

1 Introduction

The *pattern matching* problem is perhaps the most fundamental problem on strings: Given a pattern P and a text T , the task is to find all occurrences of P in T . However, in most applications, finding all *exact* occurrences of a pattern is not enough: Think of human spelling mistakes or DNA sequencing errors, for example. In this work, we focus on *approximate* pattern matching, where we are interested in finding substrings of the text that are “similar” to the pattern. While various similarity measures are imaginable, we study the two most commonly encountered metrics in this context: the *Hamming distance* and the *edit distance*.

Hamming Distance

Recall that the Hamming distance of two (equal-length) strings is the number of positions where the strings differ. Now, given a text T of length n , a pattern P of length m , and an integer threshold $k > 0$, we want to compute the *k-mismatch occurrences* of P in T , that is, all length- m substrings of T that are at Hamming distance at most k from P . This *pattern matching with mismatches* problem has been extensively studied. In the late 1980s, Abrahamson [2] and Kosaraju [33] independently proposed an FFT-based $\mathcal{O}(n\sqrt{m \log m})$ -time algorithm for computing the Hamming distance of P and all the length- m fragments of T . While their algorithms can be used to solve the pattern matching with mismatches problem, the first algorithm to benefit from the threshold k was given by Landau and Vishkin [34] and slightly improved by Galil and Giancarlo [20]: Based on so-called “kangaroo jumping”, they obtained an $\mathcal{O}(nk)$ -time algorithm, which is faster than $\mathcal{O}(n\sqrt{m \log m})$ even for moderately large k . Amir et al. [4] developed two algorithms with running time $\mathcal{O}(n\sqrt{k \log k})$ and $\tilde{\mathcal{O}}(n + k^3n/m)$, respectively; the latter algorithm was then improved upon by Clifford et al. [15], who presented an $\tilde{\mathcal{O}}(n + k^2n/m)$ -time solution. Subsequently, Gawrychowski and Uznański [22] provided a smooth trade-off between the running times $\tilde{\mathcal{O}}(n\sqrt{k})$ and $\tilde{\mathcal{O}}(n + k^2n/m)$ by designing an $\tilde{\mathcal{O}}(n + kn/\sqrt{m})$ -time algorithm. Very recently, Chan et al. [13] removed most of the polylog n factors in the latter solution at the cost of (Monte-Carlo) randomization. Furthermore, Gawrychowski and Uznański [22] showed that a significantly faster “combinatorial” algorithm would have (unexpected) consequences for the complexity of Boolean matrix multiplication. Pattern matching with mismatches on strings is thus well understood in the standard setting. Nevertheless, in the settings where the strings are not given explicitly, a similar understanding is yet to be obtained. One of the main contributions of this work is to improve the upper bounds for two such settings, obtaining algorithms with running times analogous to the algorithm of Clifford et al. [15].

Edit Distance

Recall that the edit distance (also known as the Levenshtein distance) of two strings S and T is the minimum number of edits required to transform S into T . Here, an edit is an insertion, a substitution, or a deletion of a single character. In the *pattern matching with edits* problem, we are given a text T , a pattern P , and an integer threshold $k > 0$, and the task is to find the starting positions of all the *k-edit* (or *k-error*) *occurrences* of P in T . Formally, we are to find all positions i in T such that the edit distance between $T[i..j]$ and P is at most k for some position j . Again, a classic algorithm by Landau and Vishkin [35] runs in $\mathcal{O}(nk)$ time. Subsequent research [44, 17] resulted in an $\mathcal{O}(n + k^4n/m)$ -time algorithm (which is faster for $k \leq \sqrt[3]{m}$). From a lower-bound perspective, we can benefit from the discovery that the classic quadratic-time algorithm for computing the edit distance of two strings is essentially optimal: Backurs and Indyk [6] recently proved that a significantly faster algorithm would yield a major breakthrough for the satisfiability problem. For pattern matching with edits, this means that there is no hope for an algorithm that is significantly faster than $\mathcal{O}(n + k^2n/m)$; however, apart from that “trivial” lower bound

and the 20-year-old conjecture of Cole and Hariharan [17] that an $\mathcal{O}(n + k^3n/m)$ -time algorithm *should be possible*, nothing is known that would close this gap. While we do not manage to tighten this gap, we do believe that the structural insights we obtain may be useful for doing so. What we do manage, however, is to significantly improve the running time of the known algorithms in two settings where T and P are not given explicitly, thereby obtaining running times that can be seen as analogous to the running time of Cole and Hariharan’s algorithm [17].

Grammar Compression

One of the settings that we consider in this paper is the *fully compressed* setting, where both the text T and the pattern P are given as straight-line programs. Compressing the text and the pattern is, in general, a natural thing to do—think of huge natural-language texts or genomic databases, which are easily compressible. While one approach to solve pattern matching in the fully compressed setting is to first decompress the strings and then run an algorithm for the standard setting, this voids most benefits of compression in the first place. Hence, there has been a long line of research with the goal of designing text algorithms directly operating on compressed strings. Naturally, such algorithms highly depend on the chosen compression method. In this work, we consider *grammar compression*, where a string T is represented using a context-free grammar that generates the singleton language $\{T\}$; without loss of generality, such a grammar is a *straight-line program* (SLP).

Straight-line programs are popular due to mathematical elegance and equivalence [42, 30, 29] (up to logarithmic factors and moderate constants) to widely-used dictionary compression schemes, including the LZ77 parsing [50] and the run-length-encoded Burrows–Wheeler transform [12]. Many more schemes, such as byte-pair encoding [46], Re-Pair [36], Sequitur [39], and further members of the Lempel–Ziv family [51, 49], to name but a few, can be expressed as straight-line programs. We refer an interested reader to [43, 41, 37, 45] to learn more about grammar compression.

Working directly with a compressed representation of a text, intuitively at least, seems to be hard in general—in fact, Abboud et al. [1] showed that, for some problems, decompress-and-solve is the best we can hope for, under some reasonable assumptions from fine-grained complexity theory. Nevertheless, Jež [27] managed to prove that exact pattern matching can be solved on grammar-compressed strings in near-linear time: Given an SLP of size n representing a string T and an SLP of size m representing a string P , we can find all exact occurrences of P in T in $\mathcal{O}((n + m) \log |P|)$ time. For fully compressed *approximate* pattern matching, no such near-linear time algorithm is known, though. While the $\tilde{\mathcal{O}}((n + |P|)k^4)$ -time algorithm by Bringmann et al. [11] for pattern matching with mismatches comes close, it works in an easier setting where only the text is compressed. We fill this void by providing the first algorithm for fully compressed pattern matching with mismatches that runs in near-linear time. Denote by $\text{Occ}_k^H(P, T)$ the set of (starting positions of) k -mismatch occurrences of P in T ; then, our result reads as follows.

▀ **Main Theorem 1.** *Let \mathcal{G}_T denote an SLP of size n generating a text T , let \mathcal{G}_P denote an SLP of size m generating a pattern P , let $k > 0$ denote an integer threshold, and set $N := |T| + |P|$.*

Then, we can compute $|\text{Occ}_k^H(P, T)|$ in time $\mathcal{O}(m \log N + nk^2 \log^2 N \log \log N)$. The elements of $\text{Occ}_k^H(P, T)$ can be reported within $\mathcal{O}(|\text{Occ}_k^H(P, T)|)$ extra time. ▀

For pattern matching with edits, near-linear time algorithms are not known even in the case that the pattern is given explicitly. Currently, the best pattern matching algorithms on an SLP-compressed text run in time $\mathcal{O}(n|P| \log |P|)$ [48] and $\mathcal{O}(n(\min\{|P|k, k^4 + |P|\} + \log |T|))$ [8]. Moreover, an $\tilde{\mathcal{O}}(n\sqrt{|P|}k^3)$ -time solution [21] is known for (weaker) LZW compression [49]. Again, we obtain a near-linear time algorithm for fully compressed pattern matching with edits. Denote by $\text{Occ}_k^E(P, T)$ the set of all starting positions of k -error occurrences of P in T ; then, our result reads as follows.

3 Faster Approximate Pattern Matching: A Unified Approach

▮ **Main Theorem 2.** *Let \mathcal{G}_T denote an SLP of size n generating a string T , let \mathcal{G}_P denote an SLP of size m generating a string P , let $k > 0$ denote an integer threshold, and set $N := |T| + |P|$.*

Then, we can compute $|\text{Occ}_k^E(P, T)|$ in time $\mathcal{O}(m \log N + n k^4 \log^2 N \log \log N)$. The elements of $\text{Occ}_k^E(P, T)$ can be reported within $\mathcal{O}(|\text{Occ}_k^E(P, T)|)$ extra time. ▮

Note that our algorithms also improve the state of the art when the pattern is given in an uncompressed form; this is because any string P admits a trivial SLP of size $\mathcal{O}(|P|)$.

Dynamic Strings

While compression handles large static data sets, a different approach is available if the data changes frequently. Several works on pattern matching in dynamic strings considered the indexing problem, assuming that the text is maintained subject to updates and the pattern is given explicitly at query time; we refer an interested reader to [24, 19, 44, 3, 40] and references therein.

Recently, Clifford et al. [16] considered the problem of maintaining a data structure for a text T and a pattern P , both of which undergo character substitutions, in order to be able to efficiently compute the Hamming distance between P and any given fragment of T . Among other results, for the case where $|T| \leq 2|P|$ and constant alphabet size, they presented a data structure with $\mathcal{O}(\sqrt{m \log m})$ time per operation, and they proved that, conditioned on the Online Boolean Matrix-Vector Multiplication (OMv) Conjecture [25], one cannot simultaneously achieve $\mathcal{O}(m^{1/2-\varepsilon})$ for the query and update time for any constant $\varepsilon > 0$.

We consider the following more general setting: We maintain an initially empty collection of strings \mathcal{X} that can be modified via the following “update” operations:

- ▬ **makestring**(U): Insert a string U to \mathcal{X} .
- ▬ **concat**(U, V): Insert UV to \mathcal{X} , for $U, V \in \mathcal{X}$.
- ▬ **split**(U, i): Insert $U[0..i)$ and $U[i..|U|)$ in \mathcal{X} , for $U \in \mathcal{X}$ and $i \in [1..|U|)$.

The strings in \mathcal{X} are *persistent*, meaning that **concat** and **split** do not destroy their arguments.

The main goal in this model (and for the dynamic setting in general) is to provide algorithms that are faster than recomputing the answer from scratch after every update. Specifically for dynamic strings, already the task of testing equality of strings in \mathcal{X} is challenging. After a long line of research [47, 38, 3], Gawrychowski et al. [23] proved the following: There is a data structure that supports equality queries in $\mathcal{O}(1)$ time, while each of the update operations takes $\mathcal{O}(\log N)$ time, where N is an upper bound on the total length of all strings in \mathcal{X} .^a In fact, as shown in [23], one can also support $\mathcal{O}(1)$ -time queries for the longest common prefix of two strings in \mathcal{X} with no increase in the update times. The data structure of [23] is Las-Vegas randomized: the answers are correct, but the update times are guaranteed only with high probability. Randomization can be avoided at the cost of extra logarithmic factors in the running times (see [38, 40]), and the same is true for our results.

We extend the data structure of Gawrychowski et al. [23] with approximate pattern matching queries:

▮ **Main Theorem 3.** *A collection \mathcal{X} of non-empty persistent strings of total length N can be maintained subject to **makestring**(U), **concat**(U, V), and **split**(U, i) operations requiring $\mathcal{O}(\log N + |U|)$, $\mathcal{O}(\log N)$, and $\mathcal{O}(\log N)$ time, respectively, so that given two strings $P, T \in \mathcal{X}$ with $|P| = m$ and $|T| = n$ and an integer threshold $k > 0$, we can compute $|\text{Occ}_k^H(P, T)|$ in time $\mathcal{O}(n/m \cdot k^2 \log^2 N)$ and $|\text{Occ}_k^E(P, T)|$ in time $\mathcal{O}(n/m \cdot k^4 \log^2 N)$.^b The elements of $\text{Occ}_k^H(P, T)$ and $\text{Occ}_k^E(P, T)$ can be reported in $\mathcal{O}(|\text{Occ}_k^H(P, T)|)$ and $\mathcal{O}(|\text{Occ}_k^E(P, T)|)$ extra time, respectively. ▮*

^a Strictly speaking, **makestring**(U) costs $\mathcal{O}(|U| + \log N)$ time.

^b All running time bounds hold with high probability (i.e., $1 - N^{-\Omega(1)}$).

In the Hamming distance case, for $k < \sqrt{m}/\log N$, the data structure of Main Theorem 3 is faster than recomputing the occurrences from scratch after each update: Recall that, in the standard setting, the fastest known algorithm for pattern matching with mismatches costs $\tilde{O}(n + kn/\sqrt{m}) = \tilde{O}(n + k\sqrt{m} \cdot n/m)$ time; in particular, the additive $\tilde{O}(n)$ term dominates the time complexity for the considered parameter range. Observe further that, for any k , Main Theorem 3 is not slower (ignoring polylog factors) than running the $\tilde{O}(n + k^2n/m)$ -time algorithm by Clifford et al. [15] after every update.

In the edit distance case, for $k < (m/\log^2 N)^{1/3}$, the data structure of Main Theorem 3 is faster than running the $\mathcal{O}(nk)$ -time Landau–Vishkin algorithm for the standard setting. Note further that, for any k , Main Theorem 3 is not slower (ignoring polylog N factors) than running after every update the $\mathcal{O}(n + k^4n/m)$ -time Cole–Hariharan algorithm, whose bottleneck for $k < m^{1/4}$ is the additive $\mathcal{O}(n)$ term.

Structure of Pattern Matching with Mismatches

As in [11], we obtain our algorithms by exploiting new structural insights for approximate pattern matching. Our contribution in this area is two-fold: We strengthen the structural result of [11] for pattern matching with mismatches, and we prove a similar result for pattern matching with edits.

Before we describe our new structural insights, let us recall the structure of exact pattern matching. Let P denote a pattern of length m and let T denote a text of length $n \leq \frac{3}{2}m$. Assume that T is trimmed so that P occurs both at the beginning and at the end of T , that is, $P = T[0..m] = T[n-m..n]$. By the length constraints, we have $P[n-m..m] = P[0..2m-n]$. Repeating this argument for the overlapping parts of P , we obtain the following well-known characterization (where X^∞ denotes the concatenation of infinitely many copies of a string X):

▣ **Fact 1.1** (folklore [10]). *Let P denote a pattern of length m and let T denote a text of length $n \leq \frac{3}{2}m$. If $T[0..m] = T[n-m..n] = P$, then there is a string Q such that $P = Q^\infty[0..m]$ and $T = Q^\infty[0..n]$, that is, both the text and the pattern are periodic with a common string period Q , and the starting positions of all exact occurrences of P in T form an arithmetic progression with difference $|Q|$.* ▣

Hence, it is justified to say that the structure of exact pattern matching is fully understood. Surprisingly, a similar characterization for approximate pattern matching was missing for a long time. Only recently, Bringmann et al. [11] proved a similar result for pattern matching with mismatches (we write $\delta_H(S, T)$ for the Hamming distance of S and T):

▣ **Theorem 1.2** ([11, Theorem 1.2], simplified). *Given a pattern P of length m , a text T of length $n \leq \frac{3}{2}m$, and a positive integer threshold $k \leq m$, at least one of the following holds:*

- ▣ *The number of k -mismatch occurrences of P in T is bounded by $\mathcal{O}(k^2)$.*
- ▣ *There is a primitive string Q of length $\mathcal{O}(m/k)$ such that $\delta_H(P, Q^\infty[0..m]) \leq 6k$.* ▣

Motivated by the absence of examples proving the tightness of their result, Bringmann et al. [11] conjectured that the bound on the number of k -mismatch occurrences in Theorem 1.2 can be improved to $\mathcal{O}(k)$. We resolve their conjecture positively by proving the following stronger variant of Theorem 1.2.


▣ **Main Theorem 4** (Compare Theorem 1.2). *Given a pattern P of length m , a text T of length $n \leq \frac{3}{2}m$, and a positive integer threshold $k \leq m$, at least one of the following holds:*

- ▣ *The number of k -mismatch occurrences of P in T is bounded by $\mathcal{O}(k)$.*
- ▣ *There is a primitive string Q of length $\mathcal{O}(m/k)$ that satisfies $\delta_H(P, Q^\infty[0..m]) < 2k$.* ▣

Examples from [11], illustrated in Figure 1, prove the asymptotic tightness of Main Theorem 4.

As in the exact pattern matching case, we can also characterize the (approximately) periodic case in more detail.

5 Faster Approximate Pattern Matching: A Unified Approach



(a) Consider a text $T := a^{3m/4}c^{3m/4}$ and a pattern $P := a^{m/2}c^{m/2}$, neither of which is approximately periodic. Then, shifting the exact occurrence of P in T by up to k positions in either direction still yields a k -mismatch occurrence. Hence, we need $\Omega(k)$ distinct k -mismatch occurrences to derive approximate periodicity of P .

(b) Consider a text T and a pattern P obtained from $a^{3m/2}$ and a^m , respectively, by substituting a to c at $k/2$ random positions. Then, all length- m fragments of T are k -mismatch occurrences of P , but, with high probability, neither T nor P is perfectly periodic. Hence, we need a relaxed periodicity notion allowing for $\Omega(k)$ mismatches.

■ **Figure 1** Examples (1) and (2) from [11].

▮ **Main Theorem 5** (Compare [11, Claim 3.1]). *Let P denote a pattern of length m , let T denote a text of length $n \leq \frac{3}{2}m$, and let $0 \leq k \leq m$ denote an integer threshold. Suppose that both $T[0..m)$ and $T[n-m..n)$ are k -mismatch occurrences of P . If there is a positive integer $d \geq 2k$ and a primitive string Q with $|Q| \leq m/8d$ and $\delta_H(P, Q^\infty[0..m)) \leq d$, then each of following holds:*

- (a) *Every k -mismatch occurrence of P in T starts at a position that is a multiple of $|Q|$.*
- (b) *The string T satisfies $\delta_H(T, Q^\infty[0..n)) \leq 3d$.*
- (c) *The set $\text{Occ}_k^H(P, T)$ can be decomposed into $\mathcal{O}(d^2)$ arithmetic progressions with difference $|Q|$.* ▮

Note that Theorem 1.2, as originally formulated in [11], includes a weaker version of Main Theorem 5. We also observe that part (c) of the new characterization is asymptotically tight, as justified by modifying the example of Figure 1b: Let P be obtained from a^m by placing c at $(k+1)/2$ random positions, and let T be obtained from $a^{3m/2}$ by placing c at $(k+1)/2$ random positions within the middle third of $a^{3m/2}$. Then, each k -mismatch occurrence must align at least one c from P with one c from T and, conversely, each such alignment results in a k -mismatch occurrence. Hence, the number of k -mismatch occurrences is $\Theta(k^2)$. Furthermore, for every q , with high probability, $\text{Occ}_k^H(P, T)$ can only be decomposed into $\Theta(k^2)$ progressions with difference q .


Structure of Pattern Matching with Edits

Having understood the structure of pattern matching with mismatches, we turn to the more complicated situation of pattern matching with edits. First, observe that the examples of Figures 1a and 1b are still valid: Any k -mismatch occurrence is also a k -error occurrence. However, as the edit distance allows insertions and deletions of characters, we can construct an example where neither P nor T is approximately periodic, yet the number of k -error occurrences is $\Omega(k^2)$; see Figure 2. In the example of Figure 2, there are still only $\mathcal{O}(k)$ regions of size $\mathcal{O}(k)$ each where k -error occurrences start. In fact, we can show that this is the worst that can happen (we write $\delta_E(S, T)$ for the edit distance of S and T):

▮ **Main Theorem 6.** *Given a pattern P of length m , a text T of length $n \leq \frac{3}{2}m$, and a positive integer threshold $k \leq m$, at least one of the following holds:*

- ▮ *The starting positions of all k -error occurrences of P in T lie in $\mathcal{O}(k)$ intervals of length $\mathcal{O}(k)$ each.*
- ▮ *There is a primitive string Q of length $\mathcal{O}(m/k)$ and integers i, j such that $\delta_E(P, Q^\infty[i..j]) < 2k$.* ▮

Again, we treat the (approximately) periodic case separately, thereby obtaining a result similar to Main Theorem 5:



■ **Figure 2** Consider a text $T := a^{n/2} \cdot (c \cdot a^{k-1})^{n/2k}$ and a pattern $P := a^{m/2} \cdot (c \cdot a^{k-1})^{m/2k}$ for $n := m + 2k^2$. Now, for every $i \in [-k \dots k]$, an $|i|$ -mismatch occurrence of P starts at position $n/2 - m/2 + i \cdot k$ in T . The remaining budget on the number of errors can be spent on shifting the starting positions, so for every $j \in [|i| - k \dots k - |i|]$, there is a k -error occurrence starting at position $n/2 - m/2 + i \cdot k + j$ in T . Overall, the number of k -error occurrences of P in T is $\Omega(k^2)$, but neither P nor T is approximately periodic.

■ **Main Theorem 7.** *Let P denote a pattern of length m , let T denote a text of length n , and let $0 \leq k \leq m$ denote an integer threshold such that $n < \frac{3}{2}m + k$. Suppose that the k -error occurrences of P in T include a prefix of T and a suffix of T . If there is a positive integer $d \geq 2k$ and a primitive string Q satisfying $|Q| \leq m/8d$ and $\delta_E(P, Q^\infty[i \dots j]) \leq d$ for some integers i, j , then each of following holds:*

- (a) *For every $p \in \text{Occ}_k^E(P, T)$, we have $p \bmod |Q| \leq 3d$ or $p \bmod |Q| \geq |Q| - 3d$.*
- (b) *The string T satisfies $\delta_E(T, Q^\infty[i' \dots j']) \leq 3d$ for some integers i' and j' .*
- (c) *The set $\text{Occ}_k^E(P, T)$ can be decomposed into $\mathcal{O}(d^3)$ arithmetic progressions with difference $|Q|$.* ■

Technical Overview

Gaining Structural Insights

To highlight the novelty of our approach, let us first outline the proof Theorem 1.2 by Bringmann et al. [11]. Consider a pattern P of length m and a text T of length $n \leq \frac{3}{2}m$. Split the pattern into $\Theta(k)$ blocks of length $\Theta(m/k)$ each and process each such block P_i as follows: Compute the shortest string period Q_i of P_i and align P with a substring of Q_i^∞ , starting from $P_i = Q_i^\infty[0 \dots |P_i|)$ and extending to both directions, with mismatches allowed. If there are $\mathcal{O}(k)$ mismatches for any block P_i , then P is approximately periodic; otherwise, there are many mismatches for every block P_i . In particular, in every k -mismatch occurrence where a block P_i is matched exactly, all but at most k of these mismatches between P and Q_i^∞ must be aligned to the corresponding mismatches between T and Q_i^∞ . Observing that, in any k -mismatch occurrence, all but at most k of the blocks must be matched exactly, this yields an $\mathcal{O}(k^2)$ bound on the number of k -mismatch occurrences of P in T .

The main shortcoming of this approach is the initial treatment of the pattern: Since the pattern P is independently aligned with Q_i^∞ for every block P_i , the same position in P may be accounted for as a mismatch for multiple blocks P_i . In particular, this happens if several adjacent blocks share the same period. This leads to an overcounting of the k -mismatch occurrences that is hard to control.

What we do instead is a more careful analysis of the pattern. Instead of creating all blocks P_i at once, we process P from left to right, as described below. Suppose that $P[j \dots m)$ is the unprocessed suffix of P . We first consider the length- $m/8k$ prefix P' of $P[j \dots m)$ and compute its shortest string period Q . If $|Q|$ exceeds a certain constant fraction of $|P'|$, we set P' aside as a *break* and continue processing $P[j + |P'| \dots m)$. Now, if P' is the $2k$ -th break that we set aside, our process stops, and we continue to work only with the breaks. If P' does not form a break, we try extending P' to a prefix R of $P[j \dots m)$ that satisfies $\delta_H(R, Q^\infty[0 \dots |R|]) = \Theta(k \cdot |R|/m)$. If such a prefix R exists, we set it aside as a *repetitive region* and continue processing $P[j + |R| \dots m)$. Now, if all the repetitive regions

collected so far have a total length of at least $3/8 \cdot m$, we stop our process and continue to work only with the repetitive regions computed so far. A repetitive region R does not exist only if $P[j..m)$ has too few mismatches with Q^∞ . In this case, we try extending $P[j..m)$ to a suffix R' of P that satisfies $\delta_H(R', \overline{Q}^\infty[0..|R'|)) = \Theta(k \cdot |R'|/m)$; where \overline{Q} is a suitable rotation of Q . If we fail again, we report that P is approximately periodic; otherwise, we continue to work with the single repetitive region R' , disregarding the previously generated repetitive regions. For this, we note that $|R'| \geq 3/8 \cdot m$ because all breaks and repetitive regions found beforehand have a total length of at most $5/8 \cdot m$.

Overall, for every pattern P , we obtain either $2k$ disjoint breaks, or disjoint repetitive regions of total length at least $3/8 \cdot m$, or a string with period $\mathcal{O}(m/k)$ at Hamming distance $\mathcal{O}(k)$ from P (see Lemma 3.6).

If the analysis results in breaks, we observe that at least k breaks need to be matched exactly in every k -mismatch occurrence of P in T . As both the length and the shortest period of each break are $\Theta(n/k)$, there are at most $\mathcal{O}(k)$ exact matches of each break in the text. Now, a simple marking argument shows that the number of k -mismatch occurrences of P in T is $\mathcal{O}(k)$ (see Lemma 3.8).

If the analysis results in repetitive regions, for each region R_i , we consider its k_i -mismatch occurrences in T with $k_i := \Theta(k \cdot |R_i|/m)$. Intuitively, this distributes the available budget of k mismatches among the repetitive regions according to their lengths. Next, we try extending each k_i -mismatch occurrence of each R_i to an approximate occurrence of P , and we assign $|R_i|$ marks to this extension. Using insights gained in the periodic case, we bound the total number of marks by $\mathcal{O}(k \cdot \sum_i |R_i|)$. Independently, we show that each k -mismatch occurrence of P has at least $\sum_i |R_i| - m/4$ marks. Using $\sum_i |R_i| \geq 3/8 \cdot m$, we finally obtain a bound of $\mathcal{O}(k)$ on the number of k -mismatch occurrences of P in T (see Lemma 3.11).

In total, this proves Main Theorem 4. For the characterization of the periodic case (Main Theorem 5), we use a reasoning similar to that in [11]. As in the theorem, assume that P has k -mismatch occurrences both as a prefix and as a suffix of T . Further, fix a threshold $d \geq 2k$ and a primitive string Q such that $\delta_H(P, Q^\infty[0..m)) \leq d$. First, we show that every k -mismatch occurrence of P in T starts at a multiple of $|Q|$. In particular, $|Q|$ divides $n - m$ and, using this observation, we bound $\delta_H(T, Q^\infty[0..n))$. Finally, to decompose $\text{Occ}_k^H(P, T)$ into $\mathcal{O}(k^2)$ arithmetic progressions, we analyze the sequence of Hamming distances between P and the length- m fragments of T starting at the multiples of $|Q|$: we observe that the number of changes in this sequence is bounded by $\mathcal{O}(d^2)$, which then yields the claim.

For pattern matching with edits, surprisingly few modifications in our arguments are necessary. In fact, the analysis of the pattern stays essentially the same. The main difference in the subsequent arguments is that we need to account for shifts of up to $\mathcal{O}(k)$ positions; this causes the increase in the bound on the number of occurrences. Unfortunately, for the periodic case of pattern matching with edits, the situation is messier. The key difficulty that we overcome is that an alignment corresponding to a specific edit distance may not be unique. In particular, due to insertions and deletions, combining (the arguments for) two disjoint substrings is not as easy as in the Hamming distance case. We address these issues by enclosing individual errors between a string and its approximate period with so-called *locked fragments*, which admit a unique canonical alignment. (A similar idea was used by Cole and Hariharan [17].) Combining this with a more involved marking scheme, we then obtain Main Theorem 7.

A Unified Approach to Approximate Pattern Matching

The proofs of our new structural insights are already essentially algorithmic. To obtain algorithms for all the considered settings at once, we proceed in two steps. In the first step, we devise meta-algorithms that only rely on a core set of abstract operations; in the second step, we implement these operations in various settings. Specifically, we introduce the PILLAR model—a novel abstract interface to handle strings represented in a setting-specific manner. For two strings S and T , the following operations are supported:

- $\text{Extract}(S, \ell, r)$: Retrieve a string $S[\ell..r]$.
- $\text{LCP}(S, T)$: Compute the length of the longest common prefix of S and T .
- $\text{LCP}^R(S, T)$: Compute the length of the longest common suffix of S and T .
- $\text{IPM}(S, T)$: Assuming that $|T| \leq 2|S|$, compute the starting positions of all exact occurrences of S in T .
- $\text{Access}(S, i)$: Retrieve the character $S[i]$.
- $\text{Length}(S)$: Compute the length $|S|$ of the string S .

Using the PILLAR-model operations, the meta-algorithms for both pattern matching with mismatches and with errors follow the same overall structure:

- First, we implement the analysis of the pattern. Here, the key difficulty is to detect repetitive regions. Our algorithm finds the *shortest* repetitive region: Starting from the prefix P' of the unprocessed suffix $P[j..m]$, we enumerate the mismatches (or errors) between $P[j..m]$ and Q^∞ . We stop when the number of mismatches (or errors) within the constructed region R reaches $\Theta(k/m \cdot |R|)$. Intuitively, this is correct because the number of mismatches (or errors) increases at most as fast as the length of $|R|$. We treat the special case when we reach the end of the pattern symmetrically. Note that computing the next mismatch between two strings is a prime application of the LCP operation. For finding a next edit, we adapt the Landau–Vishkin algorithm [35], based on LCP operations as well.
 - Next, we deal with the periodic case. This turns out to be the main difficulty. For the Hamming distance, implementing the proof of Main Theorem 5 is rather straightforward. However, for the edit distance case, the more complicated proof of Main Theorem 7 gets complemented with even more sophisticated algorithms. Hence, we do not discuss them in this outline.
 - Finding the occurrences in the presence of $2k$ breaks is easy: We first use IPM operations to locate exact occurrences of the breaks in the text and then perform a straightforward marking step; for the Hamming distance, we lose an $\mathcal{O}(\log \log k)$ factor for sorting marks.
 - Finding the occurrences in the presence of repetitive regions is implemented similarly; the key difference is that we use our algorithm for the periodic case to find approximate occurrences of repetitive regions.
- Overall, this approach then yields the main technical results of this work (stated below for strings of arbitrary lengths):

▣ **Main Theorem 8.** *Given a pattern P of length m , a text T of length n , and a positive integer $k \leq m$, we can compute (a representation of) the set $\text{Occ}_k^H(P, T)$ using $\mathcal{O}(n/m \cdot k^2 \log \log k)$ time plus $\mathcal{O}(n/m \cdot k^2)$ PILLAR operations. ▣*

For pattern matching with edits, the number of PILLAR-model operations matches the time cost of non-PILLAR-model operations; hence the simplified theorem statement.

▣ **Main Theorem 9.** *Given a pattern P of length m , a text T of length n , and a positive integer $k \leq m$, we can compute (a representation of) the set $\text{Occ}_k^E(P, T)$ using $\mathcal{O}(n/m \cdot k^4)$ time in the PILLAR model. ▣*

Finally, we show how to implement the PILLAR model in the settings that we consider:

- As a toy example, we start with the standard setting. Here, implementing the PILLAR-model operations boils down to collecting known tools on strings.
- For the fully compressed setting, we heavily rely on the recompression technique by Jež [27, 28] (especially for internal pattern matching queries) and on other works on straight-line programs [8, 26].
- Finally, for the dynamic setting, we use the data structure by Gawrychowski et al. [23] (for LCP and LCP^R operations). Furthermore, we reuse some tools from the fully compressed setting, because the data structure of [23] actually works with (a form of) straight-line programs.

As the primitive operations of the PILLAR model are rather simple, we believe that they can be efficiently implemented in further settings not considered here.

Contents

1	Introduction	1
2	Preliminaries	10
2.1	The PILLAR Model	11
3	Improved Structural Insights into Pattern Matching with Mismatches	13
3.1	Characterization of the Periodic Case	13
3.2	The Non-Periodic Case	16
4	Algorithm: Pattern Matching with Mismatches in the PILLAR Model	20
4.1	Auxiliary PILLAR Model Operations for Pattern Matching with Mismatches	20
4.2	Computing Structure in the Pattern	21
4.3	Computing Occurrences in the Periodic Case	23
4.4	Computing Occurrences in the Non-Periodic Case	26
4.5	A PILLAR Model Algorithm for Pattern Matches with Mismatches	28
5	Structural Insights into Pattern Matching with Edits	29
5.1	Characterization of the Periodic Case	29
5.2	Bounding the Number of Occurrences in the Non-Periodic Case	41
6	Algorithm: Pattern Matching with Edits in the PILLAR Model	45
6.1	Auxiliary PILLAR Model Operations for Pattern Matching with Edits	45
6.2	Computing Structure in the Pattern	48
6.3	Computing Occurrences in the Periodic Case	51
6.4	Computing Occurrences in the Non-Periodic Case	60
6.5	A PILLAR Model Algorithm for Pattern Matching with Edits	62
7	Implementing the PILLAR Model: Faster Approximate Pattern Matching	63
7.1	Implementing the PILLAR Model in the Standard Setting	63
7.2	Implementing the PILLAR Model in the Fully Compressed Setting	64
7.3	Implementing the PILLAR Model in the Dynamic Setting	69

2 Preliminaries

Sets and Arithmetic Progressions

For $n \in \mathbb{Z}_{\geq 0}$, we write $[n]$ to denote the set $\{0, \dots, n-1\}$. Further, for $i, j \in \mathbb{Z}$, we write $[i..j]$ to denote $\{i, \dots, j\}$ and $[i..j)$ to denote $\{i, \dots, j-1\}$; the sets $(i..j]$ and $(i..j)$ are defined similarly.

For $a, d, \ell \in \mathbb{Z}$ with $\ell > 0$, the set $\{a + j \cdot d \mid j \in [0.. \ell)\}$ is an *arithmetic progression* with starting value a , difference d , and length ℓ . Whenever we use arithmetic progressions in an algorithm, we store them as a triple (a, d, ℓ) consisting of the first value, the difference, and the length.

For a set $X \subseteq \mathbb{Z}$, we write kX to denote the set containing all elements of X multiplied by k , that is, $kX := \{k \cdot x \mid x \in X\}$. Similarly, we define $\lfloor X/k \rfloor := \{\lfloor x/k \rfloor \mid x \in X\}$ and $k\lfloor X/k \rfloor := \{k \cdot \lfloor x/k \rfloor \mid x \in X\}$.

Strings

We write $T = T[0]T[1]\cdots T[n-1]$ to denote a *string* of length $|T| = n$ over an alphabet Σ . The elements of Σ are called *characters*. We write ε to denote the *empty string*.

For a string T , we denote the *reverse string* of T by T^R , that is, $T^R := T[n-1]T[n-2]\cdots T[0]$. For two positions $i \leq j$ in T , we write $T[i..j+1) := T[i..j) := T[i]\cdots T[j]$ for the *fragment* of T that starts at position i and ends at position j . We set $T[i..j) := \varepsilon$ whenever $j < i$.

A *prefix* of a string T is a fragment that starts at position 0 (that is, a prefix is a fragment of the form $T[0..j)$ for some $j \geq 0$). A *suffix* of a string T is a fragment that ends at position $|T| - 1$ (that is, a suffix is a fragment of the form $T[i..|T|)$ for some $i \leq |T|$). We denote the length of the *longest common prefix* (*longest common suffix*) of two strings U and V , defined as the longest string that occurs as a prefix (suffix) of both U and V , by $\text{lcp}(U, V)$ (respectively, $\text{lcp}^R(U, V)$).

A string P of length $m \in [0..|T|)$ is a *substring* of a string T (denoted $P \preceq T$) if there is a fragment $T[i..i+m)$ matching P . In this case, we say that there is an *exact occurrence* of P at position i in T , or, more simply, that P *exactly occurs* in T .

For two strings U and V , we write UV or $U \cdot V$ to denote their concatenation. We also write $U^k := U \cdots U$ to denote the concatenation of k copies of the string U . Furthermore, U^∞ denotes an infinite string obtained by concatenating infinitely many copies of U . A string T is called *primitive* if it cannot be expressed as $T = U^k$ for a string U and an integer $k > 1$.

A positive integer p is called a *period* of a string T if $T[i] = T[i+p]$ for all $i \in [0..|T| - p)$. We refer to the smallest period as *the period* $\text{per}(T)$ of the string. The string $T[0.. \text{per}(T))$ is called the *string period* of T . We call a string *periodic* if its period is at most half of its length.

For a string T , we define the following *rotation* operations: The operation $\text{rot}(\cdot)$ takes as input a string, and moves its last character to the front; that is, $\text{rot}(T) := T[n-1]T[0..n-2]$. The inverse operation $\text{rot}^{-1}(\cdot)$ takes as input a string and moves its initial character to the end; that is, $\text{rot}^{-1}(T) := T[1..n-1]T[0]$. Note that a primitive string T does not match any of its non-trivial rotations, that is, we have $T = \text{rot}^j(T)$ if and only if $j \equiv 0 \pmod{|T|}$.

Finally, the *run-length encoding* (RLE) of a string T is a decomposition of T into maximal blocks such that each block is a power of a single character. (For instance, the RLE of the string `aaabbabbbb` is `a3b2ab4`.) Note that each block of the RLE can be represented in $\mathcal{O}(1)$ space.

Hamming Distance and Pattern Matching with Mismatches

For two strings S and T of the same length n , we define the set of *mismatches* between S and T as $\text{Mis}(S, T) := \{i \in [n] \mid S[i] \neq T[i]\}$. Now, the *Hamming distance* of S and T is defined as the number of mismatches between S and T , that is, $\delta_H(S, T) := |\text{Mis}(S, T)|$.

11 Faster Approximate Pattern Matching: A Unified Approach

It is easy to verify that the Hamming distance satisfies the triangle inequality:

▣ **Fact 2.1** (Triangle inequality for Hamming distance). *Any strings A , B , and C of the same length satisfy $\delta_H(A, C) + \delta_H(C, B) \geq \delta_H(A, B) \geq |\delta_H(A, C) - \delta_H(C, B)|$.* ▣

As we are often concerned with the Hamming distance of a string S and a prefix of T^∞ for a string T , we write $\text{Mis}(S, T^*) := \text{Mis}(S, T^\infty[0..|S|])$ and $\delta_H(S, T^*) := |\text{Mis}(S, T^*)|$.

Now, for a string P (also called a *pattern*), a string T (also called a *text*), and an integer $k \geq 0$ (also called a *threshold*), we say that there is a *k-mismatch occurrence* of P in T at position i if $\delta_H(P, T[i..i+|P|]) \leq k$. We write $\text{Occ}_k^H(P, T)$ to denote the set of all positions of *k-mismatch occurrences* of P in T , that is, $\text{Occ}_k^H(P, T) := \{i \mid \delta_H(P, T[i..i+|P|]) \leq k\}$. Lastly, we define the *pattern matching with mismatches* problem.

▣ **Problem 2.2** (Pattern matching with mismatches). *Given a pattern P , a text T , and a threshold k , compute the set $\text{Occ}_k^H(P, T)$.* ▣

Note that, depending on the use case (especially if the set $\text{Occ}_k^H(P, T)$ is relatively large), we may only want to compute the size $|\text{Occ}_k^H(P, T)|$ or a space-efficient representation of the set $\text{Occ}_k^H(P, T)$, e.g., as the union of disjoint arithmetic progressions.

Edit Distance and Pattern Matching with Edits

The *edit distance* (also known as *Levenshtein distance*) between two strings S and T , denoted $\delta_E(S, T)$, is the minimum number of character insertions, deletions, and substitutions required to transform S into T .

Again, it is easy to verify that the edit distance satisfies the triangle inequality:

▣ **Fact 2.3** (Triangle inequality for edit distance). *Any strings A , B , and C of the same length satisfy $\delta_E(A, C) + \delta_E(C, B) \geq \delta_E(A, B) \geq |\delta_E(A, C) - \delta_E(C, B)|$.* ▣

Further, similarly to the Hamming distance, we write $\delta_E(S, T^*) := \min\{\delta_E(S, T^\infty[0..j]) \mid j \in \mathbb{Z}_{\geq 0}\}$ to denote the minimum edit distance between a string S and any prefix of a string T^∞ . Further, we write $\delta_E(S, T^*) := \min\{\delta_E(S, T^\infty[i..j]) \mid i, j \in \mathbb{Z}_{\geq 0}, i \leq j\}$ to denote the minimum edit distance between S and any substring of T^∞ , and we set $\delta_E(S, T) := \min\{\delta_E(S, T^\infty[i..j|T|]) \mid i, j \in \mathbb{Z}_{\geq 0}, i \leq j|T|\}$.

Now, for a string P (also called a *pattern*), a string T (also called a *text*), and an integer $k \geq 0$ (also called a *threshold*), we say that there is a *k-error occurrence* of P in T at position i if $\delta_E(P, T[i..j]) \leq k$ for some $j \geq i$. We write $\text{Occ}_k^E(P, T)$ to denote the set of all positions of *k-error occurrences* of P in T , that is, $\text{Occ}_k^E(P, T) := \{i \mid \exists j \geq i, \delta_E(P, T[i..j]) \leq k\}$. Lastly, we define the *pattern matching with edits* problem.

▣ **Problem 2.4** (Pattern matching with edits). *Given a pattern P , a text T , and a threshold k , compute the set $\text{Occ}_k^E(P, T)$.* ▣

Again, we may only want to compute the size $|\text{Occ}_k^E(P, T)|$ or a space-efficient representation of $\text{Occ}_k^E(P, T)$.

2.1 The PILLAR Model

In order to unify the implementations of our approach for approximate pattern matching in different settings, we introduce the PILLAR model. The PILLAR model captures certain primitive operations which can be implemented efficiently in all considered settings (see Section 7 for the actual implementations).

Thus, in the algorithmic sections of this work, Sections 4 and 6, we bound the running times in terms of the number of PILLAR operations; if this value is asymptotically smaller than the time complexity of the remaining computations, we also specify the extra running time.

In the PILLAR model, we are given a family of strings \mathcal{X} for preprocessing. The elementary objects are fragments $X[\ell..r]$ of strings $X \in \mathcal{X}$. Each such fragment S is represented via a *handle*, which is how S can be passed as input to PILLAR operations. Initially, the model provides a handle to each $X \in \mathcal{X}$, interpreted as $X[0..|X|]$. Handles to other fragments can be obtained through an **Extract** operation:

- **Extract**(S, ℓ, r): Given a fragment S and positions $0 \leq \ell \leq r \leq |S|$, extract the (sub)fragment $S[\ell..r]$. If $S = X[\ell'..r']$ for $X \in \mathcal{X}$, then $S[\ell..r]$ is defined as $X[\ell' + \ell.. \ell' + r]$.

Furthermore, the following primitive operations are supported in the PILLAR model:

- **LCP**(S, T): Compute the length of the longest common prefix of S and T .
- **LCP^R**(S, T): Compute the length of the longest common suffix of S and T .
- **IPM**(P, T): Assuming that $|T| \leq 2|P|$, compute $\text{Occ}(P, T)$ (represented as an arithmetic progression with difference $\text{per}(P)$).
- **Access**(S, i): Assuming $i \in [|S|]$, retrieve the character $S[i]$.
- **Length**(S): Retrieve the length $|S|$ of the string S .

We now provide a small toolbox that is to be used in Sections 4 and 6. First, we note that **LCP**(S, T), **LCP^R**(S, T), or **IPM**(S, T) can be used to check whether two strings S and T are equal.

▣ **Lemma 2.5** (Equality, [31, Fact 2.5.2]). *Given strings S and T , we can check whether S and T are equal in $\mathcal{O}(1)$ time in the PILLAR model.* ▣

A more involved operation allows checking if a given string is periodic and, if so, computing the period.

▣ **Lemma 2.6** (Period(S), [32, 31]). *Given a string S , we can compute $\text{per}(S)$ or declare that $\text{per}(S) > |S|/2$ in $\mathcal{O}(1)$ time in the PILLAR model.* ▣

Next, we introduce an operation that checks cyclic equivalence and retrieves the witness shifts.

▣ **Lemma 2.7** (Rotations(S, T), [32, 31]). *Given strings S and T , we can find all integers j such that $T = \text{rot}^j(S)$ in $\mathcal{O}(1)$ time in the PILLAR model. The output is represented as an arithmetic progression.* ▣

Our subsequent goal is to generalize the primitive LCP operation to also support fragments of infinite powers of strings. We start with a special case, and then we cover the general case.

▣ **Lemma 2.8** (LCP(S, Q^∞), [31, Fact 2.5.2]; see also [5]). *Given strings S and Q , we can compute LCP(S, Q^∞) in $\mathcal{O}(1)$ time in the PILLAR model.* ▣

▣ **Corollary 2.9** (LCP($S, Q^\infty[\ell..r]$)). *Given strings S and Q and integers $0 \leq \ell \leq r$, we can compute LCP($S, Q^\infty[\ell..r]$) in $\mathcal{O}(1)$ time in the PILLAR model.*

Proof. We first compute **LCP**($S, Q[\ell \bmod |Q|..|Q|]$) using a primitive operation. If we reach the end of the string Q , we continue with an **LCP**($S[|Q| - \ell \bmod |Q|..|S|], Q^\infty$) query, implemented using Lemma 2.8. This yields **LCP**($S, Q^\infty[\ell..)$), so we cap the obtained value with $r - \ell$ to retrieve **LCP**($S, Q^\infty[\ell..r]$). ▣

A similar procedure yields an analogous generalization of the primitive **LCP^R** operation.

▣ **Corollary 2.10** (LCP^R($S, Q^\infty[\ell..r]$)). *Given strings S and Q and integers $0 \leq \ell \leq r$, we can compute LCP^R($S, Q^\infty[\ell..r]$) in $\mathcal{O}(1)$ time in the PILLAR model.* ▣

Lastly, we discuss an operation finding all exact occurrences of a given string P in a given string T .

▮ **Lemma 2.11** ($\text{ExactMatches}(P, T)$). *Let T denote a string of length n and let P denote a string of length m . We can compute the set $\text{Occ}(P, T)$ using $\mathcal{O}(n/\text{per}(P))$ time and $\mathcal{O}(n/m)$ PILLAR operations.*

Proof. We perform an $\text{IPM}(P, T_i)$ query with $T_i := T[i m .. \min\{n, (i + 2)m - 1\}]$ for each $i \in [\lfloor n/m \rfloor]$; that is a total of $\mathcal{O}(n/m)$ PILLAR operations. Each occurrence $j \in \text{Occ}(P, T)$ corresponds to a single occurrence of P in a single T_i , namely, $j \bmod m \in \text{Occ}(P, T_{\lfloor j/m \rfloor})$, and vice versa. Furthermore, each $\text{IPM}(P, T_i)$ query returns an arithmetic progression with difference $\text{per}(P)$, which thus consists of $\mathcal{O}(m/\text{per}(P))$ elements. Hence, the total number of elements of all arithmetic progressions is $\mathcal{O}(n/\text{per}(P))$. ▮

We conclude this section with introducing the concept of a generator.

▮ **Definition 2.12** (Generator of a set). *For an (ordered) set S , an $(\mathcal{O}(P), \mathcal{O}(Q))$ -time generator of S is a data structure that after $\mathcal{O}(P)$ -time initialization in the PILLAR model, supports the following operation:*

- ▬ **Next:** *In the i th call of Next, return the i th smallest element of the set S (or \perp if $i > |S|$), using $\mathcal{O}(Q(i))$ time in the PILLAR model.* ▮

Note that a generator is not specific to the PILLAR model. We use generators to obtain positions where two strings differ—either by a mismatch or by an edit. Consult Sections 4.1 and 6.1 for the details, as well as for other PILLAR operations specific to pattern matching with mismatches or edits, respectively.

3 Improved Structural Insights into Pattern Matching with Mismatches

In this section, we provide insight into the structure of k -mismatch occurrences of a pattern P in a text T . In particular, we improve the result of [11] and show the following asymptotically tight characterization (which is Main Theorem 4 with explicit constants and without the restriction on the length of T).

▮ **Theorem 3.1.** *Given a pattern P of length m , a text T of length n , and a threshold $k \in [1..m]$, at least one of the following holds:*

- ▬ *The number of k -mismatch occurrences of P in T is bounded by $|\text{Occ}_k^H(P, T)| \leq 576 \cdot n/m \cdot k$.*
- ▬ *There is a primitive string Q of length $|Q| \leq m/128k$ that satisfies $\delta_H(P, Q^*) < 2k$.* ▮


3.1 Characterization of the Periodic Case

In order to prove Theorem 3.1, we first discuss the (approximately) periodic case, that is, the case when we have $\delta_H(P, Q^*) < 2k$. In particular, we prove the following theorem, which strengthens [11, Claim 3.1].

▮ **Theorem 3.2** (Compare Main Theorem 5 and [11, Claim 3.1]). *Let P denote a pattern of length m , let T denote a text of length $n \leq \frac{3}{2}m$, and let $k \in [0..m]$ denote a threshold. Suppose that both $T[0..m)$ and $T[n - m..n)$ are k -mismatch occurrences of P (that is, $\{0, n - m\} \subseteq \text{Occ}_k^H(P, T)$). If there is a positive integer $d \geq 2k$ and a primitive string Q with $|Q| \leq m/8d$ and $\delta_H(P, Q^*) \leq d$, then each of following holds:*

- (a) *Every position in $\text{Occ}_k^H(P, T)$ is a multiple of $|Q|$.*
- (b) *The string T satisfies $\delta_H(T, Q^*) \leq 3d$.*
- (c) *The set $\text{Occ}_k^H(P, T)$ can be decomposed into $3d(d + 1)$ arithmetic progressions with difference $|Q|$.*
- (d) *If $\delta_H(P, Q^*) = d$, then $|\text{Occ}_k^H(P, T)| \leq 6d$.* ▮

Before proving Theorem 3.2, we characterize the values $\delta_H(T[j|Q|..j|Q| + m), P)$ under an extra assumption that $\delta_H(T, Q^*)$ is small as well; this assumption is dropped in Theorem 3.2.



■ **Figure 3** In both strings, all blocks apart from the last ones are of length q . For each string $X \in \{P, T\}$, we show the characters at positions in $\text{Mis}(X, Q^*)$ only. At position jq in T , we place $\mu_j(\tau_2, \pi_1) + \mu_j(\tau_4, \pi_2) + \mu_j(\tau_5, \pi_3) = 1 + 2 + 1 = 4$ marks. We have $\delta_H(P, Q^*) = |\{\pi_1, \pi_2, \pi_3\}| = 3$ and $\delta_H(T[jq..jq+m], Q^*) = |\{\tau_2, \tau_3, \tau_4, \tau_5\}| = 4$. Using Claim 3.4, we obtain that $h_j = 3$; the three corresponding mismatches are indicated by asterisks.

□ **Lemma 3.3.** *Let P denote a pattern of length m and let T denote a text of length $n \leq 3/2 m$. Further, let Q denote a string of length q and set $d := \delta_H(P, Q^*)$ and $d' := \delta_H(T, Q^*)$. Then, the sequence of values $h_j := \delta_H(T[jq..jq+m], P)$ for $0 \leq j \leq (n-m)/q$ contains at most $d'(2d+1)$ entries h_j with $h_j \neq h_{j+1}$ and, unless $d = 0$, at most $2d'$ entries h_j with $h_j \leq d/2$.*

Proof. For every $\tau \in \text{Mis}(T, Q^*)$ and $\pi \in \text{Mis}(P, Q^*)$, let us put $(2 - \delta_H(P[\pi], T[\tau]))$ marks at position $\tau - \pi$ in T , if it exists. For each $0 \leq j \leq (n-m)/q$, let $\mu_j(\tau, \pi)$ denote the number of marks placed at position jq due to the mismatches τ in T and π in P , that is,

$$\mu_j(\tau, \pi) := \begin{cases} 2 - \delta_H(P[\pi], T[\tau]) & \text{if } \pi \in \text{Mis}(P, Q^*) \text{ and } \tau = jq + \pi \in \text{Mis}(T, Q^*), \\ 0 & \text{otherwise.} \end{cases}$$

Further, define $\mu_j := \sum_{\tau, \pi} \mu_j(\tau, \pi)$ as the total number of marks at position jq .

Next, for every $0 \leq j \leq (n-m)/q$, we relate the Hamming distance $h_j := \delta_H(T[jq..jq+m], P)$ to the number of marks μ_j at position jq and the Hamming distances $\delta_H(T[jq..jq+m], Q^*)$ and $\delta_H(P, Q^*)$; consult Figure 3 for an illustration.

□ **Claim 3.4.** For each $0 \leq j \leq (n-m)/q$, we have $h_j = \delta_H(P, Q^*) + \delta_H(T[jq..jq+m], Q^*) - \mu_j$.

Proof. We show the following equivalent statement:

$$|\text{Mis}(T[jq..jq+m], P)| = |\text{Mis}(P, Q^*)| + |\text{Mis}(T[jq..jq+m], Q^*)| - \sum_{\tau, \pi} \mu_j(\tau, \pi). \quad (1)$$

By construction, $\mu_j(\tau, \pi) = 0$ whenever $\tau \neq \pi + jq$. Hence, we can prove (1) by showing that for every position $\pi \in [0..m)$ in P and every position $\tau := jq + \pi$ in T , the following equation holds:

$$\delta_H(T[\tau], P[\pi]) = \delta_H(P[\pi], Q^\infty[\pi]) + \delta_H(T[\tau], Q^\infty[\tau]) - \mu_j(\tau, \pi).$$

We proceed by case distinction on whether $\pi \in \text{Mis}(P, Q^*)$ and whether $\tau \in \text{Mis}(T, Q^*)$.

■ If $\pi \notin \text{Mis}(P, Q^*)$ and $\tau \notin \text{Mis}(T, Q^*)$, then we have $P[\pi] = Q^\infty[\pi] = Q^\infty[\tau] = T[\tau]$ and thus

$$\delta_H(T[\tau], P[\pi]) = 0 = 0 + 0 - 0 = \delta_H(P[\pi], Q^\infty[\pi]) + \delta_H(T[\tau], Q^\infty[\tau]) - \mu_j(\tau, \pi).$$

■ If $\pi \in \text{Mis}(P, Q^*)$ and $\tau \notin \text{Mis}(T, Q^*)$, then we have $P[\pi] \neq Q^\infty[\pi] = Q^\infty[\tau] = T[\tau]$ and thus

$$\delta_H(T[\tau], P[\pi]) = 1 = 1 + 0 - 0 = \delta_H(P[\pi], Q^\infty[\pi]) + \delta_H(T[\tau], Q^\infty[\tau]) - \mu_j(\tau, \pi).$$

■ If $\pi \notin \text{Mis}(P, Q^*)$ and $\tau \in \text{Mis}(T, Q^*)$, then we have $P[\pi] = Q^\infty[\pi] = Q^\infty[\tau] \neq T[\tau]$ and thus

$$\delta_H(T[\tau], P[\pi]) = 1 = 0 + 1 - 0 = \delta_H(P[\pi], Q^\infty[\pi]) + \delta_H(T[\tau], Q^\infty[\tau]) - \mu_j(\tau, \pi).$$

15 Faster Approximate Pattern Matching: A Unified Approach

- If $\pi \in \text{Mis}(P, Q^*)$ and $\tau \in \text{Mis}(T, Q^*)$, then we have $P[\pi] \neq Q^\infty[\pi] = Q^\infty[\tau] \neq T[\tau]$ and thus

$$\delta_H(T[\tau], P[\pi]) = 1 + 1 - (2 - \delta_H(T[\tau], P[\pi])) = \delta_H(P[\pi], Q^\infty[\pi]) + \delta_H(T[\tau], Q^\infty[\tau]) - \mu_j(\tau, \pi).$$

Combining the equations obtained for every pair of positions π and τ , we derive (1). \square

In particular, Claim 3.4 yields

$$h_{j+1} - h_j = |\text{Mis}(T, Q^*) \cap [jq + m \dots (j+1)q + m]| - |\text{Mis}(T, Q^*) \cap [jq \dots (j+1)q]| - \mu_{j+1} + \mu_j.$$

Hence, in order for h_{j+1} not to equal h_j , at least one of the four terms on the right-hand side of the equation above must be non-zero. Let us analyze when this is possible. To that end, we first observe that the set $\text{Mis}(T, Q^*) \cap [jq + m \dots (j+1)q + m]$ contains only elements $\tau \in \text{Mis}(T, Q^*)$ with $\tau \geq m$, and that the set $\text{Mis}(T, Q^*) \cap [jq \dots (j+1)q]$ only contains elements $\tau \in \text{Mis}(T, Q^*)$ with $\tau < n - m$. Using $n \leq \frac{3}{2}m$, we observe that h_{j+1} can be different from h_j due to the first or second term at most d' times. Further, each non-zero value in one of the terms μ_{j+1} and μ_j can be attributed to a marked position (jq or $(j+1)q$, respectively). The total number of marked positions is at most dd' , so h_{j+1} can be different from h_j due one of the terms μ_{j+1} or μ_j at most $2dd'$ times. In total, we conclude that the number of entries h_j with $h_j \neq h_{j+1}$ is at most $d'(2d+1)$.

Next, observe that $\mu_j \leq 2|\text{Mis}(T, Q^*) \cap [jq \dots jq + m]| = 2\delta_H(T[jq \dots jq + m], Q^*)$, and therefore

$$h_j = \delta_H(P, Q^*) + \delta_H(T[jq \dots jq + m], Q^*) - \mu_j \geq d - \mu_j/2.$$

Consequently, $h_j \leq d/2$ yields $\mu_j \geq d$, that is, that there are at least d marks at position jq . Given that the total number of marks is at most $2dd'$, the number of entries h_j with $h_j \leq d/2$ is at most $2d'$, assuming that $d > 0$. \blacksquare

Now, we drop the assumption that $\delta_H(T, Q^*)$ is small and prove Theorem 3.2.

Theorem 3.2 (Compare Main Theorem 5 and [11, Claim 3.1]). *Let P denote a pattern of length m , let T denote a text of length $n \leq \frac{3}{2}m$, and let $k \in [0 \dots m]$ denote a threshold. Suppose that both $T[0 \dots m]$ and $T[n - m \dots n]$ are k -mismatch occurrences of P (that is, $\{0, n - m\} \subseteq \text{Occ}_k^H(P, T)$). If there is a positive integer $d \geq 2k$ and a primitive string Q with $|Q| \leq m/8d$ and $\delta_H(P, Q^*) \leq d$, then each of following holds:*

- (a) *Every position in $\text{Occ}_k^H(P, T)$ is a multiple of $|Q|$.*
- (b) *The string T satisfies $\delta_H(T, Q^*) \leq 3d$.*
- (c) *The set $\text{Occ}_k^H(P, T)$ can be decomposed into $3d(d+1)$ arithmetic progressions with difference $|Q|$.*
- (d) *If $\delta_H(P, Q^*) = d$, then $|\text{Occ}_k^H(P, T)| \leq 6d$.*

Proof. Consider any position $\ell \in \text{Occ}_k^H(P, T)$. By the definition of a k -mismatch occurrence, we have $\delta_H(T[\ell \dots \ell + m], P) \leq k \leq d/2$. Combining this inequality with $\delta_H(P, Q^*) \leq d$ via the triangle inequality yields $\delta_H(T[\ell \dots \ell + m], Q^*) \leq \frac{3}{2}d$. Similarly, for the position $0 \in \text{Occ}_k^H(P, T)$, we obtain $\delta_H(T[0 \dots m], Q^*) \leq \frac{3}{2}d$, which lets us compare the overlapping parts of Q^∞ . Replacing strings by superstrings and applying the triangle inequality yields

$$\begin{aligned} \delta_H(Q^\infty[\ell \dots m], Q^\infty[0 \dots m - \ell]) &\leq \delta_H(T[\ell \dots m], Q^\infty[\ell \dots m]) + \delta_H(T[\ell \dots m], Q^\infty[0 \dots m - \ell]) \\ &\leq \delta_H(T[0 \dots m], Q^\infty[0 \dots m]) + \delta_H(T[\ell \dots \ell + m], Q^\infty[0 \dots m]) \\ &= \delta_H(T[0 \dots m], Q^*) + \delta_H(T[\ell \dots \ell + m], Q^*) \\ &\leq 3d. \end{aligned}$$

Towards a proof by contradiction, suppose that ℓ is not a multiple of $|Q|$. As Q is primitive, we have

$$3d \geq \delta_H(Q^\infty[\ell..m], Q^\infty[0..m-\ell]) \geq \left\lfloor \frac{m-\ell}{|Q|} \right\rfloor \geq \left\lfloor \frac{m/2}{m/8d} \right\rfloor = 4d,$$

where the second bound follows from $\ell \leq m/2$ and $|Q| \leq m/8d$. This contradiction yields Claim (a).

In order to prove Claim (b), we observe that $n-m \in \text{Occ}_k^H(P, T)$ is a multiple of $|Q|$. Consequently,

$$\delta_H(T, Q^*) = \delta_H(T[0..n-m], Q^*) + \delta_H(T[n-m..n], Q^*) \leq \delta_H(T[0..m], Q^*) + \frac{3}{2}d \leq 3d,$$

which concludes the proof of Claim (b).

For a proof of Claims (c) and (d), we apply Lemma 3.3. Due to Claim (a), each position in $\text{Occ}_k^H(P, T)$ corresponds to an entry h_j with $h_j \leq k$. In particular, each block of consecutive entries $h_j, \dots, h_{j'}$ not exceeding k yields an arithmetic progression (with difference $|Q|$) in $\text{Occ}_k^H(P, T)$. The number of entries h_j with $h_j \leq k < h_{j+1}$ or $h_j > k \geq h_{j+1}$ is in total at most $3d(2d+1)$, so the number of arithmetic progressions is at most $1 + 1/2 \cdot 3d(2d+1) \leq 3d(d+1)$, which proves Claim (c).

For Claim (d), we observe that if $d = \delta_H(P, Q^*)$, then each position in $\text{Occ}_k^H(P, T)$ corresponds to an entry h_j with $h_j \leq k \leq d/2$; thus $|\text{Occ}_k^H(P, T)| \leq 2 \cdot 3d = 6d$. \blacksquare

Corollary 3.5. *Let P denote a pattern of length m , let T denote a text of length n , and let $k \in [0..m]$ denote a threshold. If there is a positive integer $d \geq 2k$ and a primitive string Q with $|Q| \leq m/8d$ and $\delta_H(P, Q^*) \leq d$, then the set $\text{Occ}_k^H(P, T)$ can be decomposed into $6 \cdot n/m \cdot d(d+1)$ arithmetic progressions with difference $|Q|$. Moreover, if $\delta_H(P, Q^*) = d$, then $|\text{Occ}_k^H(P, T)| \leq 12 \cdot n/m \cdot d$.*

Proof. Partition the string T into $\lfloor 2n/m \rfloor$ blocks $T_0, \dots, T_{\lfloor 2n/m \rfloor - 1}$ of length less than $\frac{3}{2}m$ each, where the i th block starts at position $\lfloor i \cdot m/2 \rfloor$; formally, we set $T_i := T[\lfloor i \cdot m/2 \rfloor .. \min\{n, \lfloor (i+3) \cdot m/2 \rfloor - 1\}]$. If $\text{Occ}_k^H(P, T_i) \neq \emptyset$, we define T'_i to be the shortest fragment of T_i containing all k -mismatch occurrences of P in T_i . As a result, T'_i satisfies the assumptions of Theorem 3.2. Hence, $\text{Occ}_k^H(P, T'_i)$ can be decomposed into $3d(d+1)$ arithmetic progressions with difference $|Q|$, and $|\text{Occ}(P, T'_i)| \leq 6d$ if $\delta_H(P, Q^*) = d$.

We conclude that $\text{Occ}_k^H(P, T'_i)$ decomposes into $6 \cdot n/m \cdot d(d+1)$ arithmetic progressions with difference $|Q|$; further, $|\text{Occ}(P, T_i)| \leq 12 \cdot n/m \cdot d$ if $\delta_H(P, Q^*) = d$. \blacksquare

3.2 The Non-Periodic Case

Having dealt with the (approximately) periodic case, we now turn to the general case. In particular, we show that whenever the string P is sufficiently far from being periodic, the number of k -mismatch occurrences of P in any string T of length $n \leq \frac{3}{2}m$ is $\mathcal{O}(k)$.

Intuitively, we proceed (and thereby prove Theorem 3.1) as follows: We first analyze the string P for useful structure that can help in bounding the number of occurrences of P in any string T . If we fail to find any special structure in P , then we conclude that the string P is close to a periodic string with a small period (compared to $|P|$)—a case that we already understand thanks to the previous subsection.

Lemma 3.6. *Given a string P of length m and a threshold $k \in [1..m]$, at least one of the following holds:*

- (a) *The string P contains $2k$ disjoint breaks B_1, \dots, B_{2k} each having period $\text{per}(B_i) > m/128k$ and length $|B_i| = \lfloor m/8k \rfloor$.*
- (b) *The string P contains r disjoint repetitive regions R_1, \dots, R_r of total length $\sum_{i=1}^r |R_i| \geq 3/8 \cdot m$ such that each region R_i satisfies $|R_i| \geq m/8k$ and has a primitive approximate period Q_i with $|Q_i| \leq m/128k$ and $\delta_H(R_i, Q_i^*) = \lceil 8k/m \cdot |R_i| \rceil$.*
- (c) *The string P has a primitive approximate period Q with $|Q| \leq m/128k$ and $\delta_H(P, Q^*) < 8k$.*

■ **Algorithm 1** A constructive proof of Lemma 3.6.

```

1  $\mathcal{B} \leftarrow \{\}; \mathcal{R} \leftarrow \{\};$ 
2 while true do
3   Consider the fragment  $P' = P[j..j + \lfloor m/8k \rfloor]$  of the next  $\lfloor m/8k \rfloor$  unprocessed characters
   of  $P$ ;
4   if  $\text{per}(P') > m/128k$  then
5      $\mathcal{B} \leftarrow \mathcal{B} \cup \{P'\};$ 
6     if  $|\mathcal{B}| = 2k$  then return breaks  $\mathcal{B}$ ;
7   else
8      $Q \leftarrow P[j..j + \text{per}(P')];$ 
9     Search for a prefix  $R$  of  $P[j..m]$  with  $|R| > |P'|$  and  $\delta_H(R, Q^*) = \lceil 8k/m \cdot |R| \rceil$ ;
10    if such  $R$  exists then
11       $\mathcal{R} \leftarrow \mathcal{R} \cup \{(R, Q)\};$ 
12      if  $\sum_{(R,Q) \in \mathcal{R}} |R| \geq 3/8 \cdot m$  then
13        return repetitive regions (and their corresponding periods)  $\mathcal{R}$ ;
14    else
15      Search for a suffix  $R'$  of  $P$  with  $|R'| \geq m - j$  and  $\delta_H(R', \text{rot}^{|R'|-m+j}(Q)^*) = \lceil 8k/m \cdot |R'| \rceil$ ;
16      if such  $R'$  exists then return repetitive region  $(R', \text{rot}^{|R'|-m+j}(Q))$ ;
17    else return approximate period  $\text{rot}^j(Q)$ ;
```

Proof. We prove the claim constructively, that is, we construct either a set \mathcal{B} of $2k$ breaks, or a set \mathcal{R} of repetitive regions, or, if we fail to construct either, we derive an approximate string period Q of the string P with the desired properties.

We process the string P from left to right as follows: If the fragment P' of the next $\lfloor m/8k \rfloor$ (unprocessed) characters of P has a long period, we have found a new break and continue (or return the found set of $2k$ breaks). Otherwise, if P' has a short string period Q , we try to extend the fragment P' (to the right) into a repetitive region. If we succeed, we have found a new repetitive region and continue (or return the found set of repetitive regions if the total length of all repetitive regions found so far is at least $3/8 \cdot m$). If we fail to construct a new repetitive region, then we conclude that the suffix of P starting with P' has an approximate period Q . We try to construct a repetitive region by extending this suffix to the left, dropping all the repetitive regions computed beforehand. If we fail again, we declare that an appropriate rotation of Q is an approximate period of the string P . Consider Algorithm 1 for a detailed description.

By construction, all breaks in the set \mathcal{B} and repetitive regions in the set \mathcal{R} returned by the algorithm are disjoint and satisfy the claimed properties. To prove that the algorithm is also correct when it fails to find a new repetitive region, we start by bounding from above the length of the processed prefix of P .

□ **Claim 3.7.** Whenever we consider a new fragment $P' = P[j..j + \lfloor m/8k \rfloor]$ of the next $\lfloor m/8k \rfloor$ unprocessed characters of P , such a fragment starts at a position $j < 5/8 \cdot m$.

Proof. Observe that whenever we consider a new fragment $P[j..j + \lfloor m/8k \rfloor]$, the string $P[0..j]$ has been partitioned into breaks and repetitive regions. The total length of breaks is less than $2k \lfloor m/8k \rfloor \leq 2/8 \cdot m$, and the total length of repetitive regions is less than $3/8 \cdot m$. Hence, $j < 5/8 \cdot m$, yielding the claim. □

Note that Claim 3.7 also shows that whenever we consider a new fragment P' of $\lfloor m/8k \rfloor$ characters, there is indeed such a fragment, that is, P' is well-defined.

Now, consider the case when, for a fragment $P' = P[j..j + \lfloor m/8k \rfloor]$ (that is not a break) and its string period $Q = P[j..j + \text{per}(P')]$, we fail to obtain a new repetitive region R . In this case, we search for a repetitive region R' of length $|R'| \geq m - j$ that is a suffix of P and has an approximate period $Q' := \text{rot}^{|R'|-m+j}(Q)$. If we indeed find such a region R' , then $|R'| \geq m - j \geq m - 5/8 \cdot m = 3/8 \cdot m$ by Claim 3.7, so R' is long enough to be reported on its own. However, if we fail to find such R' , we need to show that $\text{rot}^j(Q)$ can be reported as an approximate period of P , that is, $\delta_H(P, \text{rot}^j(Q)^*) < 8k$.

We first derive $\delta_H(P[j..m], Q^*) < \lceil 8k/m \cdot (m - j) \rceil$. For this, we inductively prove that the values $\Delta_\rho := \lceil 8k/m \cdot \rho \rceil - \delta_H(P[j..j + \rho], Q^*)$ for $\rho \in [|P'|..m - j]$ are all at least 1. In the base case of $\rho = |P'|$, we have $\Delta_\rho = 1 - 0$ because Q is the string period of P' . To carry out an inductive step, suppose that $\Delta_{\rho-1} \geq 1$ for some $\rho \in [|P'|..m - j]$. Notice that $\Delta_\rho \geq \Delta_{\rho-1} - 1 \geq 0$: The first term in the definition of Δ_ρ has not decreased compared to $\Delta_{\rho-1}$, and the second term $\delta_H(P[j..j + \rho], Q^*)$ may have increased by at most one. Moreover, $\Delta_\rho \neq 0$ because $R = P[j..j + \rho]$ could not be reported as a repetitive region. Since $\Delta_\rho \in \mathbb{Z}$, we conclude that $\Delta_\rho \geq 1$. This inductive reasoning ultimately shows that $\Delta_{m-j} > 0$, that is, $\delta_H(P[j..m], Q^*) < \lceil 8k/m \cdot (m - j) \rceil$.

A symmetric argument holds for the values $\Delta'_\rho := \lceil 8k/m \cdot \rho \rceil - \delta_H(P[m - \rho..m], \text{rot}^{\rho-m+j}(Q)^*)$ for $\rho \in [m - j..m]$ because no repetitive region R' was found as an extension of $P[j..m]$ to the left. Thus, $\delta_H(P, \text{rot}^j(Q)^*) < 8k$, that is, $\text{rot}^j(Q)$ is an approximate period of P . \blacksquare

In the next steps, we discuss how to exploit the structure obtained by Lemma 3.6. First, we discuss the case that a string P contains $2k$ disjoint breaks.

\blacksquare **Lemma 3.8.** *Let P denote a pattern of length m , let T denote a text of length n , and let $k \in [1..m]$ denote a threshold. Suppose that P contains $2k$ disjoint breaks B_1, \dots, B_{2k} each satisfying $\text{per}(B_i) \geq m/128k$. Then, $|\text{Occ}_k^H(P, T)| \leq 256 \cdot n/m \cdot k$.*

Proof. For every break $B_i = P[b_i..b_i + |B_i|]$ we mark a position j in T if $j + b_i \in \text{Occ}(B_i, T)$.

\square **Claim 3.9.** We place at most $256 \cdot n/m \cdot k^2$ marks in total.

Proof. Fix a break B_i and notice that the positions in $\text{Occ}(B_i, T)$ are at distance at least $\text{per}(B_i)$ from each other. Hence, for the break B_i , we place at most $128 \cdot n/m \cdot k$ marks in T . In total, we therefore place at most $2k \cdot 128n/m \cdot k = 256 \cdot n/m \cdot k^2$ marks in T . \square

Next, we show that every k -mismatch occurrence of P in T starts at a position with at least k marks.

\square **Claim 3.10.** Each position $\ell \in \text{Occ}_k^H(P, T)$ has at least k marks in T .

Proof. Fix $\ell \in \text{Occ}_k^H(P, T)$. Out of the $2k$ breaks, at least k breaks are matched exactly, as not matching a break exactly incurs at least one mismatch. If a break B_i is matched exactly, then we have $\ell + b_i \in \text{Occ}(B_i, T)$. Hence, we have placed a mark at position ℓ . Thus, there is a mark at position ℓ for every break B_i matched exactly in the corresponding occurrence of P in T . In total, there are at least k marks at position ℓ in T . \square

By Claims 3.9 and 3.10, we have $|\text{Occ}_k^H(P, T)| \leq (256 \cdot n/m \cdot k^2)/k = 256 \cdot n/m \cdot k$. \blacksquare

Secondly, we discuss how to use repetitive regions in the string P to bound $|\text{Occ}_k^H(P, T)|$.

\blacksquare **Lemma 3.11.** *Let P denote a pattern of length m , let T denote a text of length n , and let $k \in [1..m]$ denote a threshold. If P contains disjoint repetitive regions R_1, \dots, R_r of total length at least $\sum_{i=1}^r |R_i| \geq 3/8 \cdot m$ such that each region R_i satisfies $|R_i| \geq m/8k$ and has a primitive approximate period Q_i with $|Q_i| \leq m/128k$ and $\delta_H(R_i, Q_i^*) = \lceil 8k/m \cdot |R_i| \rceil$, then $|\text{Occ}_k^H(P, T)| \leq 576 \cdot n/m \cdot k$.*

19 Faster Approximate Pattern Matching: A Unified Approach

Proof. Set $m_R := \sum_{i=1}^r |R_i|$. For each repetitive region $R_i = P[r_i \dots r_i + |R_i|)$, set $k_i := \lfloor 4k/m \cdot |R_i| \rfloor$, and place $|R_i|$ marks at each position j with $j + r_i \in \text{Occ}_{k_i}^H(R_i, T)$.

□ **Claim 3.12.** The total number of marks placed is at most $192 \cdot n/m \cdot k \cdot m_R$.

Proof. We use Corollary 3.5 to bound $|\text{Occ}_{k_i}^H(R_i, T)|$. For this, we set $d_i := \delta_H(R_i, Q_i^*)$ and notice that $d_i = \lceil 8k/m \cdot |R_i| \rceil \leq 16 \cdot k/m \cdot |R_i|$ since $|R_i| \geq m/8k$. Moreover, $d_i \geq 2k_i$ and $|Q_i| \leq m/128k \leq |R_i|/8d_i$ due to $d_i \leq 16 \cdot k/m \cdot |R_i|$. Hence, the assumptions of Corollary 3.5 are satisfied. Consequently, $|\text{Occ}_{k_i}^H(R_i, T)| \leq 12 \cdot n/|R_i| \cdot d_i \leq 192 \cdot n/m \cdot k$; the last inequality holds as $d_i \leq 16 \cdot k/m \cdot |R_i|$.

The total number of marks placed due to R_i is therefore bounded by $192 \cdot n/m \cdot k \cdot |R_i|$. Across all repetitive regions, this sums up to $192 \cdot n/m \cdot k \cdot m_R$, yielding the claim. □

Next, we show that every k -mismatch occurrence of P in T , starts at a position with many marks.

□ **Claim 3.13.** Each $\ell \in \text{Occ}_k^H(P, T)$ has at least $m_R - m/4$ marks.

Proof. Let us fix $\ell \in \text{Occ}_k^H(P, T)$ and denote $k'_i := \delta_H(R_i, T[\ell + r_i \dots \ell + r_i + |R_i|))$ to be the number of mismatches incurred by repetitive region R_i . Further, let $I := \{i \in [1 \dots r] \mid k'_i \leq k_i\} = \{i \in [1 \dots r] \mid k'_i \leq 4k/m \cdot |R_i|\}$ denote the set of indices of all repetitive regions that have k_i -mismatch occurrences at the corresponding positions in T . By construction, for each $i \in I$, we have placed $|R_i|$ marks at position ℓ . Hence, the total number of marks at position ℓ is at least $\sum_{i \in I} |R_i| = m_R - \sum_{i \notin I} |R_i|$. It remains to bound the term $\sum_{i \notin I} |R_i|$. Using the definition of I , we obtain

$$\sum_{i \notin I} |R_i| = \sum_{i \notin I} \frac{4mk}{4mk} \cdot |R_i| = \frac{m}{4k} \cdot \sum_{i \notin I} (4 \cdot |R_i|/m \cdot k) < \frac{m}{4k} \cdot \sum_{i \notin I} k'_i \leq \frac{m}{4k} \cdot \sum_{i=1}^r k'_i \leq \frac{m}{4},$$

where the last bound holds as, in total, all repetitive regions incur at most $\sum_{i=1}^r k'_i \leq k$ mismatches (since all repetitive regions are pairwise disjoint). Hence, the number of marks placed is at least $m_R - m/4$, completing the proof of the claim. □

In total, by Claims 3.12 and 3.13, the number of k -mismatch occurrences of P in T is at most

$$|\text{Occ}_k^H(P, T)| \leq \frac{192 \cdot n/m \cdot k \cdot m_R}{m_R - m/4} = \frac{192 \cdot n/m \cdot k}{1 - m/(4m_R)}.$$

As this bound is a decreasing function in m_R , the assumption $m_R \geq 3/8 \cdot m$ yields the upper bound

$$|\text{Occ}_k^H(P, T)| \leq \frac{192 \cdot n/m \cdot k \cdot 3/8 \cdot m}{3/8 \cdot m - m/4} = 576 \cdot n/m \cdot k,$$

completing the proof. □

Finally, we consider the case that P is approximately periodic, but not too close to the periodic string in scope.

□ **Lemma 3.14.** *Let P denote a string of length m , let T denote a string of length n , and let $k \in [1 \dots m]$ denote threshold. If there is a primitive string Q of length at most $|Q| \leq m/128k$ that satisfies $2k \leq \delta_H(P, Q^*) \leq 8k$, then $|\text{Occ}_k^H(P, T)| \leq 96 \cdot n/m \cdot k$.*

Proof. We apply Corollary 3.5 with $d = \delta_H(P, Q^*)$. As $2k \leq d \leq 8k$ yields $|Q| \leq m/128k \leq m/8d$, the assumptions of Corollary 3.5 are met. Consequently, $|\text{Occ}_k^H(P, T)| \leq 12 \cdot n/m \cdot d \leq 96 \cdot n/m \cdot k$. □

■ **Algorithm 2** A generator for the set $\text{Mis}(S, Q^*)$.

```

1 MismGenerator( $S, Q^*$ )
2   return  $\mathbf{G} \leftarrow \{S \leftarrow S; Q \leftarrow Q; i \leftarrow 0\}$ ;
3 Next( $\mathbf{G} = \{S; Q; i\}$ )
4   if  $i \geq |S|$  then return  $\perp$ ;
5    $\pi \leftarrow \text{LCP}(S[i..|S|], Q^\infty[i..])$ ;
6    $i \leftarrow i + \pi + 1$ ;
7   if  $i > |S|$  then return  $\perp$ ;
8   else return  $i - 1$ ;

```

Gathering Lemmas 3.6, 3.8, 3.11, and 3.14, we are now ready to prove Theorem 3.1, which we repeat here for convenience.

■ **Theorem 3.1.** *Given a pattern P of length m , a text T of length n , and a threshold $k \in [1..m]$, at least one of the following holds:*

- *The number of k -mismatch occurrences of P in T is bounded by $|\text{Occ}_k^H(P, T)| \leq 576 \cdot n/m \cdot k$.*
- *There is a primitive string Q of length $|Q| \leq m/128k$ that satisfies $\delta_H(P, Q^*) < 2k$.*

Proof. We apply Lemma 3.6 on the string P and proceed depending on the structure found in P .

If the string P contains $2k$ disjoint breaks B_1, \dots, B_{2k} (in the sense of Lemma 3.6), we apply Lemma 3.8 and obtain that $|\text{Occ}_k^H(P, T)| \leq 256 \cdot n/m \cdot k$.

If the string P contains r disjoint repetitive regions R_1, \dots, R_r (again, in the sense of Lemma 3.6), we apply Lemma 3.11 and obtain that $|\text{Occ}_k^H(P, T)| \leq 576 \cdot n/m \cdot k$.

Otherwise, Lemma 3.6 guarantees that there is a primitive string Q of length at most $|Q| \leq m/128k$ that satisfies $\delta_H(P, Q^*) < 8k$. If $\delta_H(P, Q^*) \geq 2k$, then Lemma 3.14 yields $|\text{Occ}_k^H(P, T)| \leq 96 \cdot n/m \cdot k$. If, however, $\delta_H(P, Q^*) < 2k$, then we are in the second alternative of the theorem statement. ■

4 Algorithm: Pattern Matching with Mismatches in the PILLAR Model

In this section, we implement the improved structural result (Theorem 3.1) from the previous section to obtain the following result.

■ **Main Theorem 8.** *Given a pattern P of length m , a text T of length n , and a positive integer $k \leq m$, we can compute (a representation of) the set $\text{Occ}_k^H(P, T)$ using $\mathcal{O}(n/m \cdot k^2 \log \log k)$ time plus $\mathcal{O}(n/m \cdot k^2)$ PILLAR operations. ■*

In general, the algorithm follows the outline given by the proof of Theorem 3.1: We first show how to implement Lemma 3.6 to preprocess the given pattern P . Then, depending on the structure of P , we (construct and) use algorithms implementing the insights from the corresponding lemmas from the previous section.

4.1 Auxiliary PILLAR Model Operations for Pattern Matching with Mismatches

We start by introducing some commonly used operations for Pattern Matching with Mismatches and show how to implement them efficiently in the PILLAR model.

▮ **Lemma 4.1** ($\text{MismatchGenerator}(S, Q^*), \text{MismatchGenerator}^R(S, *Q)$). *For every pair of strings S and Q , the sets $\text{Mis}(S, Q^*)$ and $\text{Mis}(S^R, (Q^R)^*)$ admit $(\mathcal{O}(1), \mathcal{O}(1))$ -time generators.*

Proof. We only develop the MismatchGenerator generator; $\text{MismatchGenerator}^R$ can be obtained similarly.

Given strings S and Q , the generator itself just stores S , Q , and an index i of the position *after* the last returned value by `Next`; initially, we set i to 0.

We implement the `Next` operation by using Corollary 2.9 to compute $\pi = \text{LCP}(S[i..|S|], Q^\infty[i..])$. If we observe that $i + \pi = |S|$, that is, if we reached the end of the string S , then we return \perp . Otherwise, we report $i + \pi$ and update the index i to $i + \pi + 1$. See Algorithm 2 for a pseudo-code.

For the correctness, we observe that due to storing the index i , we are able to retrieve the suffixes of S and Q to be compared, so the correctness follows.

For the running time, we observe that the creation of a generator is only bookkeeping, which takes constant time. Further, the `Next` operation uses one call to the primitive LCP operation and a single call to the LCP operation from Corollary 2.9, which uses $\mathcal{O}(1)$ PILLAR operations. Thus in total, the `Next` operation also uses $\mathcal{O}(1)$ PILLAR operations, completing the proof. ▮

▮ **Corollary 4.2** ($\text{Mismatches}(S, Q^*)$). *Given strings S and Q , we can compute the set $\text{Mis}(S, Q^*)$, using $\mathcal{O}(\delta_H(S, Q^*) + 1)$ primitive operations in the PILLAR model.*

Proof. We use a MismatchGenerator from Lemma 4.1 and call its `Next` operation until the `Next` operation returns \perp . The claim follows. ▮

▮ **Lemma 4.3** ($\text{Verify}(S, T, k)$). *Let S and T denote strings of length m each, and let $k \leq m$ denote a positive integer. Using $\mathcal{O}(k)$ PILLAR operations, we can check whether $\delta_H(S, T) \leq k$.*

Proof. We use a MismatchGenerator from Lemma 4.1 and call its `Next` operation until either the `Next` operation returns \perp (in which case we return `true`) or until we obtain the $(k + 1)$ st mismatch between S and T (in which case we return `false`). The claim follows. ▮

4.2 Computing Structure in the Pattern

In this section, we show how to implement Lemma 3.6. While the proof of Lemma 3.6 is already constructive, we still need to fill in some implementation details.

▮ **Lemma 4.4** ($\text{Analyze}(P, k)$: Implementation of Lemma 3.6). *Let P denote a string of length m and let $k \leq m$ denote a positive integer. Then, there is an algorithm that computes one of the following:*

- (a) $2k$ disjoint breaks $B_1, \dots, B_{2k} \preceq P$ each having period $\text{per}(B_i) > m/128k$ and length $|B_i| = \lfloor m/8k \rfloor$;
- (b) disjoint repetitive regions $R_1, \dots, R_r \preceq P$ of total length $\sum_{i=1}^r |R_i| \geq 3/8 \cdot m$ such that each region R_i satisfies $|R_i| \geq m/8k$ and is constructed along with a primitive approximate period Q_i such that $|Q_i| \leq m/128k$ and $\delta_H(R_i, Q_i^*) = \lceil 8k/m \cdot |R_i| \rceil$; or
- (c) a primitive approximate period Q of P with $|Q| \leq m/128k$ and $\delta_H(P, Q^*) < 8k$.

The algorithm uses $\mathcal{O}(k)$ time plus $\mathcal{O}(k)$ PILLAR operations.

Proof. Our implementation follows Algorithm 1 from the proof of Lemma 3.6: Recall that P is processed from left to right and split into breaks and repetitive regions. In each iteration, the algorithm first considers a fragment of length $\lfloor m/8k \rfloor$. This fragment either becomes the next break (if its shortest period is long enough) or is extended to the right to a repetitive region (otherwise). Having constructed sufficiently many breaks or repetitive regions of sufficiently large total length, the algorithm stops. Processing the string P

■ **Algorithm 3** A PILLAR model implementation of Algorithm 1.

```

1 Analyze( $P, k$ )
2    $j \leftarrow 0; r \leftarrow 1; b \leftarrow 1;$ 
3   while true do
4      $j' \leftarrow j + \lfloor m/8k \rfloor;$ 
5     if Period( $P[j..j']$ )  $> m/128k$  then
6        $B_b \leftarrow P[j..j'];$ 
7       if  $b = 2k$  then return breaks  $B_1, \dots, B_{2k};$ 
8        $b \leftarrow b + 1; j \leftarrow j';$ 
9     else
10       $q \leftarrow \text{Period}(P[j..j']);$ 
11       $Q_r \leftarrow P[j..j + q]; \delta \leftarrow 0;$ 
12      generator  $\mathbf{G} \leftarrow \text{MismatchGenerator}(P[j..m], Q_r^*);$ 
13      while  $\delta < 8k/m \cdot (j' - j)$  and  $(\pi \leftarrow \text{Next}(\mathbf{G})) \neq \perp$  do
14         $j' \leftarrow j + \pi + 1; \delta \leftarrow \delta + 1;$ 
15      if  $\delta \geq 8k/m \cdot (j' - j)$  then
16         $R_r \leftarrow P[j..j'];$ 
17        if  $\sum_{i=1}^r |R_i| \geq 3/8 \cdot m$  then
18          return repetitive regions  $R_1, \dots, R_r$  with periods  $Q_1, \dots, Q_r;$ 
19         $r \leftarrow r + 1; j \leftarrow j';$ 
20      else
21         $Q \leftarrow Q_r; j'' \leftarrow j;$ 
22        generator  $\mathbf{G}' \leftarrow \text{MismatchGenerator}^R(P[0..j], *Q);$ 
23        while  $\delta < 8k/m \cdot (m - j'')$  and  $(\pi \leftarrow \text{Next}(\mathbf{G}')) \neq \perp$  do
24           $j'' \leftarrow \pi; \delta \leftarrow \delta + 1;$ 
25         $Q \leftarrow P[j + (j'' - j) \bmod q .. j + (j'' - j) \bmod q + q];$            //  $Q \leftarrow \text{rot}^{j-j''}(Q)$ 
26        if  $\delta \geq 8k/m \cdot (m - j'')$  then
27          return repetitive region  $P[j''..m]$  with period  $Q$ 
28        else return approximate period  $Q;$ 

```

in this manner guarantees disjointness of breaks and repetitive regions. As in the proof of Lemma 3.6, a slightly different approach is needed if the algorithm encounters the end of P while growing a repetitive region. If this happens, the region is also extended to the left. This way, the algorithm either obtains a single repetitive region (which is not necessarily disjoint with the previously created ones, so it is returned on its own) or learns that the whole string P is close to being periodic.

Next, we fill in missing details of the implementation of the previous steps in the PILLAR model. To that end, first note that the PILLAR model includes a **Period** operation of checking if the period of a string S satisfies $\text{per}(S) \leq |S|/2$ and computing $\text{per}(S)$ in case of a positive answer. Since our threshold $m/128k$ satisfies $\lfloor m/128k \rfloor \leq \lfloor m/8k \rfloor/2$, no specific work is required to obtain the period of an unprocessed fragment of $\lfloor m/8k \rfloor$ characters of P .

To compute a repetitive region starting from a fragment $P' = P[j..j + \lfloor m/8k \rfloor]$ with string period $Q = P'[0.. \text{per}(P')]$, we use a $\text{MismatchGenerator}(P[j..m], Q^*)$ generator from Lemma 4.1: We extend P' up to the next mismatch between P' and Q^∞ until we either reach the end of P or the number $\delta = \delta_H(P', Q^*)$ reaches the bound $8k/m \cdot |P'|$. If we reach the end of P , we similarly extend $P' = P[j..m]$ to the left using a $\text{MismatchGenerator}^R(P[0..j], *Q)$ generator from Lemma 4.1: Again, we

always extend P' up to the next mismatch until we reach the start of P or the number $\delta = \delta_H(P', \bar{Q}^*)$ reaches the bound $8k/m \cdot |P'|$ (where $\bar{Q} = \text{rot}^{|P'|-m+j}(Q)$ is the corresponding cyclic rotation of Q). If we reach the start of the string, we return a suitable cyclic rotation of Q ; otherwise we found a long repetitive region, which we then return. Consider Algorithm 3 for a detailed pseudo-code of the implementation.

For the correctness, since our algorithm follows the proof of Lemma 3.6, we only need to show that our implementation of finding repetitive regions correctly implements the corresponding step in Algorithm 1. However, this is easy, as with each extension of P' , the number δ may increase by at most 1. As we start with $\delta = \delta_H(P[j..j + \lfloor m/8k \rfloor], Q^*) = 0$, we thus never skip over a repetitive region. Further, the fragment $P' = P[j..j + \lfloor m/8k \rfloor]$ by construction contains at least two repetitions of the period Q , so we can obtain each cyclic rotation of Q as a fragment of P . In particular we indeed compute a cyclic rotation of Q in Line 25 of Algorithm 3. Consequently, Algorithm 3 indeed correctly implements Algorithm 1.

For the running time analysis, observe that each iteration of the outer while loop processes at least $\lfloor m/8k \rfloor$ characters of P , so there are at most $\mathcal{O}(k)$ iterations of the outer while loop. In each iteration, we perform one **Period** operation, a constant number of **Access** operations, and at most $8k/m \cdot (j' - j)$ calls to the generator **MismGenerator**. Each of these calls uses $\mathcal{O}(1)$ **PILLAR** operations, which is $\mathcal{O}(8k/m \cdot m) = \mathcal{O}(k)$ in total across all iterations. Similarly, we bound the running time of the calls to the generator **MismGenerator^R**: As we find at most $8k/m \cdot m = 8k$ mismatches, **MismGenerator^R** uses at most $\mathcal{O}(k)$ operations. Overall, Algorithm 3 thus uses $\mathcal{O}(k)$ **PILLAR** operations.

The remaining running time is bounded by $\mathcal{O}(k)$ in the same way, completing the proof. ■

4.3 Computing Occurrences in the Periodic Case

▮ **Lemma 4.5** (**FindRotation**(k, Q, S)). *Let k denote a positive integer, let Q denote a primitive string, and let S denote a string with $|S| \geq (2k + 1)|Q|$. Then, we can compute a unique integer $j \in [0..|Q|)$ such that $\delta_H(S, \text{rot}^j(Q)^*) \leq k$, or report \perp if no such integer exists, using $\mathcal{O}(k)$ time plus $\mathcal{O}(k)$ **PILLAR** operations.*

Proof. For every $0 \leq i \leq 2k$, define $S_i := S[i|Q]..(i+1)|Q|$. We compute the majority of S_0, \dots, S_{2k} (using Lemma 2.5 for checking equality of fragments). If no majority exists, then we return \perp . Otherwise, we set \bar{Q} to be the majority string of S_0, \dots, S_{2k} and check if $\delta_H(S, \bar{Q}^*) \leq k$ using a **MismGenerator** from Lemma 4.1. If this test succeeds, we use a **Rotations** operation to retrieve all $j \in [0..|Q|)$ with $\bar{Q} = \text{rot}^j(Q)$ and return any such j . If the test fails or if no such j is found, then we return \perp .

For the correctness, observe that if $\delta_H(S, \bar{Q}^*) \leq k$, then at least $k + 1$ fragments S_i match \bar{Q} exactly, so \bar{Q} must be the majority of S_0, \dots, S_{2k} . Moreover, since Q is primitive, there is at most one $j \in [0..|Q|)$ with $\bar{Q} = \text{rot}^j(Q)$.

For the running time, note that we can compute the majority of $\mathcal{O}(k)$ elements with a classic linear-time algorithm by Boyer and Moore [9] using $\mathcal{O}(k)$ equality tests; as (by Lemma 2.5) each equality test takes $\mathcal{O}(1)$ time in the **PILLAR** model, we obtain the claimed running time and hence the claim. ■

▮ **Lemma 4.6** (**FindRelevantFragment**(P, T, d, Q)). *Let P denote a pattern of length m and let T denote a text of length $n \leq \frac{3}{2}m$. Further, let d denote a positive integer and let Q denote a primitive string that satisfies $|Q| \leq m/8d$ and $\delta_H(P, Q^*) \leq d$.*

*Then, using $\mathcal{O}(d)$ time plus $\mathcal{O}(d)$ **PILLAR** operations, we can report a fragment $T' = T[\ell..r]$ such that $\delta_H(T', Q^*) \leq 3d$ and, for every $k \leq d/2$, the set $\text{Occ}_k^H(P, T') = \{p - \ell \mid p \in \text{Occ}_k^H(P, T)\}$ contains only multiples of $|Q|$.*

Proof. We start by using a call to **FindRotation** from Lemma 4.5 to find the unique integer j such that $\delta_H(T[n - m..m], \text{rot}^j(Q)^*) \leq \frac{3}{2}d$. If no such j exists, then we return the empty string ε . Otherwise, we

■ **Algorithm 4** A PILLAR algorithm computing a *relevant* fragment T : a fragment T' such that all k -mismatch occurrences (for any $k \leq d/2$) of P in T start at a position in T' which is a multiple of $|Q|$.

```

1 FindRelevantFragment( $P, T, d, Q$ )
2    $j \leftarrow \text{FindRotation}(\lfloor \frac{3}{2}d \rfloor, Q, T[n-m..m]);$ 
3   if  $j = \perp$  then return  $\varepsilon$ ;
4    $\delta \leftarrow 0; r \leftarrow n - m + j$ ;
5   generator  $\mathbf{G} \leftarrow \text{MismatchGenerator}(T[n-m+j..n], Q^*)$ ;
6   while  $\delta \leq \frac{3}{2}d$  and  $(\pi \leftarrow \text{Next}(\mathbf{G})) \neq \perp$  do
7      $r \leftarrow n - m + j + \pi$ ;
8      $\delta \leftarrow \delta + 1$ ;
9   if  $\delta \leq \frac{3}{2}d$  then  $r \leftarrow m$ ;
10   $\delta' \leftarrow 0; \ell \leftarrow n - m + j; \ell' \leftarrow (n - m + j) \bmod |Q|$ ;
11  generator  $\mathbf{G}' \leftarrow \text{MismatchGenerator}^R(T[\ell'..n-m+j], Q^*)$ ;
12  while  $\delta' \leq \frac{3}{2}d$  and  $(\pi \leftarrow \text{Next}(\mathbf{G}')) \neq \perp$  do
13     $\ell \leftarrow \ell' + |Q| \cdot \lceil (\pi + 1) / |Q| \rceil$ ;
14     $\delta' \leftarrow \delta' + 1$ ;
15  if  $\delta' \leq \frac{3}{2}d$  then  $\ell \leftarrow \ell'$ ;
16  return  $T[\ell..r]$ ;
```

proceed by computing the rightmost position r such that $\delta_H(T[n-m+j..r], Q^*) \leq \frac{3}{2}d$ and the leftmost position ℓ (with $\ell \equiv (n-m+j) \pmod{|Q|}$) such that $\delta_H(T[\ell..n-m+j], Q^*) \leq \frac{3}{2}d$; afterwards, we return the fragment $T[\ell..r]$. Consider Algorithm 4 for implementation details.

For the correctness, first observe that $\delta_H(T[n-m..m], \text{rot}^{p-n+m}(Q)) \leq \frac{3}{2}d$ for each $p \in \text{Occ}_k^H(P, T)$: By triangle inequality (Fact 2.1), we have

$$\delta_H(T[p..p+m], Q^*) \leq k + \delta_H(P, Q^*) \leq \frac{3}{2}d.$$

Since $p \leq n-m$ and $p+m \geq m$, this yields $\delta_H(T[n-m..m], \text{rot}^{p-n+m}(Q)) \leq \frac{3}{2}d$. Moreover, $|T[n-m..m]| = 2m-n \geq m/2 \geq 4d|Q| \geq (2 \cdot \lfloor \frac{3}{2}d \rfloor + 1)|Q|$, so the call to **FindRotation** is valid.

Hence, if the call to **FindRotation** returns \perp , then $\text{Occ}_k^H(P, T) = \emptyset$ (for each $k \leq d/2$). Otherwise, each position $p \in \text{Occ}_k^H(P, T)$ satisfies $p \equiv n-m+j \equiv \ell \pmod{|Q|}$. Moreover, we have

$$\begin{aligned} \delta_H(T[n-m+j..p+m], Q^*) &\leq \delta_H(T[p..p+m], Q^*) \leq \frac{3}{2}d, \quad \text{and} \\ \delta_H(T[p..n-m+j], Q^*) &\leq \delta_H(T[p..p+m], Q^*) \leq \frac{3}{2}d. \end{aligned}$$

Hence, the fragment $T' = T[\ell..r]$ contains all k -mismatch occurrences of P in T (for any $k \leq d/2$), and all these occurrences start at multiples of $|Q|$ in T' . Moreover, $\delta_H(T', Q^*) = \delta_H(T[\ell..n-m+j], Q^*) + \delta_H(T[n-m+j..r], Q^*) \leq 3d$.

For the running time (and the number of PILLAR operations used), the call to **FindRotation** uses $\mathcal{O}(d)$ time plus $\mathcal{O}(d)$ PILLAR operations; the same is true for the usage of **MismatchGenerator** and **MismatchGenerator**^R. Thus, the algorithm uses $\mathcal{O}(d)$ time plus $\mathcal{O}(d)$ PILLAR operations in total, completing the proof. ■

■ **Lemma 4.7** (DistancesRLE(P, T, Q): Implementation of Lemma 3.3). *Let P denote a pattern of length m and let T denote a text of length $n \leq \frac{3}{2}m$. Further, let d denote a positive integer and let Q denote a string that satisfies $\delta_H(P, Q^*) = \mathcal{O}(d)$ and $\delta_H(T, Q^*) = \mathcal{O}(d)$.*

Then, using $\mathcal{O}(d^2 \log \log d)$ time plus $\mathcal{O}(d)$ PILLAR operations, we can compute a run-length encoded sequence of $h_j := \delta_H(T[j|Q|..j|Q|+m], P)$ for $0 \leq j \leq (n-m)/|Q|$.

■ **Algorithm 5** A PILLAR algorithm for Lemma 3.3

```

1 DistancesRLE( $P, T, Q$ )
  // Marking phase
2  $M \leftarrow \{\}$ ;
3 foreach  $\tau \in \text{Mismatches}(T, Q^*)$  do
4    $M \leftarrow M \cup \{(\tau - m, 1), (\tau, -1)\}$ ;
5   foreach  $\pi \in \text{Mismatches}(P, Q^*)$  do
6      $M \leftarrow M \cup \{(\tau - \pi - 1, \delta_H(P[\pi], T[\tau]) - 2), (\tau - \pi, 2 - \delta_H(P[\pi], T[\tau]))\}$ ;

  // Sliding-window phase
7 sort  $M$ ;
8  $h \leftarrow |\text{Mismatches}(P, Q^*)|$ ;
9 foreach  $(i', w) \in M$  with  $i' < 0$  do  $h \leftarrow h + w$ ;
10  $i \leftarrow 0$ ;
11 foreach  $(i', w) \in M$  with  $0 \leq i' < n - m$  sorted by  $i'$  do
12   Output a block of  $\lceil (i' + 1)/q \rceil - \lceil i/q \rceil$  values  $h$ ;
13    $i \leftarrow i' + 1$ ;
14    $h \leftarrow h + w$ ;
15 Output a block of  $\lceil (n - m + 1)/q \rceil - \lceil i/q \rceil$  values  $h$ ;
```

Proof. Observe that Claim 3.4 already gives rise to an algorithm: Starting with h_0 , we can obtain the value h_{j+1} from h_j by adding the value

$$h_{j+1} - h_j = |\text{Mis}(T, Q^*) \cap [jq + m \dots (j+1)q + m]| - |\text{Mis}(T, Q^*) \cap [jq \dots (j+1)q]| - \mu_{j+1} + \mu_j,$$

where μ_{j+1} and μ_j are defined as in Lemma 3.3.

We implement this idea in two steps: In the first step, we compute the values μ_j (using marking) and the positions of mismatches in $\text{Mis}(T, Q^*)$ (using `Mismatches` from Corollary 4.2). In the second step, we use a sliding-window approach (with the positions computed in the first step interpreted as events) to output the sequence of values of h_j . Consider the pseudo-code (Algorithm 5) for implementation details.

For the correctness, in the marking phase, the algorithm constructs a multiset M of pairs (i, w) (where i can be interpreted as a position in T and w as the weight) so that $h_j - \delta_H(P, Q^*) = w(M, j|Q|)$, where $w(M, i) = \sum_{\{(i', w) \in M \mid i' < i\}} w$ denotes the sum of weights of pairs (i', w) with $i' < i$.

Specifically, for each $\tau \in \text{Mis}(T, Q^*)$, the algorithm first inserts to M pairs $(\tau - m, 1)$ and $(\tau, -1)$. As a result, for each position i with $0 \leq i \leq n - m$, we have $w(M, i) = |\text{Mis}(T, Q^*) \cap [i \dots i + m]|$. In particular, if $i = j|Q|$, then $w(M, i) = \delta_H(T[j|Q| \dots j|Q| + m], Q^*)$. Next, for each $\tau \in \text{Mis}(T, Q^*)$ and each $\pi \in \text{Mis}(P, Q^*)$, the algorithm inserts to M pairs $(\tau - \pi - 1, \delta_H(P[\pi], T[\tau]) - 2)$ and $(\tau - \pi, 2 - \delta_H(P[\pi], T[\tau]))$. As a result, the values $w(M, i)$ with $i \neq \tau - \pi$ are not altered, whereas $w(M, i)$ for $i = \tau - \pi$ is decreased by the number of marks placed in the proof of Lemma 3.3 at position $i = \tau - \pi$ of T due to positions τ in T and π in P . Consequently, we have $w(M, j|Q|) = \delta_H(T[j|Q| \dots j|Q| + m], Q^*) - \mu_j$, which yields $h_j = \delta_H(P, Q^*) + w(M, j|Q|)$ due to Claim 3.4.

Hence, in order to construct the sequence h_j , the algorithm sorts the pairs in M and determines the partial sums $w(M, i)$. In each block of $[i \dots i']$ of equal partial sums, the algorithm reports a block with all $\lceil (i' + 1)/q \rceil - \lceil i/q \rceil$ entries h_j for $j|Q| \in [i \dots i']$, which is indeed correct.

The running time is $\mathcal{O}(d^2 \log \log d)$ (dominated by sorting M , which consists of $\mathcal{O}(d^2)$ integer pairs) plus $\mathcal{O}(d)$ PILLAR operations (for the calls to `Mismatches`(P, Q^*) and `Mismatches`(T, Q^*) and for accessing the mismatching positions of P and T), thus completing the proof. ■

■ **Algorithm 6** A PILLAR model algorithm for Lemma 3.8

```

1 BreakMatches( $P, T, \{B_1 = P[b_1..b_1 + |B_1|], \dots, B_{2k} = P[b_{2k}..b_{2k} + |B_{2k}|]\}, k$ )
2   multi-set  $M \leftarrow \{\}$ ;  $\text{Occ}_k^H(P, T) \leftarrow \{\}$ ;
3   for  $i \leftarrow 1$  to  $2k$  do
4     foreach  $\tau \in \text{ExactMatches}(B_i, T)$  do
5        $M \leftarrow M \cup \{\tau - b_i\}$ ; // Mark position  $\tau - b_i$  in  $T$ 
6     sort  $M$ ;
7     foreach  $\pi \in [0..n - m]$  that appears at least  $k$  times in  $M$  do
8       if  $\text{Verify}(P, T[\pi.. \pi + m], k)$  then  $\text{Occ}_k^H(P, T) \leftarrow \text{Occ}_k^H(P, T) \cup \{\pi\}$ ;
9   return  $\text{Occ}_k^H(P, T)$ ;
```

■ **Lemma 4.8** ($\text{PeriodicMatches}(P, T, k, d, Q)$: Implementation of Corollary 3.5). *Let P denote a pattern of length m and let T denote a text of length n . Further, let $k \leq m$ denote a non-negative integer, let $d \geq 2k$ denote a positive integer, and let Q denote a primitive string Q that satisfies $|Q| \leq n/8d$ and $\delta_H(P, Q^*) \leq d$.*

There is an algorithm that computes the set $\text{Occ}_k^H(P, T)$, represented as $\mathcal{O}(n/m \cdot d^2)$ arithmetic progressions with difference $|Q|$ (or as $\mathcal{O}(n/m \cdot d)$ individual positions if $\delta_H(P, Q^) = d$). The algorithm uses $\mathcal{O}(n/m \cdot d^2 \log \log d)$ time plus $\mathcal{O}(n/m \cdot d)$ PILLAR operations.*

Proof. First, we split the string T into $\lfloor 2n/m \rfloor$ blocks $T_i := T[\lfloor i \cdot m/2 \rfloor .. \min\{n, \lfloor (i+3) \cdot m/2 \rfloor - 1\}]$ for $0 \leq i < \lfloor 2n/m \rfloor$. For each block T_i , we call $\text{FindRelevantFragment}(P, T_i, d, Q)$ from Lemma 4.6 to obtain a fragment $T'_i = T[\ell_i .. r_i)$ containing all k -mismatch occurrences of P in T_i . Next, we call $\text{DistancesRLE}(P, T'_i, Q)$ from Lemma 4.7, yielding a run-length encoded sequence of values $h_t := \delta_H(T'_i[t|Q|..t|Q|+m], P)$ for $0 \leq t \leq (|T'_i| - m)/|Q|$. For each run $h_t = \dots = h_{t'} \leq k$, we add the arithmetic progression $\{\ell_i + j \cdot |Q| : j \in [t..t']\}$ to $\text{Occ}_k^H(P, T)$. In the end, we return the set $\text{Occ}_k^H(P, T)$.

For the correctness, note that we essentially follow the proof of Corollary 3.5. In particular, each k -mismatch occurrence of P in T is contained in exactly one of the fragments T_i . By Lemma 4.6, we see that T'_i contains all the k -mismatch occurrences of P in T_i . Moreover, as $\text{Occ}_k^H(P, T'_i)$ only contains multiples of $|Q|$, each $p \in \text{Occ}_k^H(P, T'_i)$ corresponds to an entry $h_j \leq k$. Consequently, all the k -mismatch occurrences of P in T are found. Furthermore, since $h_j = \delta_H(T[\ell_i + j|Q|.. \ell_i + j|Q| + m], P) \leq k$ holds whenever $\ell_i + j|Q|$ is reported, there are no false positives.

If $\delta_H(P, Q^*) = d$, then for each i , the number of entries h_j with $h_j \leq k$ is $\mathcal{O}(d)$ by Lemma 3.3, so the corresponding positions $\ell_i + j|Q|$ can be added to $\text{Occ}_k^H(P, T)$ individually.

The bounds on the overall running time follow from Lemmas 4.6 and 4.7 due to $\delta_H(P, Q^*) \leq d$ and since $\delta_H(T_i, Q^*) \leq 3d$ holds for each i by Lemma 4.7. ■

4.4 Computing Occurrences in the Non-Periodic Case

■ **Lemma 4.9** ($\text{BreakMatches}(P, T, \{B_1, \dots, B_{2k}\}, k)$: Implementation of Lemma 3.8). *Let P denote a string of length m having $2k$ disjoint breaks $B_1, \dots, B_{2k} \preceq P$ each satisfying $\text{per}(B_i) \geq m/128k$. Further, let T denote a string of length $n \leq \frac{3}{2}m$.*

Then, we can compute the set $\text{Occ}_k^H(P, T)$ using $\mathcal{O}(k^2 \log \log k)$ time plus $\mathcal{O}(k^2)$ PILLAR operations.

Proof. The implementation of (the marking in the proof of) Lemma 3.8 is straightforward: For each break $B_i = P[b_i..b_i + |B_i|]$, we use a call to $\text{ExactMatches}(B_i, T)$ from Lemma 2.11 to find all exact occurrences $\text{Occ}(B_i, T)$. For each occurrence $\pi \in \text{Occ}(B_i, T)$, we mark position $\pi - b_i$ in T . Having

■ **Algorithm 7** A PILLAR model algorithm for Lemma 3.11

```

1 RepetitiveMatches( $P, T, \{(R_1 = P[r_1 \dots r_1 + |R_1|], Q_1), \dots, (R_r = P[r_r \dots r_r + |R_r|], Q_r)\}, k$ )
2   multi-set  $M \leftarrow \{\}$ ;  $\text{Occ}_k^H(P, T) \leftarrow \{\}$ ;
3   for  $i \leftarrow 1$  to  $r$  do
4     foreach  $\tau \in \text{PeriodicMatches}(R_i, T, \lfloor 4 \cdot k/m \cdot |R_i| \rfloor, \lceil 8 \cdot k/m \cdot |R_i| \rceil, Q_i)$  do
5        $M \leftarrow M \cup \{(\tau - r_i, |R_i|)\}$ ; // Place  $|R_i|$  marks at position  $\tau - r_i$  in  $T$ 
6     sort  $M$  by positions;
7     foreach  $\pi \in [0 \dots n - m]$  appearing at least  $\sum_{(\pi, v) \in M} v \geq \sum_{i=1}^r |R_i| - m/4$  times in  $M$  do
8       if Verify( $P, T[\pi \dots \pi + m], k$ ) then  $\text{Occ}_k^H(P, T) \leftarrow \text{Occ}_k^H(P, T) \cup \{\pi\}$ ;
9   return  $\text{Occ}_k^H(P, T)$ ;

```

placed all marks, we run **Verify** from Lemma 4.3 for every position $\pi \in [0 \dots n - m]$ in T that has at least k marks. In the end, we return all positions where **Verify** confirmed a k -mismatch occurrence. See Algorithm 6 for a pseudo-code.

For the correctness, note that we placed the marks as in the proof of Lemma 3.8; in particular, by Claim 3.10, any $\pi \in \text{Occ}_k^H(P, T)$ has at least k marks. As we verify every possible candidate using **Verify**, we report no false positives, and thus the algorithm is correct.

We continue with analyzing the number of PILLAR operations used. As every break B_i has period $\text{per}(B_i) > m/128k$, every call to **ExactMatches** uses $\mathcal{O}(k)$ basic PILLAR operations; thus, all calls to **ExactMatches** use $\mathcal{O}(k^2)$ basic operations in total. As there are at most $\mathcal{O}(k^2/k) = \mathcal{O}(k)$ positions that we verify, and every call to **Verify** uses $\mathcal{O}(k)$ PILLAR operations, the verifications use $\mathcal{O}(k^2)$ PILLAR operations in total.

Finally, for the running time, by Claim 3.9, we place at most $\mathcal{O}(k^2)$ marks in T , so the marking step uses $\mathcal{O}(k^2)$ operations in total. Further, finding all positions in T with at least k marks can be done via a linear scan over the multiset M of all marks after sorting M , which can be done in time $\mathcal{O}(k^2 \log \log k)$. Overall, Algorithm 6 runs in time $\mathcal{O}(k^2 \log \log k)$ plus $\mathcal{O}(k^2)$ PILLAR operations. ■

■ **Lemma 4.10** (**RepetitiveMatches**($P, T, \{(R_1, Q_1), \dots, (R_r, Q_r)\}, k$): Implementation of Lemma 3.11). *Let P denote a string of length m , let T denote a string of length $n \leq \frac{3}{2}m$, and let $k \leq m$ denote a positive integer. Suppose that P contains disjoint repetitive regions R_1, \dots, R_r of total length at least $\sum_{i=1}^r |R_i| \geq 3/8 \cdot m$ such that each region R_i satisfies $|R_i| \geq m/8k$ and has a primitive approximate period Q_i with $|Q_i| \leq m/128k$ and $\delta_H(R_i, Q_i^*) = \lceil 8k/m \cdot |R_i| \rceil$.*

Then, we can compute the set $\text{Occ}_k^H(P, T)$ using $\mathcal{O}(k^2 \log \log k)$ time plus $\mathcal{O}(k^2)$ PILLAR operations.

Proof. As in the proof of Lemma 3.11, set $m_R := \sum_{i=1}^r |R_i| \geq 3/8 \cdot m$ and define for every $1 \leq i \leq r$ the values $k_i := \lfloor 4 \cdot k/m \cdot |R_i| \rfloor$ and $d_i := \lceil 8 \cdot k/m \cdot |R_i| \rceil = \lceil \text{Mis}(R_i, Q_i^*) \rceil$. Further, write $R_i = P[r_i \dots r_i + |R_i|]$.

We implement the marking of the proof of Lemma 3.11: for every repetitive region R_i , we call **PeriodicMatches**(R_i, T, k_i, d_i, Q_i) from Lemma 4.8 to obtain the set $\text{Occ}_{k_i}^H(R_i, T)$. Next, for each position $\tau \in \text{Occ}_{k_i}^H(R_i, T)$, we place $|R_i|$ marks at position $\tau - r_i$. Note that for performance reasons, instead of placing $|R_i|$ unweighted marks, we place a single mark of weight $|R_i|$ at position $\tau - r_i$.

Having placed all marks, we run **Verify** from Lemma 4.3 for every position $\pi \in [0 \dots n - m]$ in T that has marks of total weight at least $m_r - m/4$. In the end, we return all positions where **Verify** confirmed a k -mismatch occurrence. See Algorithm 7 for a pseudo-code.

For the correctness, first note that in every call to **PeriodicMatches** from Lemma 4.8, we have $16k/m \cdot |R_i| \geq d_i = \lceil 8k/m \cdot |R_i| \rceil = \delta_H(R_i, Q_i^*) \geq 2k_i$, so $|Q_i| \leq m/128k \leq |R_i|/8d_i$; hence, we can indeed

■ **Algorithm 8** A PILLAR model algorithm for Theorem 3.1

```

1 MismatchOccurrences( $P, T, k$ )
2   ( $B_1, \dots, B_{2k}$  or  $(R_1, Q_1), \dots, (R_r, Q_r)$  or  $Q$ )  $\leftarrow$  Analyze( $P, k$ );
3    $\text{Occ}_k^H(P, T) \leftarrow \{\}$ ;
4   for  $i \leftarrow 0$  to  $\lfloor 2n/m \rfloor$  do
5      $T_i \leftarrow T[\lfloor i \cdot m/2 \rfloor .. \min\{n, \lfloor (i+3) \cdot m/2 \rfloor - 1\}]$ ;
6     if breaks  $B_1, \dots, B_{2k}$  exist then
7        $\text{Occ}_k^H(P, T_i) \leftarrow$  BreakMatches( $P, T_i, \{B_1, \dots, B_{2k}\}, k$ );
8     else if repetitive regions  $(R_1, Q_1), \dots, (R_r, Q_r)$  exist then
9        $\text{Occ}_k^H(P, T_i) \leftarrow$  RepetitiveMatches( $P, T_i, \{(R_1, Q_1), \dots, (R_r, Q_r)\}, k$ );
10    else  $\text{Occ}_k^H(P, T_i) \leftarrow$  PeriodicMatches( $P, T_i, k, 8k, Q$ );
11     $\text{Occ}_k^H(P, T) \leftarrow \text{Occ}_k^H(P, T) \cup \{\ell + im/2 : \ell \in \text{Occ}_k^H(P, T_i)\}$ ;
12  return  $\text{Occ}_k^H(P, T)$ ;

```

call `PeriodicMatches` in this case. Further, note that we placed the marks as in the proof of Lemma 3.11; in particular, by Claim 3.13, any $\pi \in \text{Occ}_k^H(P, T)$ has at least $m_R - m/4$ marks. As we verify every possible candidate using `Verify`, we report no false positives, and thus the algorithm is correct.

For the number of PILLAR operations, observe that during the marking step, for every repetitive region R_i , we call `PeriodicMatches` once, and the call uses $\mathcal{O}(n/|R_i| \cdot d_i) = \mathcal{O}(m/|R_i| \cdot k/m \cdot |R_i|) = \mathcal{O}(k)$ PILLAR operations. Hence, the marking step uses $\mathcal{O}(r \cdot k) = \mathcal{O}(k^2)$ PILLAR operations in total. Next, during the verification step, by Claims 3.12 and 3.13, we call `Verify` at most $\mathcal{O}(k)$ times. As each call to `Verify` uses $\mathcal{O}(k)$ PILLAR operations, the verification step in total uses $\mathcal{O}(k^2)$ PILLAR operations. Overall, Algorithm 7 uses $\mathcal{O}(k^2)$ PILLAR operations.

Finally, for the running time, with similar calculations as for the number of PILLAR operations, we see that the marking step, including calls to `PeriodicMatches`, takes time $\sum_i \mathcal{O}(n/|R_i| \cdot d_i^2 \log \log d_i) = \sum_i \mathcal{O}(|R_i|/m \cdot k^2 \log \log k) = \mathcal{O}(k^2 \log \log k)$. Further, for every R_i , we place at most $|\text{Occ}_{k_i}^H(R_i, T)|$ (weighted) marks, which can be bounded using Corollary 3.5 by $|\text{Occ}_{k_i}^H(R_i, T)| = \mathcal{O}(n/|R_i| \cdot d_i) = \mathcal{O}(k)$. Thus, we place $|M| = \mathcal{O}(k^2)$ (weighted) marks in total. Therefore, we can sort M (by positions) in time $\mathcal{O}(k^2 \log \log k)$; afterwards, we can find the elements with total weight at least $m_R - m/4$ via a linear scan over M in time $\mathcal{O}(k^2)$. Hence, Algorithm 7 runs in $\mathcal{O}(k^2 \log \log k)$ overall time, completing the proof. ■

4.5 A PILLAR Model Algorithm for Pattern Matches with Mismatches

■ **Main Theorem 8.** *Given a pattern P of length m , a text T of length n , and a positive integer $k \leq m$, we can compute (a representation of) the set $\text{Occ}_k^H(P, T)$ using $\mathcal{O}(n/m \cdot k^2 \log \log k)$ time plus $\mathcal{O}(n/m \cdot k^2)$ PILLAR operations.*

Proof. First, we split T into overlapping parts $T_1, \dots, T_{\lfloor 2n/m \rfloor}$ of length less than $\frac{3}{2}m$ each. In order to compute $\text{Occ}_k^H(P, T_i)$ for each i , we follow the structure of the proof of Theorem 3.1: We first call `Analyze(P, k)` from Lemma 4.4. If the call to `Analyze(P, k)` yields $2k$ disjoint breaks B_1, \dots, B_{2k} in P , then we call `BreakMatches($P, T_i, \{B_1, \dots, B_{2k}\}, k$)` from Lemma 4.9. If the call to `Analyze(P, k)` yields disjoint repetitive regions R_1, \dots, R_r (and corresponding approximate periods Q_1, \dots, Q_r), then we call `RepetitiveMatches($P, T_i, \{(R_1, Q_1), \dots, (R_r, Q_r)\}, k$)` from Lemma 4.10. Finally, if the call to `Analyze(P, k)` yields an approximate period Q , then we call `PeriodicMatches($P, T_i, k, 8k, Q$)` from Lemma 4.8. The resulting set $\text{Occ}_k^H(P, T)$ is obtained by combining the sets $\text{Occ}_k^H(P, T_i)$. Consider Algorithm 8 for a visualization as pseudo-code.

For the correctness, first observe that we do not lose any occurrences by splitting the string T , since every length- m fragment of T is contained in one of the fragments T_i . Second, by Lemma 4.4 and due to $|T_i| \leq \frac{3}{2}m$, the parameters in the calls to `BreakMatches` and `RepetitiveMatches` each satisfy the requirements. Lastly, if we use `PeriodicMatches`, notice that again by Lemma 4.4 the string Q satisfies $\delta_H(P, Q^*) \leq 8k$ and $|Q| \leq m/128k \leq m/(8 \cdot 8k)$; hence, we can indeed call `PeriodicMatches` in this case.

For the number of PILLAR operations used, the call to `Analyze` uses $\mathcal{O}(k)$ PILLAR operations, each call to `BreakMatches` and `RepetitiveMatches` uses $\mathcal{O}(k^2)$ PILLAR operations, and each call to `PeriodicMatches` uses $\mathcal{O}(k)$ PILLAR operations. As there are at most $\mathcal{O}(n/m)$ calls to `BreakMatches`, `RepetitiveMatches`, and `PeriodicMatches`, we can bound the total number of PILLAR operations used by $\mathcal{O}(n/m \cdot k^2)$.

Similarly, for the running time, the call to `Analyze` takes $\mathcal{O}(k)$ time, whereas each call to `BreakMatches`, `RepetitiveMatches`, and `PeriodicMatches` takes $\mathcal{O}(k^2 \log \log k)$ time. Again, since there are at most $\mathcal{O}(n/m)$ calls to `BreakMatches`, `RepetitiveMatches`, and `PeriodicMatches` each, and combining the sets $\text{Occ}_k^H(P, T_i)$ to $\text{Occ}_k^H(P, T)$ can be implemented in total time $\mathcal{O}(n/m \cdot k^2)$, we can bound the total running time by $\mathcal{O}(n/m \cdot k^2 \log \log k)$, thus completing the proof. \blacksquare

5 Structural Insights into Pattern Matching with Edits

In this section, we develop insight into the structure of k -error occurrences of a pattern P in a text T . We prove the following result, which is analogous to Theorem 3.1 and is Main Theorem 6 with explicit constants.

\blacksquare **Theorem 5.1** (Compare Theorem 3.1). *Given a pattern P of length m , a text T of length n , and a positive integer $k \leq m$, then at least one of the following holds.*

- *The k -error occurrences of P in T satisfy $|\text{Occ}_k^E(P, T)/k| \leq 642045 \cdot n/m \cdot k$.*
- *There is a primitive string Q of length $|Q| \leq m/128k$ that satisfies $\delta_E(P, *Q^*) < 2k$.* \blacksquare

Similarly to Section 3, we start with an analysis of the (approximately) periodic case.

5.1 Characterization of the Periodic Case


\blacksquare **Theorem 5.2** (Compare Theorem 3.2). *Let P denote a pattern of length m , let T denote a text of length n , and let $k \leq m$ denote a non-negative integer such that $n < \frac{3}{2}m + k$. Suppose that the k -error occurrences of P in T include a prefix of T and a suffix of T . If there are a positive integer $d \geq 2k$ and a primitive string Q with $|Q| \leq m/8d$ and $\delta_E(P, Q^*) = \delta_E(P, *Q^*) \leq d$, then each of following holds:*

- (a) *For every $p \in \text{Occ}_k^E(P, T)$, we have $p \bmod |Q| \leq 3d$ or $p \bmod |Q| \geq |Q| - 3d$.*
- (b) *The string T satisfies $\delta_E(T, *Q^*) \leq 3d$.*
- (c) *If $\delta_E(P, *Q^*) = d$, then $|\text{Occ}_k^E(P, T)/d| \leq 304d$.*
- (d) *The set $\text{Occ}_k^E(P, T)$ can be decomposed into $617d^3$ arithmetic progressions with difference $|Q|$.* \blacksquare

\blacksquare **Lemma 5.3.** *Let k denote a positive integer, let Q denote a primitive string, and let S denote a string of length $|S| \geq (2k + 1)|Q|$. If there are integers $\ell \leq r$ and $\ell' \leq r'$ such that $\delta_E(S, Q^\infty[\ell \dots r]) \leq k$ and $\delta_E(S, Q^\infty[\ell' \dots r']) \leq k$, then there are integers j, j' and a decomposition $S = S_L \cdot S_R$ that satisfy*

$$\begin{aligned} \delta_E(S, Q^\infty[\ell \dots r]) &= \delta_E(S_L, Q^\infty[\ell \dots j|Q|]) + \delta_E(S_R, Q^\infty[j|Q| \dots r]) \quad \text{and} \\ \delta_E(S, Q^\infty[\ell' \dots r']) &= \delta_E(S_L, Q^\infty[\ell' \dots j'|Q|]) + \delta_E(S_R, Q^\infty[j'|Q| \dots r']). \end{aligned}$$

Furthermore, if $|Q| = 1$, then the assumption $|S| \geq (2k + 1)|Q|$ is not required.



■ **Figure 4** The setting in Lemma 5.3. Asterisks denote edit operations on the respective strings in their optimal alignment with S . S_i is denoted by a shaded rectangle, $U = Q[|Q| - p..|Q|)$ and $V = Q[0..|Q| - p)$.

Proof. If $|Q| = 1$, we can set $S_L := S$, $S_R := \varepsilon$, $j := r$, and $j' := r'$.

Now assume that we have $|S| \geq (2k + 1)|Q|$ and define $S_i := S[i|Q|..(i + 1)|Q|)$ for $0 \leq i \leq 2k$. Further, fix optimal alignments between S and $Q^\infty[l..r)$ and between S and $Q^\infty[l'..r')$.

Observe that at least one of the fragments S_i is aligned without errors in *both* alignments. Let us fix such a fragment S_i and observe that S_i is a length- $|Q|$ substring of Q^∞ , so $S_i = \text{rot}^p(Q)$ for some $p \in [0..|Q|)$. An illustration is provided in Figure 4. Based on this value, we set $S_L := S[0..i|Q| + p)$ and $S_R := S[i|Q| + p..|S|)$.

Next, consider the fragment Q' of $Q^\infty[l..r)$ that is aligned to S_i in the alignment fixed earlier. The fragment Q' matches $\text{rot}^p(Q)$. As Q is primitive, Q' is thus of the form $Q' = Q^\infty[j|Q| - p..(j + 1)|Q| - p)$ for some integer j . Consequently,

$$\delta_E(S, Q^\infty[l..r)) = \delta_E(S_L, Q^\infty[l..j|Q|)) + \delta_E(S_R, Q^\infty[j|Q|..r)).$$

A similar argument shows that for some integer j' , we also have

$$\delta_E(S, Q^\infty[l'..r')) = \delta_E(S_L, Q^\infty[l'..j'|Q|)) + \delta_E(S_R, Q^\infty[j'|Q|..r')).$$

This completes the proof. ■

■ **Lemma 5.4.** *Let T denote a text of length n , let k denote a positive integer, and let Q denote a primitive string. Suppose that $\delta_E(T[0..q], Q^\infty[x..y)) \leq k$ and $\delta_E(T[p..n], Q^\infty[x'..y')) \leq k$ holds for some integers $p \leq q$, $x \leq y$, and $x' \leq y'$. If $|Q| = 1$ or $q - p \geq (2k + 1)|Q|$, then $\delta_E(T, Q^\infty[x''..y)) = \delta_E(T, Q^\infty[x..y'')) \leq 2k$ for some $x'' \equiv x' \pmod{|Q|}$ and $y'' \equiv y' \pmod{|Q|}$, and $(p + x - x' + 2k) \bmod |Q| \leq 4k$.*

Proof. Observe that, for some integer $z \in [x..y]$, we have


$$\delta_E(T[0..q], Q^\infty[x..y)) = \delta_E(T[0..p], Q^\infty[x..z)) + \delta_E(T[p..q], Q^\infty[z..y)).$$

Similarly, for some integer $z' \in [x'..y']$, we have

$$\delta_E(T[p..n], Q^\infty[x'..y')) = \delta_E(T[p..q], Q^\infty[x'..z')) + \delta_E(T[q..n], Q^\infty[z'..y')).$$

Now, Lemma 5.3 applied to $S := T[p..q)$ yields an integer $r \in [p..q]$ and integers j, j' such that (see also Figure 5)

$$\begin{aligned} \delta_E(T[p..q], Q^\infty[z..y)) &= \delta_E(T[p..r], Q^\infty[z..j|Q|)) + \delta_E(T[r..q], Q^\infty[j|Q|..y)), \text{ and} \\ \delta_E(T[p..q], Q^\infty[x'..z')) &= \delta_E(T[p..r], Q^\infty[x'..j'|Q|)) + \delta_E(T[r..q], Q^\infty[j'|Q|..z')). \end{aligned}$$



■ **Figure 5** The setting in Lemma 5.4. T is at edit distance at most $2k$ from $Q^\infty[x..j|Q|)Q^\infty[j'|Q|..y'] = Q^\infty[x + (j' - j)|Q|..y'] = Q^\infty[x..y' + (j - j')|Q|)$.

This implies that

$$\begin{aligned} \delta_E(T[0..r], Q^\infty[x..j|Q|)) &= \delta_E(T[0..p], Q^\infty[x..z)) + \delta_E(T[p..r], Q^\infty[z..j|Q|)) \leq k, \text{ and} \\ \delta_E(T[r..n], Q^\infty[j'|Q|..y']) &= \delta_E(T[r..q], Q^\infty[j'|Q|..z']) + \delta_E([q..n], Q^\infty[z'..y']) \leq k. \end{aligned}$$

Combining the equations yields

$$\begin{aligned} \delta_E(T, Q^\infty[x + (j' - j)|Q|..y']) &= \delta_E(T, Q^\infty[x..y' + (j - j')|Q|)) \\ &\leq \delta_E(T[0..r], Q^\infty[x..j|Q|)) + \delta_E(T[r..n], Q^\infty[j'|Q|..y']) \leq 2k. \end{aligned}$$

Moreover, we deduce $|j|Q| - x - r| \leq k$ and $|j'|Q| - x' - r + p| \leq k$, which yields $|p + x - x' - (j - j')|Q|| \leq 2k$, and therefore $(p + x - x' + 2k) \bmod |Q| \leq 4k$. ■

■ **Definition 5.5.** Let S denote a string and let Q denote a primitive string. We say that a fragment L of S is locked (with respect to Q) if at least one of the following holds:


- For some integer α , we have $\delta_E(L, *Q^*) = \delta_E(L, Q^\alpha)$.
- The fragment L is a suffix of S and $\delta_E(L, *Q^*) = \delta_E(L, Q^*)$.
- The fragment L is a prefix of S and $\delta_E(L, *Q^*) = \delta_E(L, *Q)$.
- We have $L = S$. ■

The notion of locked fragments was also used in [17]. In order to develop some intuition, let us consider the following example: A string $U = Q^{k+1}SQ^{k+1}$ such that $\delta_E(U, *Q^*) \leq k$ and Q is primitive. Then, in any optimal alignment of U with a substring of Q^∞ at least one of the leading (or trailing) $k + 1$ occurrences of Q in U is matched exactly and hence also all occurrences preceding it (or succeeding it). Thus U is locked with respect to Q .

Next, we show that we can identify short locked fragments covering all errors with respect to $*Q^*$. Intuitively, our strategy is to start with at most $k|Q|$ -length fragments of S that contain all the errors and extend or/and merge them (in a sense similar to that of the intuitive example provided above), so that the resulting fragments contain sufficiently many copies of Q

■ **Lemma 5.6.** Let S denote a string and let Q denote a primitive string. There are disjoint locked fragments $L_1, \dots, L_\ell \preceq S$ with $\delta_E(L_i, *Q^*) > 0$ such that

$$\delta_E(S, *Q^*) = \sum_{i=1}^{\ell} \delta_E(L_i, *Q^*) \quad \text{and} \quad \sum_{i=1}^{\ell} |L_i| \leq (5|Q| + 1)\delta_E(S, *Q^*).$$



■ **Figure 6** A partition $S_0^{(0)} \dots S_4^{(0)}$ of a string S is shown ($s^{(0)} = 4$), in which all fragments apart from $S_0^{(0)}$ are interesting. A merge of the fragments $S_0^{(0)}$ and $S_1^{(0)}$ yields the shown partition $S_0^{(1)} \dots S_3^{(1)}$ of S .

Proof. Let us choose integers $x \leq y$ so that $\delta_E(S, *Q^*) = \delta_E(S, Q^\infty[x..y])$. Without loss of generality, we may assume that $x \in [0..|Q|)$. If $y \leq |Q|$, then $|S| \leq |Q| + \delta_E(S, *Q^*)$; thus setting the whole string S as the only locked fragment satisfies the claimed conditions. Hence, we may assume that $y > |Q|$.

An arbitrary optimum alignment of S and $Q^\infty[x..y]$ yields a partition $S = S_0^{(0)} \dots S_{s^{(0)}}^{(0)}$ with $s^{(0)} = \lfloor (y-1)/|Q| \rfloor$ such that $\delta_E(S, Q^\infty[x..y]) = \sum_{i=0}^{s^{(0)}} \Delta_i^{(0)}$, where

$$\Delta_i^{(0)} = \begin{cases} \delta_E(S_0^{(0)}, Q[x..|Q|]) & \text{if } i = 0, \\ \delta_E(S_i^{(0)}, Q) & \text{if } 0 < i < s^{(0)}, \\ \delta_E(S_{s^{(0)}}^{(0)}, Q[0..y - s^{(0)}|Q|]) & \text{if } i = s^{(0)}. \end{cases}$$

See Figure 6 for an illustration.

We start with this partition and then coarsen it by exhaustively applying the merging rules specified below, where each rule is applied only if the previous rules cannot be applied. In each case, we re-index the unchanged fragments $S_i^{(t)}$ to obtain a new partition $S = S_0^{(t+1)} \dots S_{s^{(t+1)}}^{(t+1)}$ and re-index the corresponding values $\Delta_i^{(t)}$ accordingly. We say that a fragment $S_i^{(t)}$ is *interesting* if $i = 0$, $i = s^{(t)}$, $S_i^{(t)} \neq Q$, or $\Delta_i^{(t)} > 0$.

1. If subsequent fragments $S_i^{(t)}$ and $S_{i+1}^{(t)}$ are both interesting, then merge $S_i^{(t)}$ and $S_{i+1}^{(t)}$, obtaining $S_i^{(t+1)} := S_i^{(t)} S_{i+1}^{(t)}$ and $\Delta_i^{(t+1)} := \Delta_i^{(t)} + \Delta_{i+1}^{(t)}$.
2. If $0 < i < s^{(t)}$ and $\Delta_i^{(t)} > 0$, then merge the subsequent fragments $S_{i-1}^{(t)} = Q$, $S_i^{(t)}$, and $S_{i+1}^{(t)} = Q$, obtaining $S_{i-1}^{(t+1)} := S_{i-1}^{(t)} S_i^{(t)} S_{i+1}^{(t)}$, and set $\Delta_{i-1}^{(t+1)} := \Delta_i^{(t)} - 1$.
3. If $0 < i = s^{(t)}$ and $\Delta_i^{(t)} > 0$, then merge the subsequent fragments $S_{i-1}^{(t)} = Q$ and $S_i^{(t)}$, obtaining $S_{i-1}^{(t+1)} := S_{i-1}^{(t)} S_i^{(t)}$, and set $\Delta_{i-1}^{(t+1)} := \Delta_i^{(t)} - 1$.
4. If $0 = i < s^{(t)}$ and $\Delta_i^{(t)} > 0$, then merge the subsequent fragments $S_i^{(t)}$ and $S_{i+1}^{(t)} = Q$, obtaining $S_i^{(t+1)} := S_i^{(t)} S_{i+1}^{(t)}$, and set $\Delta_i^{(t+1)} := \Delta_i^{(t)} - 1$.

Let $S = S_0 \dots S_s$ denote the obtained final partition. We select as locked fragments all the fragments S_i with $\delta_E(S_i, *Q^*) > 0$. Below, we show that this selection satisfies the desired properties. We start by proving that we indeed picked locked fragments.

□ **Claim 5.7.** Each fragment $S_i^{(t)}$ of each partition $S = S_0^{(t)} \dots S_{s^{(t)}}^{(t)}$ satisfies at least one of the following:

- $\delta_E(S_i^{(t)}, Q^\alpha) \leq \delta_E(S_i^{(t)}, *Q^*) + \Delta_i^{(t)}$ for some integer α ;
- $i = s^{(t)}$ and $\delta_E(S_i^{(t)}, Q^*) \leq \delta_E(S_i^{(t)}, *Q^*) + \Delta_i^{(t)}$;
- $i = 0$ and $\delta_E(S_i^{(t)}, *Q) \leq \delta_E(S_i^{(t)}, *Q^*) + \Delta_i^{(t)}$;
- $i = 0 = s^{(t)}$.

Proof. We proceed by induction on t . The base case follows from the definition of the values $\Delta_i^{(t)}$.

As for the inductive step, we assume that the claim holds for all fragments $S_i^{(t)}$ and we prove that it holds for all fragments $S_i^{(t+1)}$. We consider several cases based on the merge rule applied.

1. For a type-1 merge of interesting fragments $S_i^{(t)}$ and $S_{i+1}^{(t)}$ into $S_i^{(t+1)}$, it suffices to prove that $S_i^{(t+1)}$ satisfies the claim.

- If $0 < i < s^{(t+1)}$, then $\delta_E(S_i^{(t)}, Q^\alpha) \leq \delta_E(S_i^{(t)}, *Q^*) + \Delta_i^{(t)}$ and $\delta_E(S_{i+1}^{(t)}, Q^{\alpha'}) \leq \delta_E(S_{i+1}^{(t)}, *Q^*) + \Delta_{i+1}^{(t)}$ hold by the inductive assumption for some integers α, α' . Consequently,

$$\begin{aligned} \delta_E(S_i^{(t+1)}, Q^{\alpha+\alpha'}) &= \delta_E(S_i^{(t)} S_{i+1}^{(t)}, Q^\alpha Q^{\alpha'}) \leq \delta_E(S_i^{(t)}, Q^\alpha) + \delta_E(S_{i+1}^{(t)}, Q^{\alpha'}) \\ &\leq \delta_E(S_i^{(t)}, *Q^*) + \Delta_i^{(t)} + \delta_E(S_{i+1}^{(t)}, *Q^*) + \Delta_{i+1}^{(t)} \leq \delta_E(S_i^{(t+1)}, *Q^*) + \Delta_i^{(t+1)}. \end{aligned}$$

- If $0 < i = s^{(t+1)}$, then $\delta_E(S_i^{(t)}, Q^\alpha) \leq \delta_E(S_i^{(t)}, *Q^*) + \Delta_i^{(t)}$ and $\delta_E(S_{i+1}^{(t)}, Q^*) \leq \delta_E(S_{i+1}^{(t)}, *Q^*) + \Delta_{i+1}^{(t)}$ hold by the inductive assumption for some integer α . Consequently,

$$\begin{aligned} \delta_E(S_i^{(t+1)}, Q^*) &= \delta_E(S_i^{(t)} S_{i+1}^{(t)}, Q^*) \leq \delta_E(S_i^{(t)}, Q^\alpha) + \delta_E(S_{i+1}^{(t)}, Q^*) \\ &\leq \delta_E(S_i^{(t)}, *Q^*) + \Delta_i^{(t)} + \delta_E(S_{i+1}^{(t)}, *Q^*) + \Delta_{i+1}^{(t)} \leq \delta_E(S_i^{(t+1)}, *Q^*) + \Delta_i^{(t+1)}. \end{aligned}$$

- The analysis of the case that $0 = i < s^{(t+1)}$ is symmetric to that of the above case – this can be seen by reversing all strings in scope.

- If $0 = i = s^{(t+1)}$, then the claim holds trivially.

2. For a type-2 merge of $S_{i-1}^{(t)}$, $S_i^{(t)}$, and $S_{i+1}^{(t)}$ into $S_{i-1}^{(t+1)}$, it suffices to prove that $S_{i-1}^{(t+1)}$ satisfies the claim.

- If $\delta_E(S_{i-1}^{(t+1)}, *Q^*) > \delta_E(S_i^{(t)}, *Q^*)$, we observe that $\delta_E(S_i^{(t)}, Q^\alpha) \leq \delta_E(S_i^{(t)}, *Q^*) + \Delta_i^{(t)}$ holds by the inductive assumption for some integer α . Consequently,

$$\begin{aligned} \delta_E(S_{i-1}^{(t+1)}, Q^{\alpha+2}) &= \delta_E(Q S_i^{(t)} Q, Q Q^\alpha Q) \leq \delta_E(S_i^{(t)}, Q^\alpha) \\ &\leq \delta_E(S_i^{(t)}, *Q^*) + \Delta_i^{(t)} \leq \delta_E(S_{i-1}^{(t+1)}, *Q^*) - 1 + \Delta_i^{(t)} = \delta_E(S_{i-1}^{(t+1)}, *Q^*) + \Delta_{i-1}^{(t+1)}. \end{aligned}$$

- If $\delta_E(S_{i-1}^{(t+1)}, *Q^*) = \delta_E(S_i^{(t)}, *Q^*)$, then let $x' \leq y'$ denote integers that satisfy $\delta_E(S_{i-1}^{(t+1)}, *Q^*) = \delta_E(S_{i-1}^{(t+1)}, Q^\infty[x'..y'])$. This also yields integers x'', y'' with $x' \leq x'' \leq y'' \leq y'$ such that

$$\delta_E(S_{i-1}^{(t+1)}, Q^\infty[x'..y']) = \delta_E(Q, Q^\infty[x'..x'']) + \delta_E(S_i^{(t)}, Q^\infty[x''..y'']) + \delta_E(Q, Q^\infty[y''..y']).$$

Due to

$$\delta_E(S_i^{(t)}, *Q^*) \leq \delta_E(S_i^{(t)}, Q^\infty[x''..y'']) \leq \delta_E(S_{i-1}^{(t+1)}, Q^\infty[x'..y']) = \delta_E(S_{i-1}^{(t+1)}, *Q^*) = \delta_E(S_i^{(t)}, *Q^*),$$

we have $\delta_E(Q, Q^\infty[x'..x'']) = 0 = \delta_E(Q, Q^\infty[y''..y'])$. As the string Q is primitive, this means that x', x'', y'', y' are all multiples of $|Q|$. Consequently,

$$\delta_E(S_{i-1}^{(t+1)}, Q^{(y'-x')/|Q|}) = \delta_E(S_{i-1}^{(t+1)}, Q^\infty[x'..y']) = \delta_E(S_{i-1}^{(t+1)}, *Q^*) \leq \delta_E(S_{i-1}^{(t+1)}, *Q^*) + \Delta_{i-1}^{(t-1)}.$$

3. For a type-3 merge of $S_{i-1}^{(t)}$ and $S_i^{(t)}$ into $S_{i-1}^{(t+1)}$, it suffices to prove that $S_{i-1}^{(t+1)}$ satisfies the claim.

- If $\delta_E(S_{i-1}^{(t+1)}, *Q^*) > \delta_E(S_i^{(t)}, *Q^*)$, we observe that $\delta_E(S_i^{(t)}, Q^*) \leq \delta_E(S_i^{(t)}, *Q^*) + \Delta_i^{(t)}$ holds by the inductive assumption. Consequently,

$$\begin{aligned} \delta_E(S_{i-1}^{(t+1)}, Q^*) &= \delta_E(Q S_i^{(t)}, Q^*) \leq \delta_E(S_i^{(t)}, Q^*) \\ &\leq \delta_E(S_i^{(t)}, *Q^*) + \Delta_i^{(t)} \leq \delta_E(S_{i-1}^{(t+1)}, *Q^*) - 1 + \Delta_i^{(t)} = \delta_E(S_{i-1}^{(t+1)}, *Q^*) + \Delta_{i-1}^{(t+1)}. \end{aligned}$$

- If $\delta_E(S_{i-1}^{(t+1)}, *Q^*) = \delta_E(S_i^{(t)}, *Q^*)$, then let $x' \leq y'$ denote integers that satisfy $\delta_E(S_{i-1}^{(t+1)}, *Q^*) = \delta_E(S_{i-1}^{(t+1)}, Q^\infty[x'..y'])$. This also yields an integer x'' with $x' \leq x'' \leq y'$ such that

$$\delta_E(S_{i-1}^{(t+1)}, Q^\infty[x'..y']) = \delta_E(Q, Q^\infty[x'..x'']) + \delta_E(S_i^{(t)}, Q^\infty[x''..y']).$$

Due to

$$\delta_E(S_i^{(t)}, *Q^*) \leq \delta_E(S_i^{(t)}, Q^\infty[x''..y']) \leq \delta_E(S_{i-1}^{(t+1)}, Q^\infty[x'..y']) = \delta_E(S_{i-1}^{(t+1)}, *Q^*) = \delta_E(S_i^{(t)}, *Q^*),$$

we have $\delta_E(Q, Q^\infty[x'..x'']) = 0$. As the string Q is primitive, this means that x', x'' are both multiples of $|Q|$. Consequently,

$$\delta_E(S_{i-1}^{(t+1)}, Q^*) = \delta_E(S_{i-1}^{(t+1)}, Q^\infty[x'..y']) = \delta_E(S_{i-1}^{(t+1)}, *Q^*) \leq \delta_E(S_{i-1}^{(t+1)}, *Q^*) + \Delta_{i-1}^{(t-1)}.$$

4. The analysis of type 4 merges is symmetrical to that of 3 merges – this can be seen by reversing all strings in scope.

This completes the proof of the inductive step. \square

Observe that if no merge rule can be applied to a partition $S = S_0^{(t)} \dots S_{s^{(t)}}^{(t)}$, then $s^{(t)} = 0$ or $\Delta_0^{(t)} = \dots = \Delta_{s^{(t)}}^{(t)} = 0$. Consequently, Claim 5.7 implies that all fragments S_i in the final partition $S = S_0 \dots S_s$ are locked.

□ Claim 5.8. For each partition $S = S_0^{(t)} \dots S_{s^{(t)}}^{(t)}$, the total length $\lambda^{(t)}$ of interesting fragments satisfies

$$\lambda^{(t)} + 2|Q| \sum_{i=0}^{s^{(t)}} \Delta_i^{(t)} \leq (5|Q| + 1) \delta_E(S, *Q^*).$$

Proof. We proceed by induction on t . In the base case of $t = 0$, each interesting fragment other than $S_0^{(0)}$ and $S_{s^{(0)}}^{(0)}$ satisfies $\Delta_i^{(0)} > 0$. Hence, the number of interesting fragments is at most $2 + \sum_{i=0}^{s^{(0)}} \Delta_i^{(0)} = 2 + \delta_E(S, *Q^*)$. Moreover, the length of each fragment $S_i^{(0)}$ does not exceed $|Q| + \Delta_i^{(0)}$. Consequently,

$$\lambda^{(0)} + 2|Q| \sum_{i=0}^{s^{(0)}} \Delta_i^{(0)} \leq (2 + \delta_E(S, *Q^*))|Q| + (2|Q| + 1) \sum_{i=0}^{s^{(0)}} \Delta_i^{(0)} \leq (5|Q| + 1) \delta_E(S, *Q^*).$$

This completes the proof in the base case.

As for the inductive step, it suffices to prove that $\lambda^{(t+1)} + 2|Q| \sum_{i=0}^{s^{(t+1)}} \Delta_i^{(t+1)} \leq \lambda^{(t)} + 2|Q| \sum_{i=0}^{s^{(t)}} \Delta_i^{(t)}$:

- For a type-1 merge (where we merge two interesting fragments), we have

$$\lambda^{(t+1)} + 2|Q| \sum_{i=0}^{s^{(t+1)}} \Delta_i^{(t+1)} = \lambda^{(t)} + 2|Q| \sum_{i=0}^{s^{(t)}} \Delta_i^{(t)}.$$

- For a type-2, type-3, or type-4 merge (where we merge a fragment with its one or two non-interesting neighbors), we have

$$\lambda^{(t+1)} + 2|Q| \sum_{i=0}^{s^{(t+1)}} \Delta_i^{(t+1)} \leq \lambda^{(t)} + 2|Q| + 2|Q| \sum_{i=0}^{s^{(t+1)}} \Delta_i^{(t+1)} = \lambda^{(t)} + 2|Q| \sum_{i=0}^{s^{(t)}} \Delta_i^{(t)}.$$

Overall, we obtain the claimed bound. \square

We conclude that the total length of interesting fragments S_i does not exceed $(5|Q| + 1)\delta_E(S, *Q^*)$.

□ **Claim 5.9.** We have $\delta_E(S, *Q^*) = \sum_{i=0}^s \delta_E(S_i, *Q^*)$.

Proof. The claim is immediate if $s = 0$; hence, assume that $s \geq 1$. Observe that the inequality $\sum_{i=0}^s \delta_E(S_i, *Q^*) \leq \delta_E(S, *Q^*)$ easily follows from disjointness of fragments S_i ; thus, we focus on proving $\delta_E(S, *Q^*) \leq \sum_{i=0}^s \delta_E(S_i, *Q^*)$.

For $0 \leq i \leq s$, let Q_i denote a substring of Q^∞ that satisfies $\delta_E(S_i, *Q^*) = \delta_E(S_i, Q_i)$. Since each S_i is locked (by Claim 5.7), we may assume that for $0 < i < s$ the substring Q_i is a power of Q , the substring Q_s is a prefix of a power of Q , and the substring Q_0 is a suffix of a power of Q . Consequently, $Q_0 \cdots Q_s$ is a substring of Q^∞ , and we have

$$\delta_E(S, *Q^*) \leq \delta_E(S_0 \cdots S_s, Q_0 \cdots Q_s) \leq \sum_{i=0}^s \delta_E(S_i, Q_i) = \sum_{i=0}^s \delta_E(S_i, *Q^*),$$

thus completing the proof. □

The locked fragments created satisfy $\delta_E(S, *Q^*) = \sum_{i=1}^\ell \delta_E(L_i, *Q^*)$ due to Claim 5.9. Moreover, since $\delta_E(S_i, *Q^*) > 0$ holds only for interesting fragments, Claim 5.8 yields $\sum_{i=1}^\ell |L_i| \leq (5|Q| + 1)\delta_E(S, *Q^*)$, completing the proof. ■

The definition and lemma that follow, as well as Theorem 5.2 (d), are not needed for our proof of the main result of this section, Theorem 5.1 – a reader interested only in that result can safely skip them. They, however, provide additional structural insights that we exploit in Section 6.

□ **Definition 5.10.** Let S denote a string, let Q denote a primitive string, and let $k \geq 0$ denote an integer. We say that a prefix L of S is k -locked (with respect to Q) if at least one of the following holds:

- For every $p \in [0..|Q|)$, if $\delta_E(L, \text{rot}^p(Q)^*) \leq k$, then $\delta_E(L, \text{rot}^p(Q)^*) = \delta_E(L, Q^\infty[|Q| - p..j|Q|])$ holds for some integer j .
- We have $L = S$. ■

□ **Lemma 5.11.** Let S denote a string, let Q denote a primitive string, and let $k \geq 0$ be an integer. There are disjoint locked fragments $L_1, \dots, L_\ell \preceq S$, such that L_1 is a k -locked prefix of S , L_ℓ is a suffix of S , $\delta_E(L_i, *Q^*) > 0$ for $1 < i < \ell$,

$$\delta_E(S, *Q^*) = \sum_{i=1}^\ell \delta_E(L_i, *Q^*), \quad \text{and} \quad \sum_{i=1}^\ell |L_i| \leq (5|Q| + 1)\delta_E(S, *Q^*) + 2(k + 1)|Q|.$$

Proof. We proceed as in the proof of Lemma 5.6 except that $\Delta_0^{(0)}$ is artificially increased by $k + 1$, the prefix S_0 in the final partition is included as L_1 among the locked fragments even if $\delta_E(S_0, *Q^*) = 0$, and the suffix S_s is included as L_ℓ among the locked fragments even if $\delta_E(S_s, *Q^*) = 0$.

It is easy to see that Claims 5.7 and 5.9 remain satisfied, whereas the upper bound in Claim 5.8 is increased by $2(k + 1)|Q|$. We only need to prove that S_0 is a k -locked prefix of S . For this, we prove the following claim using induction.

□ **Claim 5.12.** For each partition $S = S_0^{(t)} \cdots S_{s^{(t)}}^{(t)}$, at least one of the following holds:

- For every $p \in [0..|Q|)$, if $\delta_E(S_0^{(t)}, \text{rot}^p(Q)^*) \leq k - \Delta_0^{(t)}$, then $\delta_E(S_0^{(t)}, Q^\infty[|Q| - p..j|Q|]) \leq \delta_E(S_0^{(t)}, \text{rot}^p(Q)^*) + \Delta_0^{(t)}$ holds for some integer j .
- We have $S_0^{(t)} = S$.

Proof. We proceed by induction on t . In the base case of $t = 0$, the claim holds trivially since $\delta_E(S_0^{(0)}, \text{rot}^p(Q)^*) \geq 0 > k - \Delta_0^{(0)}$ holds for every p due to $\Delta_0^{(0)} \geq k + 1$.

As for the induction step, we assume that the claim holds for $S_0^{(t)}$ and we prove that it holds for $S_0^{(t+1)}$. The claim holds trivially if the merge rule applied did not affect $S_0^{(t)}$. Given that $S_0^{(t)}$ is interesting by definition, the merges that might affect $S_0^{(t)}$ are of type 1 (if $S_1^{(t)}$ is interesting) or 4 (otherwise).

1. Consider a type-1 merge of $S_0^{(t)}$ and $S_1^{(t)}$. If $s^{(t)} = 1$, then $S_0^{(t+1)} = S$ satisfies the claim trivially. Hence, we may assume that $1 < s^{(t)}$ so that Claim 5.7 yields $\delta_E(S_1^{(t)}, Q^\alpha) \leq \delta_E(S_1^{(t)}, *Q^*) + \Delta_1^{(t)}$ for some integer α . Let us fix $p \in [0..|Q|)$ with $\delta_E(S_0^{(t+1)}, \text{rot}^p(Q)^*) \leq k - \Delta_0^{(t+1)}$. Due to $\Delta_0^{(t+1)} \geq \Delta_0^{(t)}$, this yields $\delta_E(S_0^{(t)}, \text{rot}^p(Q)^*) \leq k - \Delta_0^{(t)}$, so the inductive assumption implies $\delta_E(S_0^{(t)}, Q^\infty[|Q| - p..j|Q|]) \leq \delta_E(S_0^{(t)}, \text{rot}^p(Q)^*) + \Delta_0^{(t)}$ for some integer j . Consequently,

$$\begin{aligned} \delta_E(S_0^{(t+1)}, Q^\infty[|Q| - p..(j + \alpha)|Q|]) &= \delta_E(S_0^{(t)} S_1^{(t)}, Q^\infty[|Q| - p..j|Q|] Q^\alpha) \\ &\leq \delta_E(S_0^{(t)}, Q^\infty[|Q| - p..j|Q|]) + \delta_E(S_1^{(t)}, Q^\alpha) \leq \delta_E(S_0^{(t)}, \text{rot}^p(Q)^*) + \Delta_0^{(t)} + \delta_E(S_1^{(t)}, *Q^*) + \Delta_1^{(t)} \\ &\leq \delta_E(S_0^{(t+1)}, \text{rot}^p(Q)^*) + \Delta_0^{(t+1)}. \end{aligned}$$

2. Consider a type-4 merge of $S_0^{(t)}$ and $S_1^{(t)}$. Let us fix $p \in [0..|Q|)$ with $\delta_E(S_0^{(t+1)}, \text{rot}^p(Q)^*) \leq k - \Delta_0^{(t+1)}$.
 - If $\delta_E(S_0^{(t+1)}, \text{rot}^p(Q)^*) > \delta_E(S_0^{(t)}, \text{rot}^p(Q)^*)$, then

$$\delta_E(S_0^{(t)}, \text{rot}^p(Q)^*) \leq \delta_E(S_0^{(t+1)}, \text{rot}^p(Q)^*) - 1 \leq k - \Delta_0^{(t+1)} - 1 = k - \Delta_0^{(t)}.$$

Hence, the inductive assumption implies $\delta_E(S_0^{(t)}, Q^\infty[|Q| - p..j|Q|]) \leq \delta_E(S_0^{(t)}, \text{rot}^p(Q)^*) + \Delta_0^{(t)}$ for some integer j . Consequently,

$$\begin{aligned} \delta_E(S_0^{(t+1)}, Q^\infty[|Q| - p..(j + 1)|Q|]) &= \delta_E(S_0^{(t)} Q, Q^\infty[|Q| - p..j|Q|] Q) \\ &\leq \delta_E(S_0^{(t)}, Q^\infty[|Q| - p..j|Q|]) \leq \delta_E(S_0^{(t)}, \text{rot}^p(Q)^*) + \Delta_0^{(t)} \leq \delta_E(S_0^{(t+1)}, \text{rot}^p(Q)^*) - 1 + \Delta_0^{(t)} \\ &= \delta_E(S_0^{(t+1)}, \text{rot}^p(Q)^*) + \Delta_0^{(t+1)}. \end{aligned}$$

- If $\delta_E(S_0^{(t+1)}, \text{rot}^p(Q)^*) = \delta_E(S_0^{(t)}, \text{rot}^p(Q)^*)$, then let y' denote an arbitrary integer that satisfies $\delta_E(S_0^{(t+1)}, \text{rot}^p(Q)^*) = \delta_E(S_0^{(t+1)}, Q^\infty[|Q| - p..y'])$. This also yields an integer y'' with $|Q| - p \leq y'' \leq y'$ such that

$$\delta_E(S_0^{(t+1)}, Q^\infty[|Q| - p..y']) = \delta_E(S_0^{(t)}, Q^\infty[|Q| - p..y'']) + \delta_E(Q, Q^\infty[y''..y']).$$

Due to

$$\begin{aligned} \delta_E(S_0^{(t)}, \text{rot}^p(Q)^*) &\leq \delta_E(S_0^{(t)}, Q^\infty[|Q| - p..y'']) \leq \delta_E(S_0^{(t+1)}, Q^\infty[|Q| - p..y']) \\ &= \delta_E(S_0^{(t+1)}, \text{rot}^p(Q)^*) = \delta_E(S_0^{(t)}, \text{rot}^p(Q)^*), \end{aligned}$$

we have $\delta_E(Q, Q^\infty[y''..y'])$. As the string Q is primitive, this means that y'', y' are both multiples of $|Q|$. Consequently,

$$\delta_E(S_0^{(t+1)}, Q^\infty[|Q| - p..y']) = \delta_E(S_0^{(t+1)}, \text{rot}^p(Q)^*) \leq \delta_E(S_0^{(t+1)}, \text{rot}^p(Q)^*) + \Delta_0^{(t+1)}.$$

This completes the proof of the inductive step. □

Given that the final partition $S = S_0^{(t)} \cdots S_{s^{(t)}}^{(t)}$ satisfies $s^{(t)} = 0$ or $\Delta_0^{(t)} = 0$, we conclude that S_0 is indeed k -locked. ■

We are now ready to prove Theorem 5.2, which we restate here for convenience.

▮ **Theorem 5.2** (Compare Theorem 3.2). *Let P denote a pattern of length m , let T denote a text of length n , and let $k \leq m$ denote a non-negative integer such that $n < \frac{3}{2}m + k$. Suppose that the k -error occurrences of P in T include a prefix of T and a suffix of T . If there are a positive integer $d \geq 2k$ and a primitive string Q with $|Q| \leq m/8d$ and $\delta_E(P, Q^*) = \delta_E(P, *Q^*) \leq d$, then each of following holds:*

- (a) *For every $p \in \text{Occ}_k^E(P, T)$, we have $p \bmod |Q| \leq 3d$ or $p \bmod |Q| \geq |Q| - 3d$.*
- (b) *The string T satisfies $\delta_E(T, *Q^*) \leq 3d$.*
- (c) *If $\delta_E(P, *Q^*) = d$, then $|\lfloor \text{Occ}_k^E(P, T)/d \rfloor| \leq 304d$.*
- (d) *The set $\text{Occ}_k^E(P, T)$ can be decomposed into $617d^3$ arithmetic progressions with difference $|Q|$.*

Proof. Consider any k -error occurrence $T[\ell..r]$ of P . By definition, $\delta_E(T[\ell..r], P) \leq k \leq d/2$. Combining this inequality with $\delta_E(P, Q^*) \leq d$ via the triangle inequality (Fact 2.3), we obtain the bound $\delta_E(T[\ell..r], Q^*) \leq \frac{3}{2}d$. In particular, this inequality is true for the k -error occurrence of P as a prefix of T . Hence, for some integer $m' \in [m - k..m + k]$, we have $\delta_E(T[0..m'], Q^*) \leq \frac{3}{2}d$, and thus also $\delta_E(T[0..\min(r, m')], Q^*) \leq \frac{3}{2}d$.

Next, we apply Lemma 5.4 on the fragment $T[0..r]$, whose prefix $T[0..\min(r, m')]$ satisfies $\delta_E(T[0..\min(r, m')], Q^*) \leq \frac{3}{2}d$ and whose suffix $T[\ell..r]$ satisfies $\delta_E(T[\ell..r], Q^*) \leq \frac{3}{2}d$. Further, if $|Q| > 1$, we also have $|T[\ell..\min(r, m')]| \geq (3d + 1)|Q|$: Due to $r - \ell \geq m - k$ and $m' \geq m - k$, we have $\min(r, m') - \ell \geq 2(m - k) - n > m/2 - 3k \geq 4d|Q| - 3k$. Hence, it suffices to prove that $(d - 1)|Q| \geq 3k - 1$. This equality holds trivially if $k = 0$. For $k \geq 1$, we have $(d - 1)|Q| \geq (2k - 1) \cdot 2 = 4k - 2 \geq 3k - 1$ due to $d \geq 2k$ and $|Q| \geq 2$. Thus, we can indeed use Lemma 5.4.

In particular, Lemma 5.4 implies $(\ell + 3d) \bmod |Q| \leq 6d$. Since $T[\ell..r]$ was an arbitrary k -error occurrence of P in T , we conclude that Claim (a) holds.

Moreover, Lemma 5.4 implies $\delta_E(T[0..r], Q^*) \leq 3d$. If we choose $T[\ell..r]$ to be a k -error occurrence of P that is a suffix of T , we have $r = n$ and therefore $\delta_E(T, Q^*) \leq 3d$, which proves Claim (b).

We proceed to the proof of Claim (c). Let L_1, \dots, L_ℓ denote locked fragments of P obtained from Lemma 5.6. Note that we thus have $\sum_{i=1}^{\ell} \delta_E(L_i, *Q^*) = \delta_E(P, *Q^*) = d$, $\ell \leq d$, and $\sum_{i=1}^{\ell} |L_i| \leq (5|Q| + 1)d$.

Moreover, let us fix an optimal alignment between T and a substring Q' of Q^* , and define $d' := \delta_E(T, Q^*)$. This yields partitions $T = T_0 \cdots T_{2d'}$ and $Q' = Q'_0 \cdots Q'_{2d'}$ such that:

- ▀ $T_i = Q'_i$ for even i ,
- ▀ $T_i \neq Q'_i$ and $|T_i|, |Q'_i| \leq 1$ for odd i .

We create a multi-set $E = \{\sum_{i' < i} |T_{i'}| : i \text{ is odd}\}$ of size d' . Its elements can be interpreted as positions in T which incur errors in an optimal alignment with Q' . In particular, we show the following:

▮ **Claim 5.13.** For every fragment $T[x..y]$, we have $\delta_E(T[x..y], *Q^*) \leq |\{e \in E \mid x \leq e < y\}|$.

Proof. It suffices to observe that the alignment between T and Q' yields an alignment between $T[x..y]$ and a fragment $Q'[x'..y']$ with $\delta_E(T[x..y], Q'[x'..y']) \leq |\{e \in E \mid x \leq e < y\}|$ edits. \square

We split \mathbb{Z} into disjoint blocks of the form $[jd..(j+1)d]$ for $j \in \mathbb{Z}$. We say that a block $[jd..(j+1)d]$ is *synchronized* if it contains a position p such that $(p + 3d) \bmod |Q| \leq 6d$. For every locked fragment $L_i = P[\ell_i..r_i]$ and every $e \in E$, we mark a synchronized block $[jd..(j+1)d]$ if

$$e \in [jd + \ell_i - k..(j+1)d - 1 + r_i + k].$$

▮ **Claim 5.14.** If $[jd..(j+1)d] \cap \text{Occ}_k^E(P, T) \neq \emptyset$, then $[jd..(j+1)d]$ has at least $d - k$ marks.

Proof. Consider a k -error occurrence of P in T starting at position $p \in [jd \dots (j+1)d]$ and fix its arbitrary optimal alignment with P . For each locked fragment L_i , let L'_i be the fragment of T aligned with L_i in this alignment. Moreover, $L'_i = T[\ell'_i \dots r'_i]$ for

$$[\ell'_i \dots r'_i] \subseteq [p + \ell_i - k \dots p + r_i + k] \subseteq [jd + \ell_i - k \dots (j+1)d - 1 + r_i + k].$$

Also, by Claim 5.13, we have

$$\delta_E(L'_i, *Q^*) \leq |\{e \in E \mid \ell'_i \leq e < r'_i\}| \leq |\{e \in E \mid jd + \ell_i - k \leq e < (j+1)d - 1 + r_i + k\}|.$$

Hence, the number μ of marks at $[jd \dots (j+1)d]$ is at least $\mu \geq \sum_{i=1}^{\ell} \delta_E(L'_i, *Q^*)$. On the other hand, by disjointness of regions L_i , we have $\sum_{i=1}^{\ell} \delta_E(L'_i, L_i) \leq k$. By the triangle inequality (Fact 2.3), this yields $\sum_{i=1}^{\ell} \delta_E(L_i, *Q^*) \leq k + \mu$. Since $\sum_{i=1}^{\ell} \delta_E(L_i, *Q^*) = \delta_E(P, *Q^*) = d$, we conclude that $\mu \geq d - k$. \square

\square Claim 5.15. The total number of marks placed is at most $152d^2$.

Proof. Let us fix $e \in E$ and a locked fragment $L_i = P[\ell_i \dots r_i]$. Recall that a mark is placed at a synchronized block $[jd \dots (j+1)d]$ if $e \in [jd + \ell_i - k \dots (j+1)d - 1 + r_i + k]$, or, in other words,

$$[jd \dots (j+1)d] \cap [e - r_i - k \dots e - \ell_i + k] \neq \emptyset.$$

The length of the interval $I_{i,e} := [e - r_i - k \dots e - \ell_i + k]$ satisfies $|I_{i,e}| = 2k + |L_i| \leq d + |L_i|$.

We now consider two cases depending on whether or not the inequality $|Q| < 9d$ is satisfied. If $|Q| < 9d$, it suffices to observe that any interval I overlaps with at most $2 + |I|/d$ blocks. Hence, $I_{i,e}$ overlaps with at most $3 + |L_i|/d$ blocks, and thus the number of marks we have placed due to e and L_i is bounded by $3 + |L_i|/d$. The overall number of marks is therefore at most

$$|E| \cdot \sum_{i=1}^{\ell} (3 + |L_i|/d) \leq 9d^2 + 3 \sum_{i=1}^{\ell} |L_i| \leq 9d^2 + 3(5|Q| + 1)d \leq 9d^2 + 135d^2 = 144d^2.$$

On the other hand, if $|Q| \geq 9d$, we utilize the fact that only synchronized blocks are marked. For this, observe that any interval I overlaps at most $2 + |I|/|Q|$ intervals of the form $[(j' - 1)|Q| - 4d \dots j'|Q| + 4d]$ each of which overlaps with at most 7 synchronized blocks (covering $[j'|Q| - 3d \dots j'|Q| + 3d]$, which is of length $6d+1$). Hence, the total number of blocks marked due to e and L_i is bounded by $7(2 + (d + |L_i|)/|Q|)$. The overall number of marks is therefore at most

$$\begin{aligned} |E| \cdot \sum_{i=1}^{\ell} 7(2 + (d + |L_i|)/|Q|) &\leq 42d^2 + 21d^2/|Q| + 21d/|Q| \cdot \sum_{i=1}^{\ell} |L_i| \\ &\leq 42d^2 + 21d^2/|Q| + 21d^2 \cdot (5|Q| + 1)/|Q| \leq 42d^2 + 21d^2/|Q| + 105d^2 + 21d^2/|Q| < 152d^2. \end{aligned}$$

This completes the proof of the claim. \square

Hence, the number of blocks with at least $d/2$ marks is at most $304d$, completing the proof of Claim (c).

We proceed to the proof of Claim (d). Let $L_1^P, \dots, L_{\ell^P}^P$ denote locked fragments of P obtained from Lemma 5.11 (so that L_1^P is a k -locked prefix of P), and let $L_1^T, \dots, L_{\ell^T}^T$ denote locked fragments of T obtained from Lemma 5.6. Denote $L_i^P = P[\ell_i^P \dots r_i^P]$ for $i \in [1 \dots \ell^P]$ and $L_j^T = T[\ell_j^T \dots r_j^T]$ for $j \in [1 \dots \ell^T]$. We say that a position $p \in [0 \dots n]$ is *marked* if $p \in [n - m - k \dots n - m + k]$ or $[p + \ell_i^P - k \dots p + r_i^P + k] \cap [\ell_j^T \dots r_j^T] \neq \emptyset$ holds for some $i \in [1 \dots \ell^P]$ and $j \in [1 \dots \ell^T]$. (The positions in

$[n-m-k \dots n-m+k]$ are marked for a technical reason that will become clear in the proof of Claim 5.17.) Furthermore, we say that p is *synchronized* if $p \bmod |Q| \leq 3d$ or $p \bmod |Q| \geq |Q| - 3d$.

Let us provide some intuition. Informally, if there is a k -occurrence in an unmarked position p , then no locked region of P can overlap a locked region of T in any corresponding optimal alignment and we can exploit this structure. On the other hand, we now show that there are only a few synchronized marked positions.

□ **Claim 5.16.** Marked positions can be decomposed into at most $10d^2$ integer intervals. Moreover, the number of synchronized marked positions is at most $547d^3$.

Proof. First, we have the interval $[n-m-k \dots n-m+k]$. Observe that each pair $i \in [1 \dots \ell^P]$ and $j \in [1 \dots \ell^T]$ yields to marking positions $p \in (\ell_j^T - r_i^P - k \dots r_j^T - \ell_i^P + k)$. Consequently, marked positions can be decomposed into $1 + \ell^P \cdot \ell^T$ integer intervals. Due to $\ell^P \leq \delta_E(P, *Q^*) + 2 \leq d + 2 \leq 3d$ and $\ell^T \leq \delta_E(T, *Q^*) \leq 3d$, the number of intervals is at most $10d^2$.

The interval of positions marked due to i and j is of length $2k + |L_i^P| + |L_j^T| - 1 \leq d + |L_i^P| + |L_j^T| - 1$. Out of any $|Q|$ consecutive positions at most $6d + 1 \leq 7d$ are synchronized. Hence, the number of synchronized positions in any such interval I is at most $7d(|I| - 7d + |Q|)/|Q|$. Consequently, the total number of synchronized marked positions does not exceed $2k + 1 \leq d^3$ (for $[n-m-k \dots n-m+k]$) plus

$$\begin{aligned}
 \sum_{i=1}^{\ell^P} \sum_{j=1}^{\ell^T} 7d \frac{|L_i^P| + |L_j^T| - 7d + |Q|}{|Q|} &\leq \frac{7d\ell^T}{|Q|} \sum_{i=1}^{\ell^P} |L_i^P| + \frac{7d\ell^P}{|Q|} \sum_{i=1}^{\ell^T} |L_i^T| + \frac{7d\ell^P\ell^T(|Q| - 7d)}{|Q|} \\
 &\leq \frac{21d^2}{|Q|} \left(\sum_{i=1}^{\ell^P} |L_i^P| + \sum_{i=1}^{\ell^T} |L_i^T| + 3d(|Q| - 7d) \right) \\
 &\leq \frac{21d^2}{|Q|} ((5|Q| + 1)d + 2(k + 1)|Q| + (5|Q| + 1)3d + 3d(|Q| - 7d)) \\
 &\leq \frac{21d^2}{|Q|} ((5|Q| + 1)d + 3d|Q| + (5|Q| + 1)3d + 3d(|Q| - 7d)) \\
 &\leq \frac{21d^2}{|Q|} (26d|Q| + 4d - 21d^2) \\
 &\leq 21 \cdot 26d^3 = 546d^3.
 \end{aligned}$$

This completes the proof. □

Next, we characterize unmarked positions $p \in \text{Occ}_k^E(P, T)$. They can be decomposed into at most $10d^2$ integer intervals by Claim 5.16 and by the fact that $p < n - m - k$. Consider any such interval I .

□ **Claim 5.17.** For any $p, p' \in I$ such that $p \equiv p' \pmod{|Q|}$ we have $p \in \text{Occ}_k^E(P, T)$ if and only if $p' \in \text{Occ}_k^E(P, T)$. In particular, $I \cap \text{Occ}_k^E(P, T)$ can be decomposed into at most $6d + 1$ arithmetic progressions with difference $|Q|$.

Proof. By our marking scheme, for any i, j , for any pair $x \in [\ell_i^P \dots r_i^P)$ and $y \in [\ell_j^T \dots r_j^T)$, we see that $|(p+x) - y| > k$. Consider an unmarked position $p \in \text{Occ}_k^E(P, T) \cap I$ and fix any alignment of some prefix $T[p \dots t)$ of $T[p \dots n)$ with P with at most k errors. Then, for all i , the fragment T_i of T aligned with L_i^P is disjoint from all locked fragments of T and hence is a substring of Q^∞ . Now, recall that L_1^P is a prefix of P and $L_{\ell^P}^P$ is a suffix of P . Hence, the locked fragments of T that are considered are exactly those L_j^T 's such that $p+k < \ell_j < r_j < p+m-k$; say that this holds for $j \in [j_1 \dots j_2]$. For all $j \in [j_1 \dots j_2]$, the

fragment P_j of P aligned with L_j^T is disjoint from all locked fragments of P and is a substring of Q^∞ . In addition, since $I = [i_1 \dots i_2]$ is an interval of unmarked positions, $T[i_1 \dots i_2 + |L_1^P| + k]$ is disjoint from all locked fragments of T and hence is equal to a substring of Q^∞ . Thus, for some r we see that:

$$\begin{aligned} \delta_E(P, T[p \dots t]) &\geq \delta_E(L_1^P, \text{rot}^r(Q)^*) + \sum_{i=2}^{\ell^P} \delta_E(L_i^P, Q_i) + \sum_{j=j_1}^{j_2} \delta_E(\ell_j^T, P_j) \\ &\geq \delta_E(L_1^P, \text{rot}^r(Q)^*) + \sum_{i=2}^{\ell^P} \delta_E(L_i^P, *Q^*) + \sum_{j=j_1}^{j_2} \delta_E(\ell_j^T, *Q^*). \end{aligned} \quad (2)$$

Note that since L_1^P is k -locked and $\delta_E(L_1^P, \text{rot}^r(Q)^*) \leq k$ there exists a j such that $\delta_E(L_1^P, \text{rot}^r(Q)^*) = \delta_E(L_1^P, Q^\infty[|Q| - r \dots j|Q|])$. Let $b = |Q^\infty[|Q| - r \dots j|Q|]|$.

Consider any position $p' \in I$ with $p' \equiv p \pmod{|Q|}$. $T[p' \dots p' + b] = T[p \dots p + b]$ since both fragments lie in $T[i_1 \dots i_2 + |L_1^P| + k]$ and start a multiple of $|Q|$ positions apart. In addition, we have $P[|L_1^P| \dots m] = Q^{\alpha_1} L_2^P Q^{\alpha_2} \dots Q^{\alpha_{\ell^P-1}} L_{\ell^P}^P$ for some non-negative integers α_i and that $T[p' + b \dots p' + m + k] = Q^{\beta_{j_1-1}} L_{j_1}^T Q^{\beta_{j_1}} \dots L_{j_2}^T Q^{\beta_{j_2}} Q'$ for some non-negative integers β_j and a prefix Q' of Q . (Note that $p' + m + k < n$ since $p' < n - m - k$.) We claim that there exists a t' such that:

$$\begin{aligned} \delta_E(P, T[p' \dots t']) &\leq \delta_E(L_1^P, T[p' \dots p' + b]) + \delta_E(P[|L_1^P| \dots m], T[p' + b \dots t']) \\ &= \delta_E(L_1^P, \text{rot}^r(Q)^*) + \sum_{i=2}^{\ell^P} \delta_E(L_i^P, *Q^*) + \sum_{j=j_1}^{j_2} \delta_E(\ell_j^T, *Q^*) \leq \delta_E(P, T[p \dots t]). \end{aligned}$$

In order to prove this, let us consider the following greedy alignment of $P[|L_1^P| \dots m]$ and $T[p' + b \dots t']$. We start at the leftmost position in both strings. We will maintain the invariant that the remainder of each string starts with either Q or a locked fragment, except possibly for the case that the remainder of P is $L_{\ell^P}^P$, in which case the remainder of T can be Q' . While we have not reached the end of $P[|L_1^P| \dots m]$ we repeat the following procedure. If both strings have a prefix equal to Q , we align those prefixes exactly. Else, the prefix of one of the strings is a locked fragment L . Let us first assume that $L \neq L_{\ell^P}^P$. Then, since p' is unmarked, $Q^\infty[0 \dots |L| + k - k']$ is a prefix of the other string, where k' is the number of edits already performed by our greedy alignment. Since L is locked, $\delta_E(L, *Q^*) = \delta_E(L, Q^\alpha)$ for some integer α . We have $|\alpha|Q| - |L| \leq k - k'$ due to the fact that otherwise (2) would imply that $\delta_E(P, T[p \dots t]) > k$, a contradiction. Hence, Q^α is a prefix of $Q^\infty[0 \dots |L| + k - k']$; we optimally align those two fragments. If $L = L_{\ell^P}^P$, an analogous argument shows that there exists a prefix Q'' of the remainder of $T[p' + b \dots t']$ such that $\delta_E(L, *Q^*) = \delta_E(L, Q'')$. Upon termination of this greedy alignment, the equality in the above equation holds.

We have thus proved that if $p \in I \cap \text{Occ}_k^E(P, T)$ then all $p' \in I$ such that $p' \equiv p \pmod{|Q|}$ are also in $\text{Occ}_k^E(P, T)$. Thus, for any fixed $0 \leq j < |Q|$, for $U = \{i \cdot |Q| + j \in I \mid i \in \mathbb{Z}\}$ either $U \subseteq \text{Occ}_k^E(P, T)$ or $U \cap \text{Occ}_k^E(P, T) = \emptyset$. By Claim (a), we can restrict our attention to synchronized positions. We can thus decompose $I \cap \text{Occ}_k^E(P, T)$ into at most $6d + 1$ arithmetic progressions. \square

Combining Claims 5.16 and 5.17 we conclude that $\text{Occ}_k^E(P, T)$ can be decomposed into at most $547d^3 + 10d^2(6d + 1) \leq 617d^3$ arithmetic progressions with difference $|Q|$, thus completing the proof. \blacksquare

Corollary 5.18 (Compare Corollary 3.5). *Let P denote a pattern of length m , let T denote a string of length n , and let $k \leq m$ denote a non-negative integer. Suppose that there are a positive integer $d \geq 2k$ and a primitive string Q with $|Q| \leq m/8d$ and $\delta_E(P, *Q^*) = d$. Then $|\text{Occ}_k^E(P, T)/d| \leq 1216 \cdot n/m \cdot d$.*

■ **Algorithm 9** A constructive proof of Lemma 5.19. Changes to Algorithm 1 are highlighted.

```

1  $\mathcal{B} \leftarrow \{\}; \mathcal{R} \leftarrow \{\};$ 
2 while true do
3   Consider the fragment  $P' = P[j..j + \lfloor m/8k \rfloor]$  of the next  $\lfloor m/8k \rfloor$  unprocessed characters
   of  $P$ ;
4   if  $\text{per}(P') > m/128k$  then
5      $\mathcal{B} \leftarrow \mathcal{B} \cup \{P'\};$ 
6     if  $|\mathcal{B}| = 2k$  then return breaks  $\mathcal{B}$ ;
7   else
8      $Q \leftarrow P[j..j + \text{per}(P')];$ 
9     Search for prefix  $R$  of  $P[j..m]$  with  $\delta_E(R, *Q^*) = \lceil 8k/m \cdot |R| \rceil$  and  $|R| > |P'|$ ;
10    if such  $R$  exists then
11       $\mathcal{R} \leftarrow \mathcal{R} \cup \{(R, Q)\};$ 
12      if  $\sum_{(R,Q) \in \mathcal{R}} |R| \geq 3/8 \cdot m$  then
13        return repetitive regions (and their corresponding periods)  $\mathcal{R}$ ;
14    else
15      Search for suffix  $R'$  of  $P$  with  $\delta_E(R', *Q^*) = \lceil 8k/m \cdot |R'| \rceil$  and  $|R'| \geq m - j$ ;
16      if such  $R'$  exists then return repetitive region  $(R', Q)$ ;
17    else return approximate period  $Q$ ;

```

Proof. Partition the string T into $\lfloor 2n/m \rfloor$ blocks $T_0, \dots, T_{\lfloor 2n/m \rfloor - 1}$ of length at most $\frac{3}{2}m + k - 1$ each, where the i th block starts at position $i \cdot m/2$, that is, $T_i := T[\lfloor i \cdot m/2 \rfloor .. \min\{n, \lfloor (i+3) \cdot m/2 \rfloor + k - 1\}]$. Observe that each k -error occurrence of P in T is contained in at least one of the fragments T_i : Specifically, T_i covers all the occurrences starting in $[\lfloor i \cdot m/2 \rfloor .. \lfloor (i+1) \cdot m/2 \rfloor]$. If $\text{Occ}_k^E(P, T_i) \neq \emptyset$, we define $T'_i := T[t'_i .. t'_i + |T'_i|]$ to be the shortest fragment of T_i containing all k -error occurrences of P in T_i . As a result, T'_i satisfies the assumptions of Theorem 5.2, so $|\text{Occ}_k^E(P, T'_i)|/k \leq 304d$. Each block $[j'k .. (j'+1)k]$ of positions in T_i corresponds to a block $[t'_i + j'k .. t'_i + (j'+1)k]$ of positions in T , which intersects at most two blocks of the form $[jk .. (j+1)k]$. In total, we conclude that $|\text{Occ}(P, T_i)|/k \leq \lfloor 2n/m \rfloor \cdot 2 \cdot 304d \leq 1216 \cdot n/m \cdot d$. ■

5.2 Bounding the Number of Occurrences in the Non-Periodic Case

■ **Lemma 5.19** (Compare Lemma 3.6). *Let P denote a string of length m and let $k \leq m$ denote a positive integer. Then, at least one of the following holds:*

- (a) *The string P contains $2k$ disjoint breaks B_1, \dots, B_{2k} each having period $\text{per}(B_i) > m/128k$ and length $|B_i| = \lfloor m/8k \rfloor$.*
- (b) *The string P contains disjoint repetitive regions R_1, \dots, R_r of total length $\sum_{i=1}^r |R_i| \geq 3/8 \cdot m$ such that each region R_i satisfies $|R_i| \geq m/8k$ and has a primitive approximate period Q_i with $|Q_i| \leq m/128k$ and $\delta_E(R_i, *Q_i^*) = \lceil 8k/m \cdot |R_i| \rceil$.*
- (c) *The string P has a primitive approximate period Q with $|Q| \leq m/128k$ and $\delta_E(P, *Q^*) < 8k$.*

Proof. We use essentially the same algorithm as in the proof of Lemma 3.6: We replace all checks for a specific Hamming distance with the corresponding counterpart for the edit distance. Further, as we are only interested in (approximate) periods under an arbitrary rotation, we do not need to explicitly

rotate the string Q in the algorithm anymore. Consider Algorithm 9 for a visualization; the changes to Algorithm 1 are highlighted.

In particular, we directly get an analogue of Claim 3.7:

□ Claim 5.20 (See Claim 3.7). Whenever we consider a new fragment $P[j..j + \lfloor m/8k \rfloor]$ of $\lfloor m/8k \rfloor$ unprocessed characters of P , such a fragment starts at a position $j < 5/8 \cdot m$. □

Again, note that Claim 5.20 also shows that whenever we consider a new fragment P' of $\lfloor m/8k \rfloor$ characters, there is indeed such a fragment, that is, P' is well-defined.

Now consider the case when, for a fragment $P' = P[j..j + \lfloor m/8k \rfloor]$ (that is not a break) and its corresponding period $Q = [j..j + \text{per}(P')]$, we fail to obtain a new repetitive region R . Recall that in this case, we search for a repetitive region R' of length $|R'| \geq m - j$ that is a suffix of P and has an approximate period Q . If we indeed find such a region R' , then $|R'| \geq m - j \geq m - 5/8 \cdot m = 3/8 \cdot m$ by Claim 5.20, so R' is long enough to be reported on its own. However, if we fail to find such R' , we need to show that Q can be reported as an approximate period of P , that is, $\delta_E(P, *Q^*) < 8k$.

Similar to Lemma 3.6, we first show that $\delta_E(P[j..m], *Q^*) < \lceil 8k/m \cdot (m - j) \rceil$. For this, we inductively prove that the values $\Delta_\rho := \lceil 8k/m \cdot \rho \rceil - \delta_E(P[j..j + \rho], *Q^*)$ for $|P'| \leq \rho \leq m - j$ are all at least 1. In the base case of $\rho = |P'|$, we have $\Delta_\rho = 1 - 0$ because Q is the string period of P' . To carry out an inductive step, suppose that $\Delta_{\rho-1} \geq 1$ for some $|P'| < \rho \leq m - j$. Notice that $\Delta_\rho \geq \Delta_{\rho-1} - 1 \geq 0$: the first term in the definition of Δ_ρ has not decreased, and the second term $\delta_E(P[j..j + \rho], *Q^*)$ may have increased by at most 1 compared to $\Delta_{\rho-1}$. Moreover, $\Delta_\rho \neq 0$ because $R = P[j..j + \rho]$ could not be reported as a repetitive region. Since Δ_ρ is an integer, we conclude that $\Delta_\rho \geq 1$. This inductive reasoning ultimately shows that $\Delta_{m-j} > 0$, that is, $\delta_E(P[j..m], *Q^*) < \lceil 8k/m \cdot (m - j) \rceil$.

A symmetric argument holds for values $\Delta'_\rho := \lceil 8k/m \cdot \rho \rceil - \delta_E(P[m - \rho..m], *Q^*)$ for $m - j \leq \rho \leq m$ because no repetitive region R' was found as an extension of $P[j..m]$ to the left. Note that in contrast to the proof of Lemma 3.6, the rotation of Q is implicit. This completes the proof that $\delta_E(P, *Q^*) < 8k$, that is, Q is an approximate period of P . ▀

□ **Lemma 5.21** (Compare Lemma 3.8). *Let P denote a pattern of length m , let T denote a text of length n , and let $k \leq m$ denote a positive integer. Suppose that P that contains $2k$ disjoint breaks $B_1, \dots, B_{2k} \preceq P$ each satisfying $\text{per}(B_i) \geq m/128k$. Then, $|\lfloor \text{Occ}_k^E(P, T) / k \rfloor| \leq 1024 \cdot n/m \cdot k$.*

Proof. The proof proceeds similarly to the proof of Lemma 3.8. The only major difference is that we obtain length- k blocks of possible starting positions instead of single starting positions. This is because the edit distance allows for deletions and insertions of characters.

Hence, we split \mathbb{Z} into disjoint blocks of the form $[jk..(j+1)k)$ for $j \in \mathbb{Z}$. Now for every break $B_i = P[b_i..b_i + |B_i|)$, we mark a block $[jk..(j+1)k)$ if

$$[(j-1) \cdot k + b_i..(j+2) \cdot k + b_i) \cap \text{Occ}(B_i, T) \neq \emptyset.$$

Similarly to the proof of Lemma 3.8, we proceed to show that we place at most $\mathcal{O}(n/m \cdot k^2)$ marks and that every k -error occurrence starts in a block with at least k marks.

□ Claim 5.22. The total number of marks placed at blocks is at most $1024 \cdot n/m \cdot k^2$.

Proof. Fix a break B_i . Notice that positions in $\text{Occ}(B_i, T)$ are at distance at least $\text{per}(B_i)$ from each other. Furthermore, note that for every occurrence in $\text{Occ}(B_i, T)$ we mark at most 4 blocks. Hence, for the break B_i , we place at most $512 \cdot n/m \cdot k$ marks. In total, the number of marks placed is thus at most $2k \cdot 512n/m \cdot k = 1024 \cdot n/m \cdot k^2$. □

Next, we show that every k -error occurrence of P in T starts in a block with at least k marks.

□ **Claim 5.23.** If $[jk \dots (j+1)k] \cap \text{Occ}_k^E(P, T) \neq \emptyset$, then $[jk \dots (j+1)k]$ has at least k marks.

Proof. Consider a k -error occurrence of P in T starting at position $\ell \in [jk \dots (j+1)k]$ and fix an arbitrary optimal alignment of it with P . Out of the $2k$ breaks, at least k breaks are matched exactly, as not matching a break exactly incurs at least one error. If a break B_i is matched exactly, then for at least one $s \in [-k \dots k]$, we have $\ell + b_i + s \in \text{Occ}(B_i, T)$. Since $jk \leq \ell < (j+1)k$, we conclude that $[(j-1) \cdot k + b_i \dots (j+2) \cdot k + b_i] \cap \text{Occ}(B_i, T) \neq \emptyset$, that is, that the block $[jk \dots (j+1)k]$ has been marked for B_i . In total, there are at least k marks for the at least k breaks matched exactly. □

By Claims 5.22 and 5.23, the number of blocks where k -error occurrences of P in T may start is $|\text{Occ}_k^E(P, T)/k| \leq (1024 \cdot n/m \cdot k^2)/k = 1024 \cdot n/m \cdot k$. ▀

□ **Lemma 5.24** (Compare Lemma 3.11). *Let P denote a pattern of length m , let T denote a text of length n , and let $k \leq m$ denote a positive integer. Suppose that P contains disjoint repetitive regions R_1, \dots, R_r of total length at least $\sum_{i=1}^r |R_i| \geq 3/8 \cdot m$ such that each region R_i satisfies $|R_i| \geq m/8k$ and has a primitive approximate period Q_i with $|Q_i| \leq m/128k$ and $\delta_E(R_i, *Q_i^*) = \lceil 8k/m \cdot |R_i| \rceil$. Then, $|\text{Occ}_k^E(P, T)/k| \leq 642045 \cdot n/m \cdot k$.*

Proof. Again, the proof is similar to its Hamming distance counterpart; as before, a major difference is that we only obtain length- k blocks of possible starting positions instead of single starting positions.

As in the proof of Lemma 5.21, we split \mathbb{Z} into disjoint blocks of the form $[jk \dots (j+1)k]$ for $j \in \mathbb{Z}$. Further, we set $m_R := \sum_r |R_i|$ and define $k_i := \lfloor 4 \cdot |R_i|/m \cdot k \rfloor$ for every $1 \leq i \leq r$.

For every repetitive region $R_i = P[r_i \dots r_i + |R_i|]$, we place $|R_i|$ marks on block $[jk \dots (j+1)k]$ if

$$[(j-1) \cdot k + r_i \dots (j+2) \cdot k + r_i] \cap \text{Occ}_{k_i}^E(R_i, T) \neq \emptyset.$$

Similarly to the proof Lemma 3.11, we proceed to show that we placed at most $\mathcal{O}(n/m \cdot k \cdot m_R)$ marks and that every k -error occurrence of P in T starts in a block with at last $m_R - m/4$ marks.

□ **Claim 5.25.** The total number of marks placed is at most $214015 \cdot n/m \cdot k \cdot m_R$,

Proof. We use Corollary 5.18 to analyze $\text{Occ}_{k_i}^E(R_i, T)$. For this, we set $d_i := \delta_E(R_i, *Q_i^*) = \lceil 8k/m \cdot |R_i| \rceil$ and notice that $d_i \leq 16k/m \cdot |R_i|$ since $|R_i| \geq m/8k$. Moreover, $d_i \geq 2k_i$ and $|Q_i| \leq m/128k \leq |R_i|/8d_i$. Hence, the assumptions of Corollary 5.18 are satisfied, so $|\text{Occ}_{k_i}^E(R_i, T)/d_i| \leq 1216 \cdot n/|R_i| \cdot d_i \leq 19456 \cdot k \cdot n/m$.

For a block $[j'd_i \dots (j'+1)d_i]$ intersecting $\text{Occ}_{k_i}^E(R_i, T)$, we mark a block $[jk \dots (j+1)k]$ only if

$$[(j-1) \cdot k + r_i \dots (j+2) \cdot k + r_i] \cap [j'd_i \dots (j'+1)d_i] \neq \emptyset,$$

which holds only if $jk \in [j'd_i - r_i - 2k \dots (j'+1)d_i - r_i + k]$. The length of the latter interval is $d_i + 3k = \lceil 8k/m \cdot |R_i| \rceil + 3k \leq 11k$, so the interval contains at most 11 multiples of k . Hence, the total number of marks placed due to R_i is bounded by $11 \cdot 19456 \cdot n/m \cdot k \cdot |R_i|$. Across all repetitive regions, this sums up to no more than $214015 \cdot n/m \cdot k \cdot m_R$, yielding the claim. □

Next, we show that every k -error occurrence of P in T starts in a block with many marks.

□ **Claim 5.26.** If $[jk \dots (j+1)k] \cap \text{Occ}_k^E(P, T) \neq \emptyset$, then $[jk \dots (j+1)k]$ has at least $m_R - m/4$ marks.

Proof. Consider a k -error occurrence of P in T starting at position $\ell \in [jk \dots (j+1)k]$ and fix an arbitrary optimal alignment of it with P . For each repetitive region R_i , let R'_i be the fragment of T aligned with R_i in this alignment. Define $k'_i = \delta_E(R_i, R'_i)$ and observe that $R'_i = T[r'_i \dots r'_i + |R'_i|)$ for some $r'_i \in [\ell + r_i - k \dots \ell + r_i + k] \subseteq [(j-1) \cdot k + r_i \dots (j+2) \cdot k + r_i]$. Consequently,

$$[(j-1) \cdot k + r_i \dots (j+2) \cdot k + r_i] \cap \text{Occ}_{k'_i}^E(R_i, T) \neq \emptyset.$$

Further, let $I := \{i \mid k'_i \leq k_i\} = \{i \mid k'_i \leq 4 \cdot |R_i|/m \cdot k\}$ denote the set of indices i for which R'_i is a k_i -error occurrence of R_i . By construction, for each $i \in I$, we have placed $|R_i|$ marks at the block $[jk \dots (j+1)k]$.

Hence, the total number of marks at the block $[jk \dots (j+1)k]$ is at least $\sum_{i \in I} |R_i| = m_R - \sum_{i \notin I} |R_i|$. It remains to bound the term $\sum_{i \notin I} |R_i|$. Using the definition of I , we obtain

$$\sum_{i \notin I} |R_i| = \frac{m}{4k} \cdot \sum_{i \notin I} (4 \cdot |R_i|/m \cdot k) < \frac{m}{4k} \cdot \sum_{i \notin I} k'_i \leq \frac{m}{4k} \cdot \sum_{i=1}^r k'_i \leq \frac{m}{4},$$

where the last bound holds because, in total, all repetitive regions incur at most $\sum_{i=1}^r k'_i \leq k$ errors (since the repetitive regions are disjoint). Hence, the number of marks placed is at least $m_r - m/4k$, completing the proof of the claim. \square

In total, by Claims 5.25 and 5.26, the number of k -error occurrences of P in T is at most

$$\text{Occ}_k^E(P, T) \leq \frac{214015 \cdot n/m \cdot k \cdot m_R}{m_R - m/4}.$$

As this bound is a decreasing function in m_R , the assumption $m_R \geq 3/8 \cdot m$ yields

$$\text{Occ}_k^E(P, T) \leq \frac{214015 \cdot n/m \cdot k \cdot 3/8 \cdot m}{3/8 \cdot m - m/4} = 642045 \cdot n/m \cdot k,$$

completing the proof. \square

Lemma 5.27 (Compare Lemma 3.14). *Let P denote a pattern of length m , let T denote a text of length n , and let $k \leq m$ denote a positive integer. If there is a primitive string Q of length at most $|Q| \leq m/128k$ that satisfies $2k \leq \delta_E(P, *Q^*) \leq 8k$, then $|\text{Occ}_k^E(P, T)/k| \leq 87551 \cdot n/m \cdot k$.*

Proof. We apply Corollary 5.18 with $d = \delta_E(P, *Q^*)$. Observe that $d \geq 2k$ and that $|Q| \leq m/128k \leq m/8d$ due to $d \leq 8k$. Hence, the assumptions of Corollary 5.18 are met.

Consequently, $|\text{Occ}_k^E(P, T)/d| \leq 1216 \cdot n/m \cdot d \leq 9728 \cdot n/m \cdot k$. Every block $[j'd \dots (j'+1)d]$ is of length at most $8k$, and thus may intersect at most 9 blocks of the form $[jk \dots (j+1)k]$. Consequently, $|\text{Occ}_k^E(P, T)/k| \leq 9 \cdot 9728 \cdot n/m \cdot k$, completing the proof. \square

Theorem 5.1 (Compare Theorem 3.1). *Given a pattern P of length m , a text T of length n , and a positive integer $k \leq m$, then at least one of the following holds.*

- The k -error occurrences of P in T satisfy $|\text{Occ}_k^E(P, T)/k| \leq 642045 \cdot n/m \cdot k$.
- There is a primitive string Q of length $|Q| \leq m/128k$ that satisfies $\delta_E(P, *Q^*) < 2k$.

Proof. The proof proceeds similarly to the proof of Theorem 3.1: We apply Lemma 5.19 on the string P and proceed depending on the structure found in P .

If the string P contains $2k$ disjoint breaks B_1, \dots, B_{2k} (in the sense of Lemma 5.19), we apply Lemma 5.21 and obtain that $|\text{Occ}_k^E(P, T)/k| \leq 1024 \cdot n/m \cdot k$.

If the string P contains disjoint repetitive regions R_1, \dots, R_r (again, in the sense of Lemma 5.19), we apply Lemma 5.24 and obtain that $|\text{Occ}_k^E(P, T)/k| \leq 642045 \cdot n/m \cdot k$.

Otherwise, Lemma 5.19 guarantees that there is a primitive string Q of length at most $|Q| \leq m/128k$ that satisfies $\delta_E(P, *Q^*) < 8k$. If $\delta_E(P, *Q^*) \geq 2k$, then Lemma 5.27 yields $|\text{Occ}_k^E(P, T)/k| \leq 87551 \cdot n/m \cdot k$. If, however, $\delta_E(P, *Q^*) < 2k$, then we are in the second alternative of the theorem statement. \blacksquare

6 Algorithm: Pattern Matching with Edits in the PILLAR Model

In this section, we discuss how to solve pattern matching with edits in the PILLAR model. Specifically, we prove the following.

Main Theorem 9. *Given a pattern P of length m , a text T of length n , and a positive integer $k \leq m$, we can compute (a representation of) the set $\text{Occ}_k^E(P, T)$ using $\mathcal{O}(n/m \cdot k^4)$ time in the PILLAR model. \blacksquare*

The overall structure of the algorithm is similar to the Hamming distance case: We first introduce useful tools for the algorithms later. Then, we implement **Analyze**, which is then followed by a discussion of the case when the pattern is periodic. Finally, we discuss the easier non-periodic case and conclude with combining the various auxiliary algorithms.

6.1 Auxiliary PILLAR Model Operations for Pattern Matching with Edits

As in the Hamming distance setting, we start the discussion of the algorithms with general tools that we use as auxiliary procedures in the remaining algorithms. Specifically, we discuss a generator that computes the “next” error between two strings. Further, we discuss a procedure to verify whether there is an occurrence of the pattern at a given (interval of) position(s) in the text.

Lemma 6.1 ($\text{EditGenerator}(S, Q)$, $\text{EditGenerator}^R(S, Q)$). *Let S denote a string and let Q denote a string (that is possibly given as a cyclic rotation $Q' = \text{rot}^j(Q)$). Then, there is an $(\mathcal{O}(1), \mathcal{O}(k))$ -time generator that in the k -th call to **Next** returns the length of the longest prefix (suffix) S' of S and the length of the corresponding prefix (suffix) Q' of Q^∞ such that $\delta_E(S', Q') \leq k$. (Note that $k \geq 0$, that is, the initial call to **Next** is the zeroth call.)*

*Further, both generators support an additional operation **Alignment**, that outputs a witness for the result returned by k -th call to **Next** that is, **Alignment** outputs a sequence of edits $((i, j)$ for a replacement, (i, \perp) for an insertion in S , and (\perp, i) for an insertion in Q). The operation **Alignment** takes $\mathcal{O}(k)$ time in the PILLAR model.*

Proof. We focus on the generator $\text{EditGenerator}(S, Q)$; the generator $\text{EditGenerator}^R(S, Q)$ can be obtained in a symmetric manner.

We construct the generator as follows: We run the dynamic programming algorithm by Landau and Vishkin [35] for one additional error (per call to **Next**) at a time, storing the dynamic programming table as a state in the generator. In particular, we maintain a sequence ℓ that after k calls to **Next** stores at a position $i \in [-k \dots k]$ the length of the longest prefix S' of S that satisfies $\delta_E(S', Q^\infty[0 \dots |S'| + i]) \leq k$ (that is, intuitively, we store how far each of the “diagonals” of the dynamic programming table extends). In each call to **Next**, we update the values stored in the sequence ℓ by using three calls to **LCP** from Corollary 2.9 (one call for each of the **insert**, **replace**, **delete** cases of the edit distance) to compute each new entry. We then obtain the result as the maximum value of the newly computed sequence.

In order to support the **Alignment** operation, we additionally store every diagonal represented as list of pairs of **insert**, **replace**, and **delete** operations and the position(s) in the strings S and Q^∞ where

■ **Algorithm 10** An analogue of Algorithm 2 for edit distance: We adapt Landau and Vishkin’s algorithm to make its execution “resumable”.

```

1 EditGenerator( $S, Q' = \text{rot}^j(Q)$ )
2   return {
3      $S \leftarrow S; Q' \leftarrow Q'; k \leftarrow 0; j \leftarrow j; \text{end} \leftarrow \text{false};$ 
4      $(\ell_{-1}^{(-1)}, \dots, \ell_1^{(-1)}) \leftarrow (-\infty, -\infty, -\infty);$ 
5      $A[-2..2] \leftarrow [(), \dots, ()];$ 
6   };
7 Next( $\mathbf{G} = \{S; Q'; k; j; (\ell_{-k-1}^{(k-1)}, \dots, \ell_{k+1}^{(k-1)}); A; \text{end}\}$ )
8   if  $k = 0$  then
9      $(\ell_{-2}^{(0)}, \dots, \ell_2^{(0)}) \leftarrow (-\infty, -\infty, \text{LCP}(S, Q^\infty), -\infty, -\infty);$ 
10    replace  $\ell^{(k-1)}$  with  $\ell^{(k)}$ ;  $k \leftarrow k + 1$ ;
11    return  $\ell_0^{(0)}$ ;
12   $(\ell_{-k-2}^{(k)}, \dots, \ell_{k+2}^{(k)}) \leftarrow (-\infty, \dots, -\infty);$ 
13   $A'[-k-2..k+2] \leftarrow [(), \dots, ()]; r \leftarrow -\infty; a_r \leftarrow -1$ ;
14  for  $i \leftarrow -k-1$  to  $k+1$  do
15    // Compute new prefix lengths as long as we did not reach the end of  $S$ 
16    if not  $\text{end}$  then
17       $\ell_{\text{insert}} \leftarrow \ell_{i-1}^{(k-1)} + 1 + \text{LCP}(S[\ell_{i-1}^{(k-1)} - i..|S|], Q'^\infty[j + \ell_i^{(k-1)} + 1..]);$ 
18       $\ell_{\text{replace}} \leftarrow \ell_i^{(k-1)} + 1 + \text{LCP}(S[\ell_i^{(k-1)} - i + 1..|S|], Q'^\infty[j + \ell_i^{(k-1)} + 1..]);$ 
19       $\ell_{\text{delete}} \leftarrow \ell_{i+1}^{(k-1)} + \text{LCP}(S[\ell_i^{(k-1)} - i + 1..|S|], Q'^\infty[j + \ell_i^{(k-1)}..]);$ 
20       $\ell_i^{(k)} \leftarrow \max(\ell_{\text{insert}}, \ell_{\text{replace}}, \ell_{\text{delete}});$ 
21    else  $\ell_i^{(k)} \leftarrow \ell_i^{(k-1)}$ ;
22     $r \leftarrow \max(r, \ell_i^{(k)})$ ;
23    if  $r = \ell_i^{(k)} - i$  then  $a_r \leftarrow \ell_i^{(k)} + i$ ;
24    if  $\text{end}$  then continue;
25    // Store witness for Alignment
26    if  $\ell_i^{(k)} = \ell_{\text{insert}}$  then
27       $A'[i] \leftarrow (A[i-1], (\perp, j + \ell_i^{(k-1)}));$ 
28    if  $\ell_i^{(k)} = \ell_{\text{replace}}$  then
29       $A'[i] \leftarrow (A[i], (\ell_i - i, j + \ell_i^{(k-1)}));$ 
30    if  $\ell_i^{(k)} = \ell_{\text{delete}}$  then
31       $A'[i] \leftarrow (A[i+1], (\ell_i^{(k-1)} - i, \perp));$ 
32    replace  $\ell^{(k-1)}$  with  $\ell^{(k)}$ ;  $k \leftarrow k + 1$ ;  $A \leftarrow A'$ ;
33    if  $r \geq |S|$  then  $\text{end} \leftarrow \text{true}$ ;
34    return  $(r, a_r)$ ;
35 Alignment( $\mathbf{G} = \{S; Q'; k; j; (\ell_{-k-1}^{(k-1)}, \dots, \ell_{k+1}^{(k-1)}); A; \text{end}\}$ )
36    $r \leftarrow -\infty; a_r \leftarrow -1$ ;
37   for  $i \leftarrow -k-1$  to  $k+1$  do
38      $r \leftarrow \max(r, \ell_i^{(k-1)})$ ;
39     if  $r = \ell_i^{(k-1)}$  then  $a_r \leftarrow i$ ;
40   return  $A[a_r]$ ;

```

the edits happened. In each call to **Next**, we also update the representations of the diagonals. (Note that, for performance reasons, we need to avoid copying whole diagonals; this can be done by storing all diagonals together as a (directed) graph.)

Consider Algorithm 10 for a visualization of the generator as pseudo-code; note that we simplified how we store the diagonals for the **Alignment** operation to improve readability.

For the correctness, we show by induction that the values computed in the array ℓ are indeed correct, that is, after k calls to the **Next** operation, for all $i \in [-k \dots k]$ we have

$$\ell_i^{(k)} = \max_r \{r \mid \delta_E(S[0..r], Q^\infty[0..r+i]) \leq k\},$$

or we reached the end of the string S (in which case the output does not change anymore after calling **Next**). For the zeroth call to next, we explicitly return

$$\ell_0^{(0)} = \text{LCP}(S, Q^\infty) = \max_r \{r \mid \delta_E(S[0..r], Q^\infty[0..r+0]) = 0\},$$

which is thus correct. Now consider the k -th call to **Next** and assume that the values computed so far are correct. In particular, for all $i \in [-k+1 \dots k-1]$, we have

$$\ell_i^{(k-1)} = \max_r \{r \mid \delta_E(S[0..r], Q^\infty[0..r+i]) \leq k-1\}.$$

Now, fix a $j \in [k \dots -k]$ and consider the longest prefix $S' = S[0..r]$ with $\delta_E(S', Q^\infty[0..r+j]) = k'$, for some $0 < k' \leq k$. By definition of the edit distance, there is an integer r' that satisfies $S'(r'..r) = Q^\infty(r'+j..r+j)$ and at least one of the following

- $k' = \delta_E(S'[0..r'], Q^\infty[0..r'+j]) + 1$ and $S'[r'] \neq Q^\infty[r']$ (when changing the character $S'[r']$ to the character $Q^\infty[r'+j]$);
- $k' = \delta_E(S'[0..r'], Q^\infty[0..r'+j-1]) + 1$ (when inserting the character $Q^\infty[r'+j]$); or
- $k' = \delta_E(S'[0..r'], Q^\infty[0..r'+j]) + 1$ (when deleting the character $Q^\infty[r'+j]$).

Note that (as $S'(r'..r) = Q^\infty(r'+j..r+j)$) we may assume that r' is maximal (that is there is no larger integer with the same properties as r'). In particular, we have

$$r' = \max(\ell_j^{(k-1)} + 1, \ell_{j-1}^{(k-1)} + 1, \ell_{j+1}^{(k-1)}),$$

and hence **Next** computes $\ell_j^{(k)}$ indeed correctly.

Using the computed values $\ell^{(k)}$, we can easily compute the length $|S'|$ of the longest prefix S' of S and the length $|Q'|$ of the corresponding prefix Q' of Q^∞ such that $\delta_E(S', Q') \leq k$: For $|S'|$, we have

$$|S'| = \max_{j,r} \{r \mid \delta_E(S[0..r], Q^\infty[0..r+y]) \leq k\} = \max_j \ell_j^{(k)},$$

as k edits only allow for up to k insertions or deletions (that is, operations that can change the shift between Q^∞ and S). Hence, the computation of $|S'|$ in **Next** is correct. For $|Q'|$, observe that if $|S'| = \ell_j^{(k)}$ for some $j \in [-k \dots k]$, by construction, we have $|Q'| = \ell_j^{(k)} + j$. Hence, also the computed value for $|Q'|$ is indeed correct.

For the correctness of **Alignment**, observe that we store information computed in **Next** (which is correct); further we always synchronize the information stored, hence also **Alignment** is correct.

For the running time, observe that in the k -th call to **Next** we call **LCP** three times for each of the $2k+3$ values $\ell_i^{(k)}$. Further, all other operations are essentially book-keeping that can be implemented in $\mathcal{O}(1)$ time. Hence in total, the k -th call to **Next** takes $\mathcal{O}(k)$ time in the **PILLAR** model.

For the **Alignment** operation, observe that we traverse the sequence $\ell^{(k)}$ exactly once, hence **Alignment** uses $\mathcal{O}(k)$ time in the **PILLAR** model, completing the proof. \blacksquare

▮ **Lemma 6.2** (Verify(P, T, k, I), [17, Section 5]). *Let P denote a string of length m , let T denote a string, and let $k \leq m$ and denote a positive integer. Further, let I denote an interval of positive integers. Using $\mathcal{O}(k(k + |I|))$ PILLAR operations, we can compute $\{(\ell, \min_r \delta_E(P, T[\ell..r])) \mid \ell \in \text{Occ}_k^E(P, T) \cap I\}$.*

Proof. Observe that the algorithm in [17, Section 5] only uses LCP operations, as it mainly implements [35]. In particular, the algorithm in [17, Section 5] uses $\mathcal{O}(k(k + |I|))$ LCP operations. The claim follows. ▮

Note that we can also call Verify(P, T, k, I) for strings $T = Q^\infty$ (for some primitive Q), as, by Corollary 2.9, we can also efficiently compute LCP(P, Q^∞).

6.2 Computing Structure in the Pattern

We proceed to discuss the implementation of Lemma 5.19, that is, the analysis of the pattern. While the algorithm itself is similar to the Hamming distance case, the analysis requires more involved arguments. We start with an auxiliary combinatorial lemma:

▮ **Lemma 6.3.** *Let S denote a string, let k denote a positive integer, and let Q denote a primitive string such that $|Q| = 1$ or $|S| \geq (2k + 1)|Q|$. Suppose that $\delta_E(S, *Q^*) = \delta_E(S, Q^\infty[x..y]) \leq k$ for integers $x \leq y$. Then, for any string S' , if $\delta_E(SS', *Q^*) \leq k$, then $\delta_E(SS', *Q^*) = \delta_E(SS', \text{rot}^{-x}(Q)^*)$, and if $\delta_E(S'S, *Q^*) \leq k$, then $\delta_E(S'S, *Q^*) = \delta_E(S'S, *\text{rot}^{-y}(Q))$.*

Proof. Suppose that $\delta_E(SS', *Q^*) = \delta_E(SS', Q^\infty[x'..z']) \leq k$ for some integers $x' \leq z'$. Then, there is a position $y' \in [x'..z']$ such that

$$\delta_E(SS', *Q^*) = \delta_E(S, Q^\infty[x'..y']) + \delta_E(S', Q^\infty[y'..z']) \leq k.$$

Due to $\delta_E(S, Q^\infty[x..y]) \leq k$ and $\delta_E(S, Q^\infty[x'..y']) \leq k$, we may apply Lemma 5.3, which yields a decomposition $S = S_L \cdot S_R$ and integers j, j' such that

$$\begin{aligned} \delta_E(S, Q^\infty[x..y]) &= \delta_E(S_L, Q^\infty[x..j|Q|]) + \delta_E(S_R, Q^\infty[j|Q|..y]) \quad \text{and} \\ \delta_E(S, Q^\infty[x'..y']) &= \delta_E(S_L, Q^\infty[x'..j'|Q|]) + \delta_E(S_R, Q^\infty[j'|Q|..y']). \end{aligned}$$

In particular, we have

$$\begin{aligned} \delta_E(S_L, Q^\infty[x..j|Q|]) + \delta_E(S_R, Q^\infty[j|Q|..y]) &= \delta_E(S, Q^\infty[x..y]) = \delta_E(S, *Q^*) \\ &\leq \delta_E(S, Q^\infty[x'..y + (j' - j)|Q|]) \leq \delta_E(S_L, Q^\infty[x'..j'|Q|]) + \delta_E(S_R, Q^\infty[j|Q|..y]), \end{aligned}$$

which implies $\delta_E(S_L, Q^\infty[x..j|Q|]) \leq \delta_E(S_L, Q^\infty[x'..j'|Q|])$. Consequently,

$$\begin{aligned} \delta_E(SS', \text{rot}^{-x}(Q)^*) &\leq \delta_E(SS', Q^\infty[x..z' + (j - j')|Q|]) \\ &\leq \delta_E(S_L, Q^\infty[x..j|Q|]) + \delta_E(S_R, Q^\infty[j'|Q|..y']) + \delta_E(S', Q^\infty[y'..z']) \\ &\leq \delta_E(S_L, Q^\infty[x'..j'|Q|]) + \delta_E(S_R, Q^\infty[j'|Q|..y']) + \delta_E(S', Q^\infty[y'..z']) \\ &= \delta_E(SS', *Q^*), \end{aligned}$$

which implies $\delta_E(SS', *Q^*) = \delta_E(SS', \text{rot}^{-x}(Q)^*)$.

The claim regarding $\delta_E(S'S, *Q^*)$ and $\delta_E(S'S, *\text{rot}^{-y}(Q))$ is symmetric. ▮

▮ **Lemma 6.4** (Analyze(P, k): Implementation of Lemma 5.19). *Let P denote a string of length m and let $k \leq m$ denote a positive integer. Then, there is an algorithm that computes one of the following:*

■ **Algorithm 11** Analyzing the pattern: A PILLAR model implementation of Algorithm 9.

```

1 Analyze( $P, k$ )
2    $j \leftarrow 0; r \leftarrow 1; b \leftarrow 1;$ 
3   while true do
4      $j' \leftarrow j + \lfloor m/8k \rfloor;$ 
5     if  $\text{Period}(P[j..j']) > m/128k$  then
6        $B_b \leftarrow P[j..j'];$ 
7       if  $b = 2k$  then return breaks  $B_1, \dots, B_{2k};$ 
8        $b \leftarrow b + 1; j \leftarrow j';$ 
9     else
10       $q \leftarrow \text{Period}(P[j..j']); Q_r \leftarrow P[j..j+q];$ 
11       $\text{generator } \mathbf{G} \leftarrow \text{EditGenerator}(P[j..m], Q_r);$ 
12       $\delta \leftarrow 0;$ 
13      while  $\delta < 8k/m \cdot (j' - j)$  and  $j' \leq m$  do
14         $(\pi, \pi') \leftarrow \text{Next}(\mathbf{G});$ 
15         $j' \leftarrow j + \pi + 1; \delta \leftarrow \delta + 1;$ 
16      if  $j' \leq m$  then
17         $R_r \leftarrow P[j..j'];$ 
18        if  $\sum_{i=1}^r |R_i| \geq 3/8 \cdot m$  then
19          return repetitive regions  $R_1, \dots, R_r$  with periods  $Q_1, \dots, Q_r;$ 
20         $r \leftarrow r + 1; j \leftarrow j';$ 
21      else
22         $Q \leftarrow Q_r;$ 
23         $\text{generator } \mathbf{G}' \leftarrow \text{EditGenerator}^R(P, \text{rot}^{-\pi'}(Q));$ 
24         $j'' = m; \delta \leftarrow 0;$ 
25        while  $(j'' \geq j$  or  $\delta < 8k/m \cdot (m - j''))$  and  $j'' \geq 0$  do
26           $(\pi, \_ ) \leftarrow \text{Next}(\mathbf{G}');$ 
27           $j'' \leftarrow m - \pi - 1; \delta \leftarrow \delta + 1;$ 
28        if  $j'' \geq 0$  then return repetitive region  $P[j''..m)$  with period  $Q;$ 
29        else return approximate period  $Q;$ 

```

- (a) $2k$ disjoint breaks $B_1, \dots, B_{2k} \preceq P$, each having period $\text{per}(B_i) > m/128k$ and length $|B_i| = \lfloor m/8k \rfloor$.
 - (b) Disjoint repetitive regions $R_1, \dots, R_r \preceq P$ of total length $\sum_{i=1}^r |R_i| \geq 3/8 \cdot m$ such that each region R_i satisfies $|R_i| \geq m/8k$ and is constructed along with a primitive approximate period Q_i such that $|Q_i| \leq m/128k$ and $\delta_E(R_i, *Q_i^*) = \lceil 8k/m \cdot |R_i| \rceil$.
 - (c) A primitive approximate period Q of P with $|Q| \leq m/128k$ and $\delta_E(P, *Q^*) < 8k$.
- The algorithm uses $\mathcal{O}(k^2)$ time plus $\mathcal{O}(k^2)$ PILLAR operations.

Proof. In a similar manner to the Hamming distance case, our implementation follows Algorithm 9 from the proof of Lemma 5.19.

Recall that P is processed from left to right and split into breaks and repetitive regions. In each iteration, the algorithm first considers a fragment of length $\lfloor m/8k \rfloor$. This fragment either becomes the next break (if its shortest period is long enough) or is extended to the right to a repetitive region (otherwise). Having constructed sufficiently many breaks or repetitive regions of sufficiently large total length, the algorithm stops. Processing the string P in this manner guarantees disjointness of breaks and repetitive regions. As in the proof of Lemma 5.19, a slightly different approach is needed if the algorithm encounters

the end of P while growing a repetitive region. If this happens, the region is also extended to the left. This way, the algorithm either obtains a single repetitive region (which is not necessarily disjoint with the previously created ones, so it is returned on its own) or learns that the whole string P is approximately periodic.

Next, we fill in missing details of the implementation of the previous steps in the PILLAR model. To that end, first note that the PILLAR model includes a **Period** operation of checking if the period of a string S satisfies $\text{per}(S) \leq |S|/2$; computing $\text{per}(S)$ in case of a positive answer. Since our threshold $m/128k$ satisfies $\lfloor m/128k \rfloor \leq \lfloor m/8k \rfloor / 2$, no specific work is required to obtain the period of an unprocessed fragment of $\lfloor m/8k \rfloor$ characters of P .

To compute a repetitive region starting from a fragment $P[j..j']$ with string period Q , we use a generator $\mathbf{G} = \text{EditGenerator}(P[j..m], Q)$ from Lemma 6.1: for subsequent values $\delta \geq 1$, we find the shortest prefix P'_δ of $P[j..m]$ such that $\delta_E(P'_\delta, Q^*) = \delta$, until no such prefix exists or $\delta \geq 8k/m \cdot |P'_\delta|$. This is possible because the $(\delta - 1)$ -st call to **Next**(\mathbf{G}) returns the length π of the longest prefix of $P[j..m]$ with $\delta_E(P[j..j + \pi], Q^*) < \delta$. If $\delta \geq 8k/m \cdot |P'_\delta|$, then we have identified a repetitive region P'_δ . Otherwise, we reach $\pi = m - j$ and retrieve π' such that $\delta_E(P[j..m], Q^*) = \delta_E(P[j..m], Q^\infty[0..\pi'])$ from the last call to **Next**(\mathbf{G}). In this case, we similarly use a generator $\mathbf{G}' = \text{EditGenerator}^R(P, \text{rot}^{-\pi'}(Q))$ from Lemma 6.1: For subsequent values $\delta \geq 1$, we find the shortest suffix P''_δ of P such that $\delta_E(P''_\delta, *\text{rot}^{-\pi'}(Q)) = \delta$, until no such suffix exists or $|P''_\delta| \geq |P[j..m]|$ and $\delta \geq 8k/m \cdot |P''_\delta|$. Again, this is possible because the $(\delta - 1)$ -st call to **Next**(\mathbf{G}') returns the length π of the longest suffix of P with $\delta_E(P[m - \pi..m], *\text{rot}^{-\pi'}(Q)) < \delta$. If we reach $\pi = m$, then we return Q as an approximate period of P ; otherwise, we return the final suffix P''_δ as a long repetitive region. Consider Algorithm 11 for implementation details.

For the correctness, since our algorithm follows the proof of Lemma 5.19, we only need to show that our implementation of finding repetitive regions correctly implements the corresponding steps in Algorithm 9.

First, we inductively prove that each considered prefix P'_δ of $P[j..m]$ satisfies $\delta_E(P'_\delta, *Q^*) = \delta \leq \lceil 8k/m \cdot |P'_\delta| \rceil$. The case of $\delta = 1$ is easy since $\delta_E(P'_1, *Q^*) \leq \delta_E(P'_1, Q^*) = 1$, since $\delta_E(P'_1, *Q^*) = 0$ would imply $\delta_E(P'_1, Q^*) = 0$ because Q is a prefix of P'_1 , and since $8k/m \cdot |P'_1| > 0$ due to $|P'_1| > 0$. Next, we prove that the claim holds for $\delta + 1$ assuming that it holds for δ . The inductive assumption guarantees $\delta_E(P'_\delta, *Q^*) = \delta_E(P'_\delta, Q^*) = \delta$. Since the algorithm proceeded to the next step, we have $\delta < 8k/m \cdot |P'_\delta|$. In particular, $|P'_\delta| \geq (2\delta + 1) \cdot m/128k \geq (2\delta + 1)|Q|$, so we can apply Lemma 6.3 to P'_δ . If $\delta_E(P'_{\delta+1}, *Q^*) \leq \delta$, then Lemma 6.3 yields $\delta_E(P'_{\delta+1}, Q^*) = \delta_E(P'_{\delta+1}, *Q^*) \leq \delta$, which contradicts the definition of $P'_{\delta+1}$. Hence, $\delta_E(P'_{\delta+1}, *Q^*) \geq \delta + 1$. Due to $\delta_E(P'_{\delta+1}, Q^*) = \delta + 1$, we have $\delta_E(P'_{\delta+1}, *Q^*) = \delta + 1$. Moreover, $\lceil 8k/m \cdot |P'_{\delta+1}| \rceil \geq \lceil 8k/m \cdot |P'_\delta| \rceil > \delta$ guarantees $\lceil 8k/m \cdot |P'_{\delta+1}| \rceil \geq \delta + 1$, which completes the inductive proof.

In particular, if we encounter a prefix P'_δ that satisfies $\delta \geq 8k/m \cdot |P'_\delta|$, then $\delta_E(P'_\delta, *Q^*) = \lceil 8k/m \cdot |P'_\delta| \rceil$. However, if no such prefix P'_δ exists, then $\delta_E(R, *Q^*) < 8k/m \cdot |R|$ holds for each non-empty prefix of $P[j..m]$ (because $R = P'_\delta$ is the shortest prefix R of $P[j..m]$ with $\delta_E(R, *Q^*) = \delta$). Thus, Line 9 of Algorithm 9 is implemented correctly.

In the following, we assume that no such prefix R exists. In particular, we have $\delta_E(P[j..m], *Q^*) < 8k/m \cdot |P[j..m]|$. Then, the last call to **Next**(\mathbf{G}) resulted in $(m - j, \pi')$ such that $\delta_E(P[j..m], Q^*) = \delta_E(P[j..m], Q^\infty[0..\pi'])$. Moreover, since P'_δ with $\delta = \delta_E(P[j..m], Q^*)$ satisfies $\delta_E(P'_\delta, *Q^*) = \delta_E(P'_\delta, Q^*) = \delta$, we have

$$\delta_E(P[j..m], *Q^*) = \delta_E(P[j..m], Q^*) = \delta_E(P[j..m], *\text{rot}^{-\pi'}(Q)) = \delta.$$

We inductively prove that each considered suffix of P''_δ of P with $|P''_\delta| > j - m$ satisfies $\delta_E(P''_\delta, *Q^*) = \delta \leq \lceil 8k/m \cdot |P''_\delta| \rceil$. Let us prove that this claim is true for $\delta + 1$ assuming that it is true for δ (unless $\delta = 0$, when there is no assumption). If $\delta < \delta_E(P[j..m], *Q^*)$, then $|P''_{\delta+1}| \leq j - m$ and the claim is void, so

we only consider $\delta \geq \delta_E(P[j..m], *Q^*)$. If $\delta > \delta_E(P[j..m], *Q^*) = \delta_E(P[j..m], *\text{rot}^{-\pi'}(Q))$, then $|P''_\delta| > j - m$ and the inductive assumption guarantees $\delta_E(R, *Q^*) = \delta_E(R, *\text{rot}^{-\pi'}(Q)) = \delta$ for $R = P''_\delta$. Otherwise, we have

$$\delta = \delta_E(P[j..m], *Q^*) = \delta_E(P[j..m], *\text{rot}^{-\pi'}(Q)),$$

in which case $\delta_E(R, *Q^*) = \delta_E(R, *\text{rot}^{-\pi'}(Q)) = \delta$ holds for $R = P[j..m]$. In either case, we also have $\delta < 8k/m \cdot |R|$, and hence $|R| \geq (2\delta + 1) \cdot m / 128k \geq (2\delta + 1)|Q|$. Therefore, we can apply Lemma 6.3 to R . If $\delta_E(P''_{\delta+1}, *Q^*) \leq \delta$, then Lemma 6.3 yields $\delta_E(P''_{\delta+1}, *\text{rot}^{-\pi'}(Q)) = \delta_E(P''_{\delta+1}, *Q^*) \leq \delta$, which contradicts the definition of $P''_{\delta+1}$. Hence, $\delta_E(P''_{\delta+1}, *Q^*) \geq \delta + 1$. Due to $\delta_E(P''_{\delta+1}, *\text{rot}^{-\pi'}(Q)) = \delta + 1$, we have $\delta_E(P''_{\delta+1}, *Q^*) = \delta + 1$. Moreover, $\lceil 8k/m \cdot |P''_{\delta+1}| \rceil \geq \lceil 8k/m \cdot |R| \rceil > \delta$ guarantees $\lceil 8k/m \cdot |P''_{\delta+1}| \rceil \geq \delta + 1$, which completes the inductive proof.

In particular, if we encounter a suffix P''_δ that satisfies $|P''_\delta| > m - j$ and $\delta \geq 8k/m \cdot |P''_\delta|$, then $\delta_E(P''_\delta, *Q^*) = \lceil 8k/m \cdot |P''_\delta| \rceil$. On the other hand, if no such suffix P''_δ exists, then $\delta_E(R, *Q^*) < 8k/m \cdot |R|$ holds for each suffix R of P of length $|R| > m - j$ (because $R = P'_\delta$ is the shortest suffix R of P with $\delta_E(R, *Q^*) = \delta$ assuming $\delta > \delta_E(P[j..m], *Q^*)$). Thus, Line 15 of Algorithm 9 is also implemented correctly.

For the running time analysis, observe that each iteration of the outer while loop processes at least $\lfloor m/8k \rfloor$ characters of P , so there are at most $\mathcal{O}(k)$ iterations of the outer while loop. In each iteration, we perform one **Period** operation, a constant number of **Access** operations, and at most $8k/m \cdot (j' - j)$ calls to the generator **EditGenerator**. These calls use $\mathcal{O}((8k/m \cdot (j' - j))^2)$ time in the **PILLAR** model, which is $\mathcal{O}(k^2)$ in total across all iterations (since the function $x \mapsto x^2$ is convex). Similarly, we bound the running time of the calls to the generator **EditGenerator^R**: As we find at most $8k/m \cdot m = 8k$ errors, **EditGenerator^R** uses at most $\mathcal{O}(k^2)$ time. Overall, Algorithm 11 thus uses $\mathcal{O}(k^2)$ time in the **PILLAR** model. \blacksquare

6.3 Computing Occurrences in the Periodic Case

We start this subsection with a subroutine to compute a *witness* that a string S has a small edit distance to a string Q^∞ .

\blacksquare **Lemma 6.5** (**FindAWitness**(k, Q, S)). *Let k denote a positive integer, let S denote a string, and let Q denote a primitive string that satisfies $|S| \geq (2k + 1)|Q|$ or $|Q| \leq 3k + 1$.*

*Then, we can compute a witness $Q^\infty[x..y]$ such that $\delta_E(S, Q^\infty[x..y]) = \delta_E(S, *Q^*) \leq k$, or report that $\delta_E(S, *Q^*) > k$. The algorithm takes $\mathcal{O}(k^2)$ time in the **PILLAR** model.*

Proof. For a set $A \subseteq \mathbb{Z}$ and an integer $p > 0$, we define $A \bmod p := \{a \bmod p \mid a \in A\}$. We first compute a (short) interval J such that $\text{Occ}_k^E(S, Q^\infty) \bmod |Q| \subseteq J \bmod |Q|$. If $|Q| \leq 3k + 1$, then we simply set $J = [0..|Q|)$. Otherwise, we proceed similarly as in the Hamming distance setting (Lemma 4.5), where we computed a majority string of the first $2k + 1$ length- $|Q|$ subsequent fragments S_1, \dots, S_{2k} of S . However, now we need to accommodate for insertions and deletions of a k -error occurrence. Hence, we first compute an auxiliary set I defined as the union of intervals $[p..p + k]$ such that for at least $k + 1$ fragments S_i of S , we have $Q = \text{rot}^j(S_i)$ for $j \in [p..p + k]$. Finally, we set $J \subseteq [0..2|Q|)$ to be a shortest interval satisfying $I \bmod |Q| \subseteq J \bmod |Q|$. Here, $J \bmod |Q|$ can be interpreted as a shortest cyclic interval (modulo $|Q|$) containing $I \bmod |Q|$.

Having computed the set J , we use **Verify** from Lemma 6.2 to determine at which starting position x in J we have an occurrence with the fewest number of errors (or to report that the number of errors is greater than k everywhere). Finally, we use an **EditGenerator** from Lemma 6.1 to compute the ending

■ **Algorithm 12** Finding a witness $Q^\infty[x..y]$ for $\delta_E(S, *Q^*) \leq k$.

```

1 FindAWitness( $k, Q, S$ )
  // Compute "correct" rotation(s) of  $Q$ 
2 if  $|Q| \leq 3k + 1$  then  $J \leftarrow [0..|Q|]$ ;
3 else
4   multi-set  $R \leftarrow \{\}$ ;
5   for  $i \leftarrow 0$  to  $2k$  do
6      $R \leftarrow R \cup \text{Rotations}(S[i|Q|..(i+1)|Q|], Q)$ ;
7    $I \leftarrow \bigcup \{[p..p+k] : p \in \mathbb{Z} \text{ and } [p..p+k] \text{ contains at least } k+1 \text{ elements of } R\}$ ;
8   Let  $J \subseteq [0..2|Q|]$  denote a shortest interval that satisfies  $I \bmod |Q| \subseteq J \bmod |Q|$ ;

  // Compute the start position of the witness
9    $Occ \leftarrow \text{Verify}(S, Q^\infty, k, J)$ ;
10  if  $Occ = \emptyset$  then return  $\perp$ ;
11  Let  $(x, k') \in Occ$  be an arbitrary element minimizing  $k'$ ;

  // Compute the end position of the witness
12  generator  $\mathbf{G} \leftarrow \text{EditGenerator}(S, \text{rot}^{-x}(Q))$ ;
13  for  $i \leftarrow 1$  to  $k'$  do Next( $\mathbf{G}$ );
14   $(\lambda, \lambda') \leftarrow \text{Next}(\mathbf{G})$ ;
15  return  $Q^\infty[x..x + \lambda']$ ;
```

position y of the occurrence of S as a prefix of $Q^\infty[x..y]$). Consider Algorithm 12 for a pseudo-code implementation.

The correctness is based on the aforementioned characterization of J :

□ **Claim 6.6.** The interval J satisfies $\text{Occ}_k^E(S, Q^\infty) \bmod |Q| \subseteq J \bmod |Q|$.

Proof. The claim trivially holds if $|Q| \leq 3k + 1$, so we assume that $|Q| > 3k + 1$. For every $i \in [0..2k]$, define $S_i := S[i|Q|..(i+1)|Q|]$. Consider an optimum alignment between S and its k -error occurrence $Q^\infty[x..y]$, and let $Q_i = Q^\infty[x_i..x_{i+1}]$ denote the fragment aligned to S_i . Consider the multi-set $R := \bigcup_i \text{Rotations}(S_i, Q)$. Next, consider the values $\delta_i := x_i - i|Q|$ for $i \in [0..2k]$. We have $\delta_0 = x$, and $\delta_{i+1} = \delta_i + |Q_i| - |S_i|$ for $i > 0$. Since

$$\sum_{i=0}^{2k} \left| |Q_i| - |S_i| \right| \leq \sum_{i=0}^{2k} \delta_E(Q_i, S_i) \leq k,$$

all values δ_i belong to an interval of the form $[p..p+k]$ for some integer p . Moreover, note that $Q_i = S_i$ holds for at least $k+1$ fragments Q_i ; these fragments satisfy $Q = \text{rot}^{\delta_i}(S_i)$ and thus contribute δ_i to R . We conclude that there is an interval $[p..p+k]$ containing x and at least $k+1$ elements of R . Consequently, we have $x \in I$. By definition of J , this yields $x \bmod |Q| \in J \bmod |Q|$. □

Now, let $Q^\infty[x..y]$ denote a witness that satisfies $\delta_E(S, *Q^*) = \delta_E(S, Q^\infty[x..y])$. By Claim 6.6, there is a matching fragment $Q^\infty[x'..y'] = Q^\infty[x..y]$ starting at $x' \in J$. Thus, we may assume without loss of generality that $x \in J$. As we verify all possible starting positions in J using **Verify** from Lemma 6.2, we correctly compute the starting position x of a witness occurrence of S in Q^∞ . Further, as we use an **EditGenerator** from Lemma 6.1, we also compute the corresponding ending position correctly.

As for the running time, we prove the following characterization of J :

□ **Claim 6.7.** The interval J satisfies $|J| \leq 3k + 1$.

Proof. The claim trivially holds if $|Q| \leq 3k + 1$, so we assume that $|Q| > 3k + 1$. Recall that the multiset R in Algorithm 12 is the union of $2k + 1$ sets $\{j \in \mathbb{Z} : \text{rot}^j(S_i) = Q\}$. As the string Q is primitive, R is the union of at most $2k + 1$ infinite arithmetic progressions with difference $|Q|$. In particular, if $[p \dots p+k]$ and $[p' \dots p'+k]$ contain at least $k+1$ elements of R each, then $([p \dots p+k] \bmod |Q|) \cap ([p' \dots p'+k] \bmod |Q|) \neq \emptyset$, and thus $[p' \dots p' + k] \bmod |Q| \subseteq [p - k \dots p + 2k] \bmod |Q|$. Since J is the union of such intervals $[p' \dots p' + k]$, we have $J \bmod |Q| \subseteq [p - k \dots p + 2k] \bmod |Q|$. By definition of I , we conclude that $|I| \leq |[p - k \dots p + 2k]| = 3k + 1$. □

Now, observe that computing the multiset R (represented as the union of infinite arithmetic progressions modulo $|Q|$) takes $\mathcal{O}(k)$ time in the PILLAR model; computing the sets I and J can be done in $\mathcal{O}(k \log \log k)$ time by sorting R (restricted to $[0 \dots |Q|)$) and a subsequent cyclic scan over R . Further, by Claim 6.7, we call **Verify** on an interval of length $\mathcal{O}(k)$; hence the call to **Verify** takes $\mathcal{O}(k^2)$ time in the PILLAR model. Finally, as we query the **EditGenerator** for up to k errors, the last step of the algorithm takes $\mathcal{O}(k^2)$ time in the PILLAR model as well. Hence in total, **FindAWitness** runs in $\mathcal{O}(k^2)$ time in the PILLAR model, completing the proof. ■

□ **Lemma 6.8** (**FindRelevantFragment**(P, T, k, d, Q)). *Let P denote a pattern of length m , let T denote a text of length n , and let $0 \leq k \leq m$ denote a threshold such that $n < \frac{3}{2}m + k$. Further, let $d \geq 2k$ denote a positive integer and let Q denote a primitive string that satisfies $|Q| \leq m/8d$ and $\delta_E(P, *Q^*) \leq d$.*

*Then, there is an algorithm that computes a fragment $T' = T[\ell \dots r]$ and an integer range I such that $\delta_E(T', *Q^*) \leq 3d$, $|\text{Occ}_k^E(P, T)| = |\text{Occ}_k^E(P, T')|$, $|I| \leq 6d + 1$, and $\text{Occ}_k^E(P, T') \bmod |Q| \subseteq I \bmod |Q|$. The algorithm runs in $\mathcal{O}(d^2)$ time in the PILLAR model,*

Proof. We start with two calls to **FindAWitness** from Lemma 4.5 in order to find a fragment $Q^\infty[x \dots y]$ such that $\delta_E(P, Q^\infty[x \dots y]) = \delta_E(P, *Q^*) \leq d$ and a fragment $Q^\infty[x' \dots y']$ such that

$$\delta_E(T[n - m + k \dots m - k], Q^\infty[x' \dots y']) = \delta_E(T[n - m + k \dots m - k], *Q^*) \leq \frac{3}{2}d.$$

If the latter fragment does not exist, then we return the empty string $T' = \varepsilon$ and the empty interval $I = \emptyset$. Otherwise, we proceed by computing the rightmost position r such that $\delta_E(T[n - m + k \dots r], \text{rot}^{-x'}(Q)^*) \leq \frac{3}{2}d$ and the leftmost position ℓ such that $\delta_E(T[\ell \dots m - k], * \text{rot}^{-y'}(Q)) \leq \frac{3}{2}d$. That is, we “extend” the fragment found in the text as much as possible. Afterwards, we return the fragment $T' = T[\ell \dots r]$ and the interval $I = [n - m + k - \ell + x - x' - 3d \dots n - m + k - \ell + x - x' + 3d]$. Consider Algorithm 13 for implementation details.

For the correctness, note that the due to the assumption on Q , the first call to **FindAWitness** is valid and returns a witness $Q^\infty[x \dots y]$ (rather than \perp). Next, consider a k -error occurrence $T[p \dots q]$ of P . By triangle inequality (Fact 2.3), we have

$$\delta_E(T[p \dots q], Q^\infty[x \dots y]) \leq k + \delta_E(P, Q^\infty[x \dots y]) \leq \frac{3}{2}d.$$

Due to $|T[p \dots q]| \geq m - k$, we have $q \geq m - k$ and $p \leq n - m + k$, which yields

$$\delta_E(T[n - m + k \dots m - k], Q^\infty[x'' \dots y'']) \leq \frac{3}{2}d$$

for some integers x'', y'' with $x \leq x'' \leq y'' \leq y$. Moreover, as in the proof of Theorem 5.2, we have $|T[n - m + k \dots m - k]| = 2(m - k) - n \geq (3d + 1)|Q|$ or $|Q| = 1$, so the second call to **FindAWitness** is valid. Thus, if the call returns \perp , then $\text{Occ}_k^E(P, T) = \emptyset$.

We henceforth assume that the call returned a witness $Q^\infty[x'..y']$. Next, we apply Lemma 6.3 for $S = T[n - m + k..m - k]$. This is indeed possible because $|S| \geq (3d + 1)|Q|$ or $|Q| = 1$. Due to

$$\delta_E(T[n - m + k..q], *Q^*) \leq \delta_E(T[n - m + k..q], Q^\infty[x''..y'']) \leq \frac{3}{2}d,$$

Lemma 6.3 yields

$$\delta_E(T[n - m + k..q], \text{rot}^{-x'}(Q)^*) = \delta_E(T[n - m + k..q], *Q^*) \leq \frac{3}{2}d.$$

Consequently, we have $q \leq r$, because r is computed correctly using `EditGenerator` from Lemma 6.1. Symmetrically, due to

$$\delta_E(T[p..m + k], *Q^*) \leq \delta_E(T[p..m + k], Q^\infty[x..y'']) \leq \frac{3}{2}d,$$

Lemma 6.3 yields

$$\delta_E(T[p..m - k], *\text{rot}^{-y'}(Q)) = \delta_E(T[p..m + k], *Q^*) \leq \frac{3}{2}d.$$

Consequently, we have $p \geq \ell$, because ℓ is computed correctly using `EditGeneratorR` from Lemma 6.1. We conclude that $T[p..q]$ is contained in $T' = T[\ell..r]$. Since $T[p..q]$ was an arbitrary k -error occurrence of P in T , this implies $|\text{Occ}_k^E(P, T')| = |\text{Occ}_k^E(P, T)|$.

Now consider the fragment $T[p..m - k]$ whose prefix $T[p..m - k]$ satisfies

$$\delta_E(T[p..m - k], Q^\infty[x..y'']) \leq \frac{3}{2}d$$

and whose suffix $T[n - m + k..m - k]$ satisfies

$$\delta_E(T[n - m + k..m - k], Q^\infty[x'..y']) \leq \frac{3}{2}d.$$

We apply Lemma 5.4 to $T[p..m - k]$; this is indeed possible because $|T[n - m + k..m - k]| \geq (3d + 1)|Q|$ or $|Q| = 1$. Lemma 5.4 now implies $(n - m + k - p + x - x' + 3d) \bmod |Q| \leq 6d$. Consequently, we have

$$p \bmod |Q| \in [n - m + k + x - x' - 3d..n - m + k + x - x' + 3d] \bmod |Q|.$$

Since $T[p..q]$ was an arbitrary k -error occurrence of P in T , this implies

$$\text{Occ}_k^E(P, T') \bmod |Q| \subseteq I \bmod |Q|.$$

Moreover, $|I| \leq 6d + 1$ holds trivially by construction.

Now consider the fragment T , whose prefix $T[\ell..m - k]$ satisfies

$$\delta_E(T[\ell..m - k], *Q^*) \leq \frac{3}{2}d$$

and whose suffix $T[n - m + k..r]$ satisfies

$$\delta_E(T[n - m + k..r], *Q^*) \leq \frac{3}{2}d.$$

Again, we apply Lemma 5.4 to T' ; this is indeed possible because $|T[n - m + k..m - k]| \geq (3d + 1)|Q|$ or $|Q| = 1$. Lemma 5.4 now implies $\delta_E(T', *Q^*) \leq 3d$, as claimed.

As for the running time in the `PILLAR` model, the calls to `FindAWitness` use $\mathcal{O}(d^2)$ time; the same is true for the usage of `EditGenerator` and `EditGeneratorR`. Thus, the algorithm takes $\mathcal{O}(d^2)$ time in the `PILLAR` model. \blacksquare

■ **Algorithm 13** A PILLAR algorithm computing a *relevant* fragment T' of T containing all k -error occurrences of P in T , and an interval I such that $\text{Occ}_k^E(P, T') \bmod |Q| \subseteq I \bmod |Q|$.

```

1 FindRelevantFragment( $P, T, k, d, Q$ )
2    $Q^\infty[x..y] \leftarrow \text{FindAWitness}(d, Q, P)$ ;
3    $Q' \leftarrow \text{FindAWitness}(\lceil \frac{3}{2}d \rceil, Q, T[n-m+k..m-k])$ ;
4   if  $Q' = \perp$  then return  $(\varepsilon, \emptyset)$ ;
5    $Q^\infty[x'..y'] \leftarrow Q'$ ;
   // Extend  $Q'$  as much as possible to the right.
6   generator  $\mathbf{G} \leftarrow \text{EditGenerator}(T[n-m+k..n], \text{rot}^{-x'}(Q))$ ;
7   for  $i \leftarrow 0$  to  $\lceil \frac{3}{2}d \rceil$  do  $(\lambda, \_ ) \leftarrow \text{Next}(\mathbf{G})$ ;
8    $r \leftarrow n - m + k + \lambda$ ;
   // Extend  $Q'$  as much as possible to the left.
9   generator  $\mathbf{G}' \leftarrow \text{EditGenerator}^R(T[0..m-k], \text{rot}^{-y'}(Q))$ ;
10  for  $i \leftarrow 0$  to  $\lceil \frac{3}{2}d \rceil$  do  $(\lambda', \_ ) \leftarrow \text{Next}(\mathbf{G}')$ ;
11   $\ell \leftarrow m - k - \lambda'$ ;
12  return  $(T[\ell..r], [n-m+k-\ell+x-x'-3d..n-m+k-\ell+x-x'+3d])$ ;
```

■ **Lemma 6.9** ($\text{Locked}(S, Q, d, k)$: Implementation of Lemma 5.11). *Let S denote a string, let Q denote a primitive string, let d denote a positive integer such that $\delta_E(S, *Q^*) \leq d$ and $|S| \geq (2d+1)|Q|$, and let k denote a non-negative integer.*

*Then, there is an algorithm that computes disjoint locked fragments $L_1, \dots, L_\ell \preceq S$ such that L_1 is a k -locked prefix of S , L_ℓ is a suffix of S , and $\delta_E(L_i, *Q^*) > 0$ for $1 < i < \ell$. Moreover, we have*

$$\delta_E(S, *Q^*) = \sum_{i=1}^{\ell} \delta_E(L_i, *Q^*) \quad \text{and} \quad \sum_{i=1}^{\ell} |L_i| \leq (5|Q| + 1)d + 2(k+1)|Q|.$$

The algorithm takes $\mathcal{O}(d^2 + k)$ time in the PILLAR model.

Proof. We implement the process from in the proof of Lemma 5.11. Our algorithm is described below; see also Algorithm 14 for implementation details.

First, we construct an optimal alignment between S and a substring of Q^∞ . For this, we first use FindAWitness of Lemma 6.5 to obtain positions $x \leq y$ such that $\delta_E(S, *Q^*) = \delta_E(S, Q^\infty[x..y])$. Then, we apply a generator $\text{EditGenerator}(S, \text{rot}^{-x}(Q))$ of Lemma 6.1 to construct an optimal alignment A between S and $Q^\infty[x..x + \pi']$ for some integer $\pi' \geq 0$ (note that we cannot guarantee $y = x + \pi'$).

Then, based on the alignment A , we construct a decomposition $S = S_0^{(0)} \dots S_{s_0}^{(0)}$ such that $S_i^{(0)}$ is aligned against

$$Q_i^{(0)} := Q^\infty[\max(x, (|Q| - 1)\lceil x/|Q| \rceil) .. \min(|Q|\lceil x/|Q| \rceil, x + \pi)]$$

in the decomposition A , and a sequence $\Delta_i^{(0)}$ such that we have $\Delta_i^{(0)} = \delta_E(S_i^{(0)}, Q_i^{(0)})$ for $i > 0$ and $\Delta_0^{(0)} = \delta_E(S_0^{(0)}, Q_0^{(0)}) + k + 1$. Since this sequence might be long, we only generate *interesting* fragments $S_i^{(0)}$ and store them, along with the values $\Delta_i^{(0)}$, in a queue F in left-to-right order. (Recall that $S_i^{(t)}$ is interesting if $i = 0$, $i = s^{(t)}$, $S_i^{(t)} \neq Q$, or $\Delta_i^{(t)} > 0$.)

The process of constructing the interesting fragments $S_i^{(0)}$ is somewhat tedious. We maintain a fragment $Q^\infty[\ell_Q .. r_Q]$, interpreted as $Q_i^{(0)}$ for increasing values of i , a fragment $S[\ell_S .. r_S]$, interpreted

■ **Algorithm 14** Computing locked fragments in a string S .

```

1 Locked( $S, Q, d, k$ )
2    $Q^\infty[x..y] \leftarrow \text{FindAWitness}(d, Q, S)$ ;
3   generator  $\mathbf{G} \leftarrow \text{EditGenerator}(S, \text{rot}^{-x}(Q))$ ;
4   do  $(\pi, \pi') \leftarrow \text{Next}(\mathbf{G})$  while  $\pi < |S|$ ;
5    $A \leftarrow \text{Alignment}(\mathbf{G})$ ;
6    $\ell_Q \leftarrow x; \ell_S \leftarrow 0$ ;
7    $r_Q \leftarrow |Q| \lceil x/|Q| \rceil; r_S \leftarrow r_Q - \ell_Q$ ;
8    $\Delta \leftarrow k + 1$ ;
9   queue  $F$ ;
10  foreach  $(s, q) \in A \cup (\pi, \pi')$  do
11    if  $s = \perp$  then  $s \leftarrow q + x + r_S - r_Q - 1$ ;
12    if  $q = \perp$  then  $q \leftarrow s - x + r_Q - r_S - 1$ ;
13    if  $x + q \geq r_Q$  then
14      push( $F, (S[\ell_S..r_S], \Delta)$ );
15       $\ell_Q \leftarrow |Q| \lceil (x + q)/|Q| \rceil; \ell_S \leftarrow r_S + \ell_Q - r_Q$ ;
16       $r_Q \leftarrow \ell_Q + |Q|; r_S \leftarrow r_S + |Q|$ ;
17       $\Delta \leftarrow 0$ ;
18       $r_S \leftarrow r_Q - x + s - q$ ;
19      if  $(s, q) \neq (\pi, \pi')$  then  $\Delta \leftarrow \Delta + 1$ ;
20  push( $F, (S[\ell_S..|S|], \Delta)$ );
21  stack  $L$ ;
22  while  $F$  is not empty do
23     $(S[\ell..r], \Delta) \leftarrow \text{front}(F)$ ; pop( $F$ );
24    while true do
25      if top( $L$ ) =  $S[\ell'..r']$  and  $r' = \ell$  then
26         $\ell \leftarrow \ell'$ ;
27        pop( $L$ );
28      else if front( $F$ ) =  $(S[\ell'..r'], \Delta')$  and  $\ell' = r$  then
29         $r \leftarrow r'$ ;
30         $\Delta \leftarrow \Delta + \Delta'$ ;
31        pop( $F$ );
32      else if  $\Delta > 0$  then
33         $\ell \leftarrow \max(0, \ell - |Q|)$ ;
34         $r \leftarrow \min(|S|, r + |Q|)$ ;
35         $\Delta \leftarrow \Delta - 1$ ;
36      else
37        push( $L, S[\ell..r]$ );
38        break;
39  return  $L$ ;
```

as a candidate for $S_i^{(0)}$, and an integer Δ , interpreted as $\Delta_i^{(0)}$. They are initialized to $Q^\infty[x..|Q|\lceil x/|Q|\rceil]$, $S[0..|Q|\lceil x/|Q|\rceil - x]$, and $k + 1$, respectively.

Next, we process pairs (s, q) corresponding to subsequent errors in the alignment A . The interpretation of the j -th pair (s, q) is that $S[0..s]$ is aligned with $Q^\infty[x..x + q]$ with j errors so that the j -th error is a substitution of $S[s]$ into $Q^\infty[x + q]$, and insertion of $Q^\infty[x + q]$, or a deletion of $S[s]$.

The first step of processing (s, q) is only performed if $Q^\infty[x..x + q]$ is not (yet) contained in $Q^\infty[\ell_Q..r_Q]$. If this is not the case, then we push $S[\ell_S..r_S]$ with budget Δ to the queue F of interesting fragments, and we update the maintained data: The fragment $Q^\infty[\ell_Q..r_Q]$ is set to be the fragment of Q^∞ matching Q and containing $Q^\infty[x + q]$; between the previous and the current value of $Q^\infty[\ell_Q..r_Q]$, there are zero or more copies of Q aligned in A without error. Hence, we skip the same number of copies of Q in S (these are the uninteresting fragments $S_i^{(0)}$) and set $S[\ell_S..r_S]$ to be the subsequent fragment of length $|Q|$. Finally, the budget Δ is reset to 0.

In the second step, we update r_S according to the type of the currently processed error: We increment r_S in case of deletion of $S[s]$ and we decrement r_S in case of insertion of $Q^\infty[x + q]$. This way, we guarantee that $|S[s..r_S]| = |Q^\infty[x + q..r_Q]|$, and that A aligns $S[\ell_S..r_S]$ with $Q^\infty[\ell_Q..r_Q]$ provided that we have already processed all errors involving $Q^\infty[\ell_Q..r_Q]$. Additionally, we increase Δ to acknowledge the currently processed error between $S[\ell_S..r_S]$ and $Q^\infty[\ell_Q..r_Q]$.

In a similar way, we process $(s, q) = (|S|, \pi')$, interpreting it as extra substitution. This time, however, we do not increase Δ (because this is not a real error). Finally, we push $S[\ell_S..|S|] = S_{s^{(0)}}^{(0)}$ with budget Δ to the queue F .

In the second phase of the algorithm, we transform the decomposition $S = S_0^{(0)} \dots S_{s^{(0)}}^{(0)}$ and the sequence $\Delta_0^{(0)} \dots \Delta_{s^{(0)}}^{(0)}$ using the four types of merge operations described in the proof of Lemma 5.6.

We maintain an invariant that a stack L contains already processed interesting fragments, all with budget equal to 0, in left-to-right order (so that $\text{top}(L)$ represents the rightmost one), while F contains fragments that have not been processed yet (and may have positive budgets) also in the left-to-right order (so that $\text{front}(F)$ represents the leftmost one). Additionally, the currently processed fragment $S[\ell..r]$ is guaranteed to be to the right of all fragments in L and to the left of all fragments in F . The fragments in L , the fragment $S[\ell..r]$, and the fragments in F form the sequence of all interesting fragments in the current decomposition $S = S_0^{(t)} \dots S_{s^{(t)}}^{(t)}$.

In each iteration of the main loop, we pop the front fragment $S[\ell..r]$ with budget Δ from the queue F and exhaustively perform merge operations involving it: We first try applying a type-1 merge with the fragment to the left (which must be $\text{top}(L)$). If this is not possible, we type applying a type-1 merge with the fragment to the right (which must be $\text{front}(F)$). If also this is not possible, then $S[\ell..r]$ is surrounded by uninteresting fragments. In this case, we perform a type-2, type-3, or 4 merge provided that $\Delta > 0$. Otherwise, we push $S[\ell..r]$ to L and proceed to the next iteration.

Finally, the algorithm returns the sequence of (locked) fragments represented in the stack L .

The correctness of the algorithm follows from Lemma 5.11; no deep insight is needed to prove that our implementation indeed follows the procedure described in the proof of Lemma 5.6 and extended in the proof of Lemma 5.11.

For the running time, the initial call to **FindAWitness** and applying the generator **G** each take $\mathcal{O}(d^2)$ time in the **PILLAR** model. As the alignment A is of size $|A| \leq d$, the **for** loop in Line 10 takes $\mathcal{O}(d)$ time and generates $\mathcal{O}(d)$ interesting locked fragments with total budget $\mathcal{O}(d + k)$. Each iteration of the **while** loop in Line 24 decreases the number of interesting locked fragments or their total budget, so there are at most $\mathcal{O}(d + k)$ iterations in total. Overall the algorithm runs in $\mathcal{O}(d^2 + k)$ time in the **PILLAR** model. \blacksquare

\blacksquare **Lemma 6.10** (**SynchedMatches**(P, T, I, d, d', k, Q)). *Let P denote a pattern of length m , let*

$0 \leq k \leq m$ denote a threshold, and let T denote a text of length $n \leq \frac{3}{2}m + k$. Further, let I denote an integer range and let Q denote a primitive string that satisfies $\delta_E(P, *Q^*) \leq d$ and $\delta_E(T, *Q^*) \leq d'$.

There is an algorithm that computes the set $\text{Occ}_k^E(P, T) \cap (I + |Q|\mathbb{Z})$ as $\mathcal{O}(|I|d'(d+k))$ arithmetic progressions. The algorithm takes $\mathcal{O}(kd'(d+k)(k + |I| + d + d'))$ time in the PILLAR model.

Proof. The algorithm resembles the proof of Theorem 5.2(d). Consult Algorithm 15 for a visualization of the algorithm as pseudo-code; in the interest of readability we use $\text{Occ}_k^E(P, T)$ instead of $\text{Occ}_k^E(P, T) \cap (I + |Q|\mathbb{Z})$ in the pseudo-code. Note that if $|I| > |Q|$ we can replace I with $[0..Q]$, since in this case $I + |Q|\mathbb{Z} = \mathbb{Z} = [0..Q] + |Q|\mathbb{Z}$. Hence, we can assume that $|I| \leq |Q|$.

We first compute $\mathcal{L}^P := \text{Locked}(P, Q, d, k)$ and $\mathcal{L}^T := \text{Locked}(T, Q, d', 0)$ using Algorithm 14. We have $\ell^P := |\mathcal{L}^P| \leq \delta_E(P, *Q^*) + 2 \leq d + 2$ and $\ell^T = |\mathcal{L}^T| \leq \delta_E(T, *Q^*) \leq d' + 2$.

Then, for each of the $\mathcal{O}(dd')$ pairs of locked fragments $L_i^P = P[\ell..r] \in \mathcal{L}^P$ and $L_j^T = T[\ell'..r'] \in \mathcal{L}^T$ we (implicitly) mark the positions in the interval $[\ell' - r - k..r' - \ell + k]$. We also mark all positions in $[n - m - k..n - m + k]$. We decompose the set of marked positions M into $\mathcal{O}(dd')$ maximal ranges $J \subseteq M$. For each such maximal range J , for each maximal range $J' \subseteq J \cap (I + |Q|\mathbb{Z})$, we call $\text{Verify}(P, T, k, J')$ and add its output to $\text{Occ}_k^E(P, T) \cap (I + |Q|\mathbb{Z})$. This guarantees that we correctly compute all elements of the set $\text{Occ}_k^E(P, T) \cap (I + |Q|\mathbb{Z}) \cap M$.

The decomposition of M into maximal ranges yields a decomposition of $[0..n - m + k] \setminus M$ into $\mathcal{O}(dd')$ maximal ranges. For each such maximal range J , we rely on the characterization of Claim 5.17 in order to compute $\text{Occ}_k^E(P, T) \cap (I + |Q|\mathbb{Z}) \cap J$. Recall that for $p, p' \in J$ with $p \equiv p' \pmod{|Q|}$ we have $p \in \text{Occ}_k^E(P, T)$ if and only if $p' \in \text{Occ}_k^E(P, T)$. Hence, it suffices to restrict our attention to the intersection of the first (at most) $|Q|$ positions of J with $I + |Q|\mathbb{Z}$. This intersection consists of at most two intervals of total size at most $|I|$. We call Verify for each of them, and for each position returned by these Verify queries, we add an arithmetic progression to $\text{Occ}_k^E(P, T) \cap (I + |Q|\mathbb{Z})$.

We now proceed to analyze the time complexity of the algorithm in the PILLAR model. The two calls to Locked require $\mathcal{O}(d^2 + k + d'^2)$ time in total in the PILLAR model due to Lemma 6.9. We then decompose M into maximal ranges, which can be implemented in $\mathcal{O}(dd' \log \log(dd'))$ time. The interval R of positions marked due to locked regions L_i^P and L_j^T is of size $|L_i^P| + |L_j^T| + 2k - 1$; the number of maximal ranges $R' \subseteq R \cap (I + |Q|\mathbb{Z})$ is at most $(|R| + 2|Q| - 2|I|)/|Q|$. Consequently, the total number of maximal ranges of size at most $|I|$ that we need to consider intervals does not exceed $2k + 1$ plus

$$\begin{aligned} \sum_{i=1}^{\ell^P} \sum_{j=1}^{\ell^T} \frac{|L_i^P| + |L_j^T| + 2|Q| - 2|I|}{|Q|} &\leq \frac{\ell^T}{|Q|} \sum_{i=1}^{\ell^P} |L_i^P| + \frac{\ell^P}{|Q|} \sum_{i=1}^{\ell^T} |L_i^T| + \frac{2\ell^P \ell^T (|Q| - |I|)}{|Q|} \\ &= \mathcal{O}((d'(d|Q| + k|Q|) + dd'|Q| + dd'|Q|)/|Q|) \\ &= \mathcal{O}(d'(d+k)). \end{aligned}$$

Each call to Verify in Line 11 of Algorithm 15 requires time $\mathcal{O}(k(k + |J'|))$ by Lemma 6.2. By the above analysis, we make $\mathcal{O}(d'(d+k))$ calls to Verify , each time for an interval of size at most $|I|$. Hence, we can upper bound the overall running time for this step by $\mathcal{O}(d'(d+k)k(k + |I|))$. Finally, the total time required by Verify queries in Line 15 of Algorithm 15 is $\mathcal{O}(dd'k(k + |I|))$ as we call Verify $\mathcal{O}(dd')$ times, each time for an interval of size $\mathcal{O}(|I|)$. Thus, the overall running time is $\mathcal{O}(d'(d+k)k(k + |I|) + d^2 + d'^2 + dd' \log \log(dd')) = \mathcal{O}(kd'(d+k)(k + |I| + d + d'))$.

The bounds obtained in the time complexity analysis also imply that our representation of $\text{Occ}_k^E(P, T) \cap (I + |Q|\mathbb{Z})$ consists of $\mathcal{O}(|I|d'(d+k))$ arithmetic progressions. \blacksquare

■ **Algorithm 15** Computing k -error occurrences in the presence of locked regions in text and pattern.

```

1 SynchedMatches( $P, T, I, k, d, d', Q$ )
2    $\mathcal{L}^P \leftarrow \text{Locked}(P, Q, d, k)$ ;
3    $\mathcal{L}^T \leftarrow \text{Locked}(T, Q, d', 0)$ ;
4    $M \leftarrow [n - m - k .. n - m + k]$ ;
5   foreach  $P[\ell .. r] \in \mathcal{L}^P$  do
6     foreach  $T[\ell' .. r'] \in \mathcal{L}^T$  do
7        $M \leftarrow M \cup (\ell' - r - k .. r' - \ell + k)$ ;
8    $M \leftarrow M \cap [0 .. n - m + k]$ ;
9   foreach maximal range  $J \subseteq M$  do
10    foreach maximal range  $J' \subseteq J \cap (I + |Q|\mathbb{Z})$  do
11       $\text{Occ}_k^E(P, T) \leftarrow \text{Occ}_k^E(P, T) \cup \{pos \mid (pos, k_{pos}) \in \text{Verify}(P, T, k, J')\}$ ;
12    foreach maximal range  $[\ell .. r] \subseteq [0 .. n - m + k] \setminus M$  do
13       $J \leftarrow [\ell .. \min(r, \ell + |Q|)]$ ;
14      foreach maximal range  $J' \subseteq J \cap (I + |Q|\mathbb{Z})$  do
15        foreach  $(pos, k_{pos}) \in \text{Verify}(P, T, k, J')$  do
16           $\text{Occ}_k^E(P, T) \leftarrow \text{Occ}_k^E(P, T) \cup ((pos + |Q|\mathbb{Z}) \cap [\ell .. r])$ ;
17    return  $\text{Occ}_k^E(P, T)$ ;

```

■ **Lemma 6.11** ($\text{PeriodicMatches}(P, T, k, d, Q)$). *Let P denote a pattern of length m and let T denote a text of length n . Further, let $0 \leq k \leq m$ denote a threshold, let $d \geq 2k$ denote a positive integer, and let Q denote a primitive string that satisfies $|Q| \leq m/8d$ and $\delta_E(P, *Q^*) \leq d$.*

There is an algorithm that computes the set $\text{Occ}_k^E(P, T)$, using $\mathcal{O}(n/m \cdot d^4)$ time in the PILLAR model.

Proof. We consider $\lfloor 2n/m \rfloor$ blocks $T_0, \dots, T_{\lfloor 2n/m \rfloor - 1}$ of T , each of length at most $\frac{3}{2}m + k - 1$, where the i -th block starts at position $i \cdot m/2$, that is,

$$T_i := T[\lfloor i \cdot m/2 \rfloor .. \min\{n, \lfloor (i+3) \cdot m/2 \rfloor + k - 1\}].$$

Observe that each k -error occurrence of P in T is contained in at least one of the fragments T_i : Specifically, T_i covers all occurrences starting in $[\lfloor i \cdot m/2 \rfloor .. \lfloor (i+1) \cdot m/2 \rfloor]$. For each block T_i , we call $\text{FindRelevantFragment}(P, T_i, k, d, Q)$ from Lemma 6.8 and obtain a fragment $T'_i = T[\ell_i .. r_i]$ containing all k -error occurrences of P in T_i and an integer range I_i . Lemma 6.8 guarantees that $\delta_E(T'_i, Q) \leq 3d$ and $|I_i| \leq 6d + 1$. Next, we call $\text{SynchedMatches}(P, T'_i, I_i, d, 3d, k, Q)$ from Lemma 6.10. The output of the call to SynchedMatches consists of $\mathcal{O}((6d+1)3d(d+k)) = \mathcal{O}(d^3)$ arithmetic progressions. For each obtained arithmetic progression, we first add $\lfloor i \cdot m/2 \rfloor$ to all of its elements, and, if $i < \lfloor 2n/m \rfloor - 1$, we intersect the resulting arithmetic progression with $[\lfloor i \cdot m/2 \rfloor .. \lfloor (i+1) \cdot m/2 \rfloor]$; finally, we add the obtained set to $\text{Occ}_k^E(P, T)$. The intersection step guarantees that each k -error occurrence is accounted for by exactly one block.

For the correctness, note that by Lemma 6.8, for each i , we have $\text{Occ}_k^E(P, T'_i) \bmod |Q| \subseteq I_i \bmod |Q|$. Hence, the call to SynchedMatches indeed computes all occurrences of P in T_i .

Each call to $\text{FindRelevantFragment}$ requires $\mathcal{O}(d^2)$ time, while each call to SynchedMatches requires time $\mathcal{O}(3kd(d+k)(k + (6d+1) + d + 3d)) = \mathcal{O}(d^4)$. The claimed overall running time follows. ■

■ **Algorithm 16** A PILLAR model algorithm for Lemma 5.21.

```

1 BreakMatches( $P, T, \{B_1 = P[b_1 \dots b_1 + |B_1|], \dots, B_{2k} = P[b_{2k} \dots b_{2k} + |B_{2k}|]\}, k$ )
2   multi-set  $M \leftarrow \{\}$ ;  $\text{Occ}_k^E(P, T) \leftarrow \{\}$ ;
3   for  $i \leftarrow 1$  to  $2k$  do
4     foreach  $\tau \in \text{ExactMatches}(B_i, T)$  do
5        $M \leftarrow M \cup \{[(\tau - b_i - k)/k]\}$ ; // Mark block  $\lfloor (\tau - b_i - k)/k \rfloor$  of  $T$ 
6        $M \leftarrow M \cup \{[(\tau - b_i)/k]\}$ ; // Mark block  $\lfloor (\tau - b_i)/k \rfloor$  of  $T$ 
7        $M \leftarrow M \cup \{[(\tau - b_i + k)/k]\}$ ; // Mark block  $\lfloor (\tau - b_i + k)/k \rfloor$  of  $T$ 
8        $M \leftarrow M \cup \{[(\tau - b_i + 2k)/k]\}$ ; // Mark block  $\lfloor (\tau - b_i + 2k)/k \rfloor$  of  $T$ 
9   sort  $M$ ;
10  foreach  $\pi \in [0 \dots n - m]$  that appears at least  $k$  times in  $M$  do
11     $\text{Occ}_k^E(P, T) \leftarrow \text{Occ}_k^E(P, T) \cup \{\text{pos} \mid (\text{pos}, k_{\text{pos}}) \in \text{Verify}(P, T, k, [\pi \cdot k \dots (\pi + 1) \cdot k])\}$ ;
12  return  $\text{Occ}_k^E(P, T)$ ;

```

6.4 Computing Occurrences in the Non-Periodic Case

■ **Lemma 6.12** ($\text{BreakMatches}(P, T, \{B_1, \dots, B_{2k}\}, k)$: Implem. of Lemma 5.21). *Let k denote a threshold and let P denote a pattern of length m having $2k$ disjoint breaks $B_1, \dots, B_{2k} \preceq P$ each satisfying $\text{per}(B_i) \geq m/128k$. Further, let T denote a text of length $n \leq \frac{3}{2}m + k$.*

Then, we can compute the set $\text{Occ}_k^E(P, T)$ using $\mathcal{O}(k^3)$ time in the PILLAR model.

Proof. We proceed similarly to Lemma 4.9: Instead of marking positions, we now mark blocks of length k ; in the end, we verify complete blocks at once using Verify from Lemma 6.2. Consider Algorithm 16 for the complete algorithm visualized as pseudo-code.

For the correctness, note that we have placed the marks as in the proof of Lemma 5.21; in particular, by Claim 5.23, any block $[jk \dots (j+1)k]$ that contains any position $\pi \in \text{Occ}_k^E(P, T)$ has at least k marks. As we verify each such block using Verify from Lemma 6.2, we report no false positives, and thus the algorithm is correct.

We continue with analyzing the running time. As every break B_i has period $\text{per}(B_i) > m/128k$, every call to ExactMatches uses $\mathcal{O}(k)$ time in the PILLAR model by Lemma 2.11; thus, all calls to ExactMatches in total take $\mathcal{O}(k^2)$ time in total. Next, by Claim 5.22, we place at most $\mathcal{O}(k^2)$ marks in T , so the marking step uses $\mathcal{O}(k^2)$ operations in total. Further, finding all positions in T with at least k marks can be done via a linear scan over the multi-set M of all marks after sorting M , which can be done in time $\mathcal{O}(k^2 \log \log k)$. Finally, as there are at most $\mathcal{O}(k^2/k) = \mathcal{O}(k)$ blocks that we verify, and every call to Verify takes time $\mathcal{O}(k^2)$ in the PILLAR model, the verifications take $\mathcal{O}(k^3)$ time in the PILLAR in total. Overall, Algorithm 16 thus takes $\mathcal{O}(k^3)$ time in the PILLAR model. ■

■ **Lemma 6.13** ($\text{RepetitiveMatches}(P, T, \{(R_1, Q_1) \dots (R_r, Q_r)\}, k)$: Implementation of Lemma 5.24). *Let P denote a pattern of length m and let $k \leq m$ denote a threshold. Further, let T denote a string of length $n \leq \frac{3}{2}m + k$. Suppose that P contains disjoint repetitive regions R_1, \dots, R_r of total length at least $\sum_{i=1}^r |R_i| \geq 3/8 \cdot m$ such that each region R_i satisfies $|R_i| \geq m/8k$ and has a primitive approximate period Q_i with $|Q_i| \leq m/128k$ and $\delta_H(R_i, Q_i^*) = \lceil 8k/m \cdot |R_i| \rceil$.*

Then, we can compute the set $\text{Occ}_k^E(P, T)$ using $\mathcal{O}(k^4)$ time in the PILLAR model.

Proof. As in the proof of Lemma 5.24, set $m_R := \sum_{i=1}^r |R_i| \geq 3/8 \cdot m$ and define for every $1 \leq i \leq r$ the values $k_i := \lfloor 4 \cdot k/m \cdot |R_i| \rfloor$ and $d_i := \lceil 8 \cdot k/m \cdot |R_i| \rceil = \lceil \text{Mis}(R_i, Q_i^*) \rceil$. Further, write $R_i = P[r_i \dots r_i + |R_i|]$.

■ **Algorithm 17** A PILLAR model algorithm for Lemma 5.24.

```

1 RepetitiveMatches( $P, T, \{(R_1 = P[r_1 \dots r_1 + |R_1|], Q_1) \dots, (R_r = P[r_r \dots r_r + |R_r|], Q_r)\}, k$ )
2   multi-set  $M \leftarrow \{\}$ ;  $\text{Occ}_k^E(P, T) \leftarrow \{\}$ ;
3   for  $i \leftarrow 1$  to  $r$  do
4     set  $M_i \leftarrow \{\}$ ;
5     foreach  $\tau \in \text{PeriodicMatches}(R_i, T, \lfloor 4 \cdot k/m \cdot |R_i| \rfloor, \lceil 8 \cdot k/m \cdot |R_i| \rceil, Q_i)$  do
6        $M_i \leftarrow M_i \cup \{(\lfloor (\tau - r_i - k)/k \rfloor, |R_i|)\}$ ; // Place  $|R_i|$  marks at bl.  $\lfloor (\tau - r_i - k)/k \rfloor$ 
7        $M_i \leftarrow M_i \cup \{(\lfloor (\tau - r_i)/k \rfloor, |R_i|)\}$ ; // Place  $|R_i|$  marks at block  $\lfloor (\tau - r_i)/k \rfloor$ 
8        $M_i \leftarrow M_i \cup \{(\lfloor (\tau - r_i + k)/k \rfloor, |R_i|)\}$ ; // Place  $|R_i|$  marks at bl.  $\lfloor (\tau - r_i + k)/k \rfloor$ 
9        $M_i \leftarrow M_i \cup \{(\lfloor (\tau - r_i + 2k)/k \rfloor, |R_i|)\}$ ; // Place  $|R_i|$  marks at bl.  $\lfloor (\tau - r_i + 2k)/k \rfloor$ 
10     $M \leftarrow M \cup M_i$ ;
11    sort  $M$  by positions;
12    foreach  $\pi \in [0 \dots n - m]$  appearing at least  $\sum_{(\pi, v) \in M} v \geq \sum_{i=1}^r |R_i| - m/4$  times in  $M$  do
13       $\text{Occ}_k^E(P, T) \leftarrow \text{Occ}_k^E(P, T) \cup \{\text{pos} \mid (\text{pos}, k_{\text{pos}}) \in \text{Verify}(P, T, k, [\pi \cdot k \dots (\pi + 1) \cdot k])\}$ ;
14    return  $\text{Occ}_k^E(P, T)$ ;
```

Again, we proceed similarly to the Hamming distance setting (Lemma 4.10). However, instead of marking positions, we now mark blocks of length k ; in the end, we then verify complete blocks at once using **Verify** from Lemma 6.2. Note that we need to ensure that we mark a block of T only at most once for each repetitive part R_i ; we do so by first computing a set of all blocks to be marked due to R_i (thereby removing duplicates) and then merging the sets computed for every R_i into a multi-set. Consider Algorithm 17 for the complete algorithm visualized as pseudo-code.

For the correctness, first note that in every call to **PeriodicMatches** from Lemma 6.11, we have

$$16k/m \cdot |R_i| \geq d_i = \lceil 8k/m \cdot |R_i| \rceil = \delta_H(R_i, Q_i^*) \geq 2k_i,$$

hence $|Q_i| \leq m/128k \leq |R_i|/8d_i$; thus, we can indeed call **PeriodicMatches** in this case. Further, note that we have placed the marks as in the proof of Lemma 5.24; in particular, by Claim 5.26, any block $[jk \dots (j+1)k]$ that contains any position $\pi \in \text{Occ}_k^E(P, T)$ has at least $m_R - m/4$ marks. As we verify every possible candidate using **Verify** from Lemma 6.2, we report no false positives, and thus the algorithm is correct.

For the running time in the PILLAR model, observe that during the marking step, for every repetitive region R_i we call **PeriodicMatches** once. In total, all calls to **PeriodicMatches** take

$$\sum_i \mathcal{O}(n/|R_i| \cdot d_i^4) = \sum_i \mathcal{O}(|R_i|/m \cdot k^4) = \mathcal{O}(k^4)$$

time in the PILLAR model. Further, for every R_i , we place at most $\mathcal{O}(\lfloor |\text{Occ}_{k_i}^E(R_i, T)/k \rfloor \rfloor)$ (weighted) marks, which can be bounded by $\mathcal{O}(\lfloor |\text{Occ}_{k_i}^E(R_i, T)/k \rfloor \rfloor) = \mathcal{O}(n/|R_i| \cdot d_i) = \mathcal{O}(k)$ using Corollary 5.18. Thus, we place $|M| = \mathcal{O}(k^2)$ (weighted) marks in total. Hence, the marking step in total takes $\mathcal{O}(k^4)$ time in the PILLAR model.

As the multi-set M contains at most $\mathcal{O}(k^2)$ (weighted) marks, we can sort M (by positions) in time $\mathcal{O}(k^2 \log \log k)$; afterwards, we can find the elements with total weight at least $m_R - m/4$ via a linear scan over M in time $\mathcal{O}(k^2)$. As there are (by Claims 5.25 and 5.26) at most $\mathcal{O}(k)$ blocks with at least $m_R - m/4$ marks, we call **Verify** at most $\mathcal{O}(k)$ times. As we call verify always on a whole block of length k at once, each call to **Verify** takes $\mathcal{O}(k^2)$ time in the PILLAR model. Hence, the verification step in total takes $\mathcal{O}(k^3)$ time in the PILLAR model.

In total, Algorithm 17 thus takes time $\mathcal{O}(k^4)$ in the PILLAR model. ■

■ **Algorithm 18** Computing k -error occurrences in the PILLAR model.

```

1 EditOccurrences( $P, T, k$ )
2   ( $B_1, \dots, B_{2k}$  or  $(R_1, Q_1), \dots, (R_r, Q_r)$  or  $Q$ )  $\leftarrow$  Analyze( $P, k$ );
3    $\text{Occ}_k^E(P, T) \leftarrow \{\}$ ;
4   if approximate period  $Q$  exists then
5     return PeriodicMatches( $P, T, k, 8k, Q$ );
6   for  $i \leftarrow 0$  to  $\lfloor 2n/m \rfloor - 1$  do
7      $T_i \leftarrow T[\lfloor i \cdot m/2 \rfloor .. \min\{n, \lfloor (i+3) \cdot m/2 \rfloor - 1 + k\}]$ ;
8     if breaks  $B_1, \dots, B_{2k}$  exist then
9        $\text{Occ}_k^E(P, T_i) \leftarrow$  BreakMatches( $P, T_i, \{B_1, \dots, B_{2k}\}, k$ );
10    else if repetitive regions  $(R_1, Q_1), \dots, (R_r, Q_r)$  exist then
11       $\text{Occ}_k^E(P, T_i) \leftarrow$  RepetitiveMatches( $P, T_i, \{(R_1, Q_1), \dots, (R_r, Q_r)\}, k$ );
12    if  $i < \lfloor 2n/m \rfloor - 1$  then
13       $V \leftarrow \{\ell + \lfloor i \cdot m/2 \rfloor \mid \ell \in \text{Occ}_k^E(P, T_i)\} \cap [\lfloor i \cdot m/2 \rfloor .. \lfloor (i+1) \cdot m/2 \rfloor]$ ;
14       $\text{Occ}_k^E(P, T) \leftarrow \text{Occ}_k^E(P, T) \cup V$ ;
15  return  $\text{Occ}_k^E(P, T)$ ;
```

6.5 A PILLAR Model Algorithm for Pattern Matching with Edits

Finally, we are ready to prove Main Theorem 9.

■ **Main Theorem 9.** *Given a pattern P of length m , a text T of length n , and a positive integer $k \leq m$, we can compute (a representation of) the set $\text{Occ}_k^E(P, T)$ using $\mathcal{O}(n/m \cdot k^4)$ time in the PILLAR model.*

Proof. We proceed, as in Main Theorem 8, by separately considering each of the three possible outcomes of **Analyze** (P, k). Consider Algorithm 18 for a visualization of the whole algorithm as pseudo-code.

If there is an approximate period Q of P we call **PeriodicMatches** (from Lemma 6.11). Else, for each of the $\lfloor 2n/m \rfloor$ blocks $T_0, \dots, T_{\lfloor 2n/m \rfloor - 1}$, where $T_i := T[\lfloor i \cdot m/2 \rfloor .. \min\{n, \lfloor (i+3) \cdot m/2 \rfloor - 1 + k\}]$, we call **BreakMatches** (from Lemma 6.12) or **RepetitiveMatches** (from Lemma 6.13), depending on the case we are in, and add the computed occurrences in $\text{Occ}_k^E(P, T)$.

The correctness in the approximately periodic case follows from Lemma 6.11 and the fact that we can indeed call **PeriodicMatches** since, due to Lemma 6.4, string Q satisfies $\delta_E(P, *Q^*) \leq 8k$ and $|Q| \leq m/128k \leq m/(8 \cdot 8k)$. In the other cases, first observe that each length- $(m+k)$ fragment of T is contained in at least one of the fragments T_i and hence we do not lose any occurrences. Second, by Lemma 6.4 and due to $|T_i| \leq \frac{3}{2}m + k$, the parameters in the calls to **BreakMatches** (from Lemma 6.12) and **RepetitiveMatches** (from Lemma 6.13) each satisfy the requirements. Finally, the intersection step in Line 13 of Algorithm 18 guarantees that we account for each k -error occurrence exactly once.

For the running time in the PILLAR model, we have that the call to **Analyze** takes $\mathcal{O}(k^2)$ time in the PILLAR model, the call to **PeriodicMatches** takes $\mathcal{O}(n/m \cdot k^4)$ time in the PILLAR model, each call to **BreakMatches** takes $\mathcal{O}(k^3)$ time in the PILLAR model, and each call to **RepetitiveMatches** takes $\mathcal{O}(k^4)$ time in the PILLAR model. Finally, as the output of our calls to **BreakMatches** and **RepetitiveMatches** is of size $\mathcal{O}(k^2)$ and is sorted, Lines 13 and 14 require $\mathcal{O}(k^2)$ time. As there are at most $\mathcal{O}(n/m)$ calls to **BreakMatches** and **RepetitiveMatches**, we can bound the total time in the PILLAR model by $\mathcal{O}(n/m \cdot k^4)$, completing the proof. ■

7 Implementing the PILLAR Model: Faster Approximate Pattern Matching

In this section we implement the PILLAR model in the static, fully compressed and dynamic settings, thereby lifting Main Theorems 8 and 9 to these settings. For each setting, we first show how to implement each of the following primitive PILLAR operations:

- **Extract**(S, ℓ, r): Retrieve a string $S[\ell..r]$.
- **LCP**(S, T): Compute the length of the longest common prefix of S and T .
- **IPM**(P, T): Assuming that $|T| \leq 2|P|$, compute $\text{Occ}(P, T)$ (represented as an arithmetic progression with difference $\text{per}(P)$).
- **Access**(S, i): Retrieve the character $S[i]$.
- **Length**(S): Compute the length $|S|$ of the string S .

Operation $\text{LCP}^R(S, T)$ can be implemented analogously to $\text{LCP}(S, T)$ by reversing all the strings in scope. We then apply our main algorithmic results (Main Theorems 8 and 9) in order to obtain efficient algorithms for approximate pattern matching.

7.1 Implementing the PILLAR Model in the Standard Setting

We start with the implementation of the PILLAR model in the standard setting. This turns out to be a straightforward application of known tools for strings.

Recall that in the PILLAR model, we are to maintain a collection \mathcal{X} of strings. Let us denote the total length of all strings in \mathcal{X} by n . In the standard setting, a handle to $S = X[\ell..r]$ is implemented as a pointer to $X \in \mathcal{X}$ (recall that X is stored explicitly) along with the indices ℓ and r . Hence, the implementations of **Extract**, **Access**, and **Length** are trivial.

Next, $\text{LCP}(S, T)$ queries can be implemented efficiently as follows. We construct the generalized suffix tree for the elements of \mathcal{X} in $\mathcal{O}(n)$ time [18] and preprocess it within the same time complexity so that we can support $\mathcal{O}(1)$ -time lowest common ancestor queries [7].

As for efficiently answering $\text{IPM}(P, T)$ queries, we build the data structure of Kociumaka et al. [32, 31], encapsulated in the following theorem, for the concatenation of the elements of \mathcal{X} .

▮ **Lemma 7.1** ([32, 31]). *For every string S of length n , there is a data structure of size $\mathcal{O}(n)$, which can be constructed in $\mathcal{O}(n)$ time and answers $\text{IPM}(P, T)$ queries in $\mathcal{O}(1)$ time for fragments P and T of S .* ▮

We summarize the above discussion in the following theorem.

▮ **Theorem 7.2.** *After an $\mathcal{O}(n)$ -time preprocessing of a collection of strings of total length n , each PILLAR operation can be performed in $\mathcal{O}(1)$ time.* ▮

This yields the following theorems, which are not new, but as a warm-up, they are perhaps somewhat instructive.

Combining Theorem 7.2 and Main Theorem 8, we obtain an algorithm for pattern matching with mismatches with the same running time as the algorithm of Clifford et al. [15], which is essentially optimal when $k = \mathcal{O}(\sqrt{m})$.^c

▮ **Theorem 7.3.** *Given a text T of length n , a pattern P of length m and a threshold k , we can compute the set $\text{Occ}_k^H(P, T)$ in time $\mathcal{O}(n + n/m \cdot k^2 \log \log k)$.* ▮

^c Strictly speaking, the algorithm in [15] runs in time $\mathcal{O}(n \text{polylog}(m) + n/m \cdot k^2 \log k)$, so our algorithm is slightly faster. However, an even better improvement in the logarithmic factors was already obtained recently in [13].

Similarly, combining Theorem 7.2 and Main Theorem 9, we obtain an algorithm for pattern matching with edits that is, again, not slower than the known algorithm [17].

▣ **Theorem 7.4.** *Given a text T of length n , a pattern P of length m and a threshold k , we can compute the set $\text{Occ}_k^E(P, T)$ in time $\mathcal{O}(n + n/m \cdot k^4)$.* ▣

▣ **Remark 7.5.** The discussion of this subsection implies that our algorithms also apply to the internal setting. That is, a string S of length n can be preprocessed in $\mathcal{O}(n)$ time, so that given fragments P and T of S , and a threshold k , we can compute $\text{Occ}_k^H(P, T)$ in time $\mathcal{O}(|T|/|P| \cdot k^2 \log \log k)$ and $\text{Occ}_k^E(P, T)$ in time $\mathcal{O}(|T|/|P| \cdot k^4)$. ▣

7.2 Implementing the PILLAR Model in the Fully Compressed Setting

Next, we focus on the fully compressed setting, where we want to solve approximate pattern matching when both the text and the pattern are given as a straight-line programs—that is, in this setting, we maintain a collection \mathcal{X} of straight-line programs and show how to implement the primitive PILLAR operations on this collection. We start with a short exposition on straight-line programs and related concepts.

Straight-Line Programs

We denote the set of non-terminals of a context-free grammar \mathcal{G} by $N_{\mathcal{G}}$ and call the elements of $\mathcal{S}_{\mathcal{G}} = N_{\mathcal{G}} \cup \Sigma$ symbols. Then, a *straight line program* (SLP) \mathcal{G} is a context-free grammar that consists of a set $N_{\mathcal{G}} = \{A_1, \dots, A_n\}$ of non-terminals, such that each $A \in N_{\mathcal{G}}$ is associated with a unique production rule $A \rightarrow f_{\mathcal{G}}(A)$, where $f_{\mathcal{G}}(A) \in \mathcal{S}_{\mathcal{G}}^*$. For SLPs given as input, we can assume without loss of generality that each production rule is of the form $A \rightarrow BC$ for some symbols B and C (that is, the given SLP is in Chomsky normal form).

Every symbol $A \in \mathcal{S}_{\mathcal{G}}$ generates a unique string, which we denote by $\text{gen}(A) \in \Sigma^*$. The string $\text{gen}(A)$ can be obtained from A by repeatedly replacing each non-terminal by its production. In addition, A is associated with its *parse tree* $\text{PT}(A)$ consisting of a root labeled with A to which zero or more subtrees are attached:

- ▀ If A is a terminal, there are no subtrees.
- ▀ If A is a non-terminal $A \rightarrow B_1 \cdots B_p$, then $\text{PT}(B_i)$ are attached in increasing order of i .


Note that if we traverse the leaves of $\text{PT}(A)$ from left to right, spelling out the corresponding non-terminals, then we obtain $\text{gen}(A)$.

The parse tree $\text{PT}_{\mathcal{G}}$ of \mathcal{G} is the parse tree of the starting symbol $A_n \in N_{\mathcal{G}}$; $\text{gen}(A_n) = S$, where S is the unique string generated by \mathcal{G} . We write $\text{gen}(\mathcal{G}) := S$. Finally, an SLP can be represented naturally as a directed acyclic graph $H_{\mathcal{G}}$ of size $|\mathcal{S}_{\mathcal{G}}|$. Consult Figure 7 for an example of an SLP, its parse tree, and the corresponding acyclic graph.

We define the *value* $\text{val}(v)$ of a node v in $\text{PT}_{\mathcal{G}}$ to be the fragment $S[a..b]$ corresponding to the leaves $S[a], \dots, S[b]$ in the subtree of v . Note that $\text{val}(v)$ is an occurrence of $\text{gen}(A)$ in $\text{gen}(\mathcal{G})$, where A is the label of v . A sequence of nodes in $\text{PT}_{\mathcal{G}}$ is a *chain* if their values are consecutive fragments in T .

Given an SLP \mathcal{G} of size n , we can compute $|\text{gen}(\mathcal{G})|$ in $\mathcal{O}(n)$ time using dynamic programming. We compute the topological order of $H_{\mathcal{G}}$ and process the nodes in the reverse order: For each node corresponding to a non-terminal A with production rule $A \rightarrow BC$, we just need to compute $|\text{gen}(B)| + |\text{gen}(C)|$.

Bille et al. [8] have shown that we can efficiently access any character in $\text{gen}(\mathcal{G})$.



■ **Figure 7** (a) An SLP \mathcal{G} generating aabaab . (b) The corresponding parse tree $\text{PT}_{\mathcal{G}}$. (c) The corresponding directed acyclic graph $H_{\mathcal{G}}$.

■ **Theorem 7.6** ([8]). *An SLP \mathcal{G} of size n , generating a string S of size N , can be preprocessed in time $\mathcal{O}(n \log(N/n))$ so that, for any $i \in [0..N)$, we can access $\text{gen}(\mathcal{G})[i]$ in $\mathcal{O}(\log N)$ time.* ■

I in [26] presented an efficient data structure for answering longest common prefix queries for suffixes of a string given by an SLP, which we encapsulate in the following theorem. This data structure is based on the recompression technique, which we discuss in the next subsection in more detail.

■ **Theorem 7.7** ([26]). *An SLP \mathcal{G} of size n , generating a string S of size N , can be preprocessed in time $\mathcal{O}(n \log(N/n))$ so that for any i and j , we can compute $\text{LCP}(S[i..N], S[j..N])$ in $\mathcal{O}(\log N)$ time.* ■

Finally, we discuss how to “concatenate” two SLPs. Given two SLPs \mathcal{G}_1 and \mathcal{G}_2 , with $\text{gen}(\mathcal{G}_1) = S_1$ and $\text{gen}(\mathcal{G}_2) = S_2$, we can construct an SLP generating S_1S_2 in $\mathcal{O}(|\mathcal{G}_1| + |\mathcal{G}_2|)$ time as follows. We first rename the non-terminals in $N_{\mathcal{G}_2}$ to make sure that they are disjoint from the non-terminals in $N_{\mathcal{G}_1}$. Next, let R_1 and R_2 denote the starting non-terminals of \mathcal{G}_1 and \mathcal{G}_2 , respectively. We construct a new SLP \mathcal{G} with $N_{\mathcal{G}} := N_{\mathcal{G}_1} \cup N_{\mathcal{G}_2} \cup \{R\}$, where R has production rule $R \rightarrow R_1R_2$. Note that this procedure can be applied to more than two strings: We first apply a global renaming, and then repeatedly “concatenate” two strings in the collection.

Let us now denote the total size of all SLPs in \mathcal{X} by n , and the total length of all strings generated by those SLPs by N . We implement the handle of a fragment $S = X[\ell..r)$ so that it consists of a pointer to the SLP $\mathcal{G} \in \mathcal{X}$ generating X along with the positions ℓ and r ; this makes **Extract** and **Length** trivial.

The above discussion on computing the length of the string generated by an SLP implies that the handles to $\text{gen}(\mathcal{G})[0..|\text{gen}(\mathcal{G})|)$ for all $\mathcal{G} \in \mathcal{X}$ can be constructed in $\mathcal{O}(n)$ time. Moreover, Theorem 7.6 implies that we can preprocess \mathcal{X} in $\mathcal{O}(n \log(N/n))$ time so that operation **Access** requires $\mathcal{O}(\log N)$ time.

For efficiently answering LCP queries, we rely on Theorem 7.7. We build I’s data structure for an SLP that generates the concatenation of all elements in the multi-set $\{\text{gen}(\mathcal{G}) \mid \mathcal{G} \in \mathcal{X}\}$. Thus, after an $\mathcal{O}(n \log(N/n))$ -time preprocessing, each LCP operation takes $\mathcal{O}(\log N)$ time.

To implement the IPM operation efficiently, we rely on the recompression technique due to Jež [27, 28], which we discuss next.

Recompression of Straight-Line Programs

We start with some additional notation. A *run-length straight line program (RLSLP)* is a straight-line program \mathcal{G} that contains two kinds of non-terminals:

- *Concatenations*: Non-terminals with production rules of the form $A \rightarrow BC$ (for symbols B and C).
- *Powers*: Non-terminals with production rules of the form $A \rightarrow B^p$ (for a symbol B and an integer $p \geq 2$).

The key idea of the *recompression* technique by Jež [27, 28] is the construction of a particular RLSP \mathcal{R} that generates the input string S . The parse tree $\text{PT}_{\mathcal{R}}$ is of depth $\mathcal{O}(\log N)$ and it can be traversed efficiently based on the underlying directed acyclic graph $H_{\mathcal{R}}$. In particular, the name of the technique stems from the fact that an SLP \mathcal{G} of size n generating a string S of length N can be efficiently recompressed to the RLSP $\mathcal{R}(\mathcal{G})$ in-place, that is, without first uncompressing \mathcal{G} ; efficiently here means in $\mathcal{O}(n \log N)$ time. As observed by I [26], the parse tree $\text{PT}_{\mathcal{R}(\mathcal{G})}$ is *locally consistent* in a certain sense. To formalize this property, he introduced the *popped sequence* of every fragment $S[a..b]$, which is a sequence of symbols labeling a certain chain of nodes of $\text{PT}_{\mathcal{R}(\mathcal{G})}$, whose values constitute $S[a..b]$.

▣ **Theorem 7.8** ([26]). *If two fragments match, then their popped sequences are equal. Moreover, each popped sequence consists of $\mathcal{O}(\log N)$ runs (maximal powers of a single symbol) and can be constructed in $\mathcal{O}(\log N)$ time. The nodes corresponding to symbols in a run share a single parent. Furthermore, the popped sequence consists of a single symbol only for fragments of length 1.* ▣

Let $F_1^{p_1} \cdots F_t^{p_t}$ denote the run-length encoding of the popped sequence of a substring U of S and set

$$L(U) := \{|\text{gen}(F_1)|, |\text{gen}(F_1^{p_1})|, |\text{gen}(F_1^{p_1} F_2^{p_2})|, \dots, |\text{gen}(F_1^{p_1} \cdots F_{t-1}^{p_{t-1}})|, |\text{gen}(F_1^{p_1} \cdots F_{t-1}^{p_{t-1}} F_t^{p_t})|\}.$$

By Theorem 7.8, the set $L(U)$ can be constructed in $\mathcal{O}(\log N)$ time given a fragment $S[a..b] = U$. Now, the following lemma from [14] allows us to efficiently implement internal pattern matching queries.

▣ **Lemma 7.9** ([14]). *Let v denote a non-leaf node of $\text{PT}_{\mathcal{R}(\mathcal{G})}$ and let $S[a..b]$ denote an occurrence of S that is contained in $\text{val}(v)$, but not contained in $\text{val}(u)$ for any child u of v . If $S[a..c]$ is the longest prefix of $S[a..b]$ contained in $\text{val}(u)$ for a child u of v , then $|S[a..c]| \in L(U)$. Symmetrically, if $S[c'+1..b]$ is the longest suffix of $S[a..b]$ contained in $\text{val}(u)$ for a child u of v , then $|S[a..c']| \in L(U)$.* ▣

Finally, we discuss how to implement the IPM operation.

▣ **Lemma 7.10**. *Given $H_{\mathcal{R}(\mathcal{G})}$, a fragment $T = S[j..j + \nu)$, and a fragment $P = S[i..i + \mu)$ with $|T| \leq 2|P|$, we can compute $\text{Occ}(P, T)$ in the time required by $\mathcal{O}(\log^2 N)$ LCP and LCP^R operations on fragments of S .*

In particular, the LCP and LCP^R operations are between pairs consisting in one of $\mathcal{O}(\log N)$ fragments of P and one of $\mathcal{O}(\log N)$ fragments of S .

Proof. We can assume that $|T| \geq |P| > 1$; otherwise it suffices to perform a constant number of letter comparisons which can be done in $\mathcal{O}(\log N)$ time. We first compute the popped sequence of P and $L(P)$ using Theorem 7.8. Let $v \in \text{PT}_{\mathcal{R}(\mathcal{G})}$ denote the lowest common ancestor of the leaves representing $S[j]$ and $S[j + \nu - 1]$; the node v can be naively computed in $\mathcal{O}(\log N)$ time by a forward search from the root of the parse tree. As T is a fragment of $\text{val}(v)$, all occurrences of P in T are also contained in $\text{val}(v)$. Our first aim is to compute all occurrences of P in T that are not contained in $\text{val}(u)$ for any child u of v ; we then appropriately recurse on the children of v that may contain sought occurrences.

Let us first analyze the case that the label of v is a concatenation symbol $A \rightarrow BC$. Write v_ℓ for the left child of v and v_r for the right child of v . Further, let $T_\ell = S[j..t)$ denote the longest prefix of T that is completely contained in $\text{val}(v_\ell)$ and, similarly, let $T_r = S[t..j + \nu)$ denote the longest suffix of T that is completely contained in $\text{val}(v_r)$. Suppose that there is a fragment $U = T[a..b)$ of T that equals P and overlaps with both $\text{val}(v_\ell)$ and $\text{val}(v_r)$. The fragment U can then be naturally decomposed

into a non-empty suffix U_ℓ of T_ℓ and a non-empty prefix U_r of T_r . Lemma 7.9 implies that $|U_\ell| \in L(P)$. It thus suffices to check for each $q \in L(P)$ whether $P[0..q]$ is a suffix of T_ℓ and $P[q..m]$ is a prefix of T_r . There are $|L(P)| = \mathcal{O}(\log N)$ choices for q , and for each of them we can perform the check using $\text{LCP}^R(P[0..q], T_\ell)$ and $\text{LCP}(P[q..m], T_r)$ operations.

We now consider the case that the label of v is a power symbol $A \rightarrow B^p$ and denote the children of v in the left-to-right order by v_1, \dots, v_p . Let $T[x..y]$ denote the overlap of T with $\text{val}(v)$.

If T overlaps with $\text{val}(v_d)$ for two children v_d of v , then we can process v as in the previous case. In the case that T overlaps with exactly three children of v , some care is needed to avoid double-counting occurrences that overlap with all of them. In particular, let these three children be v_x, v_{x+1} , and v_{x+2} . We consider separately:

- occurrences that overlap with both $\text{val}(v_x)$ and $\text{val}(v_{x+1})$ by setting T_ℓ to be the longest prefix of T that is completely contained in $\text{val}(v_x)$ and T_r to be the longest suffix of T that is completely contained in $\text{val}(v_{x+1})\text{val}(v_{x+2})$, and
- occurrences that overlap with both $\text{val}(v_{x+1})$ and $\text{val}(v_{x+2})$, but not with $\text{val}(v_x)$, by setting T_ℓ to be $\text{val}(v_{x+1})$ and $T_r = S[t..j + \nu]$ to be the longest suffix of T that is completely contained in $\text{val}(v_{x+2})$.

We can thus assume that T overlaps with $\text{val}(v_d)$ for more than three children v_d of v . In that case, for all d , we have $\text{val}(v_d) < \mu$ and hence no occurrence of P in $T[x..y]$ can be completely contained in $\text{val}(v_d)$. We set $T_\ell := \text{val}(v_1)$ and $T_r := \text{val}(v_2) \cdots \text{val}(v_p)$. Using $|L(P)|$ many LCP operations and $|L(P)|$ many LCP^R operations, we can compute the set Y of occurrences of P in S (but not necessarily in T) which can be decomposed into a prefix U_ℓ that is a suffix of T_ℓ and a suffix U_r that is a prefix of T_r .

Then, by the periodicity of $\text{val}(v) = \text{gen}(B)^p$, the desired set of occurrences is

$$Z := \{i + j \cdot |\text{gen}(B)| : i \in Y, j \in [0..p-1]\} \cap [x..y - \mu].$$

Note that Z trivially decomposes into $\mathcal{O}(\log N)$ arithmetic progressions; these arithmetic progressions can be replaced by a single arithmetic progression with difference $\text{per}(P)$ in $\mathcal{O}(\log N)$ time.

If the overlap of the value of each child of v with T has length less than μ , we terminate the algorithm. Otherwise, if v has two children whose values have an overlap of length μ with T , we check whether either of them is equal to P using a single LCP operation and terminate the algorithm. Finally, at most one of v 's children has length greater than μ . In that case, we repeat the above procedure for this child. As the depth of $\text{PT}_{\mathcal{R}(\mathcal{G})}$ is $\mathcal{O}(\log N)$, the overall running time is upper bounded by the time required for $\mathcal{O}(\log^2 N)$ LCP and LCP^R operations.

In the end, we have $\text{Occ}(P, T)$ represented by at most one arithmetic progression and $\mathcal{O}(\log^2 N)$ single occurrences. We postprocess this representation in $\mathcal{O}(\log^2 N)$ time, in order to represent $\text{Occ}(P, T)$ with a single arithmetic progression with difference $\text{per}(P)$.

Note that, in each level, the LCP (resp. LCP^R) queries we perform are between $\mathcal{O}(1)$ fragments of S and each $P[0..q]$ (resp. $P[q..m]$) for $q \in L(P)$. This observation implies the last claim of the statement of this lemma, concluding its proof. ■

■ **Remark 7.11.** Essentially the same proof of the above lemma has recently appeared in [29]. ■

We can get an $\mathcal{O}(\log^3 N)$ -time implementation of IPM by employing Lemma 7.10, and answering each LCP or LCP^R query in $\mathcal{O}(\log N)$ time. However, the structure of the queries allows for a more efficient implementation.

■ **Lemma 7.12.** *We can preprocess an SLP \mathcal{G} of size n , generating a string S of length N , in $\mathcal{O}(n \log N)$ time, so that $\text{IPM}(P, T)$ queries for fragments P and T of S can be answered in $\mathcal{O}(\log^2 N \log \log N)$ time.* ■

Proof. We first build, in $\mathcal{O}(n \log(N/n))$ time, the data structures of Theorems 7.6 and 7.7 for performing **Access**, **LCP**, and LCP^R operations in $\mathcal{O}(\log N)$ time. Then, we recompress \mathcal{G} to an RLSLP $\mathcal{R}(\mathcal{G})$ in $\mathcal{O}(n \log N)$ time.

Lemma 7.10 then reduces the task at hand to answering $\mathcal{O}(\log^2 N)$ **LCP** and LCP^R queries between pairs consisting in one of $\mathcal{O}(\log N)$ fragments of P and one of $\mathcal{O}(\log N)$ fragments of S . Let us focus on efficiently answering all such **LCP** queries; LCP^R queries can be answered analogously.

We sort all fragments in scope using $\mathcal{O}(\log N \log \log N)$ comparisons, implementing each comparison in $\mathcal{O}(\log N)$ time using an **LCP** operation, followed by two **Access** operations. This step thus takes $\mathcal{O}(\log^2 N \log \log N)$ time. Then, we construct an array A of size $\mathcal{O}(\log N)$ such that $A[i]$ stores the length of the longest common prefix of the i -th and the $(i+1)$ -st elements in our sorted list. After preprocessing array A in $\mathcal{O}(\log N)$ time so that range minimum queries over it can be answered in constant time [7], we answer each of the $\mathcal{O}(\log^2 N)$ **LCP** queries in $\mathcal{O}(1)$ time. \blacksquare

In total, we have thus proved the following result.

\blacksquare **Theorem 7.13.** *Given a collection of SLPs of total size n , generating strings of total length N , each **PILLAR** operation can be performed in $\mathcal{O}(\log^2 N \log \log N)$ time after an $\mathcal{O}(n \log N)$ -time preprocessing. \blacksquare*

In the next subsection (cf. Remark 7.16), we discuss how to perform each **PILLAR** operation in $\mathcal{O}(\log^2 N)$ after an $\mathcal{O}(n \log N)$ -time preprocessing at the cost of randomization.

Approximate Pattern Matching in Fully Compressed Strings

We are now ready to present efficient algorithms for approximate pattern matching in the fully compressed setting. We choose to state our results using our deterministic implementation of the **PILLAR** model, that is Theorem 7.13.

We are given an SLP \mathcal{G}_T of size n with $T = \text{gen}(\mathcal{G}_T)$, an SLP \mathcal{G}_P of size m with $P = \text{gen}(\mathcal{G}_P)$ and a threshold k and are required to compute the k -mismatch or k -error occurrences of P in T .

Set $N := |T| + |P|$ and $\mathcal{X} := \{\mathcal{G}_T, \mathcal{G}_P\}$. The overall structure of our algorithm is as follows: We first preprocess the collection \mathcal{X} in $\mathcal{O}((n+m) \log N)$ time according to Theorem 7.13. Next, we traverse \mathcal{G}_T and compute for every non-terminal A of \mathcal{G}_T the approximate occurrences of P in T that “cross” A . Depending on the setting, we combine Theorem 7.13 with Main Theorem 8 or Main Theorem 9 to compute the occurrences. Finally, we combine the computed occurrences using dynamic programming.

Formally, for each non-terminal $A \in N_{\mathcal{G}_T}$, with production rule $A \rightarrow BC$, we wish to compute all approximate occurrences of P in the string

$$\text{gen}(B)[|\text{gen}(B)| - |P| + 1 \dots |\text{gen}(B)|] \text{gen}(C)[0 \dots |P| - 1],$$

which is indeed a fragment of $\text{gen}(\mathcal{G}_T)$ and is of length $2|P| - 2$. These approximate occurrences can be computed in time:

- $\mathcal{O}(k^2 \log^2 N \log \log N)$ in the Hamming distance case by combining Theorem 7.13 and Main Theorem 8;
 - $\mathcal{O}(k^4 \log^2 N \log \log N)$ in the edit distance case, by combining Theorem 7.13 and Main Theorem 9.
- Other approximate occurrences in $\text{gen}(A)$ lie entirely in $\text{gen}(B)$ or $\text{gen}(C)$; hence they are computed when considering B and C . (Compare [11, Theorem 4.1] for a similar algorithm.)

Now, the number of approximate occurrences of P in T (that is $|\text{Occ}_k^H(P, T)|$ or $|\text{Occ}_k^E(P, T)|$) can be computed by dynamic programming that is analogous to the dynamic programming used to compute the length of a string generated by an SLP. Further, all approximate occurrences can be reported in time proportional to their number by performing a traversal of $\text{PT}_{\mathcal{G}}$, avoiding to explore subtrees that correspond to fragments of T that do not contain any approximate occurrences.

We hence obtain the following algorithm for pattern matching with mismatches in the fully compressed setting.

▣ **Main Theorem 1.** *Let \mathcal{G}_T denote an SLP of size n generating a text T , let \mathcal{G}_P denote an SLP of size m generating a pattern P , let $k > 0$ denote an integer threshold, and set $N := |T| + |P|$.*

Then, we can compute $|\text{Occ}_k^H(P, T)|$ in time $\mathcal{O}(m \log N + n k^2 \log^2 N \log \log N)$. The elements of $\text{Occ}_k^H(P, T)$ can be reported within $\mathcal{O}(|\text{Occ}_k^H(P, T)|)$ extra time. ▣

Similarly, we obtain the following algorithm for pattern matching with edits in the fully compressed setting.

▣ **Main Theorem 2.** *Let \mathcal{G}_T denote an SLP of size n generating a string T , let \mathcal{G}_P denote an SLP of size m generating a string P , let $k > 0$ denote an integer threshold, and set $N := |T| + |P|$.*

Then, we can compute $|\text{Occ}_k^E(P, T)|$ in time $\mathcal{O}(m \log N + n k^4 \log^2 N \log \log N)$. The elements of $\text{Occ}_k^E(P, T)$ can be reported within $\mathcal{O}(|\text{Occ}_k^E(P, T)|)$ extra time. ▣

7.3 Implementing the PILLAR Model in the Dynamic Setting

Lastly, we consider the dynamic setting. In particular, we consider the dynamic maintenance of a collection of non-empty persistent strings \mathcal{X} that is initially empty and undergoes updates specified by the following operations:

- ▣ **makestring**(U): Insert a non-empty string U to \mathcal{X} .
- ▣ **concat**(U, V): Insert UV to \mathcal{X} , for $U, V \in \mathcal{X}$.
- ▣ **split**(U, i): Insert $U[0..i)$ and $U[i..|U|)$ in \mathcal{X} , for $U \in \mathcal{X}$ and $i \in [0..|U|)$.

Let N denote an upper bound on the total length of all strings in \mathcal{X} throughout the execution of the algorithm. Gawrychowski et al. [23] presented a data structure that efficiently maintains such a collection and allows for efficient longest common prefix queries.

▣ **Theorem 7.14** ([23]). *A collection \mathcal{X} of non-empty persistent strings of total length N can be dynamically maintained with update operations **makestring**(U), **concat**(U, V), **split**(U, i) requiring time $\mathcal{O}(\log N + |U|)$, $\mathcal{O}(\log N)$, and $\mathcal{O}(\log N)$,^d respectively, so that **LCP**(U, V) queries for $U, V \in \mathcal{X}$ can be answered in time $\mathcal{O}(1)$. ▣*

LCP and **LCP** ^{R} operations for arbitrary fragments of elements of \mathcal{X} can be answered in $\mathcal{O}(\log N)$ time (w.h.p.) by first performing a constant number of **split** operations to add the corresponding fragments to the collection and then asking an **LCP** query between them.

The lengths of the strings in \mathcal{X} can be maintained explicitly. Upon a **makestring** operation we naively compute the length of U , while upon a **concat** or a **split** operation, we can compute the lengths of the strings that are inserted in \mathcal{X} in constant time from the arguments of the operation.

For each string of the collection \mathcal{X} , the data structure of [23] maintains an RLSLP stemming from recompression that is of depth $\mathcal{O}(\log N)$ w.h.p. Given a string $X \in \mathcal{X}$, a pointer to the root of the parse tree of X can be retrieved in $\mathcal{O}(1)$ time.

Each **Extract**(X, ℓ, r) operation can be performed using at most two **split** operations in $\mathcal{O}(\log N)$ time w.h.p.

^d These running times hold w.h.p.

We now show that $\text{Access}(X, i)$ for $X \in \mathcal{X}$ can be performed efficiently. Although the parse trees of the strings in \mathcal{X} are not maintained explicitly, given a pointer to some node v in the parse tree of X , we can retrieve in $\mathcal{O}(1)$ time the endpoints a, b of the fragment $\text{val}(v) = X[a..b]$, the degree of v , a pointer to the parent of v , and a pointer to the j -th child of v , provided that such a child exists. Thus, an $\text{Access}(X, i)$ operation can be implemented in time proportional to the height of the parse tree, that is, in $\mathcal{O}(\log N)$ time w.h.p.

We are left with showing that IPM operations can be performed efficiently. Let us remark that the RLSPs of all $X \in \mathcal{X}$ maintained by the data structure underlying Theorem 7.14 are locally consistent with each other, that is Theorem 7.8 is also true for fragments of different strings $X_1, X_2 \in \mathcal{X}$. Thus, Lemmas 7.9 and 7.10 also hold in this setting. Combining Lemma 7.10 and Theorem 7.14 we get that $\text{IPM}(P, T)$ queries can be answered in $\mathcal{O}(\log^2 N)$ time (w.h.p.) by performing $\mathcal{O}(\log N)$ `split` operations and $\mathcal{O}(\log^2 N)$ LCP queries (cf. the last statement of Lemma 7.10). Note that the only other component of the proof of Lemma 7.10 is a forward search from the root of the relevant parse tree, which can be efficiently performed given the available pointers.

We summarize the above discussion in the following theorem.

■ **Theorem 7.15.** *A collection \mathcal{X} of non-empty persistent strings of total length N can be dynamically maintained with operations `makestring`(U), `concat`(U, V), `split`(U, i) requiring time $\mathcal{O}(\log N + |U|)$, $\mathcal{O}(\log N)$ and $\mathcal{O}(\log N)$, respectively, so that PILLAR operations can be performed in time $\mathcal{O}(\log^2 N)$.^e ■*

■ **Remark 7.16.** Given an SLP \mathcal{G} of size n , generating a string S of size N , we can efficiently implement the PILLAR operations through dynamic strings. Let us start with an empty collection \mathcal{X} of dynamic strings. Using $\mathcal{O}(n)$ `makestring`(a) operations, for $a \in \Sigma$, and $\mathcal{O}(n)$ `concat` operations (one for each non-terminal of \mathcal{G}), we can insert S to \mathcal{X} in $\mathcal{O}(n \log N)$ time w.h.p. Then, we can perform each PILLAR operation in $\mathcal{O}(\log^2 N)$ time w.h.p., due to Theorem 7.15, thus outperforming Theorem 7.13 at the cost of randomization. ■

Combining Theorem 7.15 and Main Theorems 8 and 9, we obtain an algorithm for approximate pattern matching for dynamic strings.

■ **Main Theorem 3.** *A collection \mathcal{X} of non-empty persistent strings of total length N can be maintained subject to `makestring`(U), `concat`(U, V), and `split`(U, i) operations requiring $\mathcal{O}(\log N + |U|)$, $\mathcal{O}(\log N)$, and $\mathcal{O}(\log N)$ time, respectively, so that given two strings $P, T \in \mathcal{X}$ with $|P| = m$ and $|T| = n$ and an integer threshold $k > 0$, we can compute $|\text{Occ}_k^H(P, T)|$ in time $\mathcal{O}(n/m \cdot k^2 \log^2 N)$ and $|\text{Occ}_k^E(P, T)|$ in time $\mathcal{O}(n/m \cdot k^4 \log^2 N)$.^f The elements of $\text{Occ}_k^H(P, T)$ and $\text{Occ}_k^E(P, T)$ can be reported in $\mathcal{O}(|\text{Occ}_k^H(P, T)|)$ and $\mathcal{O}(|\text{Occ}_k^E(P, T)|)$ extra time, respectively. ■*

^e All running time bounds hold w.h.p.

^f All running time bounds hold with high probability (i.e., $1 - N^{-\Omega(1)}$).

References

- 1 Amir Abboud, Arturs Backurs, Karl Bringmann, and Marvin Künnemann. Fine-grained complexity of analyzing compressed data: Quantifying improvements over decompress-and-solve. In Chris Umans, editor, *58th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2017*, pages 192–203. IEEE Computer Society, 2017. doi:10.1109/FOCS.2017.26.
- 2 Karl R. Abrahamson. Generalized string matching. *SIAM Journal on Computing*, 16(6):1039–1051, 1987. doi:10.1137/0216067.
- 3 Stephen Alstrup, Gerth Stølting Brodal, and Theis Rauhe. Pattern matching in dynamic texts. In David B. Shmoys, editor, *11th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2000*, pages 819–828. SIAM, 2000. URL: <http://dl.acm.org/citation.cfm?id=338219.338645>.
- 4 Amihod Amir, Moshe Lewenstein, and Ely Porat. Faster algorithms for string matching with k mismatches. *Journal of Algorithms*, 50(2):257–275, 2004. doi:10.1016/S0196-6774(03)00097-X.
- 5 Maxim Babenko, Paweł Gawrychowski, Tomasz Kociumaka, Ignat Kolesnichenko, and Tatiana Starikovskaya. Computing minimal and maximal suffixes of a substring. *Theoretical Computer Science*, 638:112–121, 2016. doi:10.1016/j.tcs.2015.08.023.
- 6 Arturs Backurs and Piotr Indyk. Edit distance cannot be computed in strongly subquadratic time (unless SETH is false). *SIAM Journal on Computing*, 47(3):1087–1097, 2018. doi:10.1137/15M1053128.
- 7 Michael A. Bender and Martin Farach-Colton. The LCA problem revisited. In *LATIN 2000: Theoretical Informatics, 4th Latin American Symposium, Punta del Este, Uruguay, April 10-14, 2000, Proceedings*, pages 88–94, 2000. doi:10.1007/10719839_9.
- 8 Philip Bille, Gad M. Landau, Rajeev Raman, Kunihiko Sadakane, Srinivasa Rao Satti, and Oren Weimann. Random Access to Grammar-Compressed Strings and Trees. *SIAM Journal on Computing*, 44(3):513–539, 2015. doi:10.1137/130936889.
- 9 Robert S. Boyer and J. Strother Moore. MJRTY—a fast majority vote algorithm. In Robert S. Boyer, editor, *Automated Reasoning: Essays in Honor of Woody Bledsoe*, Automated Reasoning Series, pages 105–117. Kluwer Academic Publishers, 1991. doi:10.1007/978-94-011-3488-0_5.
- 10 Dany Breslauer and Zvi Galil. Finding all periods and initial palindromes of a string in parallel. *Algorithmica*, 14(4):355–366, 1995. doi:10.1007/BF01294132.
- 11 Karl Bringmann, Marvin Künnemann, and Philip Wellnitz. Few Matches or Almost Periodicity: Faster Pattern Matching with Mismatches in Compressed Texts. In Timothy M. Chan, editor, *30th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019*, pages 1126–1145. SIAM, 2019. doi:10.1137/1.9781611975482.69.
- 12 Michael Burrows and David J. Wheeler. A block-sorting lossless data compression algorithm. Technical Report 124, Digital Equipment Corporation, Palo Alto, California, 1994.
- 13 Timothy M. Chan, Shay Golan, Tomasz Kociumaka, Tsvi Kopelowitz, and Ely Porat. Approximating text-to-pattern hamming distances. In Julia Chuzhoy, editor, *52nd Annual ACM Symposium on Theory of Computing, STOC 2020*, pages 643–656. ACM, 2020. doi:10.1145/3357713.3384266.
- 14 Panagiotis Charalampopoulos, Tomasz Kociumaka, Manal Mohamed, Jakub Radoszewski, Wojciech Rytter, and Tomasz Walen. Internal dictionary matching. In Pinyan Lu and Guochuan Zhang, editors, *30th International Symposium on Algorithms and Computation, ISAAC 2019, December 8-11, 2019, Shanghai University of Finance and Economics, Shanghai, China*, volume 149 of *LIPICs*, pages 22:1–22:17. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPICs.ISAAC.2019.22.
- 15 Raphaël Clifford, Allyx Fontaine, Ely Porat, Benjamin Sach, and Tatiana Starikovskaya. The k -mismatch problem revisited. In Robert Krauthgamer, editor, *27th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016*, pages 2039–2052. SIAM, 2016. doi:10.1137/1.9781611974331.ch142.
- 16 Raphaël Clifford, Allan Grønlund, Kasper Green Larsen, and Tatiana A. Starikovskaya. Upper and lower bounds for dynamic data structures on strings. In Rolf Niedermeier and Brigitte Vallée, editors, *35th Symposium on Theoretical Aspects of Computer Science, STACS 2018*, pages 22:1–22:14. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2018. doi:10.4230/LIPICs.STACS.2018.22.

- 17 Richard Cole and Ramesh Hariharan. Approximate String Matching: A Simpler Faster Algorithm. *SIAM Journal on Computing*, 31(6):1761–1782, 2002. doi:10.1137/S0097539700370527.
- 18 Martin Farach. Optimal suffix tree construction with large alphabets. In *38th Annual IEEE Symposium on Foundations of Computer Science, FOCS '97, Miami Beach, Florida, USA, October 19-22, 1997*, pages 137–143, 1997. doi:10.1109/SFCS.1997.646102.
- 19 Paolo Ferragina and Roberto Grossi. Optimal on-line search and sublinear time update in string matching. *SIAM Journal on Computing*, 27(3):713–736, 1998. doi:10.1137/S0097539795286119.
- 20 Zvi Galil and Raffaele Giancarlo. Improved string matching with k mismatches. *SIGACT News*, 17(4):52–54, 1986. doi:10.1145/8307.8309.
- 21 Pawel Gawrychowski and Damian Straszak. Beating $\mathcal{O}(nm)$ in approximate LZW-compressed pattern matching. In Leizhen Cai, Siu-Wing Cheng, and Tak Wah Lam, editors, *24th International Symposium on Algorithms and Computation, ISAAC 2013*, pages 78–88. Springer, 2013. doi:10.1007/978-3-642-45030-3_8.
- 22 Pawel Gawrychowski and Przemyslaw Uznański. Towards unified approximate pattern matching for Hamming and L_1 distance. In Ioannis Chatzigiannakis, Christos Kaklamanis, Dániel Marx, and Donald Sannella, editors, *45th International Colloquium on Automata, Languages, and Programming, ICALP 2018*, volume 107 of *LIPICs*, pages 62:1–62:13. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2018. doi:10.4230/LIPICs.ICALP.2018.62.
- 23 Paweł Gawrychowski, Adam Karczmarz, Tomasz Kociumaka, Jakub Łacki, and Piotr Sankowski. Optimal dynamic strings. In Artur Czumaj, editor, *29th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018*, pages 1509–1528. SIAM, 2018. arXiv:1511.02612, doi:10.1137/1.9781611975031.99.
- 24 Ming Gu, Martin Farach, and Richard Beigel. An efficient algorithm for dynamic text indexing. In Daniel Dominic Sleator, editor, *5th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 1994*, pages 697–704. ACM/SIAM, 1994.
- 25 Monika Henzinger, Sebastian Krinninger, Danupon Nanongkai, and Thatchaphol Saranurak. Unifying and strengthening hardness for dynamic problems via the online matrix-vector multiplication conjecture. In Ronitt Rubinfeld, editor, *47th Annual ACM Symposium on Theory of Computing, STOC 2015*, pages 21–30. ACM, 2015. doi:10.1145/2746539.2746609.
- 26 Tomohiro I. Longest common extensions with recompression. In Juha Kärkkäinen, Jakub Radoszewski, and Wojciech Rytter, editors, *28th Annual Symposium on Combinatorial Pattern Matching, CPM 2017*, volume 78 of *LIPICs*, pages 18:1–18:15. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2017. doi:10.4230/LIPICs.CPM.2017.18.
- 27 Artur Jeż. Faster fully compressed pattern matching by recompression. *ACM Transactions on Algorithms*, 11(3):20:1–20:43, 2015. doi:10.1145/2631920.
- 28 Artur Jeż. Recompression: A simple and powerful technique for word equations. *Journal of the ACM*, 63(1):4:1–4:51, 2016. doi:10.1145/2743014.
- 29 Dominik Kempa and Tomasz Kociumaka. Resolution of the Burrows–Wheeler transform conjecture. In Sandy Irani, editor, *61st Annual IEEE Symposium on Foundations of Computer Science, FOCS 2020*. IEEE Computer Society, 2020.
- 30 Dominik Kempa and Nicola Prezza. At the roots of dictionary compression: string attractors. In Monika Henzinger, editor, *50th Annual ACM Symposium on Theory of Computing, STOC 2018*, pages 827–840. ACM, 2018. doi:10.1145/3188745.3188814.
- 31 Tomasz Kociumaka. *Efficient Data Structures for Internal Queries in Texts*. PhD thesis, University of Warsaw, October 2018. URL: <https://www.mimuw.edu.pl/~kociumaka/files/phd.pdf>.
- 32 Tomasz Kociumaka, Jakub Radoszewski, Wojciech Rytter, and Tomasz Walen. Internal pattern matching queries in a text and applications. In Piotr Indyk, editor, *26th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015*, pages 532–551. SIAM, 2015. doi:10.1137/1.9781611973730.36.
- 33 S.R. Kosaraju. Efficient string matching. Manuscript, 1987.
- 34 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.

- 35 Gad M. Landau and Uzi Vishkin. Fast parallel and serial approximate string matching. *Journal of Algorithms*, 10(2):157–169, 1989. doi:10.1016/0196-6774(89)90010-2.
- 36 N.J. Larsson and A. Moffat. Off-line dictionary-based compression. *Proceedings of the IEEE*, 88(11):1722–1732, 2000. doi:10.1109/5.892708.
- 37 Markus Lohrey. Algorithmics on SLP-compressed strings: A survey. *Groups Complexity Cryptology*, 4(2):241–299, 2012.
- 38 Kurt Mehlhorn, R. Sundar, and Christian Uhrig. Maintaining Dynamic Sequences under Equality Tests in Polylogarithmic Time. *Algorithmica*, 17(2):183–198, 1997. doi:10.1007/BF02522825.
- 39 Craig G Nevill-Manning and Ian H Witten. Compression and explanation using hierarchical grammars. *The Computer Journal*, 40(2 and 3):103–116, 1997.
- 40 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. doi:10.1016/j.dam.2019.01.014.
- 41 Roberto Radicioni and Alberto Bertoni. Grammatical compression: compressed equivalence and other problems. *Discrete Mathematics and Theoretical Computer Science*, 12(4):109, 2010.
- 42 Wojciech Rytter. Application of Lempel-Ziv factorization to the approximation of grammar-based compression. *Theoretical Computer Science*, 302(1-3):211–222, 2003. doi:10.1016/S0304-3975(02)00777-6.
- 43 Wojciech Rytter. Grammar compression, LZ-encodings, and string algorithms with implicit input. In Josep Díaz, Juhani Karhumäki, Arto Lepistö, and Donald Sannella, editors, *31st International Colloquium on Automata, Languages, and Programming, ICALP 2004*, pages 15–27. Springer, 2004. doi:10.1007/978-3-540-27836-8_5.
- 44 Süleyman Cenk Sahinalp and Uzi Vishkin. Efficient approximate and dynamic matching of patterns using a labeling paradigm (extended abstract). In Martin Tompa, editor, *37th Annual IEEE Symposium on Foundations of Computer Science, FOCS 1996*, pages 320–328. IEEE Computer Society, 1996. doi:10.1109/SFCS.1996.548491.
- 45 Hiroshi Sakamoto. Grammar compression: Grammatical inference by compression and its application to real data. In *12th International Conference on Grammatical Inference, ICGI 2014*, pages 3–20. JMLR.org, 2014.
- 46 Yusuxke Shibata, Takuya Kida, Shuichi Fukamachi, Masayuki Takeda, Ayumi Shinohara, Takeshi Shinohara, and Setsuo Arikawa. Byte pair encoding: A text compression scheme that accelerates pattern matching. Technical report, Technical Report DOI-TR-161, Department of Informatics, Kyushu University, 1999.
- 47 Rajamani Sundar and Robert E Tarjan. Unique binary-search-tree representations and equality testing of sets and sequences. *SIAM Journal on Computing*, 23(1):24–44, 1994.
- 48 Alexander Tiskin. Threshold approximate matching in grammar-compressed strings. In Jan Holub and Jan Žďárek, editors, *Prague Stringology Conference, PSC 2014*, pages 124–138, 2014.
- 49 T. Welch. A technique for high-performance data compression. *Computer*, 17:8–19, 1984.
- 50 Jacob Ziv and Abraham Lempel. A universal algorithm for sequential data compression. *IEEE Transactions on Information Theory*, 23(3):337–343, 1977.
- 51 Jacob Ziv and Abraham Lempel. Compression of individual sequences via variable-rate coding. *IEEE Transactions on Information Theory*, 24(5):530–536, 1978.