left.\left(\$ \$_{1} X, \$_{2} Z\right)\right)-1=\min (k+$ $1, \operatorname{ed}(X, Z))$ in $\tilde{\mathcal{O}}\left(k^{3}\right)$ time and $\tilde{\mathcal{O}}\left(k^{2}\right)$ space.

## 6 Edit Distance Sketches

### 6.1 CGK embedding

In this section, we prove Proposition 3.16 based on the CGK embedding introduced in [10]. Recall that the Hamming distance between the embeddings of two strings $X, Y \in \Sigma \leq n$ is bounded in terms of the edit distance ed $(X, Y)$, which allows using Hamming distance sketches to approximate edit distance.

Definition 6.1 (CGK embedding [10]). Consider an alphabet $\Sigma$, a sentinel character $\perp \notin \Sigma$, and a 2-independent family $\mathcal{H}$ of hash functions $h: \Sigma \rightarrow\{0,1\}$. For an integer $n \in \mathbb{Z}_{+}$, a (uniformly random) sequence $R \in \mathcal{H}^{3 n}$, and a string $S \in \Sigma^{\leq n}$, the CGK walk $W_{\mathrm{CGK}}(S)=$ $\left(s_{t}\right)_{t=1}^{3 n+1}$ and the CGK embedding $\operatorname{CGK}(S) \in(\Sigma \cup\{\perp\})^{3 n}$ are defined by the following algorithm:

```
Input: An integer \(n \in \mathbb{Z}_{+}\), a string \(S \in \Sigma^{\leq n}\).
Randomness: A sequence \(R \in \mathcal{H}^{3 n}\) of 2 -independent hash functions \(\Sigma \rightarrow\{0,1\}\).
Output: The CGK walk \(W_{\text {CGK }}(S)=\left(s_{t}\right)_{t=1}^{3 n+1}\) and the CGK embedding CGK \((S)\).
\(s_{1}:=1 ;\)
for \(t:=1\) to \(3 n\) do
    if \(s_{t} \leq|S|\) then
        \(\operatorname{CGK}(S)[t]:=S\left[s_{t}\right] ;\)
        \(s_{t+1}:=s_{t}+R_{t}\left(S\left[s_{t}\right]\right) ;\)
    else
        \(\operatorname{CGK}(S)[t]:=\perp\);
        \(s_{t+1}:=s_{t} ;\)
```

Algorithm 2: The CGK algorithm

Recall from Definition 3.14 that an $m$-step walk over $S$ is complete if $s_{m+1}=|S|+1$.
Fact 6.2 ([10, Theorem 4.1]). Each $S \in \Sigma^{\leq n}$ satisfies $\operatorname{Pr}_{R}\left[W_{\text {CGK }}(S)\right.$ is complete $] \geq 1-e^{-\Omega(n)}$.
The following result summarizes the central property of the CGK embedding.
Fact 6.3 ([10, Theorem 4.3]). For every $X, Y \in \Sigma^{\leq n}$ and every constant $c>0$, the embeddings $\operatorname{CGK}(X), \operatorname{CGK}(Y)$ satisfy $\operatorname{Pr}_{R \in \mathcal{H}^{3 n}}\left[\operatorname{hd}(\operatorname{CGK}(X), \operatorname{CGK}(Y))>c \cdot \operatorname{ed}(X, Y)^{2}\right]<\frac{12}{\sqrt{c}}$.

If $W_{\mathrm{CGK}}(X)$ and $W_{\mathrm{CGK}}(Y)$ are complete, Definition 3.15 yields an edit-distance alignment of $X, Y$, which we call the CGK alignment of $X, Y$. that they induce an edit-distance alignment

Fact 6.4 (see also [10, Theorem 4.2]). If $W_{\mathrm{CGK}}(X)$ and $W_{\mathrm{CGK}}(Y)$ are complete for some $X, Y \in$ $\Sigma \leq n$ and $R \in \mathcal{H}^{3 n}$, then the CGK alignment of $X, Y$ belongs to $\operatorname{GA}_{\text {hd }(\operatorname{CGK}(X), \operatorname{CGK}(Y))}(X, Y)$.
Proof. Let $\mathcal{A}$ be the CGK alignment of $X$, $Y$, i.e., the zip alignment of $W_{\mathrm{CGK}}(X)=\left(x_{t}\right)_{t=1}^{3 n+1}$ and $W_{\mathrm{CGK}}(Y)=\left(y_{t}\right)_{t=1}^{3 n+1}$. Consider $\left(x_{t}, y_{t}\right) \in \mathcal{B}(\mathcal{A})$, with $t \in[1 \ldots 3 n+1]$ chosen so that $t=3 n+1$ or $\left(x_{t+1}, y_{t+1}\right) \neq\left(x_{t}, y_{t}\right)$. If $t=3 n+1$, then $\left(x_{t}, y_{t}\right)=(|X|+1,|Y|+1)$ by the assumption that $W_{\mathrm{CGK}}(X)$ and $W_{\mathrm{CGK}}(Y)$ are complete. If $\left(x_{t+1}, y_{t+1}\right)=\left(x_{t}+1, y_{t}+1\right)$, then we have $\operatorname{CGK}(X)[t]=$ $X\left[x_{t}\right] \neq Y\left[y_{t}\right]=\operatorname{CGK}(Y)[t]$. The remaining possibility $x_{t+1}-x_{t} \neq y_{t+1}-y_{t}$ holds only in the following three cases: if $x_{t}=|X|+1$ and $y_{t} \leq|Y|$ (when $\left.\operatorname{CGK}(X)[t]=\perp \neq \operatorname{CGK}(Y)[t]=Y\left[y_{t}\right]\right)$, if $x_{t} \leq|X|$ and $y_{t}=|Y|+1$ (when $X\left[x_{t}\right]=\operatorname{CGK}(X)[t] \neq \perp=\operatorname{CGK}(Y)[t]$, or if $x_{t} \leq|X|$, $y_{t} \leq|Y|$, and $R_{t}\left(X\left[x_{t}\right]\right) \neq R_{t}\left(Y\left[y_{t}\right]\right)$ (when $\operatorname{CGK}(X)[t]=X\left[x_{t}\right] \neq Y\left[y_{t}\right]=\operatorname{CGK}(Y)[t]$. Overall, we have $X\left[x_{t}\right] \neq Y\left[y_{t}\right]$ whenever $\left(x_{t}, y_{t}\right) \in[1 \ldots|X|] \times[1 \ldots|Y|]$, and $\operatorname{CGK}(X)[t] \neq \operatorname{CGK}(Y)[t]$ whenever $t \in[1 \ldots 3 n]$. Consequently, $\mathcal{A} \in \operatorname{GA}_{\text {hd }(\operatorname{CGK}(X), \operatorname{CGK}(Y))}(X, Y)$.

The final property of the CGK alignment required in this work is that its width is within $\mathcal{O}(\operatorname{ed}(X, Y))$ with good probability. For this, we first prove a fact about random walks:

Fact 6.5. Let $\left(w_{i}\right)_{i \geq 0}$ be an unbiased lazy random walk (that is, $w_{0}=0$, and, for every $i \geq 1$, we have $\operatorname{Pr}\left[w_{i+1}=w_{i} \mid w_{0}, \ldots, w_{i}\right]=\frac{1}{2}, \operatorname{Pr}\left[w_{i+1}=w_{i}+1 \mid w_{0}, \ldots, w_{i}\right]=\operatorname{Pr}\left[w_{i+1}=w_{i}-1 \mid\right.$ $\left.\left.w_{0}, \ldots, w_{i}\right]=\frac{1}{4}\right)$. Then, for every $m, \ell \in \mathbb{Z}_{+}$, we have $\operatorname{Pr}\left[\max _{i=0}^{m}\left|w_{i}\right| \geq \ell\right] \leq \frac{m}{\ell^{2}}$.

Proof. By the reflection principle [44, Excercise 2.10], we have $\operatorname{Pr}\left[\max _{i=0}^{m}\left|w_{i}\right| \geq \ell\right] \leq 2 \operatorname{Pr}\left[\left|w_{m}\right| \geq\right.$ $\ell]=2 \operatorname{Pr}\left[w_{m}^{2} \geq \ell^{2}\right]$. By Markov's inequality,

$$
2 \operatorname{Pr}\left[w_{m}^{2} \geq \ell^{2}\right] \leq \frac{2 \mathbb{E}\left[w_{m}^{2}\right]}{\ell^{2}}=\frac{2 \mathbb{E}\left[\left(w_{1}-w_{0}\right)^{2}+\cdots+\left(w_{m}-w_{m-1}\right)^{2}\right]}{\ell^{2}}=\frac{m}{\ell^{2}}
$$

Corollary 6.6. For every constant $\delta \in(0,1)$, there exists a constant $c$ such that for sufficiently large $n$, all strings $X, Y \in \Sigma^{\leq n}$, and their $C G K$ alignment $\mathcal{A}$, the probability over $R \in \mathcal{H}^{3 n}$ that $W_{\mathrm{CGK}}(X), W_{\mathrm{CGK}}(Y)$ are complete, $\mathcal{A} \in \mathrm{GA}_{c \cdot \mathrm{ed}(X, Y)^{2}}(X, Y)$, and width $(\mathcal{A}) \leq c \cdot \operatorname{ed}(X, Y)$ is at least $1-\delta$.

Proof. By Fact 6.2, $\operatorname{Pr}\left[W_{\text {CGK }}(S)\right.$ is incomplete $] \leq \frac{\delta}{4}$ holds for every $S \in \Sigma \leq n$ and sufficiently large $n=\Omega\left(\log \frac{1}{\delta}\right)$. By Fact 6.3 , there is a constant $c$ such that $\operatorname{Pr}[\operatorname{hd}(\operatorname{CGK}(X), \operatorname{CGK}(Y))>$ $\left.c \cdot \operatorname{ed}(X, Y)^{2}\right] \leq \frac{\delta}{4}$. We may take arbitrarily large $c$, so we shall also assume that $c \geq \frac{4}{\delta}$. Moreover, by Fact 6.4 , if $W_{\mathrm{CGK}}(X)$ and $W_{\mathrm{CGK}}(Y)$ are complete and $\operatorname{hd}(\mathrm{CGK}(X), \mathrm{CGK}(Y)) \leq c \cdot \operatorname{ed}(X, Y)^{2}$, then $\mathcal{A} \in \mathrm{GA}_{c \cdot \text { ed }(X, Y)^{2}}(X, Y)$.

To bound width $(\mathcal{A})$, let us analyze the CGK walks $W_{\text {CGK }}(X)=\left(x_{t}\right)_{t=1}^{3 n+1}$ and $W_{\text {CGK }}(Y)=$ $\left(y_{t}\right)_{t=1}^{3 n+1}$. Let $T$ be the first step such that $x_{T}=|X|+1, y_{T}=|Y|+1$, or $T=3 n+1$. Note that it suffices to bound $\operatorname{Pr}\left[\max _{t=1}^{T}\left|x_{t}-y_{t}\right| \geq c k\right]$, where $k=\operatorname{ed}(X, Y)$.

Let $A=\{1\} \cup\left\{t+1: t \in[1 \ldots T)\right.$ and $\left.X\left[x_{t}\right] \neq Y\left[y_{t}\right]\right\}$. Observe that $x_{t}-y_{t}=x_{t-1}-y_{t-1}$ holds for $t \in[1 \ldots T] \backslash A$, so we may focus on bounding $\operatorname{Pr}\left[\max _{t \in A}\left|x_{t}-y_{t}\right| \geq c k\right]$. Let $A=\left\{t_{0}, \ldots, t_{a}\right\}$ with $t_{0}<\cdots<t_{a}$, and let $d_{i}=x_{t_{i}}-y_{t_{i}}$ for $i \in[0 \ldots a]$. Observe that $\left(d_{i}\right)_{i=0}^{a}$ is an $a$-step unbiased lazy random walk and that $a \leq \operatorname{hd}(\operatorname{CGK}(X), \operatorname{CGK}(Y))$. Moreover, if we extend $\left(d_{i}\right)_{i=0}^{a}$ to an infinite unbiased lazy random walk $\left(d_{i}\right)_{i=0}^{\infty}$, then Fact 6.5 yields $\operatorname{Pr}\left[\max _{i=0}^{\left\lfloor c k^{2}\right\rfloor}\left|d_{i}\right| \geq c k\right] \leq$ $\frac{c k^{2}}{c^{2} k^{2}}=\frac{1}{c} \leq \frac{\delta}{4}$. Therefore,

$$
\operatorname{Pr}\left[\max _{i=0}^{a}\left|d_{i}\right| \geq c k\right] \leq \operatorname{Pr}\left[\operatorname{hd}(\operatorname{CGK}(X), \operatorname{CGK}(Y))>c k^{2}\right]+\operatorname{Pr}\left[{\underset{i x}{i=0}}_{\left\lfloor c k^{2}\right\rfloor}^{\max }\left|d_{i}\right| \geq c k\right]
$$

Consequently, if $\operatorname{CGK}(X)$ and $\operatorname{CGK}(Y)$ are complete, $\operatorname{hd}(\operatorname{CGK}(X), \operatorname{CGK}(Y)) \leq c k^{2}$, and the random walk satisfies $\operatorname{Pr}\left[\max _{i=0}^{\left\lfloor c k^{2}\right\rfloor}\left|d_{i}\right|<c k\right]$, then the alignment $\mathcal{A}$ satisfies the lemma. The total probability of the complementary events is at most $\delta$, so this completes the proof.

To complete the proof of Proposition 3.16 (repeated below), it remains to reduce the number of random bits using Nisan's pseudorandom generator [50].

Proposition 3.16. For every constant $\delta \in(0,1)$, there exists a constant $c$ and an algorithm W that, given an integer $n$, a seed $r$ of $\mathcal{O}\left(\log ^{2} n\right)$ random bits, and a string $S \in \Sigma \leq n$, outputs a 3n-step complete walk $\mathrm{W}(n, r, S)$ over $S$ satisfying the following property for all $X, Y \in \Sigma \leq n$ and the zip alignment $\mathcal{A}_{\mathrm{W}}$ of $\mathrm{W}(n, r, X)$ and $\mathrm{W}(n, r, Y)$ :

$$
\operatorname{Pr}_{r}\left[\mathcal{A}_{\mathrm{W}} \in \mathrm{GA}_{c \cdot e d}(X, Y)^{2}(X, Y) \text { and } \text { width }\left(\mathcal{A}_{\mathrm{W}}\right) \leq c \cdot \operatorname{ed}(X, Y)\right] \geq 1-\delta .
$$

Moreover, W is an $\mathcal{O}\left(\log ^{2} n\right)$-bit streaming algorithm that costs $\mathcal{O}(n \log n)$ time and reports any element $s_{t} \in[1 \ldots|S|]$ of $\mathrm{W}(n, r, S)$ while processing the corresponding character $S\left[s_{t}\right]$.

Proof. Let $c_{6.6}$ and $n_{6.6}$ be the constant and threshold (respectively) of Corollary 6.6 for $\delta_{6.6}=\frac{\delta}{2}$.

If $n<\max \left(\frac{10}{\delta}, n_{6.6}\right)$, we set $\mathrm{W}(n, r, S)=(\min (t,|S|+1))_{t=1}^{3 n+1}$ for all $S \in \Sigma^{\leq n}$ so that $\mathrm{W}(n, r, S)$ is trivially a $3 n$-step complete walk over $S$. Now, consider strings $X, Y \in \Sigma^{\leq n}$ and the zip alignment $\mathcal{A}_{\mathrm{W}}$ of $\mathrm{W}(n, r, X)$ and $\mathrm{W}(n, r, Y)$. Observe that $\mathcal{A}_{\mathrm{W}} \in \mathrm{GA}_{0}(X, Y)$ if $X=Y$
and $\mathcal{A}_{\mathrm{w}} \in \mathrm{GA}_{n}(X, Y)$ otherwise. Moreover, $\operatorname{width}\left(\mathcal{A}_{\mathrm{W}}\right)=||X|-|Y|| \leq \operatorname{ed}(X, Y)$. Consequently, the claimed conditions are (deterministically) satisfied for $c \geq n_{6.6}$. The construction algorithm uses $\mathcal{O}(\log n)$ bits and $\mathcal{O}(n)$ time.

If $n \geq \max \left(\frac{10}{\delta}, n_{6.6}\right)$, we first develop an algorithm $\mathrm{W}^{\prime}$ that uses $\Theta(n \log \sigma)$ random bits, interpreted as a sequence $R \in \mathcal{H}^{3 n}$. Specifically, each hash function $R_{t}$ is specified by an $\lceil\log \sigma\rceil$-bit integer $h_{t}$ so that $R_{t}(a)$ counts (modulo 2) the set bits in $a$ xor $h_{t}$.

We set $\mathrm{W}^{\prime}(n, R, S)=\left(\max \left(t-3 n+|S|, s_{t}\right)\right)_{t=1}^{3 n+1}$ based on the CGK walk $W_{\mathrm{CGK}}(S)=\left(s_{t}\right)_{t=1}^{3 n+1}$ for all $S \in \Sigma^{\leq n}$. Observe that this modification guarantees that $\mathrm{W}^{\prime}(n, R, S)$ is a complete walk over $S$ and that $\mathrm{W}^{\prime}(n, R, S)=W_{\mathrm{CGK}}(S)$ if $W_{\mathrm{CGK}}(S)$ is already complete. Consequently, by Corollary 6.6, the claimed conditions are satisfied with probability at least $1-\delta_{6.6}=1-\frac{\delta}{2}$ for $c \geq c_{6.6}$. The construction algorithm uses $\mathcal{O}(\log n)$ bits, $\operatorname{costs} \mathcal{O}(n)$ time, reads the random bits $\left(h_{t}\right)_{t=1}^{3 n+1}$ from left to right, and outputs the elements of $\mathrm{W}^{\prime}(n, R, S)$ when required.

In order to use Nisan's pseudorandom generator [50], we need to argue that we only care about properties testable using $\mathcal{O}(\log n)$-bit algorithms with one-way access to the randomness $R$. As in [10], our testers are given two (read-only) strings $X, Y \in \Sigma^{\leq n}$ and a stream of random bits representing $R$ (the sequence $\left(h_{t}\right)_{t=1}^{3 n}$ ). Observe that the following properties of the zip alignment $\mathcal{A}^{\prime}$ of $\mathrm{W}^{\prime}(n, R, X)=\left(x_{t}\right)_{t=1}^{3 n+1}$ and $\mathrm{W}^{\prime}(n, R, Y)=\left(y_{t}\right)_{t=1}^{3 n+1}$ are testable in $\mathcal{O}(\log n)$ bits and $\mathcal{O}(n)$ time:

- whether $\mathcal{A}^{\prime}$ is greedy;
- whether $\operatorname{cost}\left(\mathcal{A}^{\prime}\right) \leq k$ for a given integer $k \in[0 \ldots n]$;
- whether $\operatorname{width}\left(\mathcal{A}^{\prime}\right) \leq w$ for a given integer $w \in[0 \ldots n]$.

All these testers simply construct triples $\left(t, x_{t}, y_{t}\right)$ for subsequent $t \in[1 \ldots 3 n+1]$.
In this setting, Nisan's pseudorandom generator [50], given a sequence $r$ of $\mathcal{O}\left(\log ^{2} n\right)$ random bits, constructs a sequence $\operatorname{PRG}(r)$ of pseudorandom bits in $\mathcal{O}(1)$ amortized time per bit, using $\mathcal{O}\left(\log ^{2} n\right)$ bits of working space. Moreover, for every tester, the probabilities of accepting a given input $\mathcal{I}$ with randomness $R$ and $\operatorname{PRG}(r)$ differ by at most $\frac{1}{n^{2}}$ [10, Theorem 5]. Consequently, setting $\mathrm{W}(n, r, S)=\mathrm{W}^{\prime}(n, \operatorname{PRG}(r), S)$, we can guarantee that the claimed condition is satisfied with probability at least $1-\frac{1}{2} \delta-\frac{(2 n+3)}{n^{2}} \geq 1-\delta$. The streaming construction algorithm takes $\mathcal{O}\left(\log ^{2} n\right)$ bits and $\mathcal{O}(n \log n)$ time, dominated by the generation of $\operatorname{PRG}(r)$.

### 6.2 Context encoding

For a string $S \in \Sigma^{*}$, define $\overline{\operatorname{maxLZ}}(S)=\max _{[\ell . r) \subseteq[1 .|S|]}|\mathrm{LZ}(\overline{S[\ell \ldots r)})|$.
Observation 6.7. Consider a string $S \in \Sigma^{*}$ and an integer $k \in \mathbb{Z}_{+}$. If $\overline{\operatorname{maxLZ}}(S[\ell \ldots r)) \leq k$ holds for a non-empty fragment $S[\ell \ldots r)$, then:
(a) $\overline{\operatorname{maxLZ}}(S[\ell . . r]) \leq k$ if and only if $\operatorname{LZ}(\overline{S[\ell . . r]}) \leq k$;
(b) $\overline{\operatorname{maxLZ}}(S[\ell+1 \ldots r)) \leq k$;
(c) $\overline{\operatorname{maxLZ}}(S[\ell \ldots]) \leq k+1$.

Proof. The first two parts follow from the fact that the size of the LZ-factorisation of a prefix of a string is bounded by the size of the LZ-factorisation of the string itself. The third claim follows from the optimality of the LZ77 parsing among all LZ-like parsings (Observation 4.1).

Definition 6.8 ((Double) Context). Consider a string $S \in \Sigma^{*}$ and an integer $k \in \mathbb{Z}_{+}$. For a position $p \in[1 . .|S|]$, we define the context $C_{k}(S, p)$ as the longest prefix $S[p \ldots q)$ of $S[p \ldots|S|]$ such that $\overline{\operatorname{maxLZ}}(S[p \ldots q)) \leq k$ and the double context $C_{k}^{2}(S, p)$ as the longest prefix $S[p \ldots q)$ of $S[p \ldots|S|]$ such that $\overline{\operatorname{maxLZ}}(S[p \ldots r)) \leq k$ and $\overline{\operatorname{maxLZ}}(S[r \ldots q)) \leq k$ for some $r \in[p \ldots q]$.

For integers $t<v$, let $\mu(t, v)$ denote the integer $w \in[t \ldots v)$ divisible by the largest power of 2 . We say that a position $S[s]$ is covered by a fragment $S[\ell \ldots r)$ if $\ell \leq s<r$.

```
Input: An integer \(k \in \mathbb{Z}_{\geq 0}\), a string \(S \in \Sigma^{*}\), and a complete walk \(\left(s_{t}\right)_{t=1}^{m+1}\) over \(S\).
Output: The string \(\mathrm{CE}_{k}(W)[1 \ldots m]\).
\(\mathrm{CE}_{k}(W):=\perp^{m}\), where \(\perp=(\mathrm{LZ}(\varepsilon), 0)\);
\(s_{0}:=0 ; \mu_{0}:=0 ; u:=0 ;\)
for \(t:=1\) to \(m\) do
    if \(s_{t} \leq|S|\) then
        \(\mu_{t}:=\mu\left(t, \min \left\{v \in[1 \ldots m+1]: s_{v}=s_{t}+\left|C_{k}\left(S, s_{t}\right)\right|\right\}\right) ;\)
        if \(\mu_{t}>\mu_{t-1}\) then
            \(\mathrm{CE}_{k}(W)[t]:=\left(\operatorname{LZ}\left(\overline{C_{k}^{2}\left(S, s_{t}\right)}\right), s_{t}-s_{u}\right) ;\)
        \(u:=t\);
return \(\mathrm{CE}_{k}(W)\);
```

Algorithm 3: Function CE

Lemma 6.9. Consider $\mathrm{CE}_{k}(W)$ constructed for an integer $k \in \mathbb{Z}_{\geq 0}$ and an m-complete walk $W=\left(s_{t}\right)_{t=1}^{m+1}$ over $S \in \Sigma^{*}$. Each position $s \in[1 . .|S|]$ satisfies

$$
\begin{aligned}
1 \leq \mid\{t \in[1 \ldots m]: & \left.\mathrm{CE}_{k}(W)[t] \neq \perp \text { and } C_{k}\left(S, s_{t}\right) \text { covers } S[s]\right\} \mid \leq \\
& \leq \mid\left\{t \in[1 \ldots m]: \mathrm{CE}_{k}(W)[t] \neq \perp \text { and } C_{k}^{2}\left(S, s_{t}\right) \text { covers } S[s]\right\} \mid \leq \mathcal{O}(\log m)
\end{aligned}
$$

Proof. Let us first bound the covering number from below. Consider an index $t \in[1 \ldots m]$ such that $s=s_{t}$, and let $t^{\prime} \in[1 \ldots m]$ be the smallest index such that $\mu_{t^{\prime}}=\mu_{t}$. Note that $s_{\mu_{t^{\prime}}}<s_{t^{\prime}}+\left|C_{k}\left(S, s_{t^{\prime}}\right)\right|$ and $s_{\mu_{t}} \geq s_{t}$ by definition of $\mu$ and monotonicity of the walk $W$. Due to $\mu_{t^{\prime}}=\mu_{t}$, this implies $s_{t^{\prime}} \leq s_{t}<s_{t^{\prime}}+\left|C_{k}\left(S, s_{t^{\prime}}\right)\right|$, i.e., that $S\left[s_{t}\right]$ is covered by $C_{k}\left(S, s_{t^{\prime}}\right)$. Due to $\mu_{t^{\prime}-1} \neq \mu_{t}$, this guarantees $\mathrm{CE}_{k}(S)\left[t^{\prime}\right] \neq \perp$.

As for the upper bound, note that the indexes $t \in[1 \ldots m]$ with $\mathrm{CE}_{k}(S)[t] \neq \perp$ have distinct values $\mu_{t}$. Let us further classify them into $\mathcal{O}(\log n)$ groups depending on the largest power of two dividing $\mu_{t}$. Consider two indexes $t<t^{\prime}$ in the same group. Since $\mu_{t}<\mu_{t^{\prime}}$ are divisible by the same largest power of two, there is a number $\nu \in\left(\mu_{t} \ldots \mu_{t^{\prime}}\right)$ divisible by a strictly larger power of two. By definition of $\mu_{t}$, we have $s_{\nu} \geq s_{t}+\left|C_{k}\left(S, s_{t}\right)\right|$ and, by definition of $\mu_{t^{\prime}}$, we have $s_{\nu}<s_{t^{\prime}}$. Consequently, $s_{t^{\prime}}>s_{t}+\left|C_{k}\left(S, s_{t}\right)\right|$, i.e., the contexts $C_{k}\left(S, s_{t}\right)$ and $C_{k}\left(S, s_{t^{\prime}}\right)$ are disjoint. By monotonicity of $\overline{\operatorname{maxLZ}}$, this also means that $s_{t}+\left|C_{k}^{2}\left(S, s_{t}\right)\right| \leq s_{t^{\prime}}+\left|C_{k}\left(S, s_{t^{\prime}}\right)\right|$. In particular, each position $s \in[1 \ldots|S|]$ is covered by at most one context $C_{k}\left(S, s_{t}\right)$ and at most two double contexts $C_{k}^{2}\left(S, s_{t}\right)$ for $t \in[1 \ldots m]$ belonging to a single group.

Let $\mathcal{A}$ be an alignment between $X, Y \in \Sigma^{*}$. We define $\mathcal{B}_{X}(\mathcal{A})=\{x:(x, y) \in \mathcal{B}(\mathcal{A})\} \cap$ $[1 \ldots|X|]$. Moreover, for two alignments $\mathcal{A}, \mathcal{A}^{\prime}$ between $X, Y \in \Sigma^{*}$, we define $\Delta_{X}\left(\mathcal{A}, \mathcal{A}^{\prime}\right)$ to contain $x \in[1 . .|X|]$ unless $(x, y) \in \mathcal{M}(\mathcal{A}) \cap \mathcal{M}\left(\mathcal{A}^{\prime}\right)$ for some $y \in[1 . .|Y|]$.

Lemma 6.10. Let $b \in \mathbb{Z}_{+}$and $\mathcal{A}, \mathcal{A}^{\prime}$ be greedy alignments of strings $X$ and $Y$. The positions in $\Delta_{X}\left(\mathcal{A}, \mathcal{A}^{\prime}\right)$ can be covered by at most $\operatorname{cost}(\mathcal{A})+\frac{1}{b} \operatorname{cost}\left(\mathcal{A}^{\prime}\right)$ contexts $C_{\text {width }(\mathcal{A})+\operatorname{width}\left(\mathcal{A}^{\prime}\right)+2 b}(X, \cdot)$.

Moreover, if $b>\operatorname{cost}\left(\mathcal{A}^{\prime}\right)$, then the positions in $\Delta_{X}\left(\mathcal{A}, \mathcal{A}^{\prime}\right)$ can be covered by contexts $C_{\text {width }(\mathcal{A})+\operatorname{width}\left(\mathcal{A}^{\prime}\right)+2 b}(X, x)$ with $x \in \mathcal{B}_{X}(\mathcal{A})$.

Proof. Let $w=\operatorname{width}(\mathcal{A})$ and $w^{\prime}=\operatorname{width}\left(\mathcal{A}^{\prime}\right)$. Without loss of generality, we may trim the longest common prefix of $X$ and $Y$; this is feasible because both $\mathcal{A}$ and $\mathcal{A}^{\prime}$ match the common prefix, so $\Delta_{X}\left(\mathcal{A}, \mathcal{A}^{\prime}\right)$ only contains positions following the prefix. Moreover, we assume that $X \neq \varepsilon$; otherwise $\Delta_{X}\left(\mathcal{A}, \mathcal{A}^{\prime}\right)=\emptyset$ and the lemma is trivial.

In the remaining case, we are guaranteed that $1 \in \mathcal{B}_{X}(\mathcal{A})$. Let $\mathcal{B}_{X}(\mathcal{A})=\left\{x_{1}, \ldots, x_{m}\right\}$, where $1=x_{1}<\cdots<x_{m}$, and let $x_{m+1}=|X|+1$. This yields a decomposition $X=X_{1} \cdots X_{m}$ into $m$ non-empty substrings $X_{i}:=X\left[x_{i} \ldots x_{i+1}\right)$. The alignment $\mathcal{A}^{\prime}$ further yields a decomposition $Y=Y_{1} \cdots Y_{m}$ into $m$ (possibly empty) substrings $Y_{i}:=Y\left[y_{i} \ldots y_{i+1}\right.$ ) so that $X_{i} \sim_{\mathcal{A}^{\prime}} Y_{i}$ and $\left|x_{i}-y_{i}\right| \leq w^{\prime}$ for $i \in[1 \ldots m]$. Moreover, the cost of $\mathcal{A}^{\prime}$ can be expressed as $\operatorname{cost}\left(\mathcal{A}^{\prime}\right)=\sum_{i=1}^{m} c_{i}^{\prime}$, where $c_{i}^{\prime}$ denotes the cost of $\mathcal{A}^{\prime}$ restricted to $X_{i}$ and $Y_{i}$.
Claim 6.11. For every $i \in[1 \ldots m]$, the set $\Delta_{X}\left(\mathcal{A}, \mathcal{A}^{\prime}\right) \cap\left[x_{i} \ldots x_{i+1}\right)$ can be covered by at most $\left\lceil\frac{1}{b}\left(c_{i}^{\prime}+1\right)\right\rceil$ contexts $C_{w+w^{\prime}+2 b}(X, \cdot)$ including $C_{w+w^{\prime}+2 b}\left(X, x_{i}\right)$.

Proof. Let $q_{i}=\max \left(\Delta_{X}\left(\mathcal{A}, \mathcal{A}^{\prime}\right) \cap\left[x_{i} \ldots x_{i+1}\right)\right)$. We may assume that $q_{i} \geq x_{i}+w+w^{\prime}$; otherwise, the claim holds trivially. We shall prove that $X\left[x_{i}+1 \ldots q_{i}-\left(w+w^{\prime}\right)\right]$ can be decomposed into at most $2 c_{i}^{\prime}+1$ phrases, each of which is a single character or has another occurrence at most $w+w^{\prime}$ positions to the right. Accounting for $X\left[x_{i}\right]$, this yields a decomposition of $X\left[x_{i} \ldots q_{i}-\right.$ $\left.\left(w+w^{\prime}\right)\right]$ into at most $2 c_{i}^{\prime}+2$ such phrases. The maxLZ measure of the concatenation of every $2 b$ subsequent phrases and $w+w^{\prime}$ following single characters does not exceed $w+w^{\prime}+2 b$. Hence, $X\left[x_{i} . . q_{i}\right]$ can be covered by $\left\lceil\frac{1}{b}\left(c_{i}^{\prime}+1\right)\right\rceil$ contexts $C_{w+w^{\prime}+2 b}(X, \cdot)$ including $C_{w+w^{\prime}+2 b}\left(X, x_{i}\right)$.

Thus, it remains to construct the decomposition of $X\left[x_{i}+1 \ldots q_{i}-\left(w+w^{\prime}\right)\right]$ into phrases. By definition of $\mathcal{B}_{X}(\mathcal{A})$, we have $X\left[x_{i}+1 \ldots x_{i+1}\right) \simeq_{\mathcal{A}} Y\left[x_{i}+1+d \ldots x_{i+1}+d\right)$ for some shift $d \in[-w \ldots w]$. We consider two cases.

In the first case, we assume that $\left(x_{i}+1, \bar{y}\right) \in \mathcal{A}^{\prime}$ holds for some $\bar{y} \geq x_{i}+1+d$. The greedy nature of $\mathcal{A}^{\prime}$ then guarantees that if $(x, y) \in \mathcal{A}^{\prime}$ with $x \in\left[x_{i}+1 \ldots q_{i}\right)$, then $y>x+d$. We decompose $X\left[x_{i}+1 \ldots q_{i}-\left(w+w^{\prime}\right)\right]$ (which is contained in $X_{i}$ ) into maximal phrases $X\left[x \ldots x^{\prime}\right]$ satisfying $X\left[x \ldots x^{\prime}\right] \simeq_{\mathcal{A}^{\prime}} Y\left[x+d^{\prime} \ldots x^{\prime}+d^{\prime}\right]$ for some $d^{\prime} \in\left(d \ldots w^{\prime}\right]$ and remaining single characters (deleted or substituted by $\mathcal{A}^{\prime}$ ). Observe that this yields at most $1+c_{i}^{\prime}$ phrases and at most $c_{i}^{\prime}$ single characters. Moreover, we have $X\left[x \ldots x^{\prime}\right] \simeq_{\mathcal{A}^{\prime}} Y\left[x+d^{\prime} \ldots x^{\prime}+d^{\prime}\right] \simeq_{\mathcal{A}}$ $X\left[x+\left(d^{\prime}-d\right) \ldots x^{\prime}+\left(d^{\prime}-d\right)\right]$ due to $x_{i}+1 \leq x \leq x+\left(d^{\prime}-d\right)$ and $x^{\prime}+\left(d^{\prime}-d\right) \leq x^{\prime}+\left(w^{\prime}+w\right) \leq q_{i}$. Hence, each phrase $X\left[x \ldots x^{\prime}\right]$ has another occurrence located $d^{\prime}-d \in\left[1 \ldots w+w^{\prime}\right]$ positions to the right.

In the second case, we assume that $\left(x_{i}+1, \bar{y}\right) \in \mathcal{A}^{\prime}$ holds for some $\bar{y} \leq x_{i}+1+d$. The greedy nature of $\mathcal{A}^{\prime}$ then guarantees that if $(x, y) \in \mathcal{A}^{\prime}$ with $x \in\left[x_{i}+1 \ldots q_{i}\right)$, then $y<x+d$. We decompose $Y\left[x_{i}+1+d \ldots q_{i}-\left(w+w^{\prime}\right)+d\right]$ (which is contained in $Y_{i}$ ) into maximal phrases $Y\left[y \ldots y^{\prime}\right]$ satisfying $Y\left[y \ldots y^{\prime}\right] \simeq_{\mathcal{A}^{\prime}} X\left[y-d^{\prime} \ldots y^{\prime}-d^{\prime}\right]$ for some $d^{\prime} \in[-w \ldots d)$ and remaining single characters (deleted or substituted by $\mathcal{A}^{\prime}$ ). Observe that this yields at most $1+c_{i}^{\prime}$ phrases and at most $c_{i}^{\prime}$ single characters. Moreover, we have $Y\left[y \ldots y^{\prime}\right] \simeq_{\mathcal{A}^{\prime}} X\left[y-d^{\prime} \ldots y^{\prime}-d^{\prime}\right] \simeq_{\mathcal{A}} Y[y+(d-$ $\left.\left.d^{\prime}\right) \ldots y+\left(d-d^{\prime}\right)\right]$ due to $x_{i}+1 \leq x_{i}+1+d-d^{\prime} \leq y-d^{\prime}$ and $y^{\prime}-d^{\prime} \leq q_{i}-\left(w+w^{\prime}\right)+d-d^{\prime} \leq q_{i}$. Hence, each phrase $Y\left[y \ldots y^{\prime}\right]$ has another occurrence located $d-d^{\prime} \in\left[1 \ldots w+w^{\prime}\right]$ characters to the right. Since $Y\left[x_{i}+1+d \ldots q_{i}+d\right]=X\left[x_{i}+1 \ldots q_{i}\right]$, this decomposition of $Y\left[x_{i}+1+d \ldots q_{i}-\left(w+w^{\prime}\right)+d\right]$ gives an analogous decomposition of $X\left[x_{i}+1 \ldots q_{i}-\left(w+w^{\prime}\right)\right]$.

The set $\Delta_{X}\left(\mathcal{A}, \mathcal{A}^{\prime}\right)$ can be covered by $\sum_{i=1}^{m}\left\lceil\frac{1}{b}\left(c_{i}^{\prime}+1\right)\right\rceil \leq m+\frac{\operatorname{cost}\left(\mathcal{A}^{\prime}\right)}{b} \leq \operatorname{cost}(\mathcal{A})+\frac{\operatorname{cost}\left(\mathcal{A}^{\prime}\right)}{b}$ contexts $C_{w+w^{\prime}+2 b}(X, \cdot)$. If $b \geq \operatorname{cost}\left(\mathcal{A}^{\prime}\right)+1$, then the contexts starting at positions in $\mathcal{B}_{X}(\mathcal{A})$ are sufficient.

Lemma 6.12. Let $\mathcal{A}_{W}$ be the zip alignment of m-complete walks $W_{X}=\left(x_{t}\right)_{t=1}^{m+1}$, $W_{Y}=\left(y_{t}\right)_{t=1}^{m+1}$ over strings $X, Y \in \Sigma^{*}$, and let $k^{\prime} \in \mathbb{Z}_{+}$. If $\mathcal{A}_{W}$ is greedy, then $\mathcal{B}_{X}\left(\mathcal{A}_{W}\right) \subseteq P_{X}$, where $P_{X}$ is the set of positions $x \in[1 \ldots|X|]$ such that $X[x]$ is covered by $C_{k^{\prime}}^{2}\left(X, x_{t}\right)$ for some $t \in[1 \ldots m]$ with $\perp \neq \mathrm{CE}_{k^{\prime}}\left(W_{X}\right)[t] \neq \mathrm{CE}_{k^{\prime}}\left(W_{Y}\right)[t]$. Moreover, for every positive integer $k \geq \operatorname{ed}(X, Y)$ such that $k^{\prime} \geq$ width $\left(\mathcal{A}_{W}\right)+5 k$ :

- every $\mathcal{A} \in \mathrm{GA}_{k}(X, Y)$ satisfies $\Delta_{X}\left(\mathcal{A}, \mathcal{A}_{W}\right) \subseteq P_{X}$,
- $\operatorname{hd}\left(\mathrm{CE}_{k^{\prime}}(X), \mathrm{CE}_{k^{\prime}}(Y)\right) \leq c\left(k+\frac{1}{k} \operatorname{cost}\left(\mathcal{A}_{W}\right)\right) \log m$ holds for a sufficiently large constant $c$.

Proof. Consider a position $x \in \mathcal{B}_{X}\left(\mathcal{A}_{W}\right)$. By Lemma 6.9, there is an index $t \in[1 \ldots m]$ such that $\mathrm{CE}_{k^{\prime}}\left(W_{X}\right)[t] \neq \perp$ and $C_{k^{\prime}}^{2}\left(X, x_{t}\right)$ covers $X[x]$. To conclude that $x \in P_{X}$, it remains to prove $\mathrm{CE}_{k^{\prime}}\left(W_{X}\right)[t] \neq \mathrm{CE}_{k^{\prime}}\left(W_{Y}\right)[t]$. For a proof by contradiction, suppose that $\mathrm{CE}_{k^{\prime}}\left(W_{X}\right)[t]=$ $\mathrm{CE}_{k^{\prime}}\left(W_{Y}\right)[t]$. This implies $C_{k^{\prime}}^{2}\left(X, x_{t}\right)=C_{k^{\prime}}^{2}\left(Y, y_{t}\right)$. Since $\mathcal{A}_{W}$ is greedy, we derive $X\left[x_{t} \ldots x_{t}+\right.$ $\left.\left|C_{k^{\prime}}^{2}\left(X, x_{t}\right)\right|\right) \simeq_{\mathcal{A}_{W}} Y\left[y_{t} . . y_{t}+\left|C_{k^{\prime}}^{2}\left(Y, y_{t}\right)\right|\right)$, which contradicts $x \in \mathcal{B}_{X}\left(\mathcal{A}_{W}\right)$.

In the remainder of the proof, we assume that $k \geq e d(X, Y)$ and $k^{\prime} \geq \operatorname{width}\left(\mathcal{A}_{W}\right)+5 k$. Next, consider a greedy alignment $\mathcal{A}$ with $\operatorname{cost}(\mathcal{A}) \leq k$ and a position $x \in \Delta_{X}\left(\mathcal{A}, \mathcal{A}_{W}\right)$. Due to $k^{\prime} \geq \operatorname{width}\left(\mathcal{A}_{W}\right)+\operatorname{width}(\mathcal{A})+4 k$, Lemma 6.10 shows that there exists a position $r \in \mathcal{B}_{X}\left(\mathcal{A}_{W}\right)$ such that $C_{k^{\prime}}(X, r)$ covers $X[x]$ and, by Lemma 6.9 , there exists a position $t \in[1 \ldots m]$ such that $\mathrm{CE}_{k^{\prime}}(X)[t] \neq \perp$ and $C_{k^{\prime}}\left(X, x_{t}\right)$ covers $X[r-1]$. Since $\overline{\operatorname{maxLZ}}\left(X\left[x_{t} \ldots r\right)\right), \overline{\operatorname{maxLZ}}(X[r \ldots x]) \leq k^{\prime}$, we conclude that both $r$ and $x$ are covered by $C_{k^{\prime}}^{2}\left(X, x_{t}\right)$. We shall prove that $\mathrm{CE}_{k^{\prime}}(X)[t] \neq$ $\mathrm{CE}_{k^{\prime}}(Y)[t]$. For a proof by contradiction, suppose that $\mathrm{CE}_{k^{\prime}}\left(W_{X}\right)[t]=\mathrm{CE}_{k^{\prime}}\left(W_{Y}\right)[t]$, which implies $C_{k^{\prime}}^{2}\left(X, x_{t}\right)=C_{k^{\prime}}^{2}\left(Y, y_{t}\right)$. Let us fix the smallest $t^{\prime} \in[t \ldots m+1]$ such that $\left(x_{t^{\prime}}, y_{t^{\prime}}\right) \in \mathcal{B}\left(\mathcal{A}_{W}\right)$. Observe that $X\left[x_{t} \ldots x_{t^{\prime}}\right) \simeq_{\mathcal{A}_{W}} Y\left[y_{t} \ldots y_{t^{\prime}}\right)$ and, by the greedy nature of $\mathcal{A}_{W}, X\left[x_{t} \ldots x_{t^{\prime}}\right)=$ $Y\left[y_{t} \ldots y_{t^{\prime}}\right)$ is the longest common prefix of $X\left[x_{t} \ldots\right]$ and $Y\left[y_{t} \ldots\right]$. At the same time, due to $x_{t}<r \in \mathcal{B}_{X}\left(\mathcal{A}_{W}\right)$, we have $x_{t^{\prime}} \leq r$, so $X\left[x_{t} \ldots x_{t^{\prime}}\right]$ is a prefix of $C_{k^{\prime}}^{2}\left(X, x_{t}\right)$. However, $X\left[x_{t} \ldots x_{t^{\prime}}\right]$ is not a prefix $Y\left[y_{t} ..\right]$, so $C_{k^{\prime}}^{2}\left(X, x_{t}\right)$ is not a prefix of $Y\left[y_{t} ..\right]$ and $C_{k^{\prime}}^{2}\left(X, x_{t}\right) \neq C_{k^{\prime}}^{2}\left(Y, y_{t}\right)$. The contradiction completes the proof.

Finally, we bound $\operatorname{hd}\left(\mathrm{CE}_{k^{\prime}}(X), \mathrm{CE}_{k^{\prime}}(Y)\right)$ using several claims. Consider a set $M$ of indices $t \in[1 \ldots m]$ such that $\mathrm{CE}_{k^{\prime}}(X)[t]$ and $\mathrm{CE}_{k^{\prime}}(Y)[t]$ differ on the first coordinate.
Claim 6.13. If $t \in M$ satisfies $\mathrm{CE}_{k^{\prime}}(X)[t] \neq \perp$, then

$$
\left[x_{t-1} \ldots x_{t}+\left|C_{k^{\prime}}^{2}\left(X, x_{t}\right)\right|\right] \cap\left(\mathcal{B}_{X}\left(\mathcal{A}_{W}\right) \cup\{|X|+1\}\right) \neq \emptyset .
$$

Proof. For a proof by contradiction, suppose that $\left[x_{t-1} \ldots x_{t}+\left|C_{k^{\prime}}^{2}\left(X, x_{t}\right)\right|\right] \cap\left(\mathcal{B}_{X}\left(\mathcal{A}_{W}\right) \cup\{|X|+\right.$ $1\})=\emptyset$. Let $L$ be the length of the longest common prefix of $X\left[x_{t} ..\right]$ and $Y\left[y_{t} ..\right]$. Note that $\left(x_{t}+L, y_{t}+L\right) \in \mathcal{B}\left(\mathcal{A}_{W}\right)$, so $x_{t}+L \in \mathcal{B}_{X}\left(\mathcal{A}_{W}\right) \cup\{|X|+1\}$. Consequently, $L>\left|C_{k^{\prime}}^{2}\left(X, x_{t}\right)\right| \geq$ $\left|C_{k^{\prime}}\left(X, x_{t}\right)\right|$, and therefore $X\left[x_{t} \ldots x_{t}+\left|C_{k^{\prime}}^{2}\left(X, x_{t}\right)\right|\right]=Y\left[y_{t} \ldots y_{t}+\left|C_{k^{\prime}}^{2}\left(Y, y_{t}\right)\right|\right]$ and $X\left[x_{t} \ldots x_{t}+\right.$ $\left.\left|C_{k^{\prime}}\left(X, x_{t}\right)\right|\right]=Y\left[y_{t} \ldots y_{t}+\left|C_{k^{\prime}}\left(Y, y_{t}\right)\right|\right]$. In particular, $C_{k^{\prime}}^{2}\left(X, x_{t}\right)=C_{k^{\prime}}^{2}\left(Y, y_{t}\right)$.

If $t=1$, then we note that $\mathrm{CE}_{k^{\prime}}(X)[t] \neq \perp \neq \mathrm{CE}_{k^{\prime}}(Y)[t]$, so $C_{k^{\prime}}^{2}\left(X, x_{t}\right)=C_{k^{\prime}}^{2}\left(Y, y_{t}\right)$ contradicts $t \in M$.

In the following, we assume that $t \geq 2$. Due to $\left(x_{t-1}, y_{t-1}\right) \notin \mathcal{B}\left(\mathcal{A}_{W}\right)$, either $\left(x_{t-1}, y_{t-1}\right)=$ $\left(x_{t}, y_{t}\right)$, or $\left(x_{t-1}, y_{t-1}\right)=\left(x_{t}-1, y_{t}-1\right)$ and $X\left[x_{t}-1\right]=Y\left[y_{t}-1\right]$. In either case, we have $X\left[x_{t-1} \ldots x_{t-1}+\left|C_{k^{\prime}}\left(X, x_{t-1}\right)\right|\right]=Y\left[y_{t-1} \ldots y_{t-1}+\left|C_{k^{\prime}}\left(Y, y_{t-1}\right)\right|\right]$ due to $X\left[x_{t-1} \ldots x_{t}+\right.$ $\left.\left|C_{k^{\prime}}\left(X, x_{t}\right)\right|\right]=Y\left[y_{t-1} \ldots y_{t}+\left|C_{k^{\prime}}\left(Y, y_{t}\right)\right|\right]$ and by Observation 6.7. Consequently, $\min \{u \in$ $\left.[1 \ldots m]: x_{u}=x_{t-1}+\left|C_{k^{\prime}}\left(X, x_{t-1}\right)\right|\right\}=\min \left\{u \in[1 \ldots m]: y_{u}=y_{t-1}+\left|C_{k^{\prime}}\left(Y, y_{t-1}\right)\right|\right\}$ and, by a similar reasoning, $\min \left\{u \in[1 \ldots m]: x_{u}=x_{t}+\left|C_{k^{\prime}}\left(X, x_{t}\right)\right|\right\}=\min \left\{u \in[1 . . m]: y_{u}=\right.$ $\left.y_{t}+\left|C_{k^{\prime}}\left(Y, y_{t}\right)\right|\right\}$. Hence, the assumption $\mathrm{CE}_{k^{\prime}}(X)[t] \neq \perp$ implies $\mathrm{CE}_{k^{\prime}}(Y)[t] \neq \perp$. Therefore, $C_{k^{\prime}}^{2}\left(X, x_{t}\right)=C_{k^{\prime}}^{2}\left(Y, y_{t}\right)$ contradicts $t \in M$.
Claim 6.14. There are $\mathcal{O}\left(\left(k+\frac{1}{k} \operatorname{cost}\left(\mathcal{A}_{W}\right)\right) \log m\right)$ positions $t \in M$ such that $\mathrm{CE}_{k^{\prime}}(X)[t] \neq \perp$.
Proof. Let $\mathcal{O}$ be an optimum greedy alignment between $X$ and $Y$. Due to $k^{\prime} \geq \operatorname{width}(\mathcal{O})+$ width $\left(\mathcal{A}_{W}\right)+4 k$, from Lemma 6.10 it follows that $\Delta_{X}\left(\mathcal{O}, \mathcal{A}_{W}\right)$ can be covered by $\mathcal{O}(\operatorname{cost}(\mathcal{O})+$ $\left.\frac{1}{k} \operatorname{cost}\left(\mathcal{A}_{W}\right)\right)=\mathcal{O}\left(k+\frac{1}{k} \operatorname{cost}\left(\mathcal{A}_{W}\right)\right)$ contexts $C_{k^{\prime}}(X, \cdot)$. Let us fix such the smallest family $\mathcal{C}$ of contexts covering $\mathcal{B}_{X}\left(\mathcal{A}_{W}\right) \subseteq \Delta_{X}\left(\mathcal{O}, \mathcal{A}_{W}\right)$.

Define a set $R_{X}=\{|X|\} \cup \bigcup_{X[q . . r] \in \mathcal{C}}\{q-1, r+1\}$ and note that $\left|R_{X}\right|=\mathcal{O}(|\mathcal{C}|)$. We shall prove that, if $t \in M$ and $\mathrm{C}_{k^{\prime}}(X)[t] \neq \perp$, then $C_{k^{\prime}}^{2}\left(X, x_{t}\right)$ covers at least one position in $R_{X}$.

This is sufficient to derive the claim because, by Lemma 6.9, every position in $X$ is covered by at most $\mathcal{O}(\log m)$ double contexts $C_{k^{\prime}}^{2}\left(X, x_{t}\right)$ with $\mathrm{CE}_{k^{\prime}}(X)[t] \neq \perp$.

By Claim 6.13, we have $\left[x_{t-1} \ldots x_{t}+\left|C_{k^{\prime}}^{2}\left(X, x_{t}\right)\right|\right] \cap\left(\mathcal{B}_{X}\left(\mathcal{A}_{W}\right) \cup\{|X|+1\}\right) \neq \emptyset$. If $\{|X|+1\} \in$ $\left[x_{t-1} \ldots x_{t}+\left|C_{k^{\prime}}^{2}\left(X, x_{t}\right)\right|\right]$, then $|X|+1=x_{t}+\left|C_{k^{\prime}}^{2}\left(X, x_{t}\right)\right|$. Hence, $C_{k^{\prime}}^{2}\left(X, x_{t}\right)$ covers the position $|X| \in R_{X}$. Thus, we may assume that $\left[x_{t-1} \ldots x_{t}+\left|C_{k^{\prime}}^{2}\left(X, x_{t}\right)\right|\right]$ contains a position in $\mathcal{B}_{X}\left(\mathcal{A}_{W}\right)$. Note that the fragment $X[q \ldots r] \in \mathcal{C}$ covering that position in $\mathcal{B}_{X}\left(\mathcal{A}_{W}\right)$ satisfies $r \geq x_{t-1}$ and $q \leq x_{t}+\left|C_{k^{\prime}}^{2}\left(X, x_{t}\right)\right|$. In particular, $r+1 \geq x_{t}$ and $q-1<x_{t}+\left|C_{k^{\prime}}^{2}\left(X, x_{t}\right)\right|$. Now, if $C_{k^{\prime}}^{2}\left(X, x_{t}\right)$ covers $q-1$ or $r+1$, then we are done. Otherwise, $q-1<x_{t}$ and $r+1 \geq x_{t}+\left|C_{k^{\prime}}^{2}\left(X, x_{t}\right)\right|$, so $q \leq x_{t}$ and $r \geq x_{t}+\left|C_{k^{\prime}}^{2}\left(X, x_{t}\right)\right|$, i.e., $C_{k^{\prime}}^{2}\left(X, x_{t}\right)$ is contained in $X[q \ldots r]$ and, since $\overline{\operatorname{maxLZ}}$ is monotone, $\overline{\operatorname{maxLZ}}\left(C_{k^{\prime}}^{2}\left(X, x_{t}\right)\right) \leq k^{\prime}$. Consequently, $C_{k^{\prime}}^{2}\left(X, x_{t}\right)=C_{k^{\prime}}\left(X, x_{t}\right)$ is a suffix of $X$, so $C_{k^{\prime}}^{2}\left(X, x_{t}\right)$ covers the position $|X| \in R_{X}$.

A symmetric argument shows that there are $\mathcal{O}\left(\left(k+\frac{1}{k} \operatorname{cost}\left(\mathcal{A}_{W}\right)\right) \log m\right)$ positions $t \in M$ such that $\mathrm{CE}_{k^{\prime}}(Y)[t] \neq \perp$. Consequently, $|M|=\mathcal{O}\left(\left(k+\frac{1}{k} \operatorname{cost}\left(\mathcal{A}_{W}\right)\right) \log m\right)$. We bound $\mathrm{hd}\left(\mathrm{CE}_{k^{\prime}}(X), \mathrm{CE}_{k^{\prime}}(Y)\right)$ using the following claim.
Claim 6.15. hd $\left(\mathrm{CE}_{k^{\prime}}(X), \mathrm{CE}_{k^{\prime}}(Y)\right) \leq 2|M|$.
Proof. Let $M^{\prime}=\left\{t \in[1 \ldots m]: \mathrm{CE}_{k^{\prime}}(X)[t] \neq \mathrm{CE}_{k^{\prime}}(Y)[t]\right\}$. Note that $M \subseteq M^{\prime}$, so the claim is equivalent to $\left|M^{\prime} \backslash M\right| \leq|M|$. For every $t \in M^{\prime} \backslash M$, we shall prove that $t$ is not the leftmost position in $M^{\prime}$ and that the preceding position in $M^{\prime}$ belongs to $M$.

The assumption $t \in M^{\prime} \backslash M$ implies $\mathrm{CE}_{k^{\prime}}(X)[t] \neq \perp \neq \mathrm{CE}_{k^{\prime}}(Y)[t]$. We define $t^{\prime}$ as the largest position in $[1 \ldots t)$ such that $\mathrm{CE}_{k^{\prime}}(X)\left[t^{\prime}\right] \neq \perp$ or $\mathrm{CE}_{k^{\prime}}(Y)\left[t^{\prime}\right] \neq \perp$. To see that $t^{\prime}$ is well defined, note that $t>1$ and $\mathrm{CE}_{k^{\prime}}(X)[1] \neq \perp$.

Now, suppose that $t^{\prime} \notin M$. Consequently, we have $\mathrm{CE}_{k^{\prime}}(X)\left[t^{\prime}\right] \neq \perp \neq \mathrm{CE}_{k^{\prime}}(Y)\left[t^{\prime}\right]$. Due to $\mathrm{CE}_{k^{\prime}}(X)\left[t^{\prime}+1 \ldots t\right)=\perp^{t-t^{\prime}-1}$, Lemma 6.9 implies that $C_{k^{\prime}}\left(X, x_{t^{\prime}}\right)$ covers $X\left[x_{t}-1\right]$, that is, $\overline{\operatorname{maxLZ}}\left(X\left[x_{t^{\prime}} \ldots x_{t}\right)\right) \leq k^{\prime}$, and hence $X\left[x_{t^{\prime}} \ldots x_{t}\right]$ is a prefix of $C_{k^{\prime}}^{2}\left(X, x_{t^{\prime}}\right)$. Moreover, $C_{k^{\prime}}^{2}\left(X, x_{t^{\prime}}\right)=$ $C_{k^{\prime}}^{2}\left(Y, y_{t^{\prime}}\right)$ implies $X\left[x_{t^{\prime}} . . x_{t}\right] \simeq_{\mathcal{A}_{W}} Y\left[y_{t^{\prime}} \ldots y_{t^{\prime}}+x_{t}-x_{t^{\prime}}\right]$ by the greedy nature of $\mathcal{A}_{W}$. As $\left(x_{t}, y_{t}\right) \in \mathcal{A}_{W}$, we conclude that $y_{t}=y_{t^{\prime}}+x_{t}-x_{t^{\prime}}$. At the same time, since $\mathrm{CE}_{k^{\prime}}(X)\left[t^{\prime}+1 \ldots t\right)=$ $\perp^{t-t^{\prime}-1}=\mathrm{CE}_{k^{\prime}}(Y)\left[t^{\prime}+1 . . t\right)$, we have $\mathrm{CE}_{k^{\prime}}(X)[t]=\left(\mathrm{LZ}\left(\overline{C_{k^{\prime}}^{2}\left(X, x_{t}\right)}\right), x_{t}-x_{t^{\prime}}\right)$ and $\mathrm{CE}_{k^{\prime}}(Y)[t]=$ $\left(\mathrm{LZ}\left(\overline{C_{k^{\prime}}^{2}\left(Y, y_{t}\right)}\right), y_{t}-y_{t^{\prime}}\right)$. Thus, $\mathrm{CE}_{k^{\prime}}(X)[t]=\mathrm{CE}_{k^{\prime}}(Y)[t]$, which contradicts $t \in M^{\prime}$.

Consequently, $t^{\prime} \in M$. In this case, $M^{\prime} \cap\left[t^{\prime} . . t\right]=\left\{t^{\prime}, t\right\}$, so $t^{\prime}$ is the position of $M^{\prime}$ preceding $t$. Our goal was to show that such a position exists and belongs to $M$, so this completes the proof of the claim.

Overall, we conclude that hd $\left(\mathrm{CE}_{k^{\prime}}(X), \mathrm{CE}_{k^{\prime}}(Y)\right) \leq 2|M|=\mathcal{O}\left(\left(k+\frac{1}{k} \operatorname{cost}\left(\mathcal{A}_{W}\right)\right) \log m\right)$.
Lemma 6.16. Given integers $n \geq k \geq 1$, a seed $r$ of $\mathcal{O}\left(\log ^{2} n\right)$ random bits, and streaming access to a string $S \in \Sigma^{\leq n}$, the string $\mathrm{CE}_{k}(\mathrm{~W}(S, n, r))$ can be computed in $\tilde{\mathcal{O}}(k)$ space and $\tilde{\mathcal{O}}(n k)$ time.

Proof. Let us start with an auxiliary subroutine:
Claim 6.17. There is a streaming algorithm that computes $\mathrm{D}\left(C_{k}(S, p)\right)$ for subsequent positions $p \in[1 . .|S|]$. The algorithm uses $\tilde{\mathcal{O}}(k)$ space and $\tilde{\mathcal{O}}(n k)$ time. Moreover, if $C_{k}(S, p)=S[p \ldots q)$, then $\mathrm{D}\left(C_{k}(S, p)\right)$ is reported while the algorithm processes $S[q]$ (or the end-of-string token if $q=|S|+1)$.

Proof. The algorithm maintains a fragment $S[p \ldots q)$ satisfying $\overline{\operatorname{maxLZ}}(S[p . . q)) \leq k$ and the encoding $\mathrm{D}(S[p \ldots q))$. Initially, $S[p \ldots q)=S[1 \ldots 1)=\varepsilon$.

In each iteration, we read $S[q]$, compute $\mathrm{D}(S[p \ldots q]$ ) (using Proposition 4.2(i)), and check whether $|\operatorname{LZ}(\overline{S[p . . q]})| \leq k$ (using Proposition 4.2(f)). By Observation 6.7(a), this condition
is equivalent to $\overline{\operatorname{maxLZ}}(S[p \ldots q]) \leq k$. If $\overline{\operatorname{maxLZ}}(S[p \ldots q]) \leq k$, we discard $\mathrm{D}(S[p \ldots q))$ and increment $q$. Otherwise, we are guaranteed that $S[p \ldots q)=C_{k}(S, p)$, so we output $\mathrm{D}(S[p \ldots q)$ ), compute $\mathrm{D}(S[p+1 \ldots q)$ ) (using Proposition $4.2(\mathrm{~g})$ ), discard $\mathrm{D}(S[p \ldots q)$ ) and $\mathrm{D}(S[p \ldots q]$ ), and increment $p$. By Observation $6.7(\mathrm{~b})$, we are guaranteed that the invariant $\overline{\operatorname{maxLZ}}(S[p \ldots q)) \leq k$ remains satisfied. In the special case of $q=|S|+1$, we proceed as if $\overline{\operatorname{maxLZ}}(S[p \ldots q])>k$.

By Observation $6.7(\mathrm{c})$, we store $\mathrm{D}(X)$ only for strings $X$ satisfying $|\overline{\operatorname{maxLZ}}(X)| \leq k+1$, so the space usage and the per-iteration running time is $\tilde{\mathcal{O}}(k)$. Each iteration increments either $p$ or $q$, so the algorithm reads $S$ in a streaming fashion and the amortized running time is $\tilde{\mathcal{O}}(k)$ per character.

We maintain two instances of the algorithm of Claim 6.17 and two instances of the algorithm of Proposition 3.16. We feed the first instance of Claim 6.17 with the input stream $S$, obtaining $\mathrm{D}\left(C_{k}(S, q)\right)$ for subsequent positions $q \in[1 \ldots|S|]$. Upon retrieving $\mathrm{D}\left(C_{k}(S, q)\right)$, we extract $S[q]$ using Proposition $4.2(\mathrm{~b})$ and forward $S[q]$ to the first instance of Proposition 3.16, which lists indices $v \in[1 \ldots 3 n]$ such that $s_{v}=q$. We also forward $S[q]$ to the second instance of Claim 6.17, obtaining $\mathrm{D}\left(C_{k}(S, p)\right)$ for all subsequent positions $p$ such that $\mathrm{D}\left(C_{k}(S, p)\right)=S[p \ldots q)$. Upon retrieving $\mathrm{D}\left(C_{k}(S, p)\right)$, we extract $S[p]$ using Proposition $4.2(\mathrm{~b})$ and forward $S[p]$ to the second instance of Proposition 3.16, which lists indices $t \in[1 \ldots 3 n]$ such that $s_{t}=q$. For each such position $t$, we have $C_{k}^{2}\left(S, s_{t}\right)=C_{k}(S, p) C_{k}(S, q)$. We compute $\mu_{t}=\mu\left(t, \min \left\{v \in[1 \ldots 3 n]: s_{v}=\right.\right.$ $q\}$ ) based on the output of the first instance of Proposition 3.16. If $\mu_{t}>\mu_{t-1}$, we construct $\mathrm{LZ}\left(\overline{C_{k}^{2}\left(S, s_{t}\right)}\right)$ using Proposition $4.2(\mathrm{i})$ and Proposition 4.2(f).

By Propositions 4.2 and 3.16 and Claim 6.17, the algorithm uses $\tilde{\mathcal{O}}(k)$ space and $\tilde{\mathcal{O}}(n k)$ time.

### 6.3 Applications of Hamming distance sketches

Let us start by reminding the fingerprints (sketches) for testing string equality.
Fact 6.18 (see e.g. [38]). There exists a fingerprint $\psi$ (parameterized by an integer $n \in \mathbb{Z}_{+}$, a threshold $\delta$ with $1 \geq \delta \geq n^{-\mathcal{O}(1)}$, an alphabet $\Sigma=\left[0 \ldots n^{\mathcal{O}(1)}\right)$, and a seed of $\mathcal{O}(\log n)$ random bits) such that:

1. The fingerprint $\psi(S)$ of a string $S \in \Sigma^{\leq n}$ takes $\mathcal{O}\left(\log \delta^{-1}\right)$ bits. Given streaming access to $S$, it can be constructed in $\mathcal{O}(|S|)$ time using $\mathcal{O}(\log n)$ bits of space.
2. For all strings $X, Y \in \Sigma^{\leq n}$, we have $\operatorname{Pr}[\psi(X)=\psi(Y)] \leq \delta$ if $X \neq Y$ (and $\psi(X)=\psi(Y)$ otherwise).

For two equal-length strings $X, Y$, the set of mismatch positions is defined as $\operatorname{MP}(X, Y)=$ $\{i \in[1 \ldots|X|]: X[i] \neq Y[i]\}$ and the mismatch information $\operatorname{MI}(X, Y)=\{(i, X[i], Y[i]): i \in$ $\mathrm{MP}(X, Y)\}$. Below, we adapt the Hamming sketches of [16] to large alphabets.

Theorem 6.19. For every constant $\delta \in(0,1)$, there exists a sketch $\mathrm{sk}_{k}^{H}$ (parameterized by integers $n \geq k \geq 1$, an alphabet $\Sigma=[0 \ldots \sigma)$, and a seed of $\mathcal{O}(\log (n \log \sigma))$ random bits) such that:

1. The sketch $\mathrm{sk}_{k}^{H}(S)$ of a string $S \in \Sigma^{n}$ takes $\mathcal{O}(k \log (n \sigma))$ bits. Given streaming access to $S$, it can be constructed in $\mathcal{O}(n \log (n \sigma) \log (n \log \sigma))$ time using $\mathcal{O}(k \log (n \sigma))$ bits of space.
2. There exists a decoding algorithm that, given $\mathrm{sk}_{k}^{H}(X)$ and $\mathrm{sk}_{k}^{H}(Y)$ for strings $X, Y \in \Sigma^{n}$, with probability at least $1-\delta$ either returns $\mathrm{MI}(X, Y)$ or certifies that $\mathrm{hd}(X, Y)>k$. The algorithm uses $\mathcal{O}\left(k \log (n \sigma) \log ^{2}(n \log \sigma)\right)$ time and $\mathcal{O}(k \log (n \sigma))$ bits of space.

Proof. The construction of [16] satisfies the required conditions provided that $\sigma=n^{\mathcal{O}(1)}$. Henceforth, we assume without loss of generality that $\sigma$ is a power of two satisfying $\sigma \geq n \log \sigma$.

We interpret each character in $\Sigma$ as a block of $b=\left\lceil\log _{n \log \sigma} \sigma\right\rceil$ characters in $[0 \ldots n \log \sigma)$. Moreover, we consider the fingerprints $\psi$ of Fact 6.18 with $\delta_{6.18}=\frac{\delta}{2 n}$ and $n_{6.18}=\sigma_{6.18}=$ $n \log \sigma$. This construction uses $\mathcal{O}\left(\log n_{6.18}\right)=\mathcal{O}(\log (n \log \sigma))$ random bits and produces fingerprints of $\mathcal{O}\left(\log \delta_{6.18}^{-1}\right)=\mathcal{O}(\log n)$ bits.

Given a string $S \in \Sigma^{n}$, we define a string $\bar{S}[1 \ldots n b]$ so that $\bar{S}[i b-j+1]=(S[i][j], \psi(S[i]))$ for $i \in[1 \ldots|S|]$ and $j \in[1 \ldots b]$. Consequently, we set $\operatorname{sk}_{k}^{H}(S):=\operatorname{sk}_{\bar{k}}^{H}(\bar{S})$ using the the sketch of [16] with $\bar{\delta}=\frac{\delta}{2}, \bar{n}=n b, \bar{k}=k b$, and $\bar{\sigma} \leq n^{\mathcal{O}(1)} \log \sigma$. Note that this is feasible since $\bar{\sigma} \leq n^{\mathcal{O}(1)} \leq \bar{n}^{\mathcal{O}(1)}$ if $n \leq \log \sigma$ and $\bar{\sigma} \leq \log ^{\mathcal{O}(1)} \sigma \leq\left(\log _{\log ^{2} \sigma} \sigma\right)^{\mathcal{O}(1)} \leq b^{\mathcal{O}(1)} \leq \bar{n}^{\mathcal{O}(1)}$ otherwise. This construction uses $\mathcal{O}(\log \bar{n})=\mathcal{O}(\log (n \log \sigma))$ further random bits and produces a sketch of $\mathcal{O}(\bar{k} \log \bar{n})=\mathcal{O}(k b \log (n b))=\mathcal{O}\left(k \log _{n \log \sigma} \sigma \log (n \log \sigma)=\mathcal{O}(k \log \sigma)\right.$ bits.

The auxiliary string $\bar{S}$ is constructed in $\mathcal{O}(\bar{n})=\mathcal{O}(n \log \sigma)$ time using $\mathcal{O}(\log (n \sigma))$ bits of space. The stream representing $\bar{S}$ is passed to the encoding algorithm of [16], which takes $\mathcal{O}\left(\bar{n} \log ^{2} \bar{n}\right)=\mathcal{O}\left(n \log _{n \log \sigma} \sigma \log ^{2}(n \log \sigma)\right)=\mathcal{O}(n \log (n \sigma) \log (n \log \sigma))$ time and $\mathcal{O}(\bar{k} \log \bar{n})=$ $\mathcal{O}(k \log \sigma)$ bits of space.

The decoding algorithm, given $\mathrm{sk}_{\bar{k}}{ }^{H}(\bar{X})$ and $\mathrm{sk}_{\overline{\bar{k}}}{ }^{H}(\bar{Y})$, runs the decoding algorithm of [16]. If the latter certifies $\operatorname{hd}(\bar{X}, \bar{Y})>\bar{Y}$, we certify that $\mathrm{hd}(X, Y)>k$. Otherwise, we interpret the output as $\mathrm{MI}(\bar{X}, \bar{Y})$. For each position $i \in[1 \ldots|X|]$ such that $(i b-b \ldots i b] \subseteq \operatorname{MP}(\bar{X}, \bar{Y})$ we retrieve $X[i][j]$ and $Y[i][j]$ for each $j \in[1 \ldots b]$ from $(i b-j+1, \bar{X}[i b-j+1], \bar{Y}[i b-j+1]) \in \operatorname{MI}(\bar{X}, \bar{Y})$, and then we output $(i, X[i], Y[i]) \in \operatorname{MI}(X, Y)$.

As for correctness, with at most $\delta_{6.18}+n \bar{\delta}=\delta$ probability loss, we may assume that the decoder of [16] is successful and that, for all $i \in[1 \ldots|X|]$, we have $\psi(X[i])=\psi(Y[i])$ if and only if $X[i]=Y[i]$. The latter assumption yields $\operatorname{MP}(\bar{X}, \bar{Y})=\bigcup_{i \in \operatorname{MP}(X, Y)}(i b-b \ldots i b]$ and thus $\operatorname{hd}(\bar{X}, \bar{Y})=b \cdot \operatorname{hd}(X, Y)$. Hence, we correctly certify $\operatorname{hd}(X, Y)>k$ if $\operatorname{hd}(\bar{X}, \bar{Y})>\bar{k}$, and we correctly reconstruct $\mathrm{MI}(X, Y)$ otherwise.

The decoder uses $\mathcal{O}\left(\bar{k} \log ^{3} \bar{n}\right)=\mathcal{O}\left(k \log _{n \log \sigma} \sigma \log ^{3}(n \log \sigma)\right)=\mathcal{O}\left(k \log \sigma \log ^{2}(n \log \sigma)\right)$ time and $\mathcal{O}(\bar{k} \log \bar{n})=\mathcal{O}(k \log \sigma)$ bits of space.

Next, consider an alphabet $\hat{\Sigma}:=\Sigma \times \mathbb{Z}_{\geq 0}$ and a function $\pi: \hat{\Sigma} \rightarrow \mathbb{Z}_{\geq 0}$ defined so that $\pi((a, v))=v$ for $(a, v) \in \hat{\Sigma}$ and $\pi(S)=\sum_{i=1}^{|S|} \pi(S[i])$ for $S \in \hat{\Sigma}^{*}$. For two strings $X, Y \in \hat{\Sigma}^{*}$ of the same length, we define the prefix mismatch information

$$
\operatorname{PMI}(X, Y)=\{(i, \pi(X[1 \ldots i)), \pi(Y[1 \ldots i))): i \in \operatorname{MP}(X, Y)\}
$$

Proposition 6.20. For every constant $\delta \in(0,1)$, there exists a sketch $\mathrm{sk}_{k}^{P}$ (parameterized by integers $n \geq k \geq 1$, an alphabet $\hat{\Sigma}=[1 \ldots \sigma] \times\left[0 \ldots n^{\mathcal{O}(1)}\right]$, and a seed of $\mathcal{O}(\log (n \log \sigma))$ random bits) such that:

1. The sketch $\operatorname{sk}_{k}^{P}(S)$ of a string $S \in \hat{\Sigma}^{n}$ takes $\mathcal{O}\left(k \log ^{2} n\right)$ bits. Given streaming access to $S$, it can be constructed in $\mathcal{O}\left(n \log n \log (n \log \sigma)+n \log ^{2} n\right)$ time using $\mathcal{O}\left(k \log ^{2} n+\right.$ $\log n \log (n \log \sigma))$ bits of space.
2. There exists a decoding algorithm that, given the sketches $\mathrm{sk}_{k}^{P}(X)$ and $\mathrm{sk}_{k}^{P}(Y)$ of strings $X, Y \in \hat{\Sigma}^{n}$ satisfying $\pi(X), \pi(Y)<n$, with probability at least $1-\delta$ either returns $\operatorname{PMI}(X, Y)$ or certifies that $\operatorname{hd}(X, Y)>k$. The algorithm uses $\mathcal{O}\left(k \log ^{4} n\right)$ time and $\mathcal{O}\left(k \log ^{2} n\right)$ bits of space.

Proof. Let us fix a complete binary tree $\mathcal{T}$ with $n$ leaves, numbered with [1.. $n$ ] in the left-toright order, and let $v_{1}, \ldots, v_{2 n-1}$ denote the nodes of $\mathcal{T}$ in the pre-order. For each node $v_{i} \in \mathcal{T}$, let $\left[p_{i} \ldots q_{i}\right) \subseteq[1 \ldots n]$ be the indices of the leaves in the subtree of $v_{i}$. Consider the fingerprints $\psi$ of Fact 6.18 parameterized by $\delta_{6.18}=\frac{\delta}{4 n-2}, n_{6.18}=\mathcal{O}(n \log (n \sigma))$, and $\sigma_{6.18}=2$ : Given a string $S \in \hat{\Sigma}^{n}$, we define a string $\mathcal{T}(S)[1 \ldots 2 n)$ so that $\mathcal{T}(S)[i]=\left(\pi\left(S\left[p_{i} \ldots q_{i}\right)\right), \psi\left(S\left[p_{i} \ldots q_{i}\right)\right)\right)$
for every $i \in[1 \ldots 2 n)$, where $\psi$ expands each character of $\hat{\Sigma}$ into a sequence of $\mathcal{O}(\log (n \sigma))$ bits. We set $\mathrm{sk}_{k}^{P}(S):=\mathrm{sk}_{k}^{H}{ }_{6.19}(\mathcal{T}(S))$, where $\mathrm{sk}_{k}^{H}{ }_{6.19}$ is the sketch of Theorem 6.19 with $\delta_{6.19}=\frac{\delta}{2}$, $n_{6.19}=2 n-1, k_{6.19}=\min \left(k\lceil\log (2 n)\rceil, n_{6.19}\right)$, and $\sigma_{6.19} \leq 2^{\mathcal{O}(\log (n / \delta 6.18))} \leq n^{\mathcal{O}(1)}$.

This construction uses $\mathcal{O}\left(\log n_{6.18}+\log \left(n_{6.19} \log \sigma_{6.19}\right)\right)=\mathcal{O}(\log (n \log \sigma))$ random bits and produces sketches of $\mathcal{O}\left(k_{6.19} \log \left(n_{6.19} \sigma_{6.19}\right)\right)=\mathcal{O}\left(k \log ^{2} n\right)$ bits.

The encoding algorithm transforms $S$ into $\mathcal{T}(S)$ and feeds $\mathcal{T}(S)$ into the encoding procedure of Theorem 6.19. Construction of $\mathcal{T}(S)$ is organized in $\mathcal{O}(\log n)$ layers, each responsible for nodes $v_{i} \in \mathcal{T}$ at a fixed level. The intervals $\left[p_{i} \ldots q_{i}\right)$ corresponding to these nodes are disjoint so, at any time, a layer produces a single character $\mathcal{T}(S)[i]$ and, by Fact 6.18 , spends $\mathcal{O}(n \log \sigma)$ amortized time and uses $\mathcal{O}(\log n)$ bits of space to process a single character $S[j]$. Since the tree $\mathcal{T}$ is linearized in the post-order fashion, all levels can read the string $S$ with a common left-to-right pass and outputting $\mathcal{T}(S)$ does not require any buffers. Overall, construction of $\mathcal{T}(S)$ takes $\mathcal{O}(n \log n \log (n \log \sigma))$ time and uses $\mathcal{O}(\log n \log (n \log \sigma))$ bits of space. The encoder of Theorem 6.19 takes $\mathcal{O}\left(n_{6.19} \log \left(n_{6.19} \sigma_{6.19}\right) \log \left(n_{6.19} \log \sigma_{6.19}\right)\right)=\mathcal{O}\left(n \log ^{2} n\right)$ time and uses $\mathcal{O}\left(k_{6.19} \log \left(n_{6.19} \sigma_{6.19}\right)\right)=\mathcal{O}\left(k \log ^{2} n\right)$ bits of space.

The decoding algorithm first retrieves $\operatorname{MI}(\mathcal{T}(X), \mathcal{T}(Y))$ from $\mathrm{sk}_{k_{6.19}}^{H}(\mathcal{T}(X))$ and $\mathrm{sk}_{6.19}^{H}(\mathcal{T}(Y))$ using the decoder of Theorem 6.19. If the output of that subroutine can be interpreted as $\operatorname{MI}(\mathcal{T}(X), \mathcal{T}(Y))$, then we use Claim 6.21 below to retrieve $\operatorname{MI}(X, Y)$. If the size of the obtained set does not exceed $k$, we return the set. In the remaining cases, we certify that $\mathrm{hd}(X, Y)>k$.
Claim 6.21. For every $X, Y \in \hat{\Sigma}^{n}$, the set $\operatorname{PMI}(X, Y)$ can be extracted from $\operatorname{MI}(\mathcal{T}(X), \mathcal{T}(Y))$ in time and space $\mathcal{O}(\mathrm{hd}(X, Y) \log n)$ with success probability at least $1-\frac{\delta}{2}$.

Proof. We assume that, for every $i \in[1 \ldots 2 n)$, the equality $\psi\left(X\left[p_{i} \ldots q_{i}\right)\right)=\psi\left(Y\left[p_{i} \ldots q_{i}\right)\right)$ implies $X\left[p_{i} \ldots q_{i}\right)=Y\left[p_{i} \ldots q_{i}\right)$. By Fact 6.18, each of the implications fails with probability at most $\delta_{6.18}$, so the overall failure probability can be bounded by $\delta_{6.18} \cdot(2 n-1)=\frac{\delta}{2}$.

Our assumption yields

$$
\mathrm{MP}(\mathcal{T}(X), \mathcal{T}(Y))=\left\{i \in[1 \ldots 2 n):\left[p_{i} \ldots q_{i}\right) \cap \mathrm{MP}(X, Y) \neq \emptyset\right\}
$$

In particular, $\operatorname{MP}(X, Y)=\left\{p_{i}: v_{i}\right.$ is a leaf and $\left.i \in \operatorname{MP}(\mathcal{T}(X), \mathcal{T}(Y))\right\}$. Thus, it remains to describe how to extract $\pi(X[1 \ldots p))$ and $\pi(Y[1 \ldots p))$ for each $p \in \mathrm{MP}(X, Y)$; by symmetry, we focus on $\pi(X[1 \ldots p))$. For this, we process subsequent nodes $v_{i}$ on the path from the root of $\mathcal{T}$ to the leaf $v_{j}$ such that $\{p\}=\left[p_{j} \ldots q_{j}\right)$ maintaining $\pi\left(X\left[1 \ldots p_{i}\right)\right)$. Note that the values $\pi\left(X\left[p_{i} \ldots q_{i}\right)\right)$ can be extracted from $\operatorname{MI}(\mathcal{T}(X), \mathcal{T}(Y))$ due to $p \in\left[p_{i} \ldots q_{i}\right)$.

If $v_{i}$ is the root, then $\pi\left(X\left[1 \ldots p_{i}\right)\right)=\pi(\varepsilon)=0$. If $v_{i}$ is the left child of $v_{i^{\prime}}$, then $p_{i}=p_{i^{\prime}}$, so $\pi\left(X\left[1 \ldots p_{i}\right)\right)=\pi\left(X\left[1 \ldots p_{i^{\prime}}\right)\right)$ has already been computed. If $v_{i}$ is the right child of $v_{i^{\prime}}$, then $q_{i}=q_{i^{\prime}}$, so $\pi\left(X\left[1 \ldots p_{i}\right)\right)+\pi\left(X\left[p_{i} \ldots q_{i}\right)\right)=\pi\left(X\left[1 \ldots q_{i}\right)\right)=\pi\left(X\left[1 \ldots q_{i^{\prime}}\right)\right)=\pi\left(X\left[1 \ldots p_{i^{\prime}}\right)\right)+$ $\pi\left(X\left[p_{i^{\prime}} . . q_{i^{\prime}}\right)\right)$. Consequently, $\pi\left(X\left[1 \ldots p_{i}\right)\right)$ can be retrieved from $\pi\left(X\left[1 \ldots p_{i^{\prime}}\right)\right)$, which has already been computed, as well as $\pi\left(X\left[p_{i} \ldots q_{i}\right)\right)$ and $\pi\left(X\left[p_{i^{\prime}} \ldots q_{i^{\prime}}\right)\right)$, which are available in $\operatorname{MI}(\mathcal{T}(X), \mathcal{T}(Y))$. When this process reaches $v_{j}$, it results in the sought value $\pi(X[1 \ldots p))=$ $\pi\left(X\left[1 \ldots p_{j}\right)\right)$.

It remains to analyze correctness of the decoding algorithm. The decoder of Theorem 6.19 fails with probability at most $\delta_{6.19}=\frac{\delta}{2}$ and the procedure of Claim 6.21 fails with probability at most $\frac{\delta}{2}$. Thus, with at most $\delta$ probability loss, we may assume that both calls are successful.

If $\operatorname{hd}(\mathcal{T}(X), \mathcal{T}(Y)) \leq k_{6.19}$, then the decoder of Theorem 6.19 retrieves $\mathrm{MI}(\mathcal{T}(X), \mathcal{T}(Y))$ and the procedure of Claim 6.21 results in $\operatorname{MI}(X, Y)$. Depending on whether hd $(X, Y)=$ $|\mathrm{MI}(X, Y)| \leq k$ or not, our decoding algorithm thus correctly returns $\mathrm{MI}(X, Y)$ or certifies that $\operatorname{hd}(X, Y)>k$.

Otherwise, the algorithm of Theorem 6.19 certifies $\operatorname{hd}(\mathcal{T}(X), \mathcal{T}(Y))>k_{6.19}$, and the whole decoding procedure certifies $\operatorname{hd}(X, Y)>k$. This is correct because of $\operatorname{hd}(\mathcal{T}(X), \mathcal{T}(Y)) \leq$ hd $(X, Y)\lceil\log (2 n)\rceil$. To see this, observe that $\left[p_{i} \ldots q_{i}\right) \cap \mathrm{MP}(X, Y)=\emptyset$ implies $X\left[p_{i} \ldots q_{i}\right)=$ $Y\left[p_{i} \ldots q_{i}\right)$ and $\mathcal{T}(X)[i]=\mathcal{T}(Y)[i]$. However, every leaf of $\mathcal{T}$ has at most $\lceil\log (2 n)\rceil$ ancestors (including itself). Consequently, for every $j \in \operatorname{MP}(X, Y)$, there are at most $\lceil\log (2 n)\rceil$ nodes $v_{i}$ such that $j \in\left[p_{i} \ldots q_{i}\right)$.

### 6.4 Edit Distance Sketches

We are now ready to show the main result of this section.
Theorem 3.17. For every constant $\delta \in(0,1)$, there is a sketch $\mathrm{sk}_{k}^{E}$ (parametrized by integers $n \geq k \geq 1$, an alphabet $\Sigma=\left[0 \ldots n^{\mathcal{O}(1)}\right)$, and a seed of $\mathcal{O}\left(\log ^{2} n\right)$ random bits) such that:
(a) The sketch $\mathrm{sk}_{k}^{E}(S)$ of a string $\Sigma^{\leq n}$ takes $\mathcal{O}\left(k^{2} \log ^{3} n\right)$ bits. Given streaming access to $S$, it can be constructed in $\tilde{\mathcal{O}}(n k)$ time using $\tilde{\mathcal{O}}\left(k^{2}\right)$ space.
(b) There exists an $\tilde{\mathcal{O}}\left(k^{2}\right)$-space decoding algorithm that, given $\mathrm{sk}_{k}^{E}(X), \mathrm{sk}_{k}^{E}(Y)$ for $X, Y \in$ $\Sigma^{\leq n}$, with probability at least $1-\delta$ outputs $\mathrm{GR}_{k}(X, Y)$ and $\min (\mathrm{ed}(X, Y), k+1)$. Retrieving $\mathrm{GR}_{k}(X, Y)$ costs $\tilde{\mathcal{O}}\left(k^{5}\right)$ time, whereas computing $\min (\mathrm{ed}(X, Y), k+1)$ costs $\tilde{\mathcal{O}}\left(k^{3}\right)$ time.

Proof. Let $c_{3.16}$ be the constant of Proposition 3.16 for $\delta_{3.16}=\frac{\delta}{3}$, and let $c_{6.12}$ be the constant of Lemma 6.12. We shall use $\mathrm{CE}_{k^{\prime}}(W)$ with $k^{\prime}=\left(c_{3.16}+5\right) k$ and $W=\mathrm{W}(n, r, S)$, where $r$ is a random seed of $\mathcal{O}\left(\log ^{2} n\right)$ bits. Observe that the alphabet of $\mathrm{CE}_{k^{\prime}}(W)$ can be interpreted as $\left[0 \ldots n^{\mathcal{O}(k)}\right) \times[0 \ldots n]$ because the $\left|\mathrm{LZ}\left(\overline{C_{k^{\prime}}^{2}(S, s)}\right)\right|=\mathcal{O}(k)$ for each $s \in[1 \ldots|S|]$. We shall use Theorem 6.19 and Proposition 6.20 with $n_{6.19}=n_{6.20}=3 n, \delta_{6.19}=\delta_{6.20}=\frac{\delta}{3}, k_{6.19}=$ $k_{6.20}=\min \left(n_{6.19},\left\lfloor c_{6.12}\left(1+c_{3.16}\right) k \log (3 n)\right\rfloor\right)$, and $\sigma_{6.19}=\sigma_{6.20}=n^{\mathcal{O}(k)}$.

The sketch $\operatorname{sk}_{k}^{E}(S)$ of a string $S \in \Sigma^{\leq n}$ consists of $|S|$ as well as the sketches sk ${ }_{k}^{H}{ }_{6.19}\left(\mathrm{CE}_{k^{\prime}}(W)\right)$ and $\mathrm{sk}_{k}^{P}{ }_{6.20}\left(\mathrm{CE}_{k^{\prime}}(W)\right)$. This construction uses
$\mathcal{O}\left(\log ^{2} n+\log \left(n_{6.19} \log \sigma_{6.19}\right)+\log \left(n_{6.20} \log \sigma_{6.20}\right)\right)=\mathcal{O}\left(\log ^{2} n+\log (n k \log n)\right)=\mathcal{O}\left(\log ^{2} n\right)$
random bits and produces sketches of bit-size

$$
\mathcal{O}\left(k_{6.19} \log \left(n_{6.19} \sigma_{6.19}\right)+k_{6.20} \log ^{2} n_{6.20}\right)=\mathcal{O}\left(k \log n \log \left(n^{\mathcal{O}(k)}\right)+k \log ^{3} n\right)=\mathcal{O}\left(k^{2} \log ^{3} n\right)
$$

The encoding algorithm uses Lemma 6.16 to transform the input stream representing $S$ to an auxiliary stream representing $\mathrm{CE}_{k^{\prime}}(W)$, which we forward to the encoders constructing $\operatorname{sk}_{k_{6.19}}^{H}\left(\mathrm{CE}_{k^{\prime}}(S)\right)$ and $\mathrm{sk}_{k_{6.20}}^{P}\left(\mathrm{CE}_{k^{\prime}}(W)\right)$. Thus, it takes $\tilde{\mathcal{O}}\left(n k^{\prime}+n_{6.19} \log \sigma_{6.19}+n_{6.20}\right)=$ $\tilde{\mathcal{O}}(n k)$ time and uses $\tilde{\mathcal{O}}\left(k^{\prime}+k_{6.19} \log \sigma_{6.19}+k_{6.20}\right)=\tilde{\mathcal{O}}\left(k^{2}\right)$ space.

Decoding Algorithm Let $\mathcal{A}_{\mathrm{W}}$ be the zip alignment of walks $W_{X}=\left(x_{t}\right)_{t=1}^{3 n+1}=\mathrm{W}(n, r, X)$ and $W_{Y}=\left(y_{t}\right)_{t=1}^{3 n+1}=\mathrm{W}(n, r, Y)$ over strings $X, Y \in \Sigma^{\leq n}$.

Given sketches $\mathrm{sk}_{k}^{E}(X), \mathrm{sk}_{k}^{E}(Y)$, we run the decoders for $\mathrm{sk}_{k_{6.19}^{H}}^{H}\left(\mathrm{CE}_{k^{\prime}}\left(W_{X}\right)\right), \mathrm{sk}_{k_{6.19}^{H}}\left(\mathrm{CE}_{k^{\prime}}\left(W_{Y}\right)\right)$ and $\mathrm{sk}_{k_{6.20}}^{P}\left(\mathrm{CE}_{k^{\prime}}\left(W_{X}\right)\right), \mathrm{sk}_{k_{6.20}}^{P}\left(\mathrm{CE}_{k^{\prime}}\left(W_{Y}\right)\right)$. We certify ed $(X, Y)>k$ if either procedure certifies $\operatorname{hd}\left(\mathrm{CE}_{k^{\prime}}\left(W_{X}\right), \mathrm{CE}_{k^{\prime}}\left(W_{Y}\right)\right)>k_{6.19}=k_{6.20}$. Otherwise, the decoder interprets the outputs as $\mathrm{MI}\left(\mathrm{CE}_{k^{\prime}}\left(W_{X}\right), \mathrm{CE}_{k^{\prime}}\left(W_{Y}\right)\right)$ and $\mathrm{PMI}\left(\mathrm{CE}_{k^{\prime}}\left(W_{X}\right), \mathrm{CE}_{k^{\prime}}\left(W_{Y}\right)\right)$, respectively. In this case, we construct $\mathrm{E}^{M}(X, Y)$ for $M=\left\{(x, y) \in \mathcal{M}_{X, Y}\left(\mathcal{A}_{\mathrm{W}}\right): x \in P_{X}\right.$ or $\left.y \in P_{Y}\right\}$, where $P_{X}$ and $P_{Y}$ are defined in the statement of Lemma 6.12. Finally, we compute $\mathrm{GR}_{k}(X, Y)$ using Proposition 5.12 or $\min (\operatorname{ed}(X, Y), k+1)$ using Corollary 5.5.

Claim 6.22. Given $\operatorname{MI}\left(\mathrm{CE}_{k^{\prime}}\left(W_{X}\right), \mathrm{CE}_{k^{\prime}}\left(W_{Y}\right)\right)$ and $\operatorname{PMI}\left(\mathrm{CE}_{k^{\prime}}\left(W_{X}\right), \mathrm{CE}_{k^{\prime}}\left(W_{Y}\right)\right)$, the encoding $\mathrm{E}^{M}(X, Y)$ can be constructed in $\tilde{\mathcal{O}}\left(k^{3}\right)$ time using $\tilde{\mathcal{O}}\left(k^{2}\right)$ space. Moreover, $\left|\operatorname{LZ}\left(X^{M} Y^{M}\right)\right|=$ $\tilde{\mathcal{O}}\left(k^{2}\right)$.

Proof. We proceed in two phases. In the first phase, we construct the compressed representation $\mathrm{D}\left(X^{\prime}\right)$ of a string $X^{\prime} \in(\Sigma \cup\{\#\})^{|X|}$ such that $X^{\prime}[x]=X[x]$ if $x \in P_{X}$ and $X^{\prime}[x]=\#$ otherwise. We initialize $X^{\prime}:=\#^{|X|}$ (via LZ $\left(\#^{|X|}\right)$ using Proposition $4.2(\mathrm{~h})$ ). Next, we iterate over positions $t \in[1 . .3 n]$ such that $\perp \neq \mathrm{CE}_{k^{\prime}}\left(W_{X}\right)[t] \neq \mathrm{CE}_{k^{\prime}}\left(W_{Y}\right)[t]$. We retrieve $x_{t}$ and $\mathrm{LZ}\left(\overline{C_{k^{\prime}}^{2}\left(X, x_{t}\right)}\right)$ from the mismatch information, convert $\mathrm{LZ}\left(\overline{C_{k^{\prime}}^{2}\left(X, x_{t}\right)}\right)$ to $\mathrm{D}\left(C_{k^{\prime}}^{2}\left(X, x_{t}\right)\right)$ (Proposition 4.2(h)), and update $\mathrm{D}\left(X^{\prime}\right)$, setting $X^{\prime}\left[x_{t} \ldots x_{t}+\left|C_{k^{\prime}}^{2}\left(X, x_{t}\right)\right|\right):=C_{k^{\prime}}^{2}\left(X, x_{t}\right)$ (Proposition 4.2(g)(i)). We also symmetrically construct the compressed representation $\mathrm{D}\left(Y^{\prime}\right)$ of a string $Y^{\prime} \in(\Sigma \cup\{\#\})^{|Y|}$ such that $Y^{\prime}[y]=Y[y]$ if $y \in P_{Y}$ and $Y^{\prime}[y]=\#$ otherwise.

In the second phase, we convert $\mathrm{D}\left(X^{\prime}\right)$ to $\mathrm{D}\left(X^{M}\right)$. Here, the goal is make sure that $X^{M}[x]=$ \# not only for $x \in P_{X}$, but also when $X[x] \simeq_{\mathcal{A}_{W}} Y[y]$ for some $y \in P_{Y}$. For this, we iterate over dummy segments $Y^{\prime}\left[y \ldots y^{\prime}\right)$. By Lemma $6.12, \mathcal{A}_{\mathrm{W}}$ matches $Y\left[y \ldots y^{\prime}\right)$ to a fragment of $X\left[x \ldots x^{\prime}\right)$. Hence, we shall identify the shift $x-y$ and set $X^{M}\left[x \ldots x^{\prime}\right):=\#^{x^{\prime}-x}$ (Proposition 4.2(h)(g)(i)). Let $\mathrm{MP}_{\mathrm{CE}}:=\mathrm{MP}\left(\mathrm{CE}_{k^{\prime}}\left(W_{X}\right), \mathrm{CE}_{k^{\prime}}\left(W_{Y}\right)\right)$. Define $u, v \in[1 . .3 n+1]$ so that $u=v=3 n+1$ if $\pi\left(\mathrm{CE}_{k^{\prime}}\left(W_{Y}\right)[1 \ldots t]\right) \neq y^{\prime}$ for all $t \in \mathrm{MP}_{\mathrm{CE}}$; otherwise, $u$ is the smallest index in MP ME with $\pi\left(\mathrm{CE}_{k^{\prime}}\left(W_{Y}\right)[1 \ldots u]\right)=y^{\prime}$, whereas $v$ is the smallest index $\mathrm{MP}_{\mathrm{CE}}$ with $\pi\left(\mathrm{CE}_{k^{\prime}}\left(W_{Y}\right)[1 \ldots v)\right)=$ $\pi\left(\mathrm{CE}_{k^{\prime}}\left(W_{Y}\right)[1 \ldots u)\right)$. We claim that $x-y=\pi\left(\mathrm{CE}_{k^{\prime}}\left(W_{X}\right)[1 \ldots v)\right)-\pi\left(\mathrm{CE}_{k^{\prime}}\left(W_{Y}\right)[1 \ldots v)\right)$.

If $y^{\prime} \leq|Y|$, then $y^{\prime}-1 \notin P_{Y}$ and $y^{\prime} \in P_{Y}$. Hence, $Y\left[y^{\prime}\right]$ is covered by $C_{k^{\prime}}^{2}\left(Y, y_{t}\right)$ for some $t \in \mathrm{MP}_{\mathrm{CE}}$ with $\mathrm{CE}_{k^{\prime}}\left(W_{X}\right)[t] \neq \perp$, whereas $Y\left[y^{\prime}-1\right]$ is not covered by $C_{k^{\prime}}^{2}\left(Y, y_{t}\right)$. Hence, there is $t \in \mathrm{MP}_{\mathrm{CE}}$ with $y_{t}=y^{\prime}$, and therefore $y_{u}=y^{\prime}$. Similarly, $y_{u}=y^{\prime}$ holds if $y^{\prime}=|Y|+1$. Let $w \in[1 \ldots 3 n]$ be the smallest index such that $\mathrm{CE}_{k^{\prime}}\left(W_{Y}\right)[w] \neq \perp$ and $C_{k^{\prime}}^{2}\left(Y, y_{t}\right)$ covers $Y\left[y^{\prime}-1\right]$ (such an index exists by Lemma 6.9). Due to $y^{\prime}-1 \in P_{Y}$, we have $\mathrm{CE}_{k^{\prime}}\left(W_{X}\right)[w]=\mathrm{C}_{k^{\prime}}\left(W_{Y}\right)[w]$, and hence $x-y=x_{w}-y_{w}$. Definition of $w$ further yields $\mathrm{CE}_{k^{\prime}}\left(W_{Y}\right)[w+1 \ldots v)=\perp^{u-w-1}$, and thus $\pi\left(\mathrm{CE}_{k^{\prime}}\left(W_{Y}\right)[1 . . v)\right)=y_{w}$. Moreover, since $\pi\left(\mathrm{CE}_{k^{\prime}}\left(W_{Y}\right)[1 \ldots t)\right)=y_{w}$ for $t \in[w+1 \ldots v]$, we have $\pi\left(\mathrm{CE}_{k^{\prime}}\left(W_{X}\right)[1 \ldots v)\right)=x_{w}$ by definition of $v$. This completes the proof that $x-y=$ $\pi\left(\mathrm{CE}_{k^{\prime}}\left(W_{X}\right)[1 . . v)\right)-\pi\left(\mathrm{CE}_{k^{\prime}}\left(W_{Y}\right)[1 . . v)\right)$.

To derive $\mathrm{E}^{M}(X, Y)$, it suffices to convert $\mathrm{D}\left(Y^{\prime}\right)$ to $\mathrm{D}\left(Y^{M}\right)$ (symmetrically), and to construct $\mathrm{D}\left(X^{M} Y^{M}\right)$ using Proposition 4.2(i). Since $\overline{\operatorname{maxLZ}}\left(C_{k^{\prime}}^{2}\left(X, x_{t}\right)\right), \overline{\operatorname{maxLZ}}\left(C_{k^{\prime}}^{2}\left(Y, y_{t}\right)\right)=\mathcal{O}(k)$ holds for all $t \in[1 . .3 n+1]$ and since $\left|\mathrm{MP}_{\mathrm{CE}}\right|=\mathcal{O}(k \log n)$, the $\overline{\operatorname{maxLZ}}(\cdot)$ measure of all intermediate strings is $\tilde{\mathcal{O}}\left(k^{2}\right)$. Consequently, the $\tilde{\mathcal{O}}(k)$ applications of Proposition 4.2 cost $\tilde{\mathcal{O}}\left(k^{3}\right)$ time and use $\tilde{\mathcal{O}}\left(k^{2}\right)$ space.

To complete the complexity analysis, observe that the decoding procedure of Theorem 6.19 uses $\mathcal{O}\left(k_{6.19} \log \left(n_{6.19} \sigma_{6.19}\right) \log ^{2}\left(n_{6.19} \log \sigma_{6.19}\right)\right)=\tilde{\mathcal{O}}\left(k^{2}\right)$ time and $\mathcal{O}\left(k_{6.19} \log \left(n_{6.19} \sigma_{6.19}\right)\right)=$ $\tilde{\mathcal{O}}\left(k^{2}\right)$ bits of space. The procedure of Proposition 6.20 uses $\mathcal{O}\left(k_{6.20} \log ^{2} n_{6.20}\right)=\mathcal{O}(k)$ bits of space and costs $\mathcal{O}\left(k_{6.20} \log ^{4} n_{6.20}\right)=\tilde{\mathcal{O}}(k)$ time. Finally, due to $\left|\operatorname{LZ}\left(X^{M} Y^{M}\right)\right|=\tilde{\mathcal{O}}\left(k^{2}\right)$ and $\left|\mathcal{B}_{k}(X, Y)\right|=\mathcal{O}\left(k^{5}\right)\left(\right.$ Lemma 3.11), the algorithm of Proposition 5.12 uses $\tilde{\mathcal{O}}\left(k^{2}\right)$ space and costs $\tilde{\mathcal{O}}\left(k^{5}\right)$ time (dominating the overall decoding complexity). If we only aim to retrieve $\min (\operatorname{ed}(X, Y), k+1)$, the algorithm of Corollary 5.5 takes $\tilde{\mathcal{O}}\left(k^{2}\right)$ time and space (in which case the overall decoding uses $\tilde{\mathcal{O}}\left(k^{2}\right)$ space and costs $\tilde{\mathcal{O}}\left(k^{3}\right)$ time.

It remains to argue that the decoding algorithm is correct. With $\delta_{3.16}=\frac{\delta}{3}$ probability loss, we may assume that $\mathcal{A}_{\mathrm{W}} \in \mathrm{GA}_{c_{3.16}(\operatorname{ed}(X, Y))^{2}}(X, Y)$ and width $\left(\mathcal{A}_{\mathrm{W}}\right) \leq c_{3.16} \mathrm{ed}(X, Y)$. With $\delta_{6.19}+\delta_{6.20}=\frac{2 \delta}{3}$ further probability loss, we may assume that the decoders of Theorem 6.19 and Proposition 6.20 are successful. If $\operatorname{ed}(X, Y)>k$, then the correctness follows from Proposition 5.12 and Corollary 5.5 because $M \subseteq \mathcal{M}_{X, Y}\left(\mathcal{A}_{\mathrm{W}}\right)$ is a non-crossing matching. Otherwise, Proposition 3.16 guarantees $\operatorname{cost}\left(\mathcal{A}_{\mathrm{W}}\right) \leq c_{3.16} k^{2}$ and $\operatorname{width}\left(\mathcal{A}_{\mathrm{W}}\right) \leq c_{3.16} k$ so, in
particular, $k^{\prime} \geq \operatorname{width}\left(\mathcal{A}_{\mathrm{W}}\right)+5 k$. By Lemma 6.12 , we thus have hd $\left(\mathrm{CE}_{k^{\prime}}\left(W_{X}\right), \mathrm{CE}_{k^{\prime}}\left(W_{X}\right)\right) \leq$ $c_{6.12}\left(1+c_{3.16}\right) k \log (3 n)$. Hence, the decoders of Theorem 6.19 and Proposition 6.20 report the mismatch information. Lemma 6.10 further implies that $\Delta_{X}\left(\mathcal{A}, \mathcal{A}_{\mathrm{W}}\right) \subseteq P_{X}$ and $\Delta_{Y}\left(\mathcal{A}, \mathcal{A}_{\mathrm{W}}\right) \subseteq P_{Y}$ for all $\mathcal{A} \in \mathrm{GA}_{k}(X, Y)$. In particular, $\mathcal{M}_{X, Y}(\mathcal{A}) \subseteq M$, and therefore $\mathcal{M}_{k}(X, Y) \subseteq M$. Consequently, the algorithms of Proposition 5.12 and Corollary 5.5 correctly compute $\mathrm{GR}_{k}(X, Y)$ and $\min (\operatorname{ed}(X, Y), k+1)$, respectively.

Next, we boost the success probability and strengthen the sketches so that we can retrieve $\mathrm{qGR}_{k}(X, Y)$ instead of $\mathrm{GR}_{k}(X, Y)$.

Corollary 6.23. There is a sketch $\mathrm{sk}_{k}^{\mathrm{q}}$ (parametrized by $\delta \in\left(0, \frac{1}{2}\right)$, integers $n \geq k \geq 1$, an alphabet $\Sigma=\left[0 \ldots n^{\mathcal{O}(1)}\right)$, and a seed of $\mathcal{O}\left(\log ^{2} n \log (1 / \delta)\right)$ random bits) such that:
(a) The sketch $\mathrm{sk}_{k}^{\mathrm{q}}(S)$ of a string $\Sigma^{\leq n}$ takes $\mathcal{O}\left(k^{2} \log ^{3} n \log (1 / \delta)\right)$ bits. Given streaming access to $S$, it can be constructed in $\tilde{\mathcal{O}}(n k \log (1 / \delta))$ time using $\tilde{\mathcal{O}}\left(k^{2} \log (1 / \delta)\right)$ space.
(b) There exists a decoding algorithm that, given $\mathrm{sk}_{k}^{\mathrm{q}}(X), \mathrm{sk}_{k}^{\mathrm{q}}(Y)$ for $X, Y \in \Sigma \leq n$, with probability at least $1-\delta$ computes $\mathrm{qGR}_{k}(X, Y)$. The algorithm takes $\tilde{\mathcal{O}}\left(k^{5} \log (1 / \delta)\right)$ time and uses $\tilde{\mathcal{O}}\left(k^{2} \log (1 / \delta)\right)$ space.
Proof. We shall use $\mu=\mathcal{O}(\log (1 / \delta))$ sketches $\mathrm{sk}_{k_{3.17}}^{E}$ of Theorem 3.17 with $\delta_{3.17}=\frac{1}{3}, n_{3.17}=$ $n+1, k_{3.17}=k+1$, and independent seeds. For each of the $\mu$ sketches $\operatorname{sk}_{k_{3.17}}^{E}$, the sketch $\mathrm{sk}_{k}^{\mathrm{q}}(S)$ contains $\mathrm{sk}_{k_{3.17}}^{E}\left(\$_{1} S\right), \mathrm{sk}_{k_{3.17}}^{E}\left(\$_{2} S\right)$, where $\$_{1}, \$_{2} \notin \Sigma$ are two distinct symbols. This construction uses $\mathcal{O}\left(\mu \log ^{2} n_{3.17}\right)=\mathcal{O}\left(\log ^{2} n \log (1 / \delta)\right)$ random bits and produces sketches of $\mathcal{O}\left(\mu k_{3.17}^{2} \log ^{3} n_{3.17}\right)=\mathcal{O}\left(k^{2} \log ^{3} n \log (1 / \delta)\right)$ bits.

The encoding algorithm calls $2 \mu$ instances of the encoding algorithm of Theorem 3.17. Hence, it uses $\tilde{\mathcal{O}}\left(\mu k_{3.17}^{2}\right)=\mathcal{O}\left(k^{2} \log (1 / \delta)\right)$ space and costs $\tilde{\mathcal{O}}\left(\mu n_{3.17} k_{3.17}\right)=\tilde{\mathcal{O}}(n k \log (1 / \delta))$ time.

The decoding algorithm, given $\operatorname{sk}_{k}^{\mathrm{q}}(X), \mathrm{sk}_{k}^{\mathrm{q}}(Y)$ for $X, Y \in \Sigma^{\leq n}$, runs the decoder of Theorem 3.17 for $\mathrm{sk}_{k_{3.17}}^{E}\left(\$_{1} X\right), \mathrm{sk}_{k_{3.17}}^{E}\left(\$_{1} Y\right)$ for each of the $\mu$ sketches $\mathrm{sk}_{k_{3.17}}^{E}$. This yields $\mu$ candidates for $\mathrm{GR}_{k_{3.17}}^{3.17}\left(\$_{1} X, \$_{2} Y\right)$, which we convert to candidates for $\mathrm{qGR}_{k}(X, Y)$ using Corollary 5.17 . Finally, we determine the majority answer among the $\mu$ candidates. The equality test uses Proposition 4.2(i)(c) to compare two candidates for $\mathrm{D}_{\tilde{O}}\left(X^{M} Y^{M}\right)$.

Recall that the decoding procedures of Theorem 3.17 use $\tilde{\mathcal{O}}\left(\mu k_{3.17}^{5}\right)=\tilde{\mathcal{O}}\left(k^{5} \log (1 / \delta)\right)$ time and $\tilde{\mathcal{O}}\left(\mu k_{3.17}^{2}\right)=\tilde{\mathcal{O}}\left(k^{2} \log (1 / \delta)\right)$ space. The applications of Corollary 5.17 and the equality tests take $\tilde{\mathcal{O}}\left(\mu k_{3.17}^{2}\right)=\tilde{\mathcal{O}}\left(k^{2} \log (1 / \delta)\right)$ time and space. The entire decoding algorithm uses $\tilde{\mathcal{O}}\left(k^{2} \log (1 / \delta)\right)$ space and $\tilde{\mathcal{O}}\left(\mu k_{3.17}^{5}\right)=\tilde{\mathcal{O}}\left(k^{5} \log (1 / \delta)\right)$ time.

As for correctness, since the $\mu$ sketches $\mathrm{sk}_{k_{3.17}}^{E}$ are independent, by the Chernoff bound, the majority answer is wrong with probability at most $\exp (-\mathcal{O}(\mu))$. Setting $\mu=\mathcal{O}(\log (1 / \delta))$ (with a sufficiently large constant factor) guarantees a success probability of $1-\delta$.

## 7 Pattern Matching with $\boldsymbol{k}$ Edits

In this section, we present solutions for pattern matching with $k$ edits in the semi-streaming and streaming settings.

### 7.1 Periodicity under Edit Distance

We start by recalling combinatorial properties of strings periodic under the edit distance.

Definition 3.2 ( $k$-periodic string). A string $X$ is $k$-periodic if there exists a primitive string $Q$ with $|Q| \leq|X| / 128 k$ such that the edit distance between $X$ and a prefix of $Q^{\infty}$ is at most $2 k$. We call $Q$ a $k$-period of $X$.
Claim 7.1. Suppose that a string $X$ is a prefix of a string $Y$, where $|X|<|Y| \leq 2|X|$. If $X$ is $k$-periodic with $k$-period $Q,|Q| \leq|X| / 128 k$, then either $Y$ is not $k$-periodic, or $Y$ is $k$-periodic with $k$-period $Q$.
Proof. Suppose by contradiction that $Y$ is $k$-periodic with $k$-period $Q^{\prime} \neq Q$. Let $q=|Q|$ and $q^{\prime}=\left|Q^{\prime}\right|$. Assume first $q \leq q^{\prime}$. Fix an alignment of the smallest cost between $Y$ and a prefix of $\left(Q^{\prime}\right)^{\infty}$. It induces an alignment $\mathcal{A}^{\prime}$ of cost at most $2 k$ between $X$ and a prefix of $\left(Q^{\prime}\right)^{\infty}$ and hence generates a partition $X=X_{1} X_{2} \ldots X_{z}$, where each $X_{i}, 1 \leq i \leq z-1$, is aligned with $Q^{\prime}$ and $X_{z}$ is aligned with a prefix of $Q^{\prime}$. From $|Q| \leq|X| / 128 k$ we obtain $|X| \geq 128 k$ and therefore

$$
z \geq(|X|-2 k) / q^{\prime} \geq \frac{128-2}{128}|X| /(|Y| / 128 k) \gg 20 k
$$

Consider fragments $X_{1} X_{2} X_{3} X_{4}, X_{5} X_{6} X_{7} X_{8}$, and so on. The total number of such fragments is at least $(20 k-3) / 4>4 k$, and at least $2 k+1$ of them are not edited under $\mathcal{A}^{\prime}$. Fix an optimal alignment $\mathcal{A}$ between $X$ and a prefix $Q^{\infty}$. The cost of $\mathcal{A}$ is bounded from above by $2 k$, and therefore there is at least one fragment $X_{4 i+1} X_{4 i+2} X_{4 i+3} X_{4 i+4}$ that is not edited. Consider one such fragment $F$. On the one hand, $F=Q^{\prime} Q^{\prime} Q^{\prime} Q^{\prime}$. On the other hand, $F=\operatorname{suff}(Q) Q^{j} \operatorname{pref}(Q)$, where $\operatorname{suff}(Q)$ and $\operatorname{pref}(Q)$ are some suffix and some prefix of $Q$, respectively.

Suppose first that $q^{\prime}$ is a multiple of $q$. In this case, $Q^{\prime}=(Q[\ell+1 \ldots q] Q[1 \ldots \ell])^{r}$, where $\ell=q-|\operatorname{suff}(Q)|$ and an integer $r$, which contradicts the fact that $Q^{\prime}$ is primitive. Otherwise, consider the copy of $Q$ that contains $F\left[q^{\prime}\right]$. Consider also a substring $Q Q Q$ of $F$ formed by the copy of $Q$ that contains $F\left[2 q^{\prime}\right]$, and the preceding and succeeding copies of $Q$. We then obtain that there is an occurrence of $Q$ in $Q Q Q$ that is not aligned with any copy of $Q$ (otherwise, $q^{\prime}$ is a multiple of $q$ ), and therefore $Q$ is not primitive, a contradiction.

The case $q>q^{\prime}$ can be treated analogously.
Note that Claim 7.1 implies in particular that a string can have at most one $k$-period.

### 7.2 Semi-streaming Algorithm

We first present a deterministic algorithm for pattern matching with $k$ edits in the semistreaming setting.

### 7.2.1 Preprocessing Stage

Consider a set $\Pi$ of $\mathcal{O}(\log m)$ prefixes $P_{i}$ of $P$ initialised to contain $P$ itself as well as the prefixes of length $2^{\ell}$ for all $0 \leq \ell \leq\lfloor\log |P|\rfloor$. Order the prefixes by lengths, and consider two consecutive prefixes $P^{\prime}, P^{\prime \prime}$. If $P^{\prime}$ is $k$-periodic with $k$-period $Q^{\prime}$ while $P^{\prime \prime}$ is not $k$-periodic, we add two more prefixes to $\Pi$. Namely, if $\ell$ be the maximum integer such that $P[1 . . \ell]$ is $k$-periodic with $k$-period $Q^{\prime}$, add to $\Pi$ the prefixes $P[1 \ldots \ell]$ and $P[1 \ldots \ell+1]$. Note that $P[1 \ldots \ell+1]$ is not $k$-periodic by Claim 7.1.

Let $\Pi=\left\{P_{1}, P_{2}, \ldots, P_{z}\right\}$ be the resulting set of prefixes. We assume that the patterns are ordered in the ascending order of their lengths. During the preprocessing step, for each $i$ such that $P_{i}$ is $k$-periodic, we compute its $k$-period $Q_{i}$. We use notation $\ell_{i}=P_{i}$ and $q_{i}=\left|Q_{i}\right|$ (if defined). Importantly, we do not have to store $Q_{i}$ explicitly, we can simply memorize its endpoints in $P_{i}$ which takes $\mathcal{O}(\log m)$ extra space in total. We also store, for each of the $\mathcal{O}(k)$ rotations $D$ of $Q_{i}$ that can be a difference of a chain of $k$-edit occurrences of $P_{i}$ (Corollary 3.5), the encodings $q \operatorname{GR}_{32 k}(D, D)$ and $\mathrm{qGR}_{30 k}\left(P\left[\ell_{i-1}+1 \ldots \ell_{i}\right], D^{\infty}\left[1 \ldots \Delta_{i}\right]\right)$, where $\Delta_{i}=\ell_{i}-\ell_{i-1}+k$.

### 7.2.2 Main Stage

The main stage of the algorithm starts after we have preprocessed the pattern. During the main stage, we receive the text as a stream, one character of a time. We exploit the following result:

Fact 7.2 (cf. [59]). Given a read-only string $X$ of length $m$ and a streaming string $Y$. There is a dynamic programming algorithm that correctly identifies all prefixes $Y^{\prime}$ of $Y$ within edit distance at most $k \leq m$ from $X$, as well as $\operatorname{ed}\left(Y^{\prime}, X\right)$ itself. The algorithm takes $\mathcal{O}(k m)$ time and $\mathcal{O}(k)$ space besides the space required to store $X$.

Chains of $k$-edit occurrences. During the main stage of the algorithm, we store the following information. Let $r$ be the newly arrived position of the text $T$. For each $2 \leq i \leq z$, consider all $k$-edit occurrences of $P_{i-1}$ in $T\left[r-\ell_{i}-k+1 \ldots r\right]$. We call such occurrences active. We denote the set of right endpoints of the active occurrences by $\mathrm{aOCC}_{k}^{E}\left(P_{i-1}, T\right)$. By Corollaries 3.3 and $3.5, \operatorname{aOCC}_{k}^{E}\left(P_{i-1}, T\right)$ forms $\mathcal{O}\left(k^{3}\right)$ chains. For each chain $\mathcal{C}$, we store the following information:

1. The leftmost position $l p$ and the size $|\mathcal{C}|$ of $\mathcal{C}$;
2. An integer ed $(\mathcal{C})$ equal to the smallest edit distance from $P_{i-1}$ to a suffix of $T[1 \ldots r]$ for every endpoint $r \in \mathcal{C}$;
3. If $|\mathcal{C}| \geq 2$, we store the shift $\Delta$ of the difference $D=Q_{i-1}\left[1+\Delta \ldots q_{i-1}\right] Q_{i-1}\left[1 \ldots q_{i-1}\right]$ of $\mathcal{C}$.

If $p^{*}$ is the first position added to $\mathcal{C}$ (it can be different from $l p$ as we will update chains to contain only active occurrences), then at the position $\left(p^{*}+1\right)$ we start running the dynamic programming algorithm for $T\left[p^{*}+1 \ldots\right]$ and $P\left[\ell_{i-1}+1 \ldots \ell_{i}\right]$ (Fact 7.2).

Furthermore, consider the moment when we detect the second position in $\mathcal{C}$ (if it exists) and hence the difference $D$ of the chain. Starting from this moment, for every newly added position $p \in \mathcal{C}$, at the position $(p+1)$ we start computing the greedy encoding $\mathrm{qGR}_{32 k}(T[p+1 \ldots p+$ $\left.\Delta_{i}\right], D^{\infty}\left[1 \ldots \Delta_{i}\right]$ ). We continue running the algorithm until either the computation is over or a new position in the chain is detected. In the end, we compute the encoding for the rightmost position in the chain.

Detecting new $k$-edit occurrences of $P_{i}$. We now explain how to detect new $k$-edit occurrences of the prefixes $P_{i}$. Let $r$ be the latest arrived position of $T$. If $i=1$, then since $k \geq 1$, $r \in \operatorname{aOCC}_{k}^{E}\left(P_{1}, T\right)$. Below we consider three possible cases for $i \geq 2: P_{i-1}$ is $k$-periodic, $P_{i}$ is not $k$-periodic; $P_{i-1}$ is not $k$-periodic; $P_{i-1}$ and $P_{i}$ are $k$-periodic.

Case 1: $P_{i-1}$ is $k$-periodic, $P_{i}$ is not $k$-periodic. By construction, in this case $\ell_{i-1}+1=$ $\ell_{i}$. The position $r \in \mathrm{aOCC}_{k}^{E}\left(P_{i}, T\right)$ iff one of the following conditions is satisfied:

1. The smallest edit distance from $P_{i-1}$ to a suffix of $T[1 \ldots r]$ is at most $k-1$ (this corresponds to the case when the last character of $P_{i}$ is deleted in an optimal alignment of a suffix of $T[1 \ldots r]$ and $P_{i}$;
2. The smallest edit distance from $P_{i-1}$ to a suffix of $T[1 \ldots r-1]$ is at most $k-1$ and $P_{i}\left[\ell_{i}\right] \neq T[r]$ (this corresponds to the case when the last character of $P_{i}$ is substituted for $T[r]$ in an optimal alignment of a suffix of $T[1 \ldots r]$ and $\left.P_{i}\right)$;
3. The smallest edit distance from $P_{i-1}$ to a suffix of $T[1 \ldots r-1]$ is at most $k$ and $P_{i}\left[\ell_{i}\right]=T[r]$ (this corresponds to the case when the last character of $P_{i}$ is matched with $T[r]$ in an optimal alignment of a suffix of $T[1 \ldots r]$ and $P_{i}$ );
4. The smallest edit distance from $P_{i}$ to a suffix of $T[1 \ldots r-1]$ is at most $k-1$ (this corresponds to the case when $T[r]$ is deleted in an optimal alignment of a suffix of $T[1 \ldots r]$ and $P_{i}$ ).

We can decide which of the conditions is satisfied, and therefore whether $r \in \mathrm{aOCC}_{k}^{E}\left(P_{i}, T\right)$ in $\mathcal{O}\left(k^{3}\right)$ time using aOCC ${ }_{k}^{E}\left(P_{i-1}, T\right)$ and $\operatorname{aOCC}_{k}^{E}\left(P_{i}, T\right)$. Moreover, we can compute the smallest edit distance from $P_{i}$ to a suffix of $T[1 \ldots r]$ if it is bounded by $k$.

For the next two cases, we will use the following simple observation that follows from Fact 2.3:

Observation 7.3. Let $\operatorname{ed}_{i-1}\left(r^{\prime}\right)$ be the smallest edit distance from $P_{i-1}$ to a suffix of $T\left[1 \ldots r^{\prime}\right]$, and define

$$
\begin{equation*}
d=\min _{\substack{r^{\prime} \in \operatorname{aocc} k\left(P_{i-1}, T\right) \\ r^{\prime} \in\left[r+1-\Delta_{i}, r+1-\Delta_{i}+2 k\right]}}\left\{\operatorname{ed}_{i-1}\left(r^{\prime}\right)+\operatorname{ed}\left(P\left[\ell_{i-1}+1 \ldots \ell_{i}\right], T\left[r^{\prime}+1 \ldots r\right]\right)\right\} \tag{2}
\end{equation*}
$$

The smallest edit distance from $P_{i}$ to $T[1 \ldots r]$ is equal to $d$ if $d \leq k$ and is larger than $k$ otherwise.
It follows that to decide whether $r \in \operatorname{aOCC}_{k}^{E}\left(P_{i}, T\right)$, it suffices to compute the value $\min \{d, k+1\}$, where $d$ is as defined above.

Case 2: $P_{i-1}$ is not $k$-periodic. In this case, $\mathrm{aOCC}_{k}^{E}\left(P_{i-1}, T\right)$ is stored as $\mathcal{O}\left(k^{3}\right)$ chains of size one. Therefore, we can find the positions $r^{\prime}$ from Eq. (2) in $\mathcal{O}\left(k^{3}\right)$ time. Moreover, for each position $r^{\prime}$, we run the dynamic programming algorithm for $T\left[r^{\prime}+1 ..\right]$ and $P\left[\ell_{i-1}+1 \ldots \ell_{i}\right]$, which outputs the edit distance between $T\left[r^{\prime}+1 \ldots r\right]$ and $P\left[\ell_{i-1}+1 \ldots \ell_{i}\right]$ if it is at most $k$. As we also know the smallest edit distance between $P_{i-1}$ and a suffix of $T\left[1 \ldots r^{\prime}\right]$, we can compute $d$ in $\mathcal{O}\left(k^{3}\right)$ time.

Case 3: $P_{i-1}$ and $P_{i}$ are $k$-periodic. We can identify all positions $r^{\prime}$ from Eq. (2) in $\mathcal{O}\left(k^{3}\right)$ time. (It suffices to check each of the $\mathcal{O}\left(k^{3}\right)$ chains that we store for $P_{i-1}$ ). We must now test each of these positions. Consider a position $r^{\prime}$ and let $\mathcal{C}$ be the chain containing it. It suffices to compute the edit distance between $P\left[\ell_{i-1}+1 \ldots \ell_{i}\right]$ and $T\left[r^{\prime}+1 \ldots r\right]$ as we already know the smallest edit distance from $P_{i-1}$ to a suffix of $T\left[1 \ldots r^{\prime}\right]$. If $|\mathcal{C}|=1$, the distance has been computed by the dynamic programming algorithm. Otherwise, we use quasi-greedy encodings. On a high level, our goal is to compute the edit distance between $\pi=P\left[\ell_{i-1} \ldots \ell_{i}\right]$ and $\tau=T\left[r^{\prime}+1 \ldots r\right]$ via a string $\mu=D^{\infty}\left[1 \ldots \Delta_{i}\right]$, where $D$ is the difference of $\mathcal{C}$ and $\Delta_{i}=\ell_{i}-\ell_{i-1}+k$.

Lemma 7.4. There is $\operatorname{ed}(\pi, \mu) \leq 26 k$.
Proof. As $P_{i-1}$ and $P_{i}$ are $k$-periodic, by Claim 7.1 we obtain that $P_{i}=P\left[1 \ldots \ell_{i}\right]$ is $k$-periodic with $k$-period $Q_{i}=Q_{i-1}$, that is, there is a prefix of $Q_{i}^{\infty}$ such that the edit distance between it and $P_{i}$ is at most $2 k$. By Fact 2.3, there is a substring $Q_{i}^{\infty}[r . . t]$ such that $\left|r-\ell_{i-1}\right| \leq 2 k$ and $\left|t-\ell_{i}\right| \leq 2 k$ and ed $\left(Q_{i}^{\infty}[r \ldots t], \pi\right) \leq 2 k$. By the triangle inequality, we obtain that ed $\left(Q_{i}^{\infty}\left[\ell_{i-1}+\right.\right.$ $\left.\left.1 . . \ell_{i}\right], \pi\right) \leq 6 k$. Let $a=\ell_{i-1}-7 k\left(\bmod q_{i}\right)$ and $b=\ell_{i-1}+7 k\left(\bmod q_{i}\right)$. By Corollary 3.5, $D$ is a rotation of $Q_{i-1}=Q_{i}$ with shift $\Delta$, where $\Delta \in[a-3 k, b+3 k]$ if $a \leq b$ and $\Delta \in[0, b+3 k] \cup[a-$ $\left.3 k, q_{i}\right)$ if $a>b$. It follows that $D^{\infty}=Q_{i}^{\infty}[s .$.$] , where \left|s-\ell_{i-1}\right| \leq 10 k$. As $\mu=D^{\infty}\left[1 . . \Delta_{i}\right]=$ $Q_{i}^{\infty}\left[s \ldots s+\Delta_{i}-1\right]$, by the triangle inequality we obtain ed $\left(\mu, Q_{i}^{\infty}\left[\ell_{i-1}+1 \ldots \ell_{i}\right]\right) \leq 20 k$. Applying the triangle inequality one more time, we obtain the claim.

Let $\mathcal{G}_{P}=\mathrm{qGR}_{30 k}(\pi, \mu)$ and $\mathcal{G}_{T}=\mathrm{qGR}$ 30k $(\mu, \tau)$. By Corollary 5.27, knowing $\mathcal{G}_{P}$ and $\mathcal{G}_{T}$ is sufficient to compute the edit distance between $\pi$ and $\tau$. Note that we do not know $\mathcal{G}_{T}$ yet, we must compute it using the available information. Let $p$ be the rightmost position in $\mathcal{C}$.

1. Recall that at the position $(p+1)$ we launched an algorithm that is computing $\mathrm{q} \mathrm{GR}_{32 k}(T[p+$ $\left.\left.1 . . p+\Delta_{i}\right], D^{\infty}\left[1 . . \Delta_{i}\right]\right)$ with a delay of $k$ characters (Corollary 5.22). We have $p+\Delta_{i} \geq$ $r^{\prime}+\Delta_{i} \geq r$. Therefore, upon reaching $r$, we can use the memory of the algorithm
to compute $\mathcal{G}_{T, 1}=\mathrm{qGR}_{32 k}\left(T[p+1 \ldots r], D^{\infty}[1 \ldots r-p]\right)$ in $\tilde{\mathcal{O}}\left(k^{5}\right)$ time and $\tilde{\mathcal{O}}\left(k^{2}\right)$ space (Lemma 5.21).
2. By the definition of chains, $T\left[r^{\prime}+1 \ldots p\right]=D^{j}$ for some integer $j$. By Lemma 5.21, we can use $\mathrm{qGR}_{32 k}(D, D)$ computed during the preprocessing step to compute $\mathcal{G}_{T, 2}=$ $\mathrm{q} \mathrm{GR}_{32 k}\left(T\left[r^{\prime}+1 \ldots p\right], D^{j}\right)$ in $\tilde{\mathcal{O}}\left(k^{5}\right)$ time and $\tilde{\mathcal{O}}\left(k^{2}\right)$ space. By applying Lemma 5.21 again, we can compute $\mathrm{qGR}_{32 k}\left(T\left[r^{\prime}+1 \ldots p\right], D^{j} D^{\infty}[1 \ldots r-p]\right)$ in $\tilde{\mathcal{O}}\left(k^{5}\right)$ time and $\tilde{\mathcal{O}}\left(k^{2}\right)$ space.
3. Finally, we compute via Corollary 5.22 the encoding $\mathrm{qGR}_{30 k}\left(\varepsilon, D^{\infty}\left[r-p+1 \ldots \Delta_{i}\right]\right)$, where $\varepsilon$ is the empty string and $\Delta_{i}-(r-p) \leq 2 k$. We then apply Observation 5.20 to compute $\mathrm{qGR}_{30 k+\left(\Delta_{i}-(r-p)\right)}\left(T\left[r^{\prime}+1 \ldots p\right], D^{j} D^{\infty}[1 \ldots r-p]\right)$ and further Lemma 5.21 to compute $\mathrm{qGR}_{30 k}\left(T\left[r^{\prime}+1 \ldots p\right], D^{\infty}\left[1 \ldots \Delta_{i}\right]\right)=\mathcal{G}_{T}$.

Updating the chains. When we detect a new $k$-edit occurrence of $P_{i}$, we must decide if it should be added to some existing chain or if we must create a new chain for this occurrence.

To this end, for each $1 \leq i \leq z$, we consider $\mathcal{O}(k)$ of its rotations of $Q_{i}$ that can be the difference of a chain of $k$-edit occurrences of $P_{i}$ in $T$ (Corollary 3.5). For each rotation $R$, we run a constant-space and linear-time deterministic pattern matching algorithm [54]. The algorithm processes the text $T$ as a stream and if there is an occurrence $T[\ell \ldots r]$ of the rotation reports it while reading $T[r]$. The algorithm uses $\mathcal{O}(1)$ space and $\mathcal{O}(1)$ amortised time per character of $T$.

Suppose that we detect a new right endpoint $r$ of a $k$-edit occurrence $T[\ell . . r]$ of $P_{i}$. We must decide whether $r$ belongs to an existing chain of $k$-edit occurrences of $P_{i}$ or starts a new one. In order to do this, we first find the chain $\mathcal{C}$ that contains $r-q_{i}+1$ if it exists by checking each chain in turn. We then check that the smallest edit distance from a suffix of $T[1 . . r]$ to $P_{i}$ equals to ed $(\mathcal{C})$ and that $T\left[r-q_{i}+1 \ldots r\right]$ is equal to the difference of the chain. (Recall that we run the exact pattern matching algorithm for each rotation of $Q_{i}$ that can be the difference of a chain so that both checks can be performed in $\mathcal{O}(1)$ time $)$.

If these conditions are not satisfied, we create a new chain that contains $r$ only. Otherwise, we add $r$ to $\mathcal{C}$ (i.e., increment the size of $\mathcal{C}$ ). To finalize the update of the chains, we must delete all $k$-edit occurrences that become inactive: for each $i$ and for each chain of $k$-edit occurrences of $P_{i}$, we "delete" the first $k$-edit occurrence if it starts before $r-\ell_{i}-k+1$. To "delete" a $k$-edit occurrence, we simply update the endpoints of the first occurrence in the chain and the total number of occurrences in the chain.

### 7.2.3 Analysis

We summarize the results of this section:
Theorem 7.5. Assume a read-only pattern $P$ of length $m$ and a streaming text $T$ of length $n$. There is a deterministic algorithm that finds the set $\mathrm{OCC}_{k}(P, T)$ using $\tilde{\mathcal{O}}\left(k^{5}\right)$ space and $\tilde{\mathcal{O}}\left(k^{6}\right)$ amortised time per character of the text $T$.

Proof. Let us first upper bound the space complexity of the algorithm. For each $i=1, \ldots, z=$ $\mathcal{O}(\log m)$, we store the set aOCC ${ }_{k}\left(P_{i}, T\right)$ as $\mathcal{O}\left(k^{3}\right)$ chains. For each chain, we launch the dynamic programming algorithm (Fact 7.2), which takes $\tilde{\mathcal{O}}\left(k^{2}\right)$ space and Corollary 5.22 that takes $\tilde{\mathcal{O}}\left(k^{2}\right)$ space. The pattern matching algorithms for the rotations of $Q_{i}$ take $\tilde{\mathcal{O}}(k)$ space in total. Finally, testing if a position of the text is the rightmost position of a $k$-edit occurrence of $P_{i}$ requires $\tilde{\mathcal{O}}\left(k^{2}\right)$ space.

We now show the time bound. Updating the chains takes $\tilde{\mathcal{O}}(1)$ time. At any time, we run $\tilde{\mathcal{O}}\left(k^{3}\right)$ instances of Corollary 5.22 that takes $\mathcal{O}\left(k^{3}\right)$ amortised time per character. To test each position $r$, we spend $\tilde{\mathcal{O}}\left(k \cdot k^{3}\right)$ time.

### 7.3 Streaming Algorithm

We now modify the algorithm for the semi-streaming model to develop a fully streaming algorithm. W.l.o.g., assume $k \leq m$ and take $\delta=1 / n^{c}$ for $c$ large enough.

### 7.3.1 Preprocessing

We define the prefixes $P_{i}$ and their periods $Q_{i}$ exactly in the same way as in Section 7.2. Recall that $\ell_{i}=\left|P_{i}\right|$ and $q_{i}=\left|Q_{i}\right|$. For every $i>1$, we store the following information, where all sketches $\mathrm{sk}^{\mathrm{q}}$ (Corollary 6.23) are parametrized by probability $\delta$, maximal length $\Delta_{i}$, the alphabet of $P$ and $T$, and a seed of $\mathcal{O}\left(\log ^{2} n \log (1 / \delta)\right)$ random bits:

1. $P\left[\ell_{i}\right]$ and the sketch $\operatorname{sk}_{k}^{\mathrm{q}}\left(P\left[\ell_{i-1} \ldots \ell_{i}\right]\right)$;
2. For each of rotation $D$ of $Q_{i}$ that can be a difference of a chain of $k$-edit occurrences of $P_{i}$, the encoding $\mathrm{qGR}_{30 k}\left(P\left[\ell_{i-1}+1 \ldots \ell_{i}\right], D^{\infty}\left[1 \ldots \Delta_{i}\right]\right)$, where $\Delta_{i}=\ell_{i}-\ell_{i-1}+k$;
3. For each of rotation $D$ of $Q_{i}$ that can be a difference of a chain of $k$-edit occurrences of $P_{i}$, the sketches $\mathrm{sk}_{32 k}^{\mathrm{q}}(D)$ and $\mathrm{sk}_{32 k}^{\mathrm{q}}\left(D\left[1 \ldots \Delta_{i}\left(\bmod q_{i}\right)\right]\right)$.

### 7.3.2 Main Stage

As in Section 7.2, for each $i$, we store $\mathrm{aOCC}_{k}^{E}\left(P_{i-1}, T\right)$ in $\mathcal{O}\left(k^{3}\right)$ chains. For each chain $\mathcal{C}$, we store the leftmost position $l p$ in it, its size $|\mathcal{C}|$, and the smallest edit distance, ed $(\mathcal{C})$, from a suffix of $T[1 \ldots l p]$ to $P_{i-1}$. If $|\mathcal{C}| \geq 2$, we also store its difference (defined by the shift of the rotation of $Q_{i-1}$ ).

If $p^{*}$ is the first position added to $\mathcal{C}$, then at the position $\left(p^{*}+1\right)$, we start running the streaming algorithm of Corollary 6.23(a) for computing the sketch $\mathrm{sk}_{k}^{\mathrm{q}}\left(T\left[p^{a} s t+1 . . p^{a} s t+\Delta_{i}\right]\right)$.

Furthermore, consider the moment when we detect the second position in $\mathcal{C}$ (if it exists) and hence the difference $D$ of the chain. Starting from this moment, for every newly added position $p \in \mathcal{C}$, at the position $(p+1)$ we start computing the quasi-greedy encoding $\mathrm{q}_{\mathrm{GR}}^{32 k}(T[p+1 \ldots p+$ $\left.\left.\Delta_{i}\right], D^{\infty}\left[1 \ldots \Delta_{i}\right]\right)$ as follows: Assume that we have computed $\mathrm{qGR}_{32 k}\left(T\left[p+1 \ldots p+\ell \cdot q_{i}\right], D^{\infty}[1 \ldots \ell\right.$. $\left.q_{i}\right]$ ). Suppose first that $(\ell+1) \cdot q_{i} \leq \Delta_{i}$. While reading $T\left[p+\ell \cdot q_{i}+1 \ldots p+(\ell+1) \cdot q_{i}\right]$, we compute $\mathrm{sk}_{32 k}^{\mathrm{q}}\left(T\left[p+\ell \cdot q_{i}+1 \ldots p+(\ell+1) \cdot q_{i}\right]\right)$ again via the algorithm of Corollary 6.23(a). We then use this sketch and $\mathrm{sk}_{32 k+1}^{\mathrm{q}}(D)$ to compute $\mathrm{qGR}_{32 k}\left(T\left[p+\ell \cdot q_{i}+1 \ldots p+(\ell+1) \cdot q_{i}\right], D\right)$ (Corollary 6.23(b)) and then $\operatorname{qGR}_{32 k}\left(T\left[p+1 \ldots p+(\ell+1) \cdot q_{i}\right], D^{\infty}\left[1 \ldots(\ell+1) \cdot q_{i}\right]\right)($ Lemma 5.21$)$. If $(\ell+1) \cdot q_{i}>\Delta_{i}$, we use the sketches $\mathrm{sk}_{32 k}^{\mathrm{q}}\left(T\left[p+\ell \cdot q_{i} \ldots p+\Delta_{i}\right]\right)$ and $\mathrm{sk}_{32 k}^{\mathrm{q}}\left(D\left[1 \ldots \Delta_{i}\left(\bmod q_{i}\right)\right]\right)$, the rest is analogous. We continue running the algorithm until either the computation is completed or a new $k$-edit occurrence in the chain has been detected. In other words, in the end we compute the encoding for the rightmost position in the chain.

Detecting new $k$-edit occurrences of $P_{i}$. We now explain how we modify the algorithm for detecting new $k$-edit occurrences of the prefixes $P_{i}$. The algorithm for Case 1 does not change. Instead of the dynamic programming algorithm in Case 2, we use $\mathrm{sk}_{k}^{\mathrm{q}}$ and use Corollary 6.23(b) and then Corollary 5.27 to compute the edit distance. It remains to explain how we modify the algorithm for Case 3.

We exploit Eq. (2) again. We can find all positions $r^{\prime}$ in $\mathcal{O}\left(k^{3}\right)$ time. To test a position $r^{\prime}$, it suffices to compute the edit distance between $P\left[\ell_{i-1}+1 \ldots \ell_{i}\right]$ and $T\left[r^{\prime}+1 \ldots r\right]$. Let $\mathcal{C}$ be the chain with difference $D$ that contains $r^{\prime}$. If $|\mathcal{C}|=1$, we use the edit distance sketch, and otherwise the quasi-greedy encodings. By Lemma 7.4 , ed $\left(P\left[\ell_{i-1} \ldots \ell_{i}\right], D^{\infty}\left[1 \ldots \Delta_{i}\right]\right) \leq 26 k$ and hence by Corollary 5.27, the edit distance between $P\left[\ell_{i-1}+1 \ldots \ell_{i}\right]$ and $T\left[r^{\prime}+1 \ldots r\right]$ can be computed from the encodings $\mathcal{G}_{P}=\mathrm{qGR} \mathrm{K}_{30 k}(\pi, \mu)$ and $\mathcal{G}_{T}=\mathrm{qGR} \mathrm{B}_{30 k}(\mu, \tau)$ for $\pi=P\left[\ell_{i-1} \ldots \ell_{i}\right]$, $\mu=D^{\infty}\left[1 \ldots \Delta_{i}\right]$, and $\tau=T\left[r^{\prime}+1 \ldots r\right] . \mathcal{G}_{P}$ was computed during the preprocessing step and we
store it explicitly. Hence, we only need to explain how to compute $\mathcal{G}_{T}$. Let $p$ be the rightmost position in the chain $\mathcal{C}$.

1. Recall that $T\left[r^{\prime}+1 \ldots p\right]=D^{j}$ for $j=\left(p-r^{\prime}\right) / q_{i}$. We first compute $\mathrm{qGR}_{32 k}(D, D)$ from $\operatorname{sk}_{32 k}^{\mathrm{q}}(D, D)$ via Corollary $6.23(\mathrm{~b})$, and then $\mathcal{G}_{T, 1}=\mathrm{qGR}_{32 k}\left(T\left[r^{\prime}+1 \ldots p\right], D^{j}\right)=$ $\mathrm{qGR} \mathrm{B}_{32 k}\left(D^{j}, D^{j}\right)$ in $\tilde{\mathcal{O}}\left(k^{5}\right)$ time and $\tilde{\mathcal{O}}\left(k^{2}\right)$ space as in Section 7.2.
2. At the position $(p+1)$ we launched the streaming algorithm computing $\mathrm{q}_{\mathrm{GR}}^{32 k}(T[p+$ $\left.\left.1 . . p+\Delta_{i}\right], D^{\infty}\left[1 \ldots \Delta_{i}\right]\right)$ with a delay of $q_{i}$ characters. We have $p+\Delta_{i} \geq r^{\prime}+\Delta_{i} \geq r$. Therefore, at a position $p+\ell \cdot q_{i}$, where $\ell=\left\lfloor(r-p+1) / q_{i}\right\rfloor$, the algorithm computes $\mathcal{G}_{T, 2}=\mathrm{qGR}_{32 k}\left(T\left[p+1 \ldots p+\ell \cdot q_{i}\right], D^{\infty}\left[1 \ldots \ell \cdot q_{i}\right]\right)$. The algorithm then continues to compute the sketch $\mathrm{sk}_{32 k}^{\mathrm{q}}\left(T\left[p+\ell \cdot q_{i}+1 . . r\right]\right)$. We use this sketch and the sketch $\mathrm{sk}_{32 k}^{\mathrm{q}}\left(D\left[1 . . \Delta_{i}\right.\right.$ $\left.\left.\left(\bmod q_{i}\right)\right]\right)$ to compute $\mathcal{G}_{T, 3}=\mathrm{qGR}_{30 k}\left(T\left[p+\ell \cdot q_{i}+1 \ldots r\right], D\left[1 \ldots \Delta_{i}\left(\bmod q_{i}\right)\right]\right)$ via Corollary 6.23(b) and Observation 5.20.
3. We finally concatenate $\mathcal{G}_{T, 1}$ and $\mathcal{G}_{T, 2}$ to obtain $\mathrm{qGR}_{32 k}\left(T\left[r^{\prime}+1 \ldots r^{\prime}+(\ell+j) \cdot q_{i}\right], D^{\infty}[1 \ldots(\ell+\right.$ $\left.\left.j) \cdot q_{i}\right]\right)$, and then $\mathrm{qGR}_{32 k}\left(T\left[r^{\prime}+1 \ldots r^{\prime}+(\ell+j) \cdot q_{i}\right], D^{\infty}\left[1 \ldots(\ell+j) \cdot q_{i}\right]\right)$ and $\mathcal{G}_{T, 3}$ to obtain $\mathcal{G}_{T}$ via Lemma 5.21 (note that the difference of lengths of strings in $\mathcal{G}_{T, 3}$ is bounded by $2 k$ ).

Updating the chains. When we detect a new $k$-edit occurrence of $P_{i}$, we must decide if it should be added to some existing chain or if we must create a new chain for this occurrence. We use the algorithm of Section 7.2, but replace the constant-space pattern matching algorithm with the streaming pattern matching algorithm [52] that for a rotation of $Q_{i}$ takes $\tilde{\mathcal{O}}(1)$ space and $\tilde{\mathcal{O}}(1)$ time per character and retrieves all its occurrences correctly with probability at least $1-\delta$.

### 7.3.3 Analysis

We summarize the results of this section:
Theorem 7.6. Given a pattern $P$ of length $m$ and a text $T$ of length $n$. There is a streaming algorithm that finds the set $\operatorname{OCC}_{k}^{E}(P, T)$ using $\tilde{\mathcal{O}}\left(k^{5}\right)$ space and $\tilde{\mathcal{O}}\left(k^{8}\right)$ amortised time per character of the text $T$. The algorithm computes $\operatorname{OCC}_{k}^{E}(P, T)$ correctly with high probability.

Proof. By Corollary 3.5, for a fixed $i$ only $\mathcal{O}(k)$ rotations of $Q_{i}$ can be a difference of a chain of occurrences of $P_{i}$. The sketch $\mathbf{s k}_{30 k+1}^{\mathrm{q}}(\cdot)$ takes $\tilde{\mathcal{O}}\left(k^{2}\right)$ space (Corollary 6.23(a)) and $\mathrm{q}^{\left(\mathrm{GR}_{30 k}(\cdot, \cdot)\right.}$ $\tilde{\mathcal{O}}\left(k^{2}\right)$ space as well (Corollary 5.17). Therefore, the information computed during the preprocessing stage occupies $\tilde{\mathcal{O}}\left(k^{3}\right)$ space. During the main stage, we store $\tilde{\mathcal{O}}\left(k^{3}\right)$ chains. For each chain, we run the algorithm of Corollary $6.23(\mathrm{a})$ that takes $\tilde{\mathcal{O}}\left(k^{2}\right)$ space. The algorithm than computes the quasi-greedy encoding (Corollary $6.23(\mathrm{~b})$ ) takes $\tilde{\mathcal{O}}\left(k^{2}\right)$ space as well. In total, the information we store for the chains occupies $\tilde{\mathcal{O}}\left(k^{5}\right)$ space. When checking for new occurrences, we apply Lemma 5.21 and Corollary 5.27 , which require an overhead of $\tilde{\mathcal{O}}\left(k^{2}\right)$ space. Finally, the streaming pattern matching algorithms for the rotations of $Q_{i}$ that can be differences of occurrences of $P_{i}$ take $\tilde{\mathcal{O}}(k)$ space in total. The space bound follows.

At any time, we run $\tilde{\mathcal{O}}\left(k^{3}\right)$ instances of the algorithm of Corollary 6.23(a) that takes $\tilde{\mathcal{O}}(k)$ amortised time per character. In addition, for every character we run $\tilde{\mathcal{O}}\left(k^{3}\right)$ instances of the algorithms of Lemma 5.21 and Corollary 5.27 taking $\tilde{\mathcal{O}}\left(k^{8}\right)$ time in total. The pattern matching algorithms for the rotations of $Q_{i}$ take $\tilde{\mathcal{O}}(k)$ time per character.

Note that the only probabilistic procedures in the algorithm are streaming pattern matching [52] and that of Corollary 6.23 (b) that computes quasi-greedy encodings. These procedures are called $\operatorname{poly}(n, k)=\operatorname{poly}(n)$ times. By choosing the constant $c$ in $\delta=1 / n^{c}$ large enough, we can guarantee that the algorithm is correct with high probability by the union bound.

## References

[1] Amir Abboud, Thomas Dueholm Hansen, Virginia Vassilevska Williams, and Ryan Williams. Simulating branching programs with edit distance and friends: Or: A polylog shaved is a lower bound made. In STOC 2016, pages 375-388, 2016.
[2] Karl R. Abrahamson. Generalized string matching. SIAM J. Comput., 16(6):1039-1051, 1987. doi:10.1137/0216067.
[3] Amihood Amir, Moshe Lewenstein, and Ely Porat. Faster algorithms for string matching with $k$ mismatches. Journal of Algorithms, 50(2):257275, 2004. URL: https://doi.org/10.1016/S0196-6774(03)00097-X, doi:10.1016/S0196-6774(03)00097-X.
[4] Ajesh Babu, Nutan Limaye, Jaikumar Radhakrishnan, and Girish Varma. Streaming algorithms for language recognition problems. Theor. Comput. Sci., 494:13-23, 2013.
[5] Arturs Backurs and Piotr Indyk. Edit distance cannot be computed in strongly subquadratic time (unless SETH is false). In STOC 2015, pages 51-58, 2015. doi:10.1145/2746539.2746612.
[6] Djamal Belazzougui and Qin Zhang. Edit distance: Sketching, streaming, and document exchange. In FOCS 2016, pages 51-60, 2016. doi:10.1109/FOCS.2016.15.
[7] Dany Breslauer and Zvi Galil. Real-time streaming string-matching. ACM Trans. Algorithms, 10(4):22:1-22:12, 2014. doi:10.1145/2635814.
[8] Karl Bringmann and Marvin Künnemann. Quadratic conditional lower bounds for string problems and dynamic time warping. In FOCS 2015, pages 79-97, 2015. doi:10.1109/FOCS.2015.15.
[9] Diptarka Chakraborty, Debarati Das, and Michal Koucký. Approximate online pattern matching in sublinear time. In FSTTCS 2019, volume 150 of LIPIcs, pages 10:1-10:15, 2019. doi:10.4230/LIPIcs.FSTTCS.2019.10.
[10] Diptarka Chakraborty, Elazar Goldenberg, and Michal Koucký. Streaming algorithms for embedding and computing edit distance in the low distance regime. In STOC 2016, pages 712-725, 2016. doi:10.1145/2897518.2897577.
[11] Timothy M. Chan, Shay Golan, Tomasz Kociumaka, Tsvi Kopelowitz, and Ely Porat. Approximating text-to-pattern Hamming distances. In STOC 2020, pages 643-656, 2020. URL: https://doi.org/10.1145/3357713.3384266, doi:10.1145/3357713.3384266.
[12] Panagiotis Charalampopoulos, Tomasz Kociumaka, and Philip Wellnitz. Faster approximate pattern matching: A unified approach. In FOCS 2020, pages 978-989, 2020. doi:10.1109/FOCS46700.2020.00095.
[13] Raphaël Clifford, Allyx Fontaine, Ely Porat, Benjamin Sach, and Tatiana Starikovskaya. Dictionary matching in a stream. In ESA 2015, volume 9294 of $L N C S$, pages 361-372, 2015. doi:10.1007/978-3-662-48350-3\_31.
[14] Raphaël Clifford, Allyx Fontaine, Ely Porat, Benjamin Sach, and Tatiana Starikovskaya. The $k$-mismatch problem revisited. In SODA 2016, pages 2039-2052, 2016. doi:10.1137/1.9781611974331.ch142.
[15] Raphaël Clifford, Markus Jalsenius, and Benjamin Sach. Cell-probe bounds for online edit distance and other pattern matching problems. In SODA 2015, pages 552-561, 2015. doi:10.1137/1.9781611973730.37.
[16] Raphaël Clifford, Tomasz Kociumaka, and Ely Porat. The streaming k-mismatch problem, 2019. doi:10.1137/1.9781611975482.68.
[17] Richard Cole and Ramesh Hariharan. Approximate string matching: A simpler faster algorithm. SIAM J. Comput., 31(6):1761-1782, 2002. doi:10.1137/S0097539700370527.
[18] Funda Ergün, Elena Grigorescu, Erfan Sadeqi Azer, and Samson Zhou. Streaming periodicity with mismatches. In APPROX 2017, pages 42:1-42:21, 2017.
[19] Funda Ergün, Elena Grigorescu, Erfan Sadeqi Azer, and Samson Zhou. Periodicity in data streams with wildcards. In CSR 2018, pages 90-105, 2018. doi:10.1007/978-3-319-90530-3_9.
[20] Funda Ergün, Hossein Jowhari, and Mert Sağlam. Periodicity in streams. In APPROX 2010, pages 545-559, 2010. doi:10.1007/978-3-642-15369-3\_41.
[21] N. François, F. Magniez, M. de Rougemont, and O. Serre. Streaming property testing of visibly pushdown languages. In ESA 2016, volume 57 of LIPIcs, pages 43:1-43:17, 2016.
[22] Moses Ganardi, Danny Hucke, Daniel König, Markus Lohrey, and Konstantinos Mamouras. Automata theory on sliding windows. In STACS 2018, volume 96 of LIPIcs, pages 31:131:14, 2018. doi:10.4230/LIPIcs.STACS.2018.31.
[23] Moses Ganardi, Danny Hucke, and Markus Lohrey. Querying regular languages over sliding windows. In FSTTCS 2016, volume 65 of LIPIcs, pages 18:1-18:14, 2016. doi:10.4230/LIPIcs.FSTTCS.2016.18.
[24] Moses Ganardi, Danny Hucke, and Markus Lohrey. Randomized sliding window algorithms for regular languages. In ICALP 2018, volume 107 of LIPIcs, pages 127:1-127:13, 2018. doi:10.4230/LIPIcs.ICALP.2018.127.
[25] Moses Ganardi, Danny Hucke, and Markus Lohrey. Sliding window algorithms for regular languages. In LATA 2018, volume 10792, pages 26-35, 2018. doi:10.1007/978-3-319-77313-1 \_2.
[26] Moses Ganardi, Danny Hucke, Markus Lohrey, and Tatiana Starikovskaya. Sliding window property testing for regular languages. In ISAAC 2019, volume 149 of LIPIcs, pages 6:16:13, 2019. doi:10.4230/LIPIcs.ISAAC.2019.6.
[27] Moses Ganardi, Artur Jeż, and Markus Lohrey. Sliding windows over contextfree languages. In MFCS 2018, volume 117 of LIPIcs, pages 15:1-15:15, 2018. doi:10.4230/LIPIcs.MFCS.2018.15.
[28] Paweł Gawrychowski, Oleg Merkurev, Arseny M. Shur, and Przemyslaw Uznański. Tight tradeoffs for real-time approximation of longest palindromes in streams. Algorithmica, 81(9):3630-3654, 2019. doi:10.1007/s00453-019-00591-8.
[29] Paweł Gawrychowski, Jakub Radoszewski, and Tatiana Starikovskaya. Quasiperiodicity in streams. In CPM, volume 128 of LIPIcs, pages 22:1-22:14, 2019. doi:10.4230/LIPIcs.CPM.2019.22.
[30] Pawel Gawrychowski and Tatiana Starikovskaya. Streaming dictionary matching with mismatches. In CPM 2019, volume 128 of LIPIcs, pages 21:1-21:15, 2019. doi:10.4230/LIPIcs.CPM.2019.21.
[31] Paweł Gawrychowski and Przemysław Uznański. Towards unified approximate pattern matching for Hamming and L_1 distance. In ICALP 2018, volume 107 of LIPIcs, pages 62:1-62:13, 2018. doi:10.4230/LIPIcs.ICALP.2018.62.
[32] Shay Golan, Tomasz Kociumaka, Tsvi Kopelowitz, and Ely Porat. The streaming kmismatch problem: Tradeoffs between space and total time. In CPM 2020, volume 161 of LIPIcs, pages 15:1-15:15, 2020. doi:10.4230/LIPIcs.CPM.2020.15.
[33] Shay Golan, Tsvi Kopelowitz, and Ely Porat. Towards optimal approximate streaming pattern matching by matching multiple patterns in multiple streams. In ICALP 2018, volume 107 of LIPIcs, pages 65:1-65:16, 2018. doi:10.4230/LIPIcs.ICALP.2018.65.
[34] Shay Golan and Ely Porat. Real-time streaming multi-pattern search for constant alphabet. In ESA 2017, volume 107 of LIPIcs, pages 41:1-41:15, 2017. doi:10.4230/LIPIcs.ESA.2017.41.
[35] Tomohiro I. Longest common extensions with recompression. In CPM 2017, volume 78 of LIPIcs, pages 18:1-18:15, 2017.
[36] Russell Impagliazzo and Ramamohan Paturi. On the complexity of $k$-SAT. Journal of Computer and System Sciences, 62(2):367-375, 2001. doi:10.1006/jcss.2000.1727.
[37] Ce Jin, Jelani Nelson, and Kewen Wu. An improved sketching bound for edit distance, 2021. doi:10.4230/LIPIcs.STACS.2021.45.
[38] Richard M. Karp and Michael O. Rabin. Efficient randomized pattern-matching algorithms. IBM J. of RED, 31(2):249-260, 1987.
[39] Dominik Kempa and Tomasz Kociumaka. Resolution of the Burrows-Wheeler transform conjecture, 2020. doi:10.1109/FOCS46700.2020.00097.
[40] S. Kosaraju. Efficient string matching. 1987.
[41] Gad M. Landau, Eugene W. Myers, and Jeanette P. Schmidt. Incremental string comparison. SIAM J. Comput., 27(2):557-582, April 1998.
[42] Gad M. Landau and Uzi Vishkin. Efficient string matching with $k$ mismatches. Theoretical Computer Science, 43:239-249, 1986. doi:10.1016/0304-3975 (86) 90178-7.
[43] Gad M. Landau and Uzi Vishkin. Fast parallel and serial approximate string matching. J. Algorithms, 10(2):157-169, 1989. doi:10.1016/0196-6774(89)90010-2.
[44] David Levin and Yuval Peres. Markov Chains and Mixing Times. American Mathematical Society, 2017. doi:10.1090/mbk/107.
[45] Frédéric Magniez, Claire Mathieu, and Ashwin Nayak. Recognizing well-parenthesized expressions in the streaming model. SIAM J. Comput., 43(6):1880-1905, 2014. doi:10.1137/130926122.
[46] William J. Masek and Michael S. Paterson. A faster algorithm computing string edit distances. Journal of Computer and System Sciences, 20(1):18-31, 1980. doi:https://doi.org/10.1016/0022-0000(80)90002-1.
[47] O. Merkurev and A. M. Shur. Searching runs in streams. In SPIRE 2019, volume 11811 of $L N C S$, pages 203-220, 2019.
[48] Oleg Merkurev and Arseny M. Shur. Searching long repeats in streams. In CPM 2019, volume 128 of LIPIcs, pages 31:1-31:14, 2019.
[49] Gonzalo Navarro. A guided tour to approximate string matching. ACM Comput. Surv., $33(1): 31-88$, March 2001. doi:10.1145/375360. 375365.
[50] Noam Nisan. Pseudorandom generators for space-bounded computation. Comb., 12(4):449461, 1992. doi:10.1007/BF01305237.
[51] Takaaki Nishimoto, Tomohiro I, Shunsuke Inenaga, Hideo Bannai, and Masayuki Takeda. Dynamic index and LZ factorization in compressed space. Discrete Applied Mathematics, 274:116-129, 2020. Stringology Algorithms. doi:https://doi.org/10.1016/j.dam.2019.01.014.
[52] Benny Porat and Ely Porat. Exact and approximate pattern matching in the streaming model. In FOCS 2009, pages 315-323, 2009. doi:10.1109/FOCS.2009.11.
[53] Jakub Radoszewski and Tatiana Starikovskaya. Streaming $k$-mismatch with error correcting and applications. Journal of Information and Computation, 271:104513, 2020. doi:10.1016/j.ic.2019.104513.
[54] Wojciech Rytter. On maximal suffixes and constant-space linear-time versions of KMP algorithm. Theoretical Computer Science, 299(1):763 - 774, 2003. doi:https://doi.org/10.1016/S0304-3975(02)00590-X.
[55] Süleyman Cenk Sahinalp and Uzi Vishkin. Efficient approximate and dynamic matching of patterns using a labeling paradigm (extended abstract). In FOCS 1996, pages 320-328, 1996. doi:10.1109/SFCS. 1996.548491.
[56] Peter H. Sellers. The theory and computation of evolutionary distances: Pattern recognition. Journal of Algorithms, 1(4):359 - 373, 1980. doi:https://doi.org/10.1016/0196-6774(80)90016-4.
[57] Tatiana Starikovskaya. Communication and streaming complexity of approximate pattern matching. In CPM 2017, volume 78 of LIPIcs, pages 13:1-13:11, 2017. doi:10.4230/LIPIcs.CPM.2017.13.
[58] Alexandre Tiskin. Semi-local string comparison: algorithmic techniques and applications, 2013. arXiv:0707.3619v21.
[59] Esko Ukkonen. Algorithms for approximate string matching. Information and Control, 64(1-3):100-118, 1985. doi:10.1016/S0019-9958(85)80046-2.


[^0]:    *Partly supported by NSF 1652303, 1909046, and HDR TRIPODS 1934846 grants, and an Alfred P. Sloan Fellowship.
    ${ }^{\dagger}$ Partly supported by the grant ANR-20-CE48-0001 from the French National Research Agency (ANR).

[^1]:    ${ }^{1}$ Hereafter, $\tilde{\mathcal{O}}(\cdot)$ hides a factor of poly $(\log n)$.

[^2]:    ${ }^{2}$ With high probability means with probability at least $1-1 / n^{c}$ for any predefined constant $c>1$.

[^3]:    ${ }^{3}$ This definition is rather complex, but it is equivalent to the standard definition given in the textbooks. We chose this particular formulation as it allowed us to introduce notions essential for this work in a rigorous way.

[^4]:    ${ }^{4}$ Pairs of strings of any lengths can be supported in the same way provided that concatenations require larger input thresholds (compared to the output threshold) to accommodate length differences.
    ${ }^{5}$ In fact, $\mathcal{O}\left(k^{2}\right)$ edit distances suffice and they can be encoded in $\mathcal{O}(k)$ space using techniques of Tiskin [58].

[^5]:    ${ }^{6}$ The authors would like to thank Panagiotis Charalampopoulos for an early write-up of the proof of this lemma, which was originally meant to be included into [12].

